

IBM Tivoli Storage Manager
for Linux
Version 6.3.4

Administrator's Reference



IBM Tivoli Storage Manager
for Linux
Version 6.3.4

Administrator's Reference



Note:

Before using this information and the product it supports, read the information in "Notices" on page 1595.

This edition applies to Version 6.3.4 of IBM Tivoli Storage Manager (product numbers 5608-E01, 5608-E02, 5608-E03), and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters. This edition replaces SC23-9777-04.

© **Copyright IBM Corporation 1993, 2013.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this publication	xi
Who should read this publication	xi
Publications	xi
Tivoli Storage Manager publications	xii
Tivoli Storage FlashCopy Manager publications	xiv
Related hardware publications	xiv
Conventions used in this publication	xiv

New for IBM Tivoli Storage Manager

Version 6.3	xvii
Server updates	xvii
New for the server in Version 6.3.4	xvii
New for the server in Version 6.3.3	xviii
New for the server in Version 6.3.1	xx
New for the server in Version 6.3.0	xx

Chapter 1. Administering the server from the command-line interface 1

Issuing commands from the administrative client	1
Starting and stopping the administrative client	2
Monitoring server activities from the administrative client	2
Monitoring removable-media mounts from the administrative client	3
Processing commands one at a time from the administrative client	3
Processing a series of commands from the administrative client	3
Formatting output from commands	4
Saving command output to a specified location	4
Administrative client options	5
Issuing commands from the Administration Center	7
Issuing commands from the server console	7
Entering administrative commands	8
Reading syntax diagrams	8
Entering long commands	12
Naming Tivoli Storage Manager objects	13
Using wildcard characters to specify object names	14
Specifying descriptions in keyword parameters	15
Controlling command processing	15
Server command processing	15
Cancelling commands	16
Performing tasks concurrently on multiple servers	17
Routing commands to a single server	17
Routing commands to multiple servers	17
Routing commands to a server group	18
Routing commands to server groups	18
Routing commands to two servers and a server group	18
Routing commands inside scripts	18
Privilege classes for commands	19
Commands requiring system privilege	19
Commands requiring policy privilege	22
Commands requiring storage privilege	23

Commands requiring operator privilege	24
Commands any administrator can issue	25

Chapter 2. Administrative commands 27

ACCEPT DATE (Accepts the current system date)	28
ACTIVATE POLICYSET (Activate a new policy set)	29
ASSIGN DEFMGMTCLASS (Assign a default management class)	31
AUDIT commands	33
AUDIT LDAPDIRECTORY (Audit an LDAP directory server)	34
AUDIT LIBRARY (Audit volume inventories in an automated library)	36
AUDIT LICENSES (Audit server storage usage)	39
AUDIT VOLUME (Verify database information for a storage pool volume)	40
BACKUP commands	46
BACKUP DB (Back up the database)	47
BACKUP DEVCONFIG (Create backup copies of device configuration information)	51
BACKUP NODE (Back up a NAS node)	53
BACKUP STGPOOL (Back up primary storage pool to copy storage pool)	58
BACKUP VOLHISTORY (Save sequential volume history information)	62
BEGIN EVENTLOGGING (Begin logging events)	64
CANCEL commands	66
CANCEL EXPIRATION (Cancel an expiration process)	67
CANCEL EXPORT (Delete a suspended export operation)	68
CANCEL PROCESS (Cancel an administrative process)	69
CANCEL REPLICATION (Cancel node replication processes)	71
CANCEL REQUEST (Cancel one or more mount requests)	72
CANCEL RESTORE (Cancel a restartable restore session)	73
CANCEL SESSION (Cancel one or more client sessions)	74
CHECKIN LIBVOLUME (Check a storage volume into a library)	75
CHECKOUT LIBVOLUME (Check a storage volume out of a library)	83
CLEAN DRIVE (Clean a drive)	89
COMMIT (Control committing of commands in a macro)	90
COPY commands	91
COPY ACTIVEDATA (Copy active backup data from a primary storage pool to an active-data pool)	92
COPY CLOPTSET (Copy a client option set)	96
COPY DOMAIN (Copy a policy domain)	97
COPY MGMTCLASS (Copy a management class)	99

COPY POLICYSET (Copy a policy set)	101	DEFINE SERVERGROUP (Define a server group).	333
COPY PROFILE (Copy a profile)	103	DEFINE SPACETRIGGER (Define the space trigger)	334
COPY SCHEDULE (Copy a client or an administrative command schedule)	105	DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)	337
COPY SCRIPT (Copy a Tivoli Storage Manager script)	109	DEFINE STGPOOL (Define a storage pool)	341
COPY SERVERGROUP (Copy a server group)	110	DEFINE SUBSCRIPTION (Define a profile subscription)	387
DEFINE commands	111	DEFINE VIRTUALFSMAPPING (Define a virtual file space mapping)	389
DEFINE ALERTTRIGGER (Define an alert trigger)	112	DEFINE VOLUME (Define a volume in a storage pool)	391
DEFINE ASSOCIATION (Associate client nodes with a schedule)	115	DELETE commands	398
DEFINE BACKUPSET (Define a backup set)	117	DELETE ALERTTRIGGER (Remove a message from an alert trigger)	400
DEFINE CLIENTACTION (Define a one-time client action)	121	DELETE ASSOCIATION (Delete the node association to a schedule)	401
DEFINE CLIENTOPT (Define an option to an option set)	127	DELETE BACKUPSET (Delete a backup set)	403
DEFINE CLOPTSET (Define a client option set name)	130	DELETE CLIENTOPT (Delete an option in an option set)	408
DEFINE COLLOGROUP (Define a collocation group).	131	DELETE CLOPTSET (Delete a client option set)	409
DEFINE COLLOCMEMBER (Define collocation group member).	133	DELETE COLLOGROUP (Delete a collocation group).	410
DEFINE COPYGROUP (Define a copy group)	135	DELETE COLLOCMEMBER (Delete collocation group member).	411
DEFINE DATAMOVER (Define a data mover)	145	DELETE COPYGROUP (Delete a backup or archive copy group)	413
DEFINE DEVCLASS (Define a device class)	148	DELETE DATAMOVER (Delete a data mover)	415
DEFINE DEVCLASS - z/OS media server (Define device class for z/OS media server)	212	DELETE DEVCLASS (Delete a device class)	416
DEFINE DOMAIN (Define a new policy domain)	239	DELETE DOMAIN (Delete a policy domain)	417
DEFINE DRIVE (Define a drive to a library)	241	DELETE DRIVE (Delete a drive from a library)	418
DEFINE EVENTSERVER (Define a server as the event server)	245	DELETE EVENT (Delete event records).	419
DEFINE GRPMEMBER (Add a server to a server group)	246	DELETE EVENTSERVER (Delete the definition of the event server)	421
DEFINE LIBRARY (Define a library).	247	DELETE FILESPACE (Delete client node data from the server)	422
DEFINE MACHINE (Define machine information for disaster recovery)	267	DELETE GRPMEMBER (Delete a server from a server group)	426
DEFINE MACHNODEASSOCIATION (Associate a node with a machine)	269	DELETE KEYRING (Delete password information in the key database)	427
DEFINE MGMTCLASS (Define a management class)	271	DELETE LIBRARY (Delete a library).	428
DEFINE NODEGROUP (Define a node group)	274	DELETE MACHINE (Delete machine information).	429
DEFINE NODEGROUPMEMBER (Define node group member).	275	DELETE MACHNODEASSOCIATION (Delete association between a machine and a node)	430
DEFINE PATH (Define a path)	276	DELETE MGMTCLASS (Delete a management class)	431
DEFINE POLICYSET (Define a policy set).	286	DELETE NODEGROUP (Delete a node group)	432
DEFINE PROFASSOCIATION (Define a profile association)	288	DELETE NODEGROUPMEMBER (Delete node group member).	433
DEFINE PROFILE (Define a profile).	294	DELETE PATH (Delete a path)	434
DEFINE RECMEDMACHASSOCIATION (Associate recovery media with a machine)	296	DELETE POLICYSET (Delete a policy set).	436
DEFINE RECOVERYMEDIA (Define recovery media)	298	DELETE PROFASSOCIATION (Delete a profile association)	437
DEFINE SCHEDULE (Define a client or an administrative command schedule)	300	DELETE PROFILE (Delete a profile).	440
DEFINE SCRIPT (Define a Tivoli Storage Manager script)	323	DELETE RECMEDMACHASSOCIATION (Delete recovery media and machine association)	442
DEFINE SERVER (Define a server for server-to-server communications).	326	DELETE RECOVERYMEDIA (Delete recovery media)	443

DELETE SCHEDULE (Delete a client or an administrative command schedule)	444	GRANT PROXYNODE (Grant proxy authority to a client node)	560
DELETE SCRIPT (Delete command lines from a script or delete the entire script)	447	HALT (Shut down the server)	561
DELETE SERVER (Delete a server definition)	448	HELP (Get help on commands and error messages)	563
DELETE SERVERGROUP (Delete a server group).	449	IDENTIFY DUPLICATES (Identify duplicate data in a storage pool)	565
DELETE SPACETRIGGER (Delete the storage pool space triggers)	450	IMPORT commands	569
DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)	451	IMPORT ADMIN (Import administrator information).	570
DELETE STGPOOL (Delete a storage pool)	453	IMPORT NODE (Import client node information).	574
DELETE SUBSCRIBER (Delete subscriptions from a configuration manager database)	454	IMPORT POLICY (Import policy information)	581
DELETE SUBSCRIPTION (Delete a profile subscription)	455	IMPORT SERVER (Import server information)	584
DELETE VIRTUALFSMAPPING (Delete a virtual file space mapping)	456	INSERT MACHINE (Insert machine characteristics information or recovery instructions)	589
DELETE VOLHISTORY (Delete sequential volume history information)	457	ISSUE MESSAGE (Issue a message from a server script)	590
DELETE VOLUME (Delete a storage pool volume)	462	LABEL LIBVOLUME (Label a library volume)	592
DISABLE commands	465	LOCK commands	598
DISABLE EVENTS (Disable events for event logging)	466	LOCK ADMIN (Lock out an administrator)	599
DISABLE REPLICATION (Prevent outbound replication processing on a server)	470	LOCK NODE (Lock out a client node)	601
DISABLE SESSIONS (Prevent new sessions from accessing Tivoli Storage Manager)	471	LOCK PROFILE (Lock a profile)	603
DISMOUNT command	473	MACRO (Invoke a macro)	605
DISMOUNT VOLUME (Dismount a volume by volume name)	474	MIGRATE STGPOOL (Migrate storage pool to next storage pool)	607
DISPLAY OBJNAME (Display a full object name)	475	MOVE commands	610
ENABLE commands	476	MOVE DATA (Move files on a storage pool volume)	611
ENABLE EVENTS (Enable server or client events for logging)	477	MOVE DRMEDIA (Move disaster recovery media offsite and back onsite)	615
ENABLE REPLICATION (Allow outbound replication processing on a server)	480	MOVE GRPMEMBER (Move a server group member)	630
ENABLE SESSIONS (Resume user activity on the server)	481	MOVE MEDIA (Move sequential access storage pool media)	631
END EVENTLOGGING (Stop logging events)	483	MOVE NODEDATA (Move data by node in a sequential access storage pool)	639
EXPIRE INVENTORY (Manually start inventory expiration processing)	485	NOTIFY SUBSCRIBERS (Notify managed servers to update profiles).	648
EXPORT commands	489	PERFORM LIBACTION (Define or delete all drives and paths for a library)	649
EXPORT ADMIN (Export administrator information).	490	PING SERVER (Test the connection between servers)	652
EXPORT NODE (Export client node information).	497	PREPARE (Create a recovery plan file)	653
EXPORT POLICY (Export policy information)	518	QUERY commands	658
EXPORT SERVER (Export server information)	525	QUERY ACTLOG (Query the activity log).	661
EXTEND DBSPACE (Increase space for the database).	543	QUERY ADMIN (Display administrator information).	667
GENERATE commands	544	QUERY ALERTTRIGGER (Query the list of defined alert triggers)	672
GENERATE BACKUPSET (Generate a backup set of Backup-Archive Client data)	545	QUERY ALERTSTATUS (Query the status of an alert)	674
GENERATE BACKUPSETTOC (Generate a table of contents for a backup set)	553	QUERY ASSOCIATION (Query client node associations with a schedule)	679
GRANT commands	555	QUERY AUDITOCAPACITY (Query client node storage utilization).	681
GRANT AUTHORITY (Add administrator authority)	556	QUERY BACKUPSET (Query a backup set)	684
		QUERY BACKUPSETCONTENTS (Query contents of a backup set)	689
		QUERY CLOPTSET (Query a client option set)	692
		QUERY COLLOGGROUP (Query a collocation group).	694

QUERY CONTENT (Query the contents of a storage pool volume)	697	QUERY PROFILE (Query a profile)	841
QUERY COPYGROUP (Query copy groups)	705	QUERY PROXYNODE (Query proxy authority for a client node)	844
QUERY DATAMOVER (Display data mover definitions)	710	QUERY PVUESTIMATE (Display processor value unit estimate)	845
QUERY DB (Display database information)	713	QUERY RECOVERYMEDIA (Query recovery media)	849
QUERY DBSPACE (Display database storage space)	716	QUERY REPLICATION (Query node replication processes)	852
QUERY DEVCLASS (Display information on one or more device classes).	718	QUERY REPLNODE (Display information about replication status for a client node)	862
QUERY DIRSPACE (Query storage utilization of FILE directories)	723	QUERY REPLRULE (Query replication rules)	866
QUERY DOMAIN (Query a policy domain)	724	QUERY REQUEST (Query one or more pending mount requests)	868
QUERY DRIVE (Query information about a drive)	727	QUERY RESTORE (Query restartable restore sessions)	869
QUERY DRMEDIA (Query disaster recovery media)	731	QUERY RPFCONTENT (Query recovery plan file contents stored on a target server)	872
QUERY DRMSTATUS (Query disaster recovery manager system parameters)	740	QUERY RPFFILE (Query recovery plan file information stored on a target server)	874
QUERY ENABLED (Query enabled events)	743	QUERY SAN (Query the devices on the SAN)	877
QUERY EVENT (Query scheduled and completed events)	745	QUERY SCHEDULE (Query schedules).	880
QUERY EVENTRULES (Query rules for server or client events)	757	QUERY SCRIPT (Query Tivoli Storage Manager scripts)	888
QUERY EVENTSERVER (Query the event server).	759	QUERY SERVER (Query a server)	891
QUERY EXPORT (Query for active or suspended export operations)	760	QUERY SERVERGROUP (Query a server group)	895
QUERY FILESPACE (Query one or more file spaces)	767	QUERY SESSION (Query client sessions)	897
QUERY LIBRARY (Query a library)	774	QUERY SHREDSTATUS (Query shredding status)	901
QUERY LIBVOLUME (Query a library volume)	777	QUERY SPACETRIGGER (Query the space triggers)	903
QUERY LICENSE (Display license information)	780	QUERY SSLKEYRINGPW (Query SSL key database file password)	905
QUERY LOG (Display information on the recovery log)	782	QUERY STATUS (Query system parameters)	906
QUERY MACHINE (Query machine information).	784	QUERY STATUSTHRESHOLD (Query status monitoring thresholds)	914
QUERY MEDIA (Query sequential access storage pool media)	787	QUERY STGPOOL (Query storage pools)	917
QUERY MGMTCLASS (Query a management class)	793	QUERY SUBSCRIBER (Display subscriber information).	928
QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)	796	QUERY SUBSCRIPTION (Display subscription information).	930
QUERY MONITORSTATUS (Query the monitoring status).	799	QUERY SYSTEM (Query the system configuration and capacity).	932
QUERY MOUNT (Display information on mounted sequential access volumes).	803	QUERY TAPEALERTMSG (Display status of SET TAPEALERTMSG command)	934
QUERY NASBACKUP (Query NAS backup images)	805	QUERY TOC (Display table of contents for a backup image)	935
QUERY NODE (Query nodes).	809	QUERY VIRTUALFSMAPPING (Query a virtual file space mapping)	938
QUERY NODEDATA (Query client data in volumes)	820	QUERY VOLHISTORY (Display sequential volume history information)	940
QUERY NODEGROUP (Query a node group)	823	QUERY VOLUME (Query storage pool volumes)	948
QUERY OCCUPANCY (Query client file spaces in storage pools)	825	QUIT (End the interactive mode of the administrative client)	956
QUERY OPTION (Query server options)	829	RECLAIM STGPOOL (Reclaim volumes in a sequential-access storage pool)	957
QUERY PATH (Display a path definition)	831	RECONCILE VOLUMES (Reconcile differences in the virtual volume definitions)	961
QUERY POLICYSET (Query a policy set)	835	REGISTER commands	964
QUERY PROCESS (Query one or more server processes)	838		

REGISTER ADMIN (Register an administrator ID)	965	SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)	1065
REGISTER LICENSE (Register a new license)	969	SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)	1066
REGISTER NODE (Register a node)	971	SET ALERTMONITOR (Set the alert monitor to on or off)	1067
REMOVE commands	988	SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)	1068
REMOVE ADMIN (Delete an administrator user ID)	989	SET ARCHIVERETENTIONPROTECTION (Activate data retention protection)	1069
REMOVE NODE (Delete a node or an associated machine node)	991	SET ARREPLRULEDEFAULT (Set the server replication rule for archive data)	1071
REMOVE REPLNODE (Remove a client node from replication)	993	SET AUTHENTICATION (Set password authentication)	1073
RENAME commands	995	SET BKREPLRULEDEFAULT (Set the server replication rule for backup data)	1074
RENAME ADMIN (Rename an administrator)	996	SET CLIENTACTDURATION (Set the duration period for the client action)	1077
RENAME FILESPACE (Rename a client file space on the server)	998	SET CONFIGMANAGER (Specify a configuration manager)	1078
RENAME NODE (Rename a node)	1001	SET CONFIGREFRESH (Set managed server configuration refresh)	1079
RENAME SCRIPT (Rename a Tivoli Storage Manager script)	1003	SET CONTEXTMESSAGING (Set message context reporting on or off)	1080
RENAME SERVERGROUP (Rename a server group)	1004	SET CPUINFOREFRESH (Refresh interval for the client workstation information scan)	1081
RENAME STGPOOL (Change the name of a storage pool)	1005	SET CROSSDEFINE (Specifies whether to cross-define servers)	1082
REPLICATE NODE (Replicate data in file spaces that belong to a client node)	1006	SET DBRECOVERY (Set the device class for automatic backups)	1083
REPLY (Allow a request to continue processing)	1013	SET DBREPORTMODE (Set the level of database reporting)	1085
RESET PASSEXP (Reset password expiration)	1014	SET DEDUPVERIFICATIONLEVEL (Set the percentage of extents to verify)	1086
RESTART EXPORT (Restart a suspended export operation)	1016	SET DEFAULTAUTHENTICATION (Set the default authentication method for REGISTER NODE and REGISTER ADMIN commands)	1088
RESTORE commands	1018	SET DRMACTIVEDATASTGPOOL (Specify the active-data pools to be managed by DRM)	1090
RESTORE NODE (Restore a NAS node)	1019	SET DRMCHECKLABEL (Specify label checking)	1091
RESTORE STGPOOL (Restore storage pool data from a copy pool or an active-data pool)	1024	SET DRMCMDFILENAME (Specify the name of a file to contain commands)	1092
RESTORE VOLUME (Restore primary volume data from a copy pool or an active-data pool)	1028	SET DRMCOPYSTGPOOL (Specify the copy storage pools to be managed by DRM)	1093
REVOKE commands	1032	SET DRMCOURIERNAME (Specify the courier name)	1094
REVOKE AUTHORITY (Remove administrator authority)	1033	SET DRMDBBACKUPEXPIREDAYS (Specify DB backup series expiration)	1095
REVOKE PROXYNODE (Revoke proxy authority for a client node)	1037	SET DRMFILEPROCESS (Specify file processing)	1097
ROLLBACK (Rollback uncommitted changes in a macro)	1038	SET DRMINSTRPREFIX (Specify the prefix for recovery instructions file names)	1098
RUN (Run a Tivoli Storage Manager script)	1039	SET DRMNOTMOUNTABLENAME (Specify the not mountable location name)	1100
SELECT (Perform an SQL query of the Tivoli Storage Manager database)	1042	SET DRMPLANPREFIX (Specify a prefix for recovery plan file names)	1101
SET commands	1052	SET DRMPLANVPOSTFIX (Specify replacement volume names)	1103
SET ACCOUNTING (Set accounting records on or off)	1055		
SET ACTLOGRETENTION (Set the retention period or the size of the activity log)	1056		
SET ALERTACTIVEDURATION (Set the duration of an active alert)	1058		
SET ALERTCLOSEDDURATION (Set the duration of a closed alert)	1059		
SET ALERTEMAIL (Set the alert monitor to email alerts to administrators)	1060		
SET ALERTEMAILFROMADDR (Set the email address of the sender)	1062		
SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)	1063		
SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)	1064		

SET DRMPRIMSTGPOOL (Specify the primary storage pools to be managed by DRM)	1104	SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)	1147
SET DRMRPFEXPIREDAYS (Set criteria for recovery plan file expiration)	1105	SET SUBFILE (Set subfile backup for client nodes)	1149
SET DRMVAULTNAME (Specify the vault name)	1107	SET SUMMARYRETENTION (Set number of days to keep data in activity summary table)	1150
SET EVENTRETENTION (Set the retention period for event records)	1108	SET TAPEALERTMSG (Set tape alert messages on or off)	1151
SET INVALIDPWLIMIT (Set the number of invalid logon attempts)	1109	SET TOCLOADRETENTION (Set load retention period for table of contents)	1152
SET LDAPPASSWORD (Set the LDAP password for the server)	1111	SETOPT (Set a server option for dynamic update)	1153
SET LDAPUSER (Define the server administrator for the LDAP directory server)	1112	SHRED DATA (Shred data)	1155
SET LICENSEAUDITPERIOD (Set license audit period)	1114	SUSPEND EXPORT (Suspend a currently running export operation)	1157
SET MAXCMDRETRIES (Set the maximum number of command retries)	1115	UNLOCK commands	1158
SET MAXSCHEDESESSIONS (Set maximum scheduled sessions)	1116	UNLOCK ADMIN (Unlock an administrator)	1159
SET MINPWLENGTH (Set minimum password length)	1118	UNLOCK NODE (Unlock a client node)	1161
SET MONITORINGADMIN (Set the name of the monitoring administrator)	1119	UNLOCK PROFILE (Unlock a profile)	1163
SET MONITOREDSEVERGROUP (Set the group of monitored servers)	1120	UPDATE commands	1164
SET PASSEXP (Set password expiration date)	1121	UPDATE ALERTTRIGGER (Update a defined alert trigger)	1165
SET QUERYSCHEDPERIOD (Set query period for polling client nodes)	1123	UPDATE ALERTSTATUS (Update the status of an alert)	1168
SET RANDOMIZE (Set randomization of scheduled start times)	1124	UPDATE ADMIN (Update an administrator)	1170
SET REGISTRATION (Set open or closed registration)	1126	UPDATE BACKUPSET (Update a retention value assigned to a backup set)	1174
SET REPLRETENTION (Set the retention period for replication records)	1128	UPDATE CLIENTOPT (Update a client option sequence number)	1179
SET REPLSERVER (Set the target replication server)	1130	UPDATE CLOPTSET (Update a client option set description)	1180
SET RETRYPERIOD (Set time between retry attempts)	1132	UPDATE COLLOCGROUP (Update a collocation group)	1181
SET SCHEDMODES (Select a central scheduling mode)	1133	UPDATE COPYGROUP (Update a copy group)	1182
SET SERVERHLADDRESS (Set the high-level address of a server)	1135	UPDATE DATAMOVER (Update a data mover)	1190
SET SERVERLLADDRESS (Set the low-level address of a server)	1136	UPDATE DEVCLASS (Update the attributes of a device class)	1192
SET SERVERNAME (Specify the server name)	1137	UPDATE DEVCLASS - z/OS media server (Update device class for z/OS media server)	1249
SET SERVERPASSWORD (Set password for server)	1138	UPDATE DOMAIN (Update a policy domain)	1275
SET SPREPLRULEDEFAULT (Set the server replication rule for space-managed data)	1139	UPDATE DRIVE (Update a drive)	1277
SET SSLKEYRINGPW (Set the SSL key ring password)	1141	UPDATE FILESPACE (Update file-space node-replication rules)	1281
SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)	1142	UPDATE LIBRARY (Update a library)	1286
SET STATUSMONITOR (Specifies whether to enable status monitoring)	1144	UPDATE LIBVOLUME (Change the status of a storage volume)	1302
SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)	1146	UPDATE MACHINE (Update machine information)	1304
		UPDATE MGMTCLASS (Update a management class)	1306
		UPDATE NODE (Update node attributes)	1309
		UPDATE NODEGROUP (Update a node group)	1325
		UPDATE PATH (Change a path)	1326
		UPDATE POLICYSET (Update a policy set description)	1335
		UPDATE PROFILE (Update a profile description)	1337
		UPDATE RECOVERYMEDIA (Update recovery media)	1338
		UPDATE REPLRULE (Update replication rules)	1340

UPDATE SCHEDULE (Update a schedule)	1342	DBDIAGLOGSIZE	1471
UPDATE SCRIPT (Update a Tivoli Storage Manager script)	1365	DBDIAGPATHFSTHRESHOLD (Specify a threshold for free space)	1472
UPDATE SERVER (Update a server defined for server-to-server communications)	1368	DBMEMPERCENT	1473
UPDATE SERVERGROUP (Update a server group description)	1372	DBMTCPPORT	1474
UPDATE SPACETRIGGER (Update the space triggers)	1373	DEDUPREQUIRESBACKUP	1475
UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)	1375	DEDUPTIER2FILESIZE	1476
UPDATE STGPOOL (Update a storage pool)	1379	DEDUPTIER3FILESIZE	1477
UPDATE VIRTUALFSMAPPING (Update a virtual file space mapping)	1419	DEVCONFIG	1478
UPDATE VOLHISTORY (Update sequential volume history information)	1421	DISABLESCHEDULES	1479
UPDATE VOLUME (Change a storage pool volume)	1423	DISPLAYLFINFO	1480
VALIDATE commands	1427	DNSLOOKUP	1481
VALIDATE LANFREE (Validate LAN-Free paths)	1428	DRIVEACQUIRERETRY	1482
VALIDATE POLICYSET (Verify a policy set)	1430	ENABLENASDEDUP	1483
VALIDATE REPLICATION (Validate replication for a client node)	1432	EVENTSERVER	1484
VARY (Bring a random access volume online or offline)	1436	EXPINTERVAL	1485
		EXPQUIET	1486
		FFDCLOGNAME	1487
		FFDCMAXLOGSIZE	1488
		FILEEXIT	1489
		FILETEXTEXIT	1490
		FSUSEDTHRESHOLD	1491
		IDLETIMEOUT	1492
		LANGUAGE	1493
		LDAPCACHEDURATION	1495
		LDAPURL	1496
		MAXSESSIONS	1498
		MESSAGEFORMAT	1499
		MIRRORLOGDIRECTORY	1500
		MOVEBATCHSIZE	1501
		MOVESIZETHRESH	1502
		MSGINTERVAL	1503
		NDMPCONTROLPORT	1504
		NDMPENABLEKEEPALIVE	1505
		NDMPKEEPIDLEMINUTES	1506
		NDMPPORTRANGE	1507
		NDMPPREFDATAINTERFACE	1508
		NOPREEMPT	1509
		NORETRIEVEDATE	1510
		NUMOPENVOLSALLOWED	1511
		QUERYAUTH	1513
		RECLAIMDELAY	1514
		RECLAIMPERIOD	1515
		REORGBEGINTIME	1516
		REORGDURATION	1517
		REPORTRETRIEVE	1518
		REPLBATCHSIZE	1519
		REPLSIZETHRESH	1520
		REQSYSAUTHOUTFILE	1521
		RESOURCETIMEOUT	1522
		RESTOREINTERVAL	1523
		RETENTIONEXTENSION	1524
		SANDISCOVERY	1525
		SANDISCOVERYTIMEOUT	1526
		SANREFRESHTIME	1527
		SEARCHMPQUEUE	1528
		SERVEDEDUPTXNLIMIT	1529
		SHMPORT	1531
		SHREDDING	1532
		SNMPHEARTBEATINTERVAL	1533
		SNMPMESSAGECATEGORY	1534
Chapter 3. Server options	1439		
Modifying server options	1439		
Types of server options	1440		
Server communication options	1440		
Server storage options	1442		
Client-server options	1443		
Date, number, time, and language options	1443		
Database options	1443		
Data transfer options	1444		
Message options	1444		
Event logging options	1445		
Security options and licensing options	1445		
Miscellaneous options	1446		
3494SHARED	1447		
ACSACCESSID	1448		
ACSLOCKDRIVE	1449		
ACSQUICKINIT	1450		
ACSTIMEOUTX	1451		
ACTIVELOGDIRECTORY	1452		
ACTIVELOGSIZE	1453		
ADMINCOMMTIMEOUT	1454		
ADMINIDLETIMEOUT	1455		
ADMINONCLIENTPORT	1456		
ALIASHALT	1457		
ALLOWREORGINDEX	1458		
ALLOWREORGTABLE	1459		
ARCHFAILOVERLOGDIRECTORY	1460		
ARCHLOGDIRECTORY	1461		
ASSISTVCRRECOVERY	1462		
AUDITSTORAGE	1463		
CHECKTAPEPOS	1464		
CLIENTDEDUPTXNLIMIT	1466		
COMMMETHOD	1468		
COMMTIMEOUT	1470		

SNMPSUBAGENT	1535
SNMPSUBAGENTHOST	1536
SNMPSUBAGENTPORT	1537
SSLFIPSMODE	1537
SSLTCPADMINPORT	1538
SSLTCPPOINT	1539
SSLTLS12	1540
TCPADMINPORT	1541
TCPBUFSIZE	1542
TCPNODELAY	1543
TCPPOINT	1544
TCPWINDOWSIZE	1545
TECBEGINEVENTLOGGING	1546
TECHOST	1547
TECPORT	1548
TECUTF8EVENT	1549
THROUGHPUTDATATHRESHOLD	1550
THROUGHPUTTIMETHRESHOLD	1551
TXNGROUPMAX	1552
UNIQUETDPTECEVENTS	1553
UNIQUETECEVENTS	1554
USEREXIT	1555
VERBCHECK	1556
VOLUMEHISTORY	1557

Chapter 4. Server utilities. 1559

DSMSERV (Start the server)	1560
The server startup script: rc.dsmserv	1561
The server startup script: dsmserv.rc	1562
DSMSERV DISPLAY DBSPACE (Display information about database storage space)	1563
DSMSERV DISPLAY LOG (Display recovery log information)	1564
DSMSERV FORMAT (Format the database and log)	1566
DSMSERV INSERTDB (Move a server database into an empty database)	1568
DSMSERV LOADFORMAT (Format a database)	1571

DSMSERV REMOVEDB (Remove a database)	1573
DSMSERV RESTORE DB (Restore the database)	1574
DSMSERV RESTORE DB (Restore a database to its most current state)	1575
DSMSERV RESTORE DB (Restore a database to a point-in-time)	1578

Appendix A. Return codes for use in IBM Tivoli Storage Manager scripts . 1583

Appendix B. Support information . 1587

Getting technical training	1587
Searching knowledge bases	1587
Searching the Internet	1588
Using IBM Support Assistant.	1588
Finding product fixes	1589
Receiving notification of product fixes.	1589
Contacting IBM Software Support	1589
Setting up a subscription and support contract	1590
Determining the business impact	1590
Describing the problem and gathering background information	1590
Submitting the problem to IBM Software Support	1591

Appendix C. Accessibility features for the Tivoli Storage Manager product family 1593

Notices 1595
Trademarks 1597

Glossary 1599

Index 1621

About this publication

IBM® Tivoli® Storage Manager is a client/server program that provides storage management solutions to customers in a multi-vendor computer environment. IBM Tivoli Storage Manager provides an automated, centrally scheduled, policy-managed backup, archive, and space-management facility for file servers and workstations.

This publication provides you with the commands and options that you can use to manage the Tivoli Storage Manager server.

Who should read this publication

This reference is intended for anyone who is registered as an administrator. A single administrator can manage Tivoli Storage Manager, or several people can share administrative responsibilities.

You should be familiar with the operating system on which the server resides and the communication protocols required for the client/server environment. You also need to understand the storage management practices of your organization, such as how you are currently backing up workstation files and how you are using storage devices.

Publications

Publications for the IBM Tivoli Storage Manager family of products are available online. The Tivoli Storage Manager product family includes IBM Tivoli Storage FlashCopy® Manager, IBM Tivoli Storage Manager for Space Management, IBM Tivoli Storage Manager for Databases, and several other storage management products from IBM Tivoli.

To search all publications, search across the appropriate Tivoli Storage Manager information center:

- Version 6.3 information center: <http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3>
- Version 6.4 information center: <http://pic.dhe.ibm.com/infocenter/tsminfo/v6r4>

You can download PDF versions of publications from the Tivoli Storage Manager information center or from the IBM Publications Center at <http://www.ibm.com/shop/publications/order/>.

Go to Tivoli Documentation Central to find information centers that contain official product documentation for current and previous versions of Tivoli products, including the Tivoli Storage Manager product family. You can find Tivoli Documentation Central at <http://www.ibm.com/tivoli/documentation>.

You can also order some related publications from the IBM Publications Center website at <http://www.ibm.com/shop/publications/order/>. The website provides information about ordering publications from countries other than the United States. In the United States, you can order publications by calling 1-800-879-2755.

Tivoli Storage Manager publications

The following tables list the publications that make up the Tivoli Storage Manager library.

Table 1. Tivoli Storage Manager server publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for AIX Installation Guide</i>	GC23-9781
<i>IBM Tivoli Storage Manager for AIX Administrator's Guide</i>	SC23-9769
<i>IBM Tivoli Storage Manager for AIX Administrator's Reference</i>	SC23-9775
<i>IBM Tivoli Storage Manager for HP-UX Installation Guide</i>	GC23-9782
<i>IBM Tivoli Storage Manager for HP-UX Administrator's Guide</i>	SC23-9770
<i>IBM Tivoli Storage Manager for HP-UX Administrator's Reference</i>	SC23-9776
<i>IBM Tivoli Storage Manager for Linux Installation Guide</i>	GC23-9783
<i>IBM Tivoli Storage Manager for Linux Administrator's Guide</i>	SC23-9771
<i>IBM Tivoli Storage Manager for Linux Administrator's Reference</i>	SC23-9777
<i>IBM Tivoli Storage Manager for Oracle Solaris Installation Guide</i>	GC23-9784
<i>IBM Tivoli Storage Manager for Oracle Solaris Administrator's Guide</i>	SC23-9772
<i>IBM Tivoli Storage Manager for Oracle Solaris Administrator's Reference</i>	SC23-9778
<i>IBM Tivoli Storage Manager for Windows Installation Guide</i>	GC23-9785
<i>IBM Tivoli Storage Manager for Windows Administrator's Guide</i>	SC23-9773
<i>IBM Tivoli Storage Manager for Windows Administrator's Reference</i>	SC23-9779
<i>IBM Tivoli Storage Manager for z/OS Media Installation and User's Guide</i>	SC27-4018
<i>IBM Tivoli Storage Manager Upgrade and Migration Guide for V5 Servers</i>	GC27-4017
<i>IBM Tivoli Storage Manager Integration Guide for Tivoli Storage Manager FastBack®</i>	SC27-2828

Table 2. Tivoli Storage Manager storage agent publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for SAN for AIX Storage Agent User's Guide</i>	SC23-9797
<i>IBM Tivoli Storage Manager for SAN for HP-UX Storage Agent User's Guide</i>	SC23-9798
<i>IBM Tivoli Storage Manager for SAN for Linux Storage Agent User's Guide</i>	SC23-9799
<i>IBM Tivoli Storage Manager for SAN for Oracle Solaris Storage Agent User's Guide</i>	SC23-9800
<i>IBM Tivoli Storage Manager for SAN for Windows Storage Agent User's Guide</i>	SC23-9553

Table 3. Tivoli Storage Manager client publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for UNIX and Linux: Backup-Archive Clients Installation and User's Guide</i>	SC23-9791

Table 3. Tivoli Storage Manager client publications (continued)

Publication title	Order number
<i>IBM Tivoli Storage Manager for Windows: Backup-Archive Clients Installation and User's Guide</i>	SC23-9792
<i>IBM Tivoli Storage Manager Using the Application Programming Interface</i>	SC23-9793
<i>IBM Tivoli Storage Manager for Space Management for UNIX and Linux: User's Guide</i>	SC23-9794
<i>IBM Tivoli Storage Manager HSM for Windows Administration Guide</i>	SC23-9795

Table 4. Tivoli Storage Manager data protection publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Installation and User's Guide</i>	GC27-4010
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for UNIX and Linux Installation and User's Guide</i>	SC27-4019
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for Windows Installation and User's Guide</i>	SC27-4020
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Installation and User's Guide</i>	GC27-4009
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino UNIX and Linux Installation and User's Guide</i>	SC27-4021
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino for Windows Installation and User's Guide</i>	SC27-4022
<i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for DB2</i>	SC33-6341
<i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for Oracle</i>	SC33-6340
<i>IBM Tivoli Storage Manager for Virtual Environments Installation and User's Guide</i>	SC27-2898
<i>IBM Tivoli Storage Manager for Microsoft SharePoint Guide</i>	N/A

Table 5. IBM Tivoli Storage Manager troubleshooting and tuning publications

Publication title	Order number
<i>IBM Tivoli Storage Manager Problem Determination Guide</i>	GC23-9789
<i>IBM Tivoli Storage Manager Optimizing Performance</i>	GC23-9788
<i>IBM Tivoli Storage Manager Client Messages and Application Programming Interface Return Codes</i>	SC27-2878
<i>IBM Tivoli Storage Manager Server Messages and Error Codes</i>	SC27-2877
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Messages</i>	GC27-4011
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Messages</i>	GC27-4012
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle Messages</i>	SC27-4014
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino Messages</i>	SC27-4015

Table 5. IBM Tivoli Storage Manager troubleshooting and tuning publications (continued)

Publication title	Order number
IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Messages	SC27-4016

Note: You can find information about IBM System Storage® Archive Manager at the Tivoli Storage Manager v6.3.0 information center.

Tivoli Storage FlashCopy Manager publications

The following table lists the publications that make up the Tivoli Storage FlashCopy Manager library.

Table 6. Tivoli Storage FlashCopy Manager publications

Publication title	Order number
IBM Tivoli Storage FlashCopy Manager for UNIX and Linux Installation and User's Guide	SC27-4005
IBM Tivoli Storage FlashCopy Manager for Windows Installation and User's Guide	SC27-4006
IBM Tivoli Storage FlashCopy Manager for VMware Installation and User's Guide	SC27-4007
IBM Tivoli Storage FlashCopy Manager Messages	GC27-4008

Related hardware publications

The following table lists related IBM hardware products publications.

For additional information on hardware, see the resource library for tape products at <http://www.ibm.com/systems/storage/tape/library.html>.

Title	Order Number
IBM TotalStorage 3494 Tape Library Introduction and Planning Guide	GA32-0448
IBM TotalStorage 3494 Tape Library Operator Guide	GA32-0449
IBM 3490E Model E01 and E11 User's Guide	GA32-0298
IBM Tape Device Drivers Installation and User's Guide	GC27-2130
IBM TotalStorage Enterprise Tape System 3590 Operator Guide	GA32-0330
IBM TotalStorage Enterprise Tape System 3592 Operator Guide	GA32-0465

Conventions used in this publication

- Command to be entered on the Linux command line:
`> dsmadm`
- Command to be entered on the command line of an administrative client:
`query devclass`

In the usage and descriptions for administrative commands, the term characters corresponds to the number of bytes available to store an item. For languages in which it takes a single byte to represent a displayable character, the character to byte ratio is 1 to 1. However, for DBCS and other multi-byte languages, the reference to characters refers only to the number of bytes available for the item and

may represent fewer actual characters.

New for IBM Tivoli Storage Manager Version 6.3

Many features in the Tivoli Storage Manager Version 6.3 server are new for previous Tivoli Storage Manager users.

Server updates

New features and other changes are available in the IBM Tivoli Storage Manager V6.3 server. Technical updates since the previous edition are marked with a vertical bar (|) in the left margin.

New for the server in Version 6.3.4

Server fix pack 6.3.4 contains several new features, in addition to fixes for problems.

The server that is included with the Tivoli Storage Manager and IBM Tivoli Storage Manager Extended Edition V6.4 products is at the V6.3.4 level. The V6.3.4 server is also available for download separately, as a fix pack for current users of V6.3.

Tivoli Storage Manager migration to V6.3.4 or later on Linux x86_64

| You can now migrate a Tivoli Storage Manager V5 server that runs on an AIX®, HP-UX, or Solaris operating system to V6.3.4 or later on a Linux x86_64 operating system.

| Depending on your hardware and software environment, this migration procedure might be useful for achieving server consolidation, load balancing, or standardization on the Linux operating system.

IBM Tivoli Storage Manager Operations Center

| The V6.4.1 IBM Tivoli Storage Manager Operations Center is a new web-based user interface for managing a storage environment.

| The V6.4.1 Operations Center includes an Overview page that shows the interaction of Tivoli Storage Manager servers and clients. You can use the Operations Center to identify potential issues at a glance, manage alerts, and access the Tivoli Storage Manager command line. The Administration Center interface is also available, but the Operations Center is the preferred monitoring interface.

Tivoli Monitoring for Tivoli Storage Manager updates

| Tivoli Monitoring for Tivoli Storage Manager V6.3.4 includes some new Cognos® reports, and features, including some methods for distributing Cognos reports to other organizations.

| The following new Cognos reports are available:

Status reports

- Client storage summary and details
- VE activity status
- VE backup type summary
- VE current occupancy summary

Trending reports

- Client storage usage trends

To allow reports to be shared, Cognos reports can be exported and imported in to other Tivoli Common Reporting instances.

The Agent Log workspace is enhanced to display whether the monitored servers are up and running.

Pruning values are now automatically configured during new installations. If you upgraded the application, you must manually configure the pruning settings to periodically remove data from the WAREHOUS database.

New for the server in Version 6.3.3

Server fix pack 6.3.3 contains several new features, in addition to fixes for problems.

The server that is included with the Tivoli Storage Manager and IBM Tivoli Storage Manager Extended Edition V6.4 products is at the V6.3.3 level. The V6.3.3 server is also available for download separately, as a fix pack for current users of V6.3.

LDAP-authenticated passwords

IBM Tivoli Storage Manager server V6.3.3 can use an LDAP directory server to authenticate passwords. LDAP-authenticated passwords give you an extra level of security by being case-sensitive, offering advanced password rule enforcement, and a centralized server on which to authenticate them.

You must have an LDAP directory server on which to authenticate a hash representation of the password. After you configure the Tivoli Storage Manager server, you can authenticate administrator and node passwords with the LDAP directory server.

The two methods of authentication are LDAP and LOCAL. LOCAL means that the password is authenticated with the Tivoli Storage Manager server.

Passwords that are authenticated with the Tivoli Storage Manager server are not case-sensitive. All passwords can be composed of characters from the following list:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9
~ ! @ # $ % ^ & * _ - + = ` | ( ) { } [ ] : ; < > , . ? /
```

Restriction: Client nodes must be at V6.4 or later to use the LDAP directory server to authenticate passwords.

Tivoli Monitoring for Tivoli Storage Manager updates

Tivoli Monitoring for Tivoli Storage Manager now includes some new Tivoli Enterprise Portal workspaces, new Cognos reports, and the option for existing IBM Tivoli Monitoring customers to install and deploy monitoring agents from a small agent package.

The new Cognos reports are:

Status reports:

- Client activity status
- Client backup currency
- Client backup status
- Client schedule status
- Client storage pool usage summary
- Current client occupancy summary
- Current storage pool summary
- Highest storage space usage
- Server schedule status
- Yesterday's missed and failed client schedules

Trending Reports

- Client activity success rate
- Client schedule success rate
- Server database growth trends
- Server storage growth trends

New Tivoli Enterprise Portal monitoring workspaces are:

- Activity summary workspace, which provides data about server and client operations for both virtual and non-virtual environments.
- Sessions workspace, which provides a view of all the active client sessions running on the specified server.

Additional VMware backup information is now provided in the Tivoli Storage Manager server activity log and summary table for Data Protection for VMware operations. This new information provides improved data collection and reporting content when you use reporting facilities such as Tivoli Common Reporting.

A new option to enable LDAP authentication, when creating and configuring monitoring agent instances.

The improved ability to exclude, or filter-out message numbers, to narrow results.

Existing IBM Tivoli Monitoring users can now install the small agent package, and remotely deploy monitoring agents without having to download the larger Tivoli Monitoring for Tivoli Storage Manager package.

Control for inactive NDMP operation connections

Transmission Control Protocol (TCP) keepalive is a mechanism by which small packets of data are sent across the network at predefined intervals. The packets prevent a long-running, inactive connection from being closed by firewall software that detects and closes inactive connections. With this release, you can enable the TCP keepalive function for control connections of network data-management protocol (NDMP).

Enhancements for expiration processing

Improvements for expiring inventory are available with this release. Node processing can now be spread across more than one thread in parallel, at the file space level. This change in process helps distribute the workload and more efficiently avoid bottlenecks for nodes that use virtual servers.

Related reference:

“EXPIRE INVENTORY (Manually start inventory expiration processing)” on page 485

New for the server in Version 6.3.1

Server fix pack 6.3.1 contains new features, in addition to fixes for problems.

Data validation during read/write operations to tape

With logical block protection, IBM Tivoli Storage Manager includes a cyclic redundancy check (CRC) value at the end of each logical block of data to be written to tape. You can specify CRC data-block validation during read and write operations, or only during write operations.

You can use logical block protection only with the following types of drives and media:

- IBM LTO5 and later
- IBM 3592 Generation 3 drives, and later, with 3592 Generation 2 media, and later

New for the server in Version 6.3.0

New features and other changes are available in the Tivoli Storage Manager V6.3 server.

Node replication

Node replication is the process of incrementally copying or *replicating* client node data from one server of Tivoli Storage Manager to another server of Tivoli Storage Manager for the purpose of disaster recovery.

The server from which client node data is replicated is called a *source replication server*. The server to which client node data is replicated is called a *target replication server*.

Node replication avoids the logistics and security exposure of physically moving tape media to a remote location. If a disaster occurs and the source replication server is unavailable, backup-archive clients of Tivoli Storage Manager can recover their data from the target replication server. If you cannot recover the source replication server, you can convert client nodes to nonreplicating nodes for store operations on the target replication server.

If you use the export and import functions of Tivoli Storage Manager to store client node data on a disaster-recovery server, you can convert the nodes to replicating nodes. When replicating data, you can also use data deduplication to reduce bandwidth and storage requirements.

Tivoli Storage Manager V6.3 servers can be used for node replication. However, you can replicate data for client nodes that are at V6.3 or earlier. You can also replicate data that was stored on a Tivoli Storage Manager V6.2 or earlier server before you upgraded it to V6.3.

You cannot replicate nodes from a Tivoli Storage Manager V6.3.3 server to a server that is running on an earlier level of Tivoli Storage Manager.

Related reference:

- “CANCEL REPLICATION (Cancel node replication processes)” on page 71
- “DEFINE SERVER (Define a server for server-to-server communications)” on page 326
- “DISABLE SESSIONS (Prevent new sessions from accessing Tivoli Storage Manager)” on page 471
- “ENABLE REPLICATION (Allow outbound replication processing on a server)” on page 480
- “ENABLE SESSIONS (Resume user activity on the server)” on page 481
- “DISABLE REPLICATION (Prevent outbound replication processing on a server)” on page 470
- “QUERY FILESPACE (Query one or more file spaces)” on page 767
- “QUERY NODE (Query nodes)” on page 809
- “QUERY REPLICATION (Query node replication processes)” on page 852
- “QUERY REPLNODE (Display information about replication status for a client node)” on page 862
- “QUERY REPLRULE (Query replication rules)” on page 866
- “QUERY SERVER (Query a server)” on page 891
- “QUERY STATUS (Query system parameters)” on page 906
- “REGISTER NODE (Register a node)” on page 971
- “REMOVE REPLNODE (Remove a client node from replication)” on page 993
- “REPLICATE NODE (Replicate data in file spaces that belong to a client node)” on page 1006
- “SET ARREPLRULEDEFAULT (Set the server replication rule for archive data)” on page 1071
- “SET BKREPLRULEDEFAULT (Set the server replication rule for backup data)” on page 1074
- “SET REPLRETENTION (Set the retention period for replication records)” on page 1128
- “SET REPLSERVER (Set the target replication server)” on page 1130
- “SET SPREPLRULEDEFAULT (Set the server replication rule for space-managed data)” on page 1139
- “UPDATE FILESPACE (Update file-space node-replication rules)” on page 1281
- “UPDATE NODE (Update node attributes)” on page 1309
- “UPDATE REPLRULE (Update replication rules)” on page 1340
- “UPDATE SERVER (Update a server defined for server-to-server communications)” on page 1368
- “VALIDATE REPLICATION (Validate replication for a client node)” on page 1432

Deduplication of NetApp file-server data

Deduplication of data that belongs to network-attached storage (NAS) file-servers is disabled by default. To enable deduplication of NetApp file-server data, use the new `ENABLENASDEDUP` server option.

Related reference:

“`ENABLENASDEDUP`” on page 1483

Database table and index reorganization

If automatic table and index reorganization is affecting server performance, you can manually schedule reorganizations.

Related reference:

“`ALLOWREORGINDEX`” on page 1458

“`ALLOWREORGTABLE`” on page 1459

“`REORGBEGINTIME`” on page 1516

“`REORGDURATION`” on page 1517

Automatic backup-archive client deployment

IBM Tivoli Storage Manager, Version 6.3 can be scheduled to automatically deploy backup-archive client software to all workstations that have the backup-archive client installed.

You can deploy backup-archive clients on Microsoft Windows operating systems from a fix pack or interim fixes for all releases at V5.4 or later. You can migrate the Backup-Archive Client to a newer version, release, modification, or fix pack level that is V5.5 and later.

You can deploy backup-archive clients on operating systems other than Windows from all releases at V5.5 or later. These Backup-Archive Clients can go to any later version, release, modification, or fix level. You can coordinate the updates to each Backup-Archive Client from the Administration Center.

Multistream database backup and restore processing

Multiple, concurrent data streams can reduce the amount of time that is required to back up or restore the database. You can specify multiple, concurrent data streams for automatic or manual database-backup operations.

During restore operations, the Tivoli Storage Manager server attempts to use the same number of data streams that you specified for the backup operation. For example, suppose that you specify four data streams for a database backup operation. During a restore operation, the server attempts to use four drives. If one drive is offline and unavailable, the server uses three drives for the restore operation.

The benefit of multiple, concurrent data streaming depends on the size of the database. In general, if the database is less than 100 GB, the amount of time that you can save is relatively small. Multiple, concurrent data streaming also uses more volumes. If the volumes are high-capacity and if you use data compression, the result can be wasted space.

Related reference:

“BACKUP DB (Back up the database)” on page 47

“QUERY DB (Display database information)” on page 713

“QUERY VOLHISTORY (Display sequential volume history information)” on page 940

“SELECT (Perform an SQL query of the Tivoli Storage Manager database)” on page 1042

“SET DBRECOVERY (Set the device class for automatic backups)” on page 1083

Tivoli Monitoring for Tivoli Storage Manager updates

IBM Tivoli Monitoring for Tivoli Storage Manager, previously referred to as the Reporting and Monitoring feature, has an improved installation wizard. Cognos is now included for custom report creation.

Updates to Tivoli Monitoring for Tivoli Storage Manager include the following items:

- Cognos Business Intelligence V8 is an integrated business intelligence suite that is provided as part of Tivoli Common Reporting. Tivoli Common Reporting is included in the Administration Center installation when you select the Tivoli Common Reporting component. See *Customizing reports with Cognos Business Intelligence*, in the *Monitoring operations* section of the *Administrator's Guide* for details. All of the information regarding client and server reports can also be found in that section.
- The installation process has been improved to include a prerequisite checker, and now performs all installation configuration tasks automatically.
- A customizable dashboard workspace has been added to display many commonly viewed items in a single view. With the default setting, the dashboard displays data about the storage space used by node; unsuccessful client and server schedules; and details about storage pools, drives, and activity log error messages.
- You can include multiple servers in a single report. Reports have been enhanced to refine the accuracy of the data being displayed.
- New Tivoli Enterprise Portal workspaces are: activity log, agent log, updates to client node status, drives, libraries, occupancy, PVU details, and replication status and details.
- New client reports are available: storage pool media details, storage summary details, replication details, replication growth, and replication summary.
- New server reports are available: activity log details, server throughput, and an updated server throughput report for data collected by agents earlier than version 6.3.

Estimation of processor value units

You can use new methods to obtain information about the number of client and server devices connected to the system, and the utilization of processor value units (PVUs) by server devices. The new methods provide information to help you assess the license compliance of the Tivoli Storage Manager system.

By using the new **QUERY PVUESTIMATE** command, you can generate reports that estimate the number of server devices and client devices managed by the Tivoli Storage Manager server. You can also view PVU information on a per-node basis. These reports are not legally binding, but provide a starting point for determining license requirements. Alternatively, you can view PVU information in the

Tivoli Monitoring for Tivoli Storage Manager updates in Version 6.3

Administration Center. The Administration Center provides summaries of client devices, server devices, and estimated PVUs, and more detailed information.

For purposes of PVU estimation, only nodes on Windows 7, Windows XP Professional, and Apple systems are classified as client devices by default. Nodes on all other platforms are classified as server devices by default. You can update the role classification by issuing the **UPDATE NODE** command. You can also view and change the role classifications in the Administration Center client node notebook.

For a detailed report, issue the **SQL SELECT * FROM PVUESTIMATE_DETAILS** command. This command extracts information at the node level. This data can be exported to a spreadsheet and modified to more accurately represent the system environment.

For more information about PVU calculations and their use for licensing purposes, see the topic describing the role of PVUs in the *Administrator's Guide*.

Prerequisite checker

Tivoli Storage Manager Version 6.3 includes a prerequisite checker, a tool that can be run before starting the Tivoli Storage Manager installation.

The prerequisite checker verifies requirements for the Tivoli Storage Manager server, the Administration Center, and Tivoli Monitoring for Tivoli Storage Manager. The prerequisite checker verifies the operating system, the amount of free disk space, the required memory for the server, and other prerequisites. The tool presents a summary of results, informs you about changes that are required in your environment before installation, and creates required directories. In this way, the prerequisite checker can help simplify the installation process.

For more information, see the section about running the prerequisite checker in the *Installation Guide*.

Storage device updates

New device support and other changes to storage devices are available in Tivoli Storage Manager Version 6.3.

Virtual tape libraries:

With enhancements available in Version 6.3, you can define a library as a virtual tape library (VTL) to Tivoli Storage Manager.

VTLs primarily use disk subsystems to internally store data. Because they do not use tape media, you can exceed the capabilities of a physical tape library when using VTL storage. Using a VTL, you can define many volumes and drives which provides for greater flexibility in the storage environment and increases productivity by allowing more simultaneous mounts and tape I/O.

Related reference:

DEFINE LIBRARY
PERFORM LIBACTION
UPDATE LIBRARY

Tivoli Monitoring for Tivoli Storage Manager updates in Version 6.3

Faster formatting of DISK and preallocated FILE volumes:

On some Linux and UNIX file systems, including ext4 on Linux and JFS2 on AIX, IBM Tivoli Storage Manager formats random-access volumes and preallocated sequential-access disk volumes nearly instantaneously. For example, Tivoli Storage Manager can format a FILE or DISK volume on JFS2 in less than one second.

The new method for formatting volumes on some Linux and UNIX systems is similar to the method that Tivoli Storage Manager has been using to format volumes on Windows operating systems.

The new method provides the following advantages:

- During volume formatting, I/O is negligible and does not affect performance. For example, if you are running a server operation and you require an additional volume, you can format the volume without affecting the performance of the current operation.
- File system storage is allocated for the full file size without writing blocks and without creating sparse files. For example, if you issue the **DEFINE VOLUME** command to create a 10 GB volume, the underlying file in the file system consumes 10 GB of storage. Because the FILE volume size does not change, file-system log files do not require constant updating.
- The possibility for file fragmentation is minimized.

Before using the new volume formatting method, Tivoli Storage Manager checks the file system to determine whether the file system supports the method. If the file system does not, Tivoli Storage Manager uses the old method.

Access to storage devices attached by FICON on a z/OS system:

The database of a Tivoli Storage Manager V5 server that is running on a z/OS[®] system can be migrated to a V6.3 server that runs on AIX or Linux on System z[®]. After the upgrade, z/OS users can continue to access data stored on tape volumes whose contents are accessed by using FICON[®] attached storage devices.

The Tivoli Storage Manager V6.3 server accesses client data by using a storage device attached to z/OS. The storage device is made available by IBM Tivoli Storage Manager for z/OS Media.

In addition, Tivoli Storage Manager for z/OS Media facilitates access to Virtual Storage Access Method (VSAM) linear data sets on z/OS by using an enhanced sequential FILE storage method.

For more information, see the section about migrating Tivoli Storage Manager V5 servers on z/OS systems to V6 in the *Tivoli Storage Manager Upgrade and Migration Guide for V5 Servers*.

Tivoli Monitoring for Tivoli Storage Manager updates in Version 6.3

Append-only mode for IBM LTO-5 drives:

The CHECKTAPEPOS server option allows the Tivoli Storage Manager server to check the validity and consistency of data block positions on tape.

Enhancements to this option enable a drive to check for data overwrite problems before each WRITE operation and allow Tivoli Storage Manager to reposition tapes to the correct location and continue to write data. Use the CHECKTAPEPOS option with IBM LTO Generation 5 drives.

Note: You can enable append-only mode for IBM LTO Generation 5 and later drives, and for any drives that support this feature.

Related reference:

QUERY OPTION
SETOPT
CHECKTAPEPOS

Persistent reserve for tape drives:

Persistent reservation allows tape drive reservations from other servers to be cleared, if, for example, a server crashes.

In Tivoli Storage Manager Version 6.3, persistent reserve is enabled for drives and driver levels that support the feature.

For additional details about persistent reserve support, see <http://www.ibm.com/support/docview.wss?uid=swg21470319>.

Related reference:

DEFINE LIBRARY
UPDATE LIBRARY

Administration Center updates

New Administration Center support is available in Tivoli Storage Manager Version 6.3.

Tivoli Integrated Portal, Version 2.1:

The Tivoli Storage Manager Administration Center uses Tivoli Integrated Portal for its graphical user interface (GUI). With Tivoli Integrated Portal V2.1, you can now monitor the Administration Center with Internet Explorer 8 and Mozilla Firefox 3.5. All browsers that you used with Tivoli Integrated Portal V1.1.1 and later can be used with this latest version.

When you install Tivoli Integrated Portal V2.1 installing Tivoli Common Reporting, embedded security service, or the time scheduling service is optional. These features can be added and registered with Tivoli Integrated Portal V2.1 at a later time.

Administration Center policy domain updates:

With enhancements to the Administration Center, you can now specify server event-based archive settings using the Policy Domain and Management Class wizards.

If you set an archive retention period for an object through the server, you can update these settings using the Administration Center Management Class notebook.

Setting an archive retention period ensures that objects are not deleted from the Tivoli Storage Manager server until policy-based retention requirements for that object are satisfied.

Analysis of client performance data:

With the new client performance monitor function, you have the capability to gather and analyze performance data about backup and restore operations for an IBM Tivoli Storage Manager client.

The client performance monitor function is accessed from the Tivoli Storage Manager Administration Center and uses data that is collected by the API. You can view performance information about processor, disk, and network utilization, and performance data that relates to data transfer rates and data compression. You can analyze data throughput rates at any time during a backup or restore operation. Also, you can use the performance information to analyze processor, disk, or network performance bottlenecks.

Server session disablement and enablement

You can now temporarily disable and enable all outbound or inbound sessions for a particular Tivoli Storage Manager server.

This feature is useful, for example, if you have a planned network outage that might affect communication between a source and a target replication server. To prevent replication failures, you can disable outbound sessions from the source replication server before the outage. After communications have been reestablished, you can resume replication by enabling outbound sessions.

Related reference:

“DISABLE SESSIONS (Prevent new sessions from accessing Tivoli Storage Manager)” on page 471

“ENABLE SESSIONS (Resume user activity on the server)” on page 481

Command-line help for subcommands

In this release, you can obtain help for Tivoli Storage Manager subcommands. For example, you can display help for the **DEFINE DEVCLASS** command for 3570 device classes and for 3590 device classes. To display command-line help for a subcommand, type `help` followed by the topic number for the command.

Topic numbers are listed in the table of contents, for example:

3.0 Administrative commands

...

3.13.10 DEFINE DEVCLASS (Define a device class)

3.13.10.1 DEFINE DEVCLASS (Define a 3570 device class)

3.13.10.2 DEFINE DEVCLASS (Define a 3590 device class)

...

To display help for the **DEFINE DEVCLASS** command for 3570 device classes, type:
help 3.13.10.1

As in previous releases, you can use this method to display help for commands that have unique names, such as **REGISTER NODE**:

3.46 REGISTER
3.46.1 REGISTER ADMIN (Register an administrator)
3.46.2 REGISTER LICENSE (Register a new license)
3.46.3 REGISTER NODE (Register a node)

To display help for the **REGISTER NODE** command, you can type:
help 3.46.1

You can also type help *commandName*, where *commandName* is the name of the server command for which you want information:

help register node

Related reference:

“HELP (Get help on commands and error messages)” on page 563

Data encryption with TLS/SSL

You can use Transport Layer Security (TLS)/Secure Sockets Layer (SSL) on HP-UX, Linux, Oracle Solaris, AIX, and Windows platforms.

With TLS/SSL industry-standard communications, you can encrypt all traffic between the backup-archive client, the administrative command-line clients, and the IBM Tivoli Storage Manager server. You can use either self-signed or vendor-acquired SSL certificates.

For Tivoli Storage Manager V6.3 and later, to use SSL with self-signed certificates, use the SSLTLS12 option after you distribute new self-signed certificates to all V6.3 backup-archive clients. You can use certificates from previous server versions, but you then cannot use TLS 1.2.

For Tivoli Storage Manager V6.3.3 server, TLS/SSL is available for LAN-free and server-to-server functions.

Chapter 1. Administering the server from the command-line interface

Tivoli Storage Manager provides several different command-line interfaces. The interface you choose depends on the tasks that you want to perform and accessibility.

Tivoli Storage Manager provides three command-line interfaces:

- Administrative command-line client
- Server console
- Administration Center command line

The administrative command-line client is a program that runs on a file server, workstation, or mainframe. It is installed as part of the Tivoli Storage Manager server installation process. With the administrative client, you can issue any server commands.

Compared to the administrative client, the capabilities of the server console are limited. For example, you cannot issue certain commands from the server console, and you cannot specify that certain commands process before other commands can be issued. (This procedure can be useful if, for example, you want to run two commands in quick succession.) Furthermore, because the server console is a command-line window on the system on which the server is installed, you must be physically located at the server system to use the console. The administrative client can be accessed remotely from a different location. You cannot route commands to other servers from the server console.

The Administration Center is a Web-based, task-oriented interface for centrally configuring and managing Tivoli Storage Manager servers. The Administration Center provides a command-line interface that you can use if necessary. For example, you might want to use the command-line interface to perform Tivoli Storage Manager functions that are limited or not supported in the Administration Center. Using the command line in the Administration Center, you can issue any server commands.

Server scripts provide for automation of common administrative tasks. A macro is a file that contains one or more Tivoli Storage Manager administrative commands. When you issue the **MACRO** command, the server processes all commands in the macro file in order, including commands contained in any nested macros. For information about scripts and macros, see the *Administrator's Guide*.

Issuing commands from the administrative client

The administrative command-line client is a program that runs on a file server, workstation, or mainframe.

Ensure that your administrative client and your server are running in compatible languages. See "LANGUAGE" on page 1493 for language and locale options. If your client and server are using different languages, the messages that Tivoli Storage Manager generates might not be understandable.

Note: Text strings that are sent from the client to the server do not depend on the server language setting. The text will be displayed properly if the administrative client runs in the same locale when sending the string and when receiving the string.

For example, assume that you update a node contact field with a value that contains national characters (update node *myNode* contact=*NLcontact_info*), and later query the node (query node *myNode* format=detailed). If the client is running in the same locale when you update as when you query, the *NLcontact_info* will display properly. If you update the node contact field when the client is running in one locale, and query the node when the client is running in a different locale, it is possible that the *NLcontact_info* will not display properly.

Starting and stopping the administrative client

Use the **DSMADMC** command to start an administrative client session.

The Tivoli Storage Manager server must be running before an administrative client can connect. For instructions about starting the server, see the *Installation Guide*.

To start an administrative client session in command-line mode, enter this command on your workstation:

```
dsmadm -id=admin -password=admin -dataonly=yes
```

Note: If you do not want to be prompted for a user ID and password, enter the **DSMADMC** command using the **-ID** and **-PASSWORD** options as shown.

Stop an administrative command-line client session by entering the following command on your workstation:

```
quit
```

Monitoring server activities from the administrative client

To monitor Tivoli Storage Manager activities, such as server migration and client logons, run the administrative client in console mode. You cannot enter any administrative commands in console mode.

To start an administrative client session in console mode, enter:

```
dsmadm -consolemode
```

You are prompted for a password if authentication is turned on for the server. If you do not want to be prompted for your user ID and password, enter the **DSMADMC** command with the **-ID** and **-PASSWORD** options.

To end an administrative client session in console mode, see Table 7.

Table 7. Keyboard break sequences

Environment	Break Sequence
UNIX-based and Linux clients	Ctrl+C
Windows	Ctrl+C, Ctrl+Break
TSO	ATTN

Monitoring removable-media mounts from the administrative client

To monitor the mounting and dismounting of removable media, run the administrative client in mount mode. When the client is running in mount mode, you cannot enter any administrative commands.

To start an administrative client session in mount mode, enter:

```
dsmdmc -mountmode
```

You are prompted for a password if authentication is turned on for the server. If you do not want to be prompted for your user ID and password, enter the DSMADMC command with the `-ID` and `-PASSWORD` options.

To end an administrative client session in mount mode, see Table 7 on page 2.

Processing commands one at a time from the administrative client

Use batch mode to enter a single administrative command. Your administrative client session automatically ends when the command has processed.

To start an administrative client session in batch mode, enter:

```
dsmdmc server_command
```

If you do not want to be prompted for your user ID and password, you can enter the DSMADMC command with the `-ID` and `-PASSWORD` options.

In batch mode, you must enter the complete command on one line. If a command does not fit on one line, enter the command by using a macro or a script. If you specify a parameter with a string of text using batch mode, enclose the text in single quotation marks (' ') in the macro. Do not use double quotation marks for commands in batch mode, because your operating system might not parse the quotation marks correctly.

Processing a series of commands from the administrative client

Use the interactive mode to process a series of administrative commands.

To start an administrative client session in interactive mode, a server session must be available. To ensure the availability of server sessions for both administrative and client node sessions, the interactive mode of the administrative client is disconnected if one or more of the following conditions is true:

- The server was stopped using the HALT command.
- Commands were not issued from the administrative client session for the length of time specified with the IDLETIMEOUT server option.
- The administrative client session was canceled with the CANCEL SESSION command.

You can use continuation characters when using interactive mode. For more information, see “Entering long commands” on page 12.

To start an administrative session in interactive mode, enter:

```
dsmdmc
```

Do not enter a server command with the DSMADMC command. Doing so will start the administrative client in batch, not interactive, mode. For example, do not enter:

```
dsmadmc server_command
```

You can automatically restart your administrative client session by entering another command each time the **tsm: servername >** prompt appears.

Formatting output from commands

Tivoli Storage Manager formats the output processed from commands according to your screen or window width.

If the width of your screen or window is not wide enough to display the output horizontally, Tivoli Storage Manager arranges and displays the information vertically.

You can format the output of QUERY commands using the DISPLAYMODE and OUTFILE administrative client options.

Saving command output to a specified location

The most common use for redirecting output is to save the output from query commands to a specified file or program. You can then browse the contents of the file or, in some cases, print the contents.

Some platforms support redirection of output using special characters such as >, >>, and |. You can save the output from a command by entering redirection characters at the end of the command. Redirection characters direct the output of a command to a file or program you specify instead of to your screen. To redirect output, leave a blank between the redirection character and the file or program name. See the following examples.

Note: When redirecting output, follow the naming conventions of the operating system running your administrative client.

If you want to:	Enter this:
Redirect the output of a QUERY DOMAIN command to a new file in batch or interactive mode	<pre>dsmadmc -id=xxx -pa=xxx query domain acctg > dominfo.acc</pre> A single greater-than sign (>) indicates that Tivoli Storage Manager redirects the output to a new file or writes over an existing file.
Append the output of a QUERY DOMAIN command to the end of an existing file in batch or interactive mode	<pre>dsmadmc -id=xxx -pa=xxx query domain acctg >> dominfo.acc</pre> Double greater-than signs (>>) indicates that Tivoli Storage Manager appends the output to the end of an existing file.
Redirect all output from an administrative client session in console mode to a program called <code>filter.exe</code>	<pre>dsmadmc -console -id=admin -password=xxx filter.exe</pre> The program can be set up to monitor the output for individual messages as they occur and take appropriate action, such as sending mail to another user.


You can also redirect all output to a specified file by specifying the `-OUTFILE` option with a destination file name. For example, enter:

```
dsmadmc -id=sullivan -password=secret -consolemode -outfile=save.out
```

Administrative client options

In all administrative client modes, you can use options to modify administrative client session responses.

DSMADMC



Here is an example of a task and how to use the administrative client options: You can enter the **DSMADMC** command with your user ID and password by using the **-ID** and **-PASSWORD** options, respectively if you do not want to be prompted for that information. To have Tivoli Storage Manager redirect all output to a file, specify the **-OUTFILE** option with a destination file name. For example, to issue the **QUERY NODE** command in batch mode with the output redirected to the **SAVE.OUT** file, enter:

```
dsmadmc -id=sullivan -password=secret -outfile=save.out query node
```

Options

Administrative client options can be specified with the **DSMADMC** command and are valid from an administrative client session only. You can type an option in uppercase letters, lowercase letters, or any combination. Uppercase letters denote the shortest acceptable abbreviation. If an option appears entirely in uppercase letters, you cannot abbreviate it.

-ALWAYSPrompt

Specifies that a command prompt is displayed if the input is from the keyboard or if it is redirected (for example, from a file). If this option is not specified and the input is redirected, the command prompt is not written.

If the input is redirected, only the command output is displayed. If this option is specified, the command prompt and the command output are displayed.

-CHECKAlias**halt**

Allows the administrative client to recognize an alias for the **HALT** command as set in the **ALIASHALT** server option. See “**ALIASHALT**” on page 1457 for details.

-COMMAdelimited

Specifies that any tabular output from a server query is to be formatted as comma-separated strings rather than in readable format. This option is intended to be used primarily when redirecting the output of an SQL query (**SELECT** command). The comma-separated value format is a standard data format, which can be processed by many common programs, including spreadsheets, databases, and report generators.

-CONsolemode

Specifies that Tivoli Storage Manager runs in console mode. All server console output is echoed to your screen, excluding items such as responses to query commands issued from the console, trace output, or any system messages that might appear on the console.

-DATAONLY=NO or YES

Specifies whether product version information and output headers display with the output. The default is NO.

NO Specifies that the product version information and output column headers display.

YES

Suppresses the product version information and output column headers.

-DISPLAymode=LIST or TABLE

Allows you to force the QUERY output to tabular or list format regardless of the command line window column width.

If you are using the -DISPLAYMODE option and you want the output to go to a file, do not specify the -OUTFILE option. Use redirection to write to the file.

-ID=userid

Specifies the administrator's user ID.

-Itemcommit

Specifies that Tivoli Storage Manager commits commands inside a script or a macro as each command is processed.

-MOUNTmode

Specifies that Tivoli Storage Manager runs in mount mode. All server removable-media mount messages are echoed to your screen.

-NEWLINEAFTERPrompt

Specifies that a newline character is written immediately after the command prompt and commands entered from the keyboard appear immediately below the prompt. If this option is not specified, commands entered from the keyboard appear immediately to the right of the prompt.

-NOConfirm

Specifies that you do not want Tivoli Storage Manager to request confirmation before processing commands that affect the availability of the server or data managed by the server.

-OUTfile

Specifies that output from a server query is formatted one line per query. This option is available in batch mode only.

-OUTfile=filename

Specifies that output from a server query is redirected to a specified file. In batch mode, output is redirected to a file you specify and the format of the output matches the format of the output on your screen.

In interactive, console, or mount mode sessions, output displays on your screen.

-PAssword=password

Specifies the administrator's password.

-Quiet

Specifies that Tivoli Storage Manager does not display standard output messages to your screen. However, when you use this option, certain error messages still appear.

-SErveraddress

Specifies the server stanza in the dsm.sys file. The client uses the server stanza

to determine the server it will connect to. The `SERVERADDRESS` option is supported by administrative clients running on UNIX, Linux, and Macintosh operating systems only.

-TABdelimited

Specifies that any tabular output from a server query is to be formatted as tab-separated strings rather than in readable format. This option is intended to be used primarily when redirecting the output of an SQL query (`SELECT` command). The tab-separated value format is a standard data format, which can be processed by many common programs, including spreadsheets, databases, and report generators.

-TCPport

Specifies a TCP/IP port address for a Tivoli Storage Manager server. The `TCPport` option is only supported by administrative clients running on Windows operating systems and is valid on the Windows administrative client command line.

-TCPserveraddress

Specifies a TCP/IP server address for a Tivoli Storage Manager server. The `TCPserveraddress` option is only supported by administrative clients running on Windows operating systems and is valid on the Windows administrative client command line.

Besides the options listed here, you can also specify any option that is in the client options file. Each option must be preceded with a hyphen and delimited with a space.

Issuing commands from the Administration Center

A command-line interface is available from all of the main server tables in the Administration Center. To access the command line, select a server, click **Select Action**, and select **Use Command Line**.

Issuing commands from the server console

Tivoli Storage Manager provides a user ID named `SERVER_CONSOLE` that allows you to issue commands and administer the server from the server console after Tivoli Storage Manager is installed. At installation, `SERVER_CONSOLE` is automatically registered as an administrator and is given system authority.

If you have system privilege, you can revoke or grant new privileges to the `SERVER_CONSOLE` user ID. However, you cannot:

- Register or update the `SERVER_CONSOLE` user ID
- Lock or unlock the `SERVER_CONSOLE` user ID
- Rename the `SERVER_CONSOLE` user ID
- Remove `SERVER_CONSOLE` user ID
- Route commands from the `SERVER_CONSOLE` user ID

Note: Not all Tivoli Storage Manager commands are supported by the server console. You cannot specify the `WAIT` parameter from the server console.

Entering administrative commands

Commands consist of command names and usually parameters and variables. Syntax diagrams depict the rules to follow when entering commands.

To display command-line help for server commands that have unique names, you can type `help commandName`, where *commandName* is the name of the server command for which you want information. For example, to display help for the **REGISTER NODE** command, type `help register node`. Command syntax and parameter descriptions are displayed in the output.

You can also type `help` followed by the topic number for the command. Topic numbers are listed in the table of contents for command-line help, for example:

```
3.0 Administrative commands
  3.46 REGISTER
    3.46.1 REGISTER ADMIN (Register an administrator)
    3.46.2 REGISTER LICENSE (Register a new license)
    3.46.3 REGISTER NODE (Register a node)
```

To display help about the **REGISTER NODE** command, type:

```
help 3.46.3
```

Use topic numbers to display command-line help for subcommands. **DEFINE DEVCLASS** is an example of a command that has subcommands. For example, you can specify the **DEFINE DEVCLASS** command for 3570 device classes and for 3590 device classes:

```
3.0 Administrative commands
  ...
  3.13.10 DEFINE DEVCLASS (Define a device class)
    3.13.10.1 DEFINE DEVCLASS (Define a 3570 device class)
    3.13.10.2 DEFINE DEVCLASS (Define a 3590 device class)
    ...
```

To display help for the **DEFINE DEVCLASS** command for 3570 device classes, type:

```
help 3.13.10.1
```

Reading syntax diagrams

To read a syntax diagram for entering a command, follow the path of the line. Read from left to right and from top to bottom.

- The **▶**— symbol indicates the beginning of a syntax diagram.
- The —**▶** symbol at the end of a line indicates that the syntax diagram continues onto the next line.
- The **▶**— symbol at the beginning of a line indicates that a syntax diagram continues from the previous line.
- The —**▶** symbol indicates the end of a syntax diagram.

Command names

The command name can consist of a single action word, such as **HALT**, or it can consist of an action word and an object for the action, such as **DEFINE DOMAIN**. You can enter the command in any column of the input line.

Enter the entire command name or the abbreviation that is specified in the syntax diagram for the command. Uppercase letters denote the shortest acceptable abbreviation. If a command appears entirely in uppercase letters, you cannot

abbreviate it. You can enter the command in uppercase letters, lowercase letters, or any combination. In this example, you can enter CMDNA, CMDNAM, or CMDNAME in any combination of uppercase and lowercase letters.

▶▶—CMDName—————▶▶

Note: Command names in descriptive text are always capitalized.

Required parameters

When a parameter is on the same line as the command name, the parameter is required. When two or more parameter values are in a stack and one of them is on the line, you *must* specify one value.

In this example, you must enter PARMNAME=A, PARMNAME=B, or PARMNAME=C. Do not include any blanks immediately before or after the equal sign (=).

▶▶—PARMName==
 ┌A
 ├B
 └C

Optional parameters

When a parameter is below the line, the parameter is optional. In this example, you can enter PARMNAME=A or nothing at all. Do not include any blanks immediately before or after the equal sign (=).

▶▶
┌—PARMName==—A

When two or more parameter values are in a stack below the line, all of them are optional. In this example, you can enter PARMNAME=A, PARMNAME=B, PARMNAME=C, or nothing at all. Do not include any blanks immediately before or after the equal sign (=).

▶▶
┌—PARMName==
 ┌A
 ├B
 └C

Defaults


Defaults are above the line. The system uses the default unless you override it. You can override the default by entering an option from the stack below the line.

In this example, PARMNAME=A is the default. You can also enter PARMNAME=A, PARMNAME=B, or PARMNAME=C. Do not include any blanks before or after the equal sign (=).



Variables

Highlighted lowercase items (*like this*) denote variables. In these examples, *var_name* represents variables::



Special characters

You must code these symbols exactly as they appear in the syntax diagram.

- * Asterisk
- :
- ,
- = Equal sign
- Hyphen
- () Parentheses
- .

Repeating values

An arrow returning to the left means that the item can be repeated. A character within the arrow means that you must separate repeated items with that character.



Repeatable choices

A stack of values followed by an arrow returning to the left means that you can select more than one value or, when permitted, repeat a single item. In this example, you can choose more than one value, with each name delimited with a comma. Do not include any blanks before or after the equal sign (=).



Footnotes

Footnotes are enclosed in parentheses.




Notes:

- 1 You can specify up to five file names.

Entering parameters

The order in which you enter parameters can be important. The following example shows a portion of the command for defining a copy storage pool:



The first two parameters in this command (*pool_name* and *device_class_name*) are required parameters. *pool_name* and *device_class_name* are also positional. That is, they must be entered in the order shown, immediately after the command name. The **POOLTYPE** parameter is a required keyword parameter. **DESCRIPTION** and **RECLAIM**, are optional keyword parameters. Keyword parameters are identified by an equal sign that specifies a specific value or a variable. Keyword parameters must follow any positional parameters in a command.

The following command entries, in which the keyword parameters are ordered differently, are both acceptable:

```
define stgpool mycopypool mydeviceclass pooltype=copy description=engineering  
reclaim=50  
define stgpool mycopypool mydeviceclass description=engineering pooltype=copy  
reclaim=50
```

The following example, in which one of the positional parameters follows a keyword parameter, is not acceptable:

```
define stgpool mycopypool pooltype=copy mydeviceclass description=engineering  
reclaim=50
```

Syntax fragments

Some diagrams, because of their length, must display parts of the syntax with fragments. The fragment name appears between vertical bars in the diagram.

The expanded fragment appears in the diagram after all other parameters or at the bottom of the diagram. A heading with the fragment name identifies the expanded fragment. Commands appearing directly on the line are required.

In this example, the fragment is named "Fragment".



Fragment:



Entering long commands

Continuation characters are useful when you want to process a command that is longer than your screen or window width. You can use continuation characters in the interactive mode of the administrative client.

Without continuation characters you can enter up to 256 characters. With continuation characters you can enter up to 1500 characters.

Note: In the MACRO command, the maximums apply *after* any substitution variables have been applied.

With continuation characters, you can do the following:

- Enter a dash at the end of the line you want to continue.

For example:

```
register admin pease mypasswd -  
contact="david, ext1234"
```

- Continue a list of values by entering a dash or a back slash, with no preceding blank spaces, after the last comma of the list that you enter on the first line. Then, enter the remaining items in the list on the next line with no preceding blank spaces. For example:

```
stgpools=stg1,stg2,stg3,-  
stg4,stg5,stg6
```

- Continue a string of values enclosed in quotation marks by entering the first part of the string enclosed in quotation marks, followed by a dash or a back slash at the end of the line. Then, enter the remainder of the string on the next line enclosed in the *same* type of quotation marks.

For example:

```
contact="david pease, bldg. 100, room 2b, san jose,"-  
"ext. 1234, alternate contact-norm pass,ext 2345"
```

Tivoli Storage Manager concatenates the two strings with no intervening blanks. You must use *only* this method to continue a quoted string of values across more than one line.

Naming Tivoli Storage Manager objects

IBM Tivoli Storage Manager restricts the number and type of characters that you can use to name objects.

The following characters are available for defining object names.

Character	Description
A–Z	Any letter, A through Z
0–9	Any number, 0 through 9
_	Underscore
.	Period
-	Hyphen
+	Plus
&	Ampersand

The following table shows the maximum length of characters permitted for naming objects.

Type of Name	Maximum Length
Administrators, client option sets, client nodes, passwords, server groups, server, names, virtual file space names	64
Restartable export identifiers	64
High-level and low-level TCP/IP (IPv4 or IPv6) addresses	64
Device classes, drives, libraries, management classes, policy domains, profiles, schedules scripts, backup sets, storage pools	30

The following characters are available for defining password names:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9
~ ! @ # $ % ^ & * _ - + = ` | ( ) { } [ ] : ; < > , . ? /
```

Passwords considered “LOCAL” are those passwords that authenticate with the Tivoli Storage Manager server and are not case-sensitive. Passwords considered “LDAP” are those passwords that authenticate with an LDAP directory server and are case-sensitive.

When you use DEFINE commands to define database, recovery log, and storage pool volumes, the naming convention for the volume name is dependent on the type of sequential access media or random access media that you are using. Refer to the specific VOLUME command for details.

Using wildcard characters to specify object names

In some commands, such as the query commands, you can use wildcard characters to create a pattern-matching expression that specifies more than one object. Using wildcard characters makes it easier to tailor a command to your needs.

The wildcard characters you use depend on the operating system from which you issue commands. For example, you can use wildcard characters such as an asterisk (*) to match any (0 or more) characters, or you can use a question mark (?) or a percent sign (%) to match exactly one character.

Table 8 provides references to wildcard characters for some operating systems. Use wildcard characters appropriate for your system.

Table 8. Wildcard characters by operating system

Operating system	Match any	Match exactly one
AIX, HP-UX, Linux, OS/2, Oracle Solaris, Windows	*	?
TSO	*	%

For example, if you want to query all the management classes whose names begin with DEV in all the policy sets in DOMAIN1, and your system uses an asterisk as the *match-any* character, you can enter:

```
query mgmtclass domain1 * dev*
```

If your system uses a question mark as the *match-exactly-one* character, and you want to query the management classes in POLICYSET1 in DOMAIN1, you can enter:

```
query mgmtclass domain1 policyset1 mc?
```

Tivoli Storage Manager displays information about management classes with names MC.

Table 9 shows additional examples of using wildcard characters to match any characters.

Table 9. Match-any character

Pattern	Matches	Does not match
ab*	ab, abb, abxxx	a, b, aa, bb
ab*rs	abrs, abtrs, abrsrs	ars, aabrs, abrss
ab*ef*rs	abefrs, abefghrs	abefr, abers

Table 10 shows additional examples of using wildcard characters to match exactly one character. The question mark (?) can be replaced by a percent sign (%) if your platform uses that character instead of (?).

Table 10. Match-exactly-one character

Pattern	Matches	Does not match
ab?	abc	ab, abab, abzzzz
ab?rs	abrs	abrs, abllrs
ab?ef?rs	abdefjrs	abefrs, abdefrs, abefjrs

Table 10. Match-exactly-one character (continued)

Pattern	Matches	Does not match
ab??rs	abcdrs, abzzrs	abrs, abjrs, abkkrs

Specifying descriptions in keyword parameters

If a description (a string of text) for a parameter begins with a single or double quotation mark, or contains any embedded blanks or equal signs, you must surround the value with either single (') or double (") quotation marks.

The opening and closing quotation marks must be the same type of quotation marks. For example, if the opening quotation is a single quotation mark, the closing quotation mark must also be a single quotation mark.

For example, to register a new client node named Louie, with a password of secret, and with his title included as contact information, enter:

```
register node louie secret contact="manager of dept. 61f"
```

The following table presents ways of entering a description for the CONTACT parameter. The value can contain quotation marks, embedded blanks, or equal signs.

For this description	Enter this
manager	contact=manager
manager's	contact="manager's" or contact='manager's'
"manager"	contact=""manager"" or contact=""'manager'""
manager's report	contact="manager's report" or contact='manager's report'
manager's "report"	contact='manager's "report"'
manager=dept. 61f	contact='manager=dept. 61f'
manager reports to dept. 61f	contact='manager reports to dept. 61f' or contact="manager reports to dept. 61f"

Controlling command processing

You can run some Tivoli Storage Manager commands sequentially or concurrently with other commands. You can also route commands from one server to other servers for processing.

Server command processing

Tivoli Storage Manager processes administrator commands either in the foreground or in the background. Commands that process in the foreground must complete before you can issue another command. When commands are processing in the background, you can issue additional commands at any time.

Most Tivoli Storage Manager commands process in the foreground. For some commands that normally process in the background (for example, BACKUP DB), you can specify the WAIT parameter (WAIT=YES) with the command so that the command processes in the foreground. You might want to process a command in the foreground rather than in the background for any of the following reasons:

- To quickly determine whether a command completed successfully. When you issue a command that processes in the foreground, Tivoli Storage Manager sends a confirmation message indicating that the command completed successfully. If you process the command in the background, you need to open operational reporting or query the activity log to determine whether the command completed successfully.
- To monitor server activities (for example, messages) on the administrative client as a command is being processed. This might be preferable to searching a long activity log after the command has completed.
- To be able to start another process immediately after a command has completed. For example, you might specify WAIT=YES for a command that takes a short time to process so that, when the processing completes, you can immediately start processing another command.
- To serialize commands in an administrative script when it is important that one command completes before another begins.

Check the individual command description to determine whether a command has a WAIT parameter.

You can cancel commands that are processed in the foreground from the server console or from another administrative client session.

Each background process is assigned a process number. Use the QUERY PROCESS command to obtain the status and process number of a background process.

Note:

- If you are defining a schedule with a command that specifies WAIT=NO (the default), and you issue QUERY EVENT to determine the status of your scheduled operation, failed operations will report an event status of COMPLETED with a return of OK. In order for the QUERY EVENT output to reflect the failed status, the WAIT parameter must be set to YES. This will run the scheduled operation in the foreground and inform you of the status when it completes.
- You cannot process commands in the foreground from the server console.

Cancelling commands

Use the CANCEL PROCESS command to cancel commands that generate background processes.

Use the QUERY PROCESS command to obtain the status and process number of a background process. If a background process is active when you cancel it, the server stops the process. Any changes that are uncommitted are rolled back. However, changes that are committed are not rolled back.

When you issue a QUERY command from the administrative client, multiple screens of output might be generated. If this occurs and additional output is not needed, you can cancel the display of output to the client workstation. Doing so does not end the processing of the command.

Performing tasks concurrently on multiple servers

Command routing allows you to route commands to one or more servers for processing and then collect the output from these servers.

To route commands to other servers, you must have the same administrator ID and password as well as the required administrative authority on each server to which the command is being routed. You cannot route commands to other servers from the server console.

After the command has completed processing on all servers, the output displays, in its entirety, for each server. For example, the output from SERVER_A displays in its entirety, followed by the output from SERVER_B. The output includes summary messages for each individual server and identifies which server processed the output. Return codes indicate whether commands processed on the servers successfully. These return codes include one of three severities: 0, ERROR, or WARNING.

Each server that is identified as the target of a routed command must first be defined using the DEFINE SERVER command. The command is automatically routed to all servers specified as members of a server group or to individual servers specified with the command. For details about setting up and managing multiple servers for command routing, see the *Administrator's Guide*.

The following examples describe how to route the QUERY STGPOOL command to one server, multiple servers, a server group, multiple server groups, or a combination of servers and server groups. Each server or server group in a list must be separated with a comma, without spaces.

Routing commands to a single server

To route the QUERY STGPOOL command to a server named ASTRO, enter:

```
astro: query stgpool
```

The colon after the server name indicates the end of the routing information. This is also called the *server prefix*. Another way to indicate the end of routing information is to use parentheses around the server name, for example:

```
(astro) query stgpool
```

Routing commands to multiple servers

To route the QUERY STGPOOL command to multiple servers named HD_QTR, MIDAS, SATURN, enter:

```
hd_qtr,midas,saturn: query stgpool
```

If the first server has not been defined to Tivoli Storage Manager, the command is routed to the next defined server in the list of servers.

You can also enter the command this way:

```
(hd_qtr,midas,saturn) query stgpool
```

Routing commands to a server group

In this example, the server group ADMIN has servers named SECURITY, PAYROLL, PERSONNEL defined as group members. The command is routed to each of these servers.

To route the QUERY STGPOOL command to the server group named ADMIN, enter:

```
admin: query stgpool
```

You can also enter the command this way:

```
(admin) query stgpool
```

Routing commands to server groups

In this example, the server group ADMIN2 has servers SERVER_A, SERVER_B, and SERVER_C defined as group members, and server group ADMIN3 has servers ASTRO, GUMBY, and CRUSTY defined as group members. The command is routed to servers SERVER_A, SERVER_B, SERVER_C, ASTRO, GUMBY, and CRUSTY.

To route the QUERY STGPOOL command to two server groups that are named ADMIN2 and ADMIN3, enter:

```
admin2,admin3: query stgpool
```

You can also enter the command this way:

```
(admin2,admin3) query stgpool
```

Routing commands to two servers and a server group

In this example, the server group DEV_GROUP has servers SALES, MARKETING, and STAFF defined as group members. The command is routed to servers SALES, MARKETING, STAFF, MERCURY, and JUPITER.

To route the QUERY STGPOOL command to a server group named DEV_GROUP and to the servers named MERCURY and JUPITER, enter:

```
dev_group,mercury,jupiter: query stgpool
```

You can also enter the command this way:

```
(dev_group,mercury,jupiter) query stgpool
```

Routing commands inside scripts

When routing commands inside scripts, you must enclose the server or server group in parentheses and omit the colon. Otherwise, the command will not be routed when the RUN command is issued, and will only be run on the server where the RUN command is issued.

For example, to route the QUERY STGPOOL command inside a script:

1. Define a script called QU_STG to route it to the DEV_GROUP server group.

```
define script qu_stg "(dev_group) query stgpool"
```
2. Run the QU_STG script:

```
run qu_stg
```

In this example, the server group DEV_GROUP has servers SALES, MARKETING, and STAFF defined as group members. The QUERY STGPOOL command is routed

to these servers.

Privilege classes for commands

The authority granted to an administrator through the privilege class determines which administrative commands that the administrator can issue.

There are four administrator privilege classes in Tivoli Storage Manager:

- System
- Policy
- Storage
- Operator

After an administrator has been registered using the REGISTER ADMIN command, the administrator can issue a limited set of commands, including all query commands. When you install Tivoli Storage Manager, the server console is defined as a system administrator named SERVER_CONSOLE and is granted system privilege.

The following sections describe each type of administrator privilege and the commands that can be issued by an administrator who has been granted the corresponding authority.

Commands requiring system privilege

An administrator with system privilege has the highest level of authority in Tivoli Storage Manager. With system privilege, an administrator can issue any administrative command and has authority to manage all policy domains and all storage pools.

Table 11 on page 20 lists the commands that administrators with system privilege can issue. In some cases administrators with lower levels of authority, for example, unrestricted storage privilege, can also issue these commands. In addition, the REQSYSAUTHOUTFILE server option can be used to specify that certain commands require system privilege if they cause Tivoli Storage Manager to write to an external file. For more information about this server option see “REQSYSAUTHOUTFILE” on page 1521.

Table 11. System privilege commands

Command Name	Command Name
AUDIT LDAPDIRECTORY	DEFINE SPACETRIGGER
AUDIT LICENSES	DEFINE STGPOOL
ACCEPT DATE	DEFINE SUBSCRIPTION
BEGIN EVENTLOGGING	DEFINE VIRTUALFSMAPPING
CANCEL EXPIRATION	DEFINE VOLUME
CANCEL PROCESS	DELETE BACKUPSET
CANCEL REPLICATION	DELETE CLIENTOPT
CANCEL REQUEST	DELETE CLOPTSET
CANCEL RESTORE	DEFINE COLLOGROUP
CLEAN DRIVE	DEFINE COLLOCMEMBER
COPY ACTIVATEDATA	DELETE DOMAIN
COPY DOMAIN	DELETE DRIVE
COPY POLICYSET	DELETE EVENTSERVER
COPY PROFILE	DELETE GRPMEMBER
COPY SCHEDULE (See note.)	DELETE LIBRARY
COPY SCRIPT	DELETE MACHINE
COPY SERVERGROUP	DELETE MACHNODEASSOCIATION
DEFINE BACKUPSET	DELETE NODEGROUP
DEFINE CLIENTACTION	DELETE NODEGROUPMEMBER
DEFINE CLIENTOPT	DELETE PROFASSOCIATION
DEFINE CLOPTSET	DELETE PROFILE
DEFINE COLLOGROUP	DELETE RECMEDMACHASSOCIATION
DEFINE COLLOCMEMBER	DELETE RECOVERYMEDIA
DEFINE DEVCLASS	DELETE SCHEDULE (See note.)
DEFINE DOMAIN	DELETE SCRIPT
DEFINE DRIVE	DELETE SERVER
DEFINE EVENTSERVER	DELETE SERVERGROUP
DEFINE GRPMEMBER	DELETE SPACETRIGGER
DEFINE LIBRARY	DELETE STGPOOL
DEFINE MACHINE	DELETE SUBSCRIBER
DEFINE MACHNODEASSOCIATION	DELETE SUBSCRIPTION
DEFINE NODEGROUP	DELETE VIRTUALFSMAPPING
DEFINE NODEGROUPMEMBER	DISABLE EVENTS
DEFINE PATH	ENABLE EVENTS
DEFINE PROFASSOCIATION	END EVENTLOGGING
DEFINE PROFILE	EXPIRE INVENTORY
DEFINE RECMEDMACHASSOCIATION	EXPORT ADMIN
DEFINE RECOVERYMEDIA	EXPORT NODE
DEFINE SCHEDULE (See note.)	EXPORT POLICY
DEFINE SCRIPT	EXPORT SERVER
DEFINE SERVER	GENERATE BACKUPSET
DEFINE SERVERGROUP	GRANT AUTHORITY

Table 11. System privilege commands (continued)

Command Name	Command Name
GRANT PROXYNODE	SET CONFIGMANAGER
IDENTIFY DUPLICATES	SET CONFIGREFRESH
IMPORT NODE	SET CONTEXTMESSAGING
IMPORT POLICY	SET CROSSDEFINE
IMPORT SERVER	SET DBRECOVERY
INSERT MACHINE	SET DEFAULTAUTHENTICATION
LABEL LIBVOLUME	SET DRMACTIVEDATASTGPOOL
LOCK ADMIN	SET DRMCHECKLABEL
LOCK PROFILE	SET DRMCMDFILENAME
MIGRATE STGPOOL	SET DRMCOPYSTGPOOL
MOVE DRMEDIA	SET DRMCOURIERNAME
MOVE MEDIA	SET DRMDBBACKUPEXPIREDAYS
MOVE GRPMEMBER	SET DRMFILPROCESS
NOTIFY SUBSCRIBERS	SET DRMINSTRPREFIX
PERFORM LIBACTION	SET DRMNOTMOUNTABLENAME
PING SERVER	SET DRMPPLANPREFIX
PREPARE	SET DRMPPLANVPOSTFIX
QUERY BACKUPSETCONTENTS	SET DRMPRIMSTGPOOL
QUERY MEDIA	SET DRMRPFEXPIREDAYS
QUERY RPFCONTENT	SET DRMVaultNAME
QUERY TOC	SET EVENTRETENTION
RECLAIM STGPOOL	SET INVALIDPWLIMIT
RECONCILE VOLUMES	SET LDAPPASSWORD
REGISTER ADMIN	SET LDAPUSER
REGISTER LICENSE	SET LICENSEAUDITPERIOD
REMOVE ADMIN	SET MAXCMDRETRIES
REMOVE REPLNODE	SET MAXSCHEDSESSIONS
RENAME ADMIN	SET MINPWLENGTH
RENAME SCRIPT	SET PASSEXP
RENAME SERVERGROUP	SET QUERYSCHEDPERIOD
RENAME STGPOOL	SET RANDOMIZE
REPLICATE NODE	SET REGISTRATION
RESET PASSEXP	SET REPLRETENTION
RESTORE NODE	SET REPLSERVER
REVOKE AUTHORITY	SET RETRYPERIOD
REVOKE PROXYNODE	SET SCHEDMODES
RUN	SET SERVERHLADDRESS
SET ACCOUNTING	SET SERVERLLADDRESS
SET ACTLOGRETENTION	SET SERVERNAME
SET ARCHIVERETENTIONPROTECTION	SET SERVERPASSWORD
SET AUTHENTICATION	SET SPREPLRULEDEFAULT
SET ARREPLRULEDEFAULT	SET SUBFILE
SET BKREPLRULEDEFAULT	SET TOCLOADRETENTION
SET CLIENTACTDURATION	

Table 11. System privilege commands (continued)

Command Name	Command Name
SETOPT	UPDATE NODEGROUP
UNLOCK ADMIN	UPDATE PATH
UNLOCK PROFILE	UPDATE PROFILE
UPDATE ADMIN	UPDATE RECOVERYMEDIA
UPDATE BACKUPSET	UPDATE REPLRULE
UPDATE CLIENTOPT	UPDATE SCHEDULE (See note.)
UPDATE CLOPTSET	UPDATE SCRIPT
UPDATE COLLOCGROUP	UPDATE SERVER
UPDATE DEVCLASS	UPDATE SERVERGROUP
UPDATE DRIVE	UPDATE SPACETRIGGER
UPDATE LIBRARY	UPDATE VIRTUALFSMAPPING
UPDATE LIBVOLUME	UPDATE VOLHISTORY
UPDATE MACHINE	VALIDATE LANFREE
	VALIDATE REPLICATION

Note: This command is restricted by the authority granted to an administrator. System privilege is required only for administrative command schedules. System or policy privilege is required for client operation schedules.

Commands requiring policy privilege

An administrator with policy privilege can issue commands that relate to policy management objects such as policy domains, policy sets, management classes, copy groups, and schedules. The policy privilege can be unrestricted, or can be restricted to specific policy domains.

Unrestricted policy privilege permits you to issue all of the administrator commands that require policy privilege. You can issue commands that affect all existing policy domains as well as any policy domains that are defined in the future. An unrestricted policy administrator cannot define, delete, or copy policy domains.

Restricted policy privilege permits you to issue administrator commands that affect one or more policy domains for which you have been explicitly granted authority. For example, the DELETE MGMTCLASS command requires you to have policy privilege for the policy domain to which the management class belongs.

Table 12 on page 23 lists the commands that an administrator with policy privilege can issue.

Table 12. Policy privilege commands

Command Name	Command Name
ACTIVATE POLICYSET	DELETE POLICYSET
ASSIGN DEFMGMTCLASS	DELETE PATH
CLEAN DRIVE	DELETE SCHEDULE (See note 2.)
BACKUP NODE	GENERATE BACKUPSET
COPY MGMTCLASS	LOCK NODE
COPY POLICYSET	QUERY BACKUPSETCONTENTS
COPY SCHEDULE (See note 2.)	REGISTER NODE
DEFINE ASSOCIATION	REMOVE NODE
DEFINE BACKUPSET	RENAME FILESPACE
DEFINE COPYGROUP	RENAME NODE
DEFINE CLIENTACTION	SET SUMMARYRETENTION
DEFINE CLIENTOPT	RESTORE NODE
DEFINE MGMTCLASS	QUERY TOC
DEFINE NODEGROUP	UNLOCK NODE
DEFINE NODEGROUPMEMBER	UPDATE BACKUPSET
DEFINE POLICYSET	UPDATE COPYGROUP
DEFINE SCHEDULE	UPDATE DOMAIN
DELETE ASSOCIATION	UPDATE MGMTCLASS
DELETE BACKUPSET	UPDATE NODE
DELETE COPYGROUP	UPDATE NODEGROUP
DELETE EVENT (See note 1.)	UPDATE POLICYSET
DELETE FILESPACE	UPDATE SCHEDULE (See note 2.)
DELETE MGMTCLASS	VALIDATE POLICYSET
DELETE NODEGROUP	
DELETE NODEGROUPMEMBER	

Notes:

1. This command can be restricted by policy domain. An administrator with unrestricted policy privilege or restricted policy privilege for a specified policy domain can issue this command.
2. This command is restricted by the authority granted to an administrator. System privilege is required only for administrative command schedules. System or policy privilege is required for client operation schedules.

Commands requiring storage privilege

An administrator with storage privilege can issue commands that allocate and control storage resources for the server. The storage privilege can be unrestricted, or can be restricted to specific storage pools.

Unrestricted storage privilege permits you to issue all of the administrator commands that require storage privilege. You can issue commands that affect all existing storage pools as well as any storage pools that are defined in the future. You can also issue commands that affect the database and the recovery log. An unrestricted storage administrator cannot define or delete storage pools.

Restricted storage privilege permits you to issue administrator commands that only affect a storage pool for which you have been granted authority. For example, the DELETE VOLUME command only affects a storage pool volume that is defined to a specific storage pool.

Table 13 lists the commands an administrator with storage privilege can issue.

Table 13. Storage privilege commands

Command Name	Command Name
AUDIT LIBRARY	DELETE SPACETRIGGER
AUDIT VOLUME (See note.)	DELETE VIRTUALFSMAPPING
BACKUP DB	DELETE VOLHISTORY
BACKUP DEVCONFIG	DELETE VOLUME (See note.)
BACKUP STGPOOL	GRANT PROXYNODE
BACKUP VOLHISTORY	LABEL LIBVOLUME
CHECKIN LIBVOLUME	MIGRATE STGPOOL
CHECKOUT LIBVOLUME	MOVE DATA (See note.)
COPY ACTIVATEDATA (See note.)	MOVE MEDIA
DEFINE COLLOGROUP	QUERY TAPEALERTMSG
DEFINE COLLOCMEMBER	RECLAIM STGPOOL
DEFINE DATAMOVER	RESTORE STGPOOL
DEFINE DEVCLASS	RESTORE VOLUME
DEFINE DRIVE	REVOKE PROXYNODE
DEFINE LIBRARY	SET TAPEALERTMSG
DEFINE PATH	UPDATE COLLOGROUP
DEFINE VIRTUALFSMAPPING	UPDATE DATAMOVER
DEFINE VOLUME (See note.)	UPDATE DEVCLASS
DEFINE SPACETRIGGER	UPDATE DRIVE
DELETE COLLOGROUP	UPDATE LIBRARY
DELETE COLLOCMEMBER	UPDATE PATH
DELETE DATAMOVER	UPDATE SPACETRIGGER
DELETE DEVCLASS	UPDATE STGPOOL (See note.)
DELETE DRIVE	UPDATE VIRTUALFSMAPPING
DELETE LIBRARY	
DELETE PATH	

Note: This command can be restricted by storage pool. An administrator with unrestricted storage privilege or restricted storage privilege for a specified storage pool can issue this command.

Commands requiring operator privilege

An administrator with operator privilege can issue commands that control the immediate operation of the server and the availability of storage media.

Table 14 on page 25 lists the commands an administrator with operator privilege can issue.

Table 14. Operator privilege commands

Command Name	Command Name
CANCEL SESSION	MOVE DRMEDIA
DISABLE SESSIONS	MOVE MEDIA
DISMOUNT VOLUME	QUERY MEDIA
ENABLE SESSIONS	REPLY
HALT	UPDATE VOLUME
	VARY

Commands any administrator can issue

A limited number of commands can be used by any administrator, even if that administrator has not been granted any specific administrator privileges.

Table 15 on page 26 lists the commands any registered administrator can issue.

Table 15. Commands issued by all administrators

Command Name	Command Name
COMMIT	QUERY NODE
HELP	QUERY NODEDATA
ISSUE MESSAGE	QUERY NODEGROUP
MACRO	QUERY OCCUPANCY
PARALLEL	QUERY OPTION
QUERY ACTLOG	QUERY PATH
QUERY ADMIN	QUERY POLICYSET
QUERY ASSOCIATION	QUERY PROCESS
QUERY AUDITOCUPANCY	QUERY PROFILE
QUERY BACKUPSET	QUERY PROXYNODE
QUERY CLOPTSET	QUERY RECOVERYMEDIA
QUERY COLLOGGROUP	QUERY REPLICATION
QUERY CONTENT	QUERY REPLNODE
QUERY COPYGROUP	QUERY REPLRULE
QUERY DATAMOVER	QUERY REQUEST
QUERY DB	QUERY RESTORE
QUERY DBSPACE	QUERY RPFIL
QUERY DEVCLASS	QUERY SCHEDULE
QUERY DIRSPACE	QUERY SCRIPT
QUERY DOMAIN	QUERY SERVER
QUERY DRIVE	QUERY SERVERGROUP
QUERY DRMEDIA	QUERY SESSION
QUERY DRMSTATUS	QUERY SPACETRIGGER
QUERY ENABLED	QUERY STATUS
QUERY EVENT	QUERY STGPOOL
QUERY EVENTRULES	QUERY SUBSCRIBER
QUERY EVENTSERVER	QUERY SUBSCRIPTION
QUERY FILESPACE	QUERY SYSTEM
QUERY LIBRARY	QUERY VIRTUALFSMAPPING
QUERY LIBVOLUME	QUERY VOLHISTORY
QUERY LICENSE	QUERY VOLUME
QUERY LOG	QUIT
QUERY MACHINE	ROLLBACK
QUERY MGMTCLASS	SELECT
QUERY MOUNT	SERIAL
QUERY NASBACKUP	

Chapter 2. Administrative commands

Administrative commands are available to manage and configure the server.

Information for each command includes:

- A description of the tasks a command performs
- The administrator privilege class required to use the command
- A syntax diagram that identifies the required and optional parameters for the command
- Descriptions of each parameter of the command
- Examples of using the command
- A list of related commands

ACCEPT DATE (Accepts the current system date)

Use this command to allow the server to begin normal processing, when the server does not start normal processing because of a discrepancy between the server date and the current date on the system.

When the server does not start normal processing because of a discrepancy between the server date and the current date, this command forces the server to accept the current date and time as valid. If the system time is valid and the server has not been run for an extended time, this command should be run to allow the server to begin normal processing.

Attention: If the system date is invalid or the server was created or run previously with an invalid system date and this command is issued, any server processing or command that uses dates can have unexpected results. File expiration can be affected, for example. When the server is started with the correct date, files backed up with future dates will not be considered for expiration until that future date is reached. Files backed up with dates that have passed will expire faster. When the server processing encounters a future date, an error message is issued. See the *Administrator's Guide* for more details.

If the server detects an invalid date or time, server sessions become disabled (as if the **DISABLE SESSIONS** command had been issued). Expiration, migration, reclamation, and volume history deletion operations are not able to continue processing.

Use the **ENABLE SESSIONS ALL** command after you issue the **ACCEPT DATE** command to re-enable sessions to start.

Privilege class

To issue this command, you must have system privilege.

Syntax

▶▶—ACcEpt Date—▶▶

Parameters

None.

Example: Accept the current system date

Allow the server to accept the current date as the valid date.

```
accept date
```

Related commands

Table 16. Command related to **ACCEPT DATE**

Command	Description
ENABLE SESSIONS	Resumes server activity following the DISABLE command or the ACCEPT DATE command.

ACTIVATE POLICYSET (Activate a new policy set)

Use this command to copy the contents of a policy set to the ACTIVE policy set for the domain. The server uses the rules in the ACTIVE policy set to manage client operations in the domain. You can define multiple policy sets for a policy domain, but only one policy set can be active. The current ACTIVE policy set is replaced by the one you specify when you issue this command. You can modify the ACTIVE policy set only by activating another policy set.

Before activating a policy set, check that the policy set is complete and valid by using the **VALIDATE POLICYSET** command.

The **ACTIVATE POLICYSET** command fails if any of the following conditions exist:

- A copy group specifies a copy storage pool as a destination.
- A management class specifies a copy storage pool as the destination for files that were migrated by a Tivoli Storage Manager for Space Management client.
- The policy set has no default management class.
- A **TOCDESTINATION** parameter is specified, and the storage pool is either a copy pool or has a data format other than NATIVE or NONBLOCK.

The ACTIVE policy set and the last activated policy set are not necessarily identical. You can modify the original policy set that you activated without affecting the ACTIVE policy set.

If the server has data retention protection enabled, the following conditions must exist:

- All management classes in the policy set to be activated must contain an archive copy group.
- If a management class exists in the active policy set, a management class with the same name must exist in the policy set to be activated.
- If an archive copy group exists in the active policy set, the corresponding copy group in the policy set to be activated must have a RETVER value at least as large as the corresponding values in the active copy group.

Attention: Retention protection only applies to archive objects.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
▶▶—ACTivate POLicyset—domain_name—policy_set_name—▶▶
```

Parameters

domain_name **(Required)**

Specifies the policy domain for which you want to activate a policy set.

policy_set_name **(Required)**

Specifies the policy set to activate.

ACTIVATE POLICYSET

Example: Activate a policy set on a specific policy domain

Activate the VACATION policy set in the EMPLOYEE_RECORDS policy domain.

```
activate policyset employee_records vacation
```

Related commands

Table 17. Commands related to **ACTIVATE POLICYSET**

Command	Description
COPY POLICYSET	Creates a copy of a policy set.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.
QUERY DOMAIN	Displays information about policy domains.
QUERY POLICYSET	Displays information about policy sets.
UPDATE POLICYSET	Changes the description of a policy set.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

ASSIGN DEFMGMTCLASS (Assign a default management class)

Use this command to specify a management class as the default management class for a policy set. You must assign a default management class for a policy set before you can activate that policy set.

To ensure that clients can always back up and archive files, choose a default management class that contains both an archive copy group and a backup copy group.

The server uses the default management class to manage client files when a management class is not otherwise assigned or appropriate. For example, the server uses the default management class when a user does not specify a management class in the include-exclude list. See the *Administrator's Guide* for details.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
►►—ASsign DEFMGmtclass—domain_name—policy_set_name—class_name—►►
```

Parameters

domain_name **(Required)**

Specifies the policy domain to which the management class belongs.

policy_set_name **(Required)**

Specifies the policy set for which you want to assign a default management class. You cannot assign a default management class to the ACTIVE policy set.

class_name **(Required)**

Specifies the management class that is to be the default management class for the policy set.

Example: Assign a default management class

Assign DEFAULT1 as the default management class for policy set SUMMER in the PROG1 policy domain.

```
assign defmgmtclass prog1 summer default1
```

Related commands

Table 18. Commands related to ASSIGN DEFMGMTCLASS

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE MGMTCLASS	Defines a management class.

ASSIGN DEFMGMTCLASS

Table 18. Commands related to ASSIGN DEFMGMTCLASS (continued)

Command	Description
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY MGMTCLASS	Displays information about management classes.
QUERY POLICYSET	Displays information about policy sets.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.
UPDATE MGMTCLASS	Changes the attributes of a management class.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

AUDIT commands

Use the **AUDIT** commands to review or examine the adequacy of the database information and the storage pool volume. The **AUDIT LDAPDIRECTORY** command deletes nodes or administrator IDs from an LDAP directory server, that do not authenticate their passwords with the LDAP directory server.

- “AUDIT LDAPDIRECTORY (Audit an LDAP directory server)” on page 34
- “AUDIT LIBRARY (Audit volume inventories in an automated library)” on page 36
- “AUDIT LICENSES (Audit server storage usage)” on page 39
- “AUDIT VOLUME (Verify database information for a storage pool volume)” on page 40

AUDIT LDAPDIRECTORY (Audit an LDAP directory server)

Use this command to audit an IBM Tivoli Storage Manager-controlled namespace on a Lightweight Directory Access Protocol (LDAP) server. The LDAP directory server and namespace are specified through one or more **LDAPURL** options.

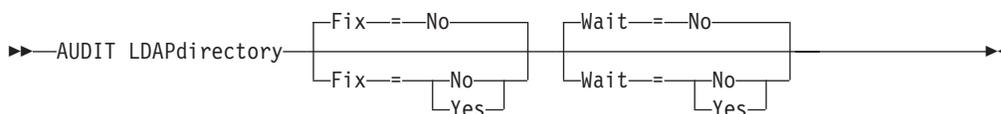
Nodes and administrator user IDs that do not authenticate their passwords with the LDAP directory server are deleted with the **AUDIT LDAPDIRECTORY FIX=YES** command. Nodes or administrator user IDs that no longer exist in the Tivoli Storage Manager database are also deleted.

Before you issue this command, ensure that the **LDAPURL** option is specified in the `dsmserv.opt` file. See the **LDAPURL** option for more information. If you specified more than one **LDAPURL** option in the `dsmserv.opt` file, each option is validated in the order in which they are placed. If the **LDAPURL** option is not specified, the command fails.

Privilege class

You must have system privileges to issue this command.

Syntax



Parameters

Fix

This optional parameter specifies how the Tivoli Storage Manager server resolves inconsistencies between the database and the external directory. The default is **NO**. You can specify the following values:

No The server reports all inconsistencies but does not change the external directory.

Yes

The server resolves any inconsistencies that it can and suggests further actions, if needed.

Important: If there are LDAP entries that are shared with other Tivoli Storage Manager servers, choosing **YES** might cause those servers to become out-of-sync.

Wait

This optional parameter specifies whether to wait for the Tivoli Storage Manager server to complete processing this command in the foreground. The default is **NO**. You can specify the following values:

No The server processes this command in the background and you can continue with other tasks while the command is processing. Messages related to the background process are shown either in the activity log file or the server console, depending on where the messages are logged.

Yes

The server processes this command in the foreground. The operation must

complete before you can continue with other tasks. Messages are shown either in the activity log file or the server console, or both, depending on where the messages are logged.

Restriction: You cannot specify WAIT=YES from the server console.

Example: Audit an LDAP directory and repair inconsistencies

Audit the LDAP directory that you specified in the LDAPURL option. The Tivoli Storage Manager server resolves some inconsistencies.

```
audit ldapdirectory fix=yes
```

```
ANR2749W Admin ADMIN1 was located in the LDAP directory server but not in the database.
ANR2749W Admin ADMIN2 was located in the LDAP directory server but not in the database.
ANR2749W Admin NODE1 was located in the LDAP directory server but not in the database.
ANR2749W Admin NODE2 was located in the LDAP directory server but not in the database.
ANR2748W Node NODE1 was located in the LDAP directory server but not in the database.
ANR2748W Node NODE2 was located in the LDAP directory server but not in the database.
ANR2745I AUDIT LDAPDIRECTORY command completed: 4 administrator entries are only in the
LDAP directory server (not in the Tivoli Storage Manager server), 0 administrator entries
are only in the Tivoli Storage Manager server (not in the LDAP directory server), 2 node
entries are only in the LDAP directory server (not in the Tivoli Storage Manager server),
0 node entries are only in the Tivoli Storage Manager server, (not in the LDAP directory
server), 6 entries were deleted from the LDAP server in total.
```

Related commands

Table 19. Commands related to AUDIT LDAPDIRECTORY

Command	Description
SET DEFAULTAUTHENTICATION	Specifies the default password authentication method for any REGISTER NODE or REGISTER ADMIN commands.
SET LDAPPASSWORD	Sets the password for the LDAPUSER.
SET LDAPUSER	Sets the user who oversees the passwords and administrators on the LDAP directory server.

AUDIT LIBRARY (Audit volume inventories in an automated library)

Use this command to audit and synchronize volume inventories in an automated library.

When the **AUDIT LIBRARY** command is issued on a library client, the client synchronizes its inventory with the inventory on the library manager. If the library client detects inconsistencies, it corrects them by changing the ownership of the volume on the library manager.

When the **AUDIT LIBRARY** command is issued on a server where the library is SCSI, 349X, or ACSLS (LIBTYPE=SCSI, LIBTYPE=349X, or LIBTYPE=ACSL), the server synchronizes its inventory with the inventory of the library device. If the server detects inconsistencies, it deletes missing volumes from its inventory.

- In SCSI libraries, the server also updates the locations of volumes in its inventory that have been moved since the last audit.
- In 349X libraries, the server also ensures that scratch volumes are in the scratch category and that private volumes are in the private category.

When the **AUDIT LIBRARY** command is issued on a server that is a library manager for the library (SHARED=YES), the server updates ownership of its volumes if it detects inconsistencies.

Regardless the type of server or type of library, issuing the **AUDIT LIBRARY** command does not automatically add new volumes to a library. To add new volumes, you must use the **CHECKIN LIBVOLUME** command.

Attention: The following precautions apply to SCSI, 349X, and ACSLS libraries only (LIBTYPE=SCSI, LIBTYPE=349X, and LIBTYPE=ACSL):

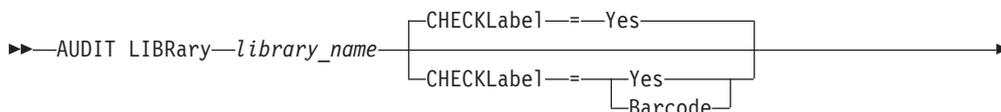
- Running the **AUDIT LIBRARY** command prevents any other library activity until the audit completes. For example, Tivoli Storage Manager will not process restore or retrieve requests that involve the library when the **AUDIT LIBRARY** command is running.
- If other activity is occurring in the library, do not issue the **AUDIT LIBRARY** command. Issuing the **AUDIT LIBRARY** command when a library is active can produce unpredictable results (for example, a hang condition) if a process currently accessing the library attempts to acquire a new tape mount.

This command creates a background process that you can cancel with the **CANCEL PROCESS** command. To display information about background processes, use the **QUERY PROCESS** command.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax





Parameters

library_name (Required)

Specifies the name of the library to audit.

CHECKLabel

Specifies how the storage volume label is checked during the audit. This parameter applies to SCSI libraries only. The parameter is ignored for other library types. The default is YES. Possible values are:

Yes

Specifies that Tivoli Storage Manager checks each volume label to verify the identity of the volume.

Barcode

Specifies that Tivoli Storage Manager uses the barcode reader to read the storage label. Using the barcode decreases the audit processing time. This parameter applies only to SCSI libraries.

Attention: If the scanner cannot read the barcode label or the barcode label is missing, Tivoli Storage Manager loads that tape in a drive to read the label.

REFRESHstate

Specifies whether the server's information about a library, which is normally obtained during initialization, is refreshed, so that any changes in configuration are reflected. By setting the REFRESHSTATE parameter to Yes, this action is completed without having to restart the server or re-define the library. The default is No. Possible values are:

No Specifies that the Tivoli Storage Manager server does not refresh the library's state when the library is audited.

Yes

Specifies that the Tivoli Storage Manager server does refresh the library's state when the **AUDIT LIBRARY** command is issued.

Example: Audit an automated library

Audit the EZLIFE automated library.

```
audit library ezlife
```

Related commands

Table 20. Commands related to **AUDIT LIBRARY**

Command	Description
CANCEL PROCESS	Cancels a background server process.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE LIBRARY	Deletes a library.
DISMOUNT VOLUME	Dismounts a sequential, removable volume by the volume name.

AUDIT LIBRARY

Table 20. Commands related to **AUDIT LIBRARY** (continued)

Command	Description
QUERY LIBRARY	Displays information about one or more libraries.
QUERY LIBVOLUME	Displays information about a library volume.
QUERY PROCESS	Displays information about background processes.
UPDATE LIBRARY	Changes the attributes of a library.

AUDIT LICENSES (Audit server storage usage)

Use this command to audit the server storage used by client nodes and to audit the server licenses. The audit determines whether the current configuration is in compliance with the license terms.

An audit creates a background process you can cancel with the **CANCEL PROCESS** command. If you halt and restart the server, an audit is run automatically as specified by the **SET LICENSEAUDITPERIOD**. To view audit results, use the **QUERY LICENSE** command.

Attention: The audit of server storage can take a lot of CPU time. You can use the **AUDITSTORAGE** server option to specify that storage is not to be audited.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—AUDIT LICENSES—◄◄

Parameters

None.

Example: Audit server licenses

Issue the **AUDIT LICENSES** command.

```
audit licenses
```

Related commands

Table 21. Commands related to **AUDIT LICENSES**

Command	Description
CANCEL PROCESS	Cancels a background server process.
QUERY AUDITOCCUPANCY	Displays the server storage utilization for a client node.
QUERY LICENSE	Displays information about licenses and audits.
QUERY PROCESS	Displays information about background processes.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER LICENSE	Registers a license with the Tivoli Storage Manager server.
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.

AUDIT VOLUME (Verify database information for a storage pool volume)

Use this command to check for inconsistencies between database information and a storage pool volume. Processing information generated during an audit is sent to the activity log and server console.

You can only audit volumes that belong to storage pools with `DATAFORMAT=NATIVE` and `DATAFORMAT=NONBLOCK`.

You cannot audit a volume if it is being deleted from a primary or copy storage pool.

While an audit process is active, clients cannot restore data from the specified volume or store new data to that volume.

If the server detects a file with errors, handling of the file will depend on the type of storage pool to which the volume belongs, whether the `FIX` option is specified on this command, and whether the file is also stored on a volume assigned to other pools.

If Tivoli Storage Manager does not detect errors for a file that was marked as damaged, the state of the file is reset so that it can be used.

The Tivoli Storage Manager server will not delete archive files that are on deletion hold. If archive retention protection is enabled, the Tivoli Storage Manager server will not delete archive files whose retention period has not expired.

To display information about the contents of a storage pool volume, use the **QUERY CONTENT** command.

To audit multiple volumes, you can use the `FROMDATE` and `TODATE` parameters. Use the `STGPOOL` parameter to audit all volumes in a storage pool. When you use the parameters `FROMDATE`, `TODATE`, or both, the server limits the audit to only the sequential media volumes that meet the date criteria, and automatically includes all online disk volumes in storage. To limit the number of volumes that may include disk volumes, use the `FROMDATE`, `TODATE`, and `STGPOOL` parameters.

If you are running a server with archive retention protection enabled, and you have data stored in storage pools which are defined with the parameter `RECLAMATIONTYPE=SNAPLOCK`, the Last Access Date on the NetApp SnapLock Filer for a volume should be equal to the End Reclaim Period date that you see when you issue a **QUERY VOLUME F=D** command on that volume. During `AUDIT VOLUME` processing, these dates are compared. If they do not match and the **AUDIT VOLUME** command is being run with the `FIX=NO` parameter, a message will be issued to you indicating that the command should be run with the `FIX=YES` parameter to resolve the inconsistency. If they do not match and the `AUDIT VOLUME` command is being run with the `FIX=YES` parameter, the inconsistencies will be resolved.

|
|
|
|
|
|

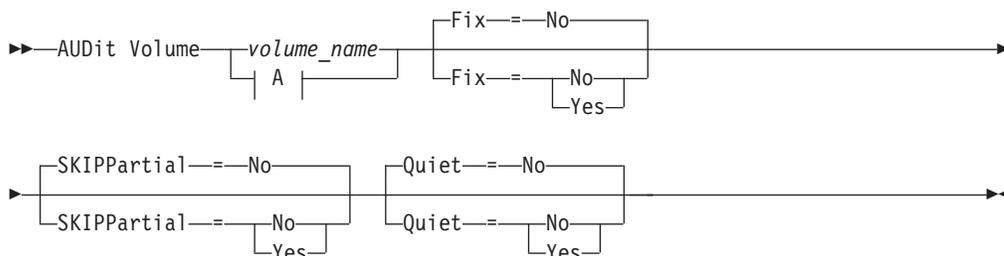
Attention: Use the FIX=Yes parameter only if your tape drive and storage area network (SAN) infrastructure is stable. Ensure that the tape heads are clean and that the tape device drivers are stable and reliable. Otherwise, you risk deleting data that is error free when you use this parameter. The Tivoli Storage Manager server cannot determine whether a tape is physically damaged or whether a tape infrastructure is unstable.

This command creates a background process that can be canceled with the **CANCEL PROCESS** command. To display information on background processes, use the **QUERY PROCESS** command.

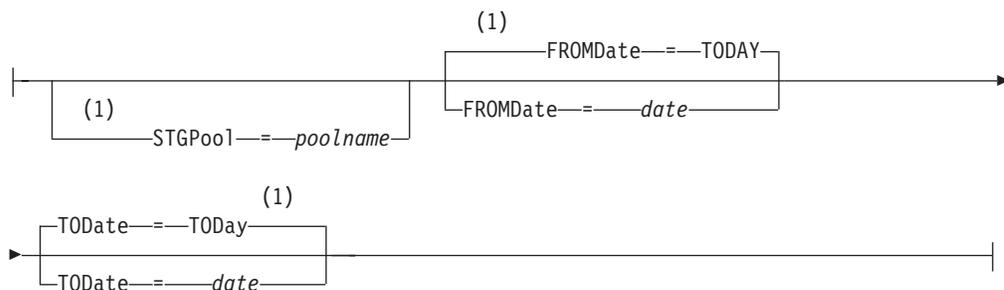
Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to which the volume is defined.

Syntax



A (at least one of these parameters must be specified):



Notes:

- 1 You cannot specify a volume name if you specify a storage pool name, FROMDATE, or TODATE.

Parameters

volume_name

Specifies the name of the storage pool volume you want to audit. This parameter is required if you do not specify a storage pool. You cannot specify a volume name together with the FROMDATE and TODATE parameters.

Fix

Specifies how the server resolves inconsistencies between the database inventory and the specified storage pool volume. This parameter is optional. The default is NO.

The actions the server performs depend on whether the volume is assigned to a primary or a copy storage pool.

Primary Storage Pool:

Note: If the **AUDIT VOLUME** command does not detect an error in a file that was previously marked as damaged, Tivoli Storage Manager resets the state of the file so that it can be used. This provides a means for resetting the state of damaged files if it is determined that the errors were caused by a correctable hardware problem such as a dirty tape head.

Fix=No

Tivoli Storage Manager reports, but does not delete, database records that refer to files with inconsistencies:

- Tivoli Storage Manager marks the file as damaged in the database. If a backup copy is stored in a copy storage pool, you can restore the file using the **RESTORE VOLUME** or **RESTORE STGPOOL** command.
- If the file is a cached copy, you must delete references to the file on this volume by issuing the **AUDIT VOLUME** command and specifying **FIX=YES**. If the physical file is not a cached copy, and a duplicate is stored in a copy storage pool, it can be restored by using the **RESTORE VOLUME** or **RESTORE STGPOOL** command.

Fix=Yes

The server fixes any inconsistencies as they are detected:

- If the physical file is a cached copy, the server deletes the database records that refer to the cached file. The primary file is stored on another volume.
- If the physical file is not a cached copy, and the file is also stored in one or more copy storage pools, the error will be reported and the physical file marked as damaged in the database. You can restore the physical file by using the **RESTORE VOLUME** or **RESTORE STGPOOL** command.
- If the physical file is not a cached copy, and the physical file is not stored in a copy storage pool, each logical file for which inconsistencies are detected are deleted from the database.
- If archive retention protection is enabled by using the **SET ARCHIVERETENTIONPROTECTION** command, a cached copy of data can be deleted if needed. Data in primary and copy storage pools can only be marked damaged and never deleted.

Do not use the **AUDIT VOLUME** command with **FIX=YES** if a restore process (**RESTORE STGPOOL** or **RESTORE VOLUME**) is running. The **AUDIT VOLUME** command could cause the restore to be incomplete.

Copy Storage Pool:

Fix=No

The server reports the error and marks the physical file copy as damaged in the database.

Fix=Yes

The server deletes any references to the physical file and any database records that point to a physical file that does not exist.

SKIPPARTIAL

Specifies whether Tivoli Storage Manager ignores skipped files, which are files that span multiple storage pool volumes. This parameter is optional. The default value is NO. When performing an audit operation on a sequential access media volume, this parameter prevents additional sequential access media mounts that may be necessary to audit any skipped files. Possible values are:

No Tivoli Storage Manager audits files that span multiple volumes.

Unless you specify SKIPPARTIAL=YES, Tivoli Storage Manager attempts to process each file stored on the volume, including files that span into and out of other volumes. To audit files that span multiple volumes, the following conditions must be true:

- For sequential access volumes, the additional sequential access volumes must have an access mode of read/write or read-only.
- For random access volumes, the additional volumes must be online.

Yes

Tivoli Storage Manager audits only files that are stored on the volume to be audited. The status of any skipped files is unknown.

Quiet

Specifies whether Tivoli Storage Manager sends detailed informational messages to the activity log and the server console about irretrievable files on the volume. This parameter is optional. The default is NO. Possible values are:

No Specifies that Tivoli Storage Manager sends detailed informational messages and a summary. Each message contains the node, file space, and client name for the file.

Yes

Specifies that Tivoli Storage Manager sends only a summary report.

FROMDate

Specifies the beginning date of the range to audit volumes. The default is the current date. All sequential media volumes meeting the time range criteria that were written to after this date are audited. The server includes all online disk volumes in storage. The server starts one audit process for each volume and runs the process serially. You cannot use this parameter if you have specified a volume. This parameter is optional. To limit the number of volumes that may include disk volumes, use the FROMDATE, TODATE, and STGPOOL parameters.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	10/15/2001
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -7 or -7. To display information beginning with volumes written a week ago, you can specify FROMDATE=TODAY-7 or FROMDATE= -7.

AUDIT VOLUME

Value	Description	Example
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

TODate

Specifies the ending date of the range for volumes to audit. All sequential media volumes meeting the time range criteria that were written to before this date are audited. The server includes all online disk volumes in storage. If you do not specify a value, the server defaults to the current date. You cannot use this parameter if you have specified a volume. This parameter is optional. To limit the number of volumes that may include disk volumes, use the FROMDATE, TODATE, and STGPOOL parameters.

You can specify the date by using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	10/15/2001 If a date is entered, all candidate volumes written on that day (ending at 11:59:59 pm) will be evaluated.
TODAY	The current date	TODAY
TODAY- <i>days</i> or - <i>days</i>	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY-1 or -1. To display information created up to yesterday, you can specify TODATE=TODAY-1 or simply TODATE= -1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

STGPool

This parameter specifies that the server only audits the volumes from the

specified storage pool. This parameter is optional. You cannot use this parameter if you have specified a volume.

Example: Verify database information for a specific storage pool volume

Verify that the database information for storage pool volume PROG2 is consistent with the data stored on the volume. Tivoli Storage Manager fixes any inconsistencies.

```
audit volume prog2 fix=yes
```

Example: Verify database information for all volumes written to during a specific date range

Verify that the database information for all eligible volumes written to from 3/20/2002 to 3/22/2002 is consistent with data stored on the volume.

```
audit volume fromdate=03/20/2002 todate=03/22/2002
```

Example: Verify database information for all volumes in a specific storage pool

Verify that the database information for all volumes in storage pool STPOOL3 is consistent with data stored on the volume for today.

```
audit volume stgpool=STPOOL3
```

Example: Verify database information for all volumes in a specific storage pool written to in the last two days

Verify that the database information for all volumes in storage pool STPOOL3 is consistent with data stored on the volume for the last two days.

```
audit volume stgpool=STPOOL3 fromdate=-1
```

Related commands

Table 22. Commands related to AUDIT VOLUME

Command	Description
CANCEL PROCESS	Cancels a background server process.
QUERY CONTENT	Displays information about files in a storage pool volume.
QUERY PROCESS	Displays information about background processes.
QUERY VOLUME	Displays information about storage pool volumes.
SET ARCHIVERETENTIONPROTECTION	Specifies whether data retention protection is activated.

BACKUP commands

Use the **BACKUP** commands to create backup copies of Tivoli Storage Manager information or objects.

- “BACKUP DB (Back up the database)” on page 47
- “BACKUP DEVCONFIG (Create backup copies of device configuration information)” on page 51
- “BACKUP NODE (Back up a NAS node)” on page 53
- “BACKUP STGPOOL (Back up primary storage pool to copy storage pool)” on page 58
- “BACKUP VOLHISTORY (Save sequential volume history information)” on page 62

BACKUP DB (Back up the database)

Use this command to back up a Tivoli Storage Manager database to sequential access volumes.

Attention: To restore a database, the server must use information from the volume history file and the device configuration file. You must make and save copies of the volume history file and the device configuration file. These files cannot be recreated.

To determine how much additional storage space a backup requires, issue the **QUERY DB** command.

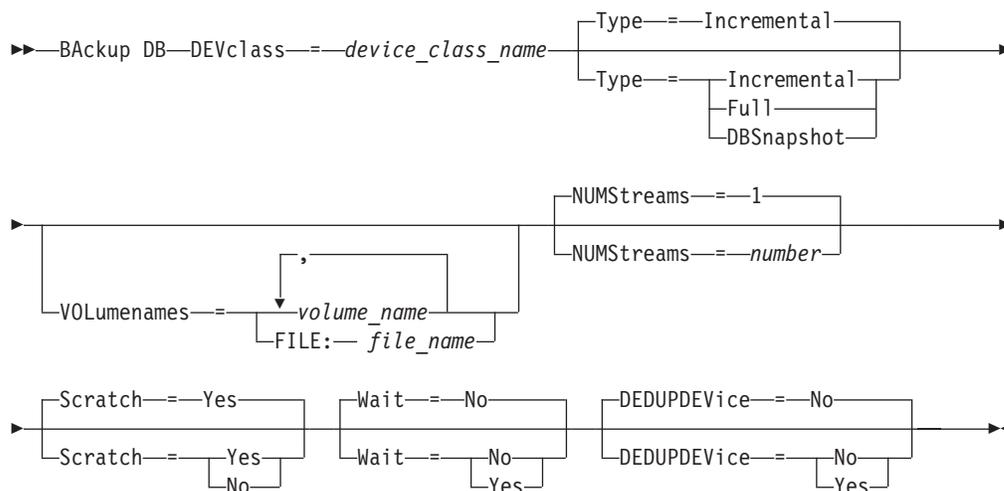
Restriction: You cannot restore a server database from a different Tivoli Storage Manager server if both of the following conditions are present:

- The database backup is stored on virtual volumes
- The connection to the Tivoli Storage Manager server is protected by SSL (Secure Sockets Layer)

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

DEVclass (Required)

Specifies the name of the sequential access device class to use for the backup. If you issue the **BACKUP DB** command, and the device class is not the one that is specified in the **SET DBRECOVERY** command, a warning message is issued. However, the backup operation continues and is not affected.

If the **SET DBRECOVERY** command has not been issued to set a device class, the **BACKUP DB** command fails.

Restriction:

- You cannot use a device class with a device type of NAS or CENTERA.

BACKUP DB

- A restore database operation fails if the source for the restore is a FILE library. A FILE library is created if the FILE device class specifies SHARED=YES.

If all drives for this device class are busy when the backup runs, Tivoli Storage Manager cancels lower priority operations, such as reclamation, to make a drive available for the backup.

Type

Specifies the type of backup to run. This parameter is optional. The default is INCREMENTAL. Possible values are:

Incremental

Specifies that you want to run an incremental backup of the Tivoli Storage Manager database. An incremental (or cumulative) backup image contains a copy of all database data that has changed since the last successful full backup operation was performed.

Full

Specifies that you want to run a full backup of the Tivoli Storage Manager database.

DBSnapshot

Specifies that you want to run a full snapshot database backup. The entire contents of a database are copied and a new snapshot database backup is created without interrupting the existing full and incremental backup series for the database.

VOLUMenames

Specifies the volumes used to back up the database. This parameter is optional. However, if you specify SCRATCH=NO, you must specify a list of volumes.

volume_name

Specifies the volumes used to back up the database. Specify multiple volumes by separating the names with commas and no intervening spaces.

FILE:*filename*

Specifies the name of a file that contains a list of volumes used to back up the database. Each volume name must be on a separate line. Blank lines and comment lines, which begin with an asterisk, are ignored.

For example, to use volumes DB0001, DB0002, and DB0003, create a file that contains these lines:

```
DB0001
DB0002
DB0003
```

Name the file appropriately. For example:

```
TAPEVOL
```

You can then specify the volumes for the command as follows:

```
VOLUMENAMES=FILE:TAPEVOL
```

NUMStreams

Specifies the number of parallel data movement streams to use when backing up the database. The minimum value is 1, and the maximum value is 4. Increasing the value causes a corresponding increase in the number of database backup sessions to be used and the number of drives to be used for the device class. If you specify a **NUMSTREAMS** value in the **BACKUP DB** command, it

overrides any value set in the **SET DBRECOVERY** command. Otherwise, the value set in the **SET DBRECOVERY** command is used. The **NUMSTREAMS** value is used for all types of database backups.

If a value is specified that is greater than the number of drives available for the device class, only the number of available drives are used. The available drives are those defined to the device class by the **MOUNTLIMIT** parameter or by the number of online drives for the specified device class. The session is displayed in the **QUERY SESSION** output.

If you increase the number of streams, more volumes are used from the corresponding device class for this operation. Using more volumes might improve the speed of the database backups, but at the cost of more volumes that are not fully used.

For details about using multiple streams for database backup and restore operations, see the *Administrator's Guide*.

Scratch

Specifies whether scratch volumes can be used for the backup. This parameter is optional. The default is YES. Possible values are:

Yes

Specifies that scratch volumes can be used.

If you specify **SCRATCH=YES** and the **VOLUMENAMES** parameter, Tivoli Storage Manager only uses scratch volumes if space is unavailable on the specified volumes.

If you do not include a list of volumes by using the **VOLUMENAMES** parameter, you must either specify **SCRATCH=YES** or use the default.

No

Specifies that scratch volumes cannot be used. If you specify volumes by using the **VOLUMENAMES** parameter and **SCRATCH=NO**, the backup fails if there is not enough space available to store the backup data on the specified volumes.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. The default is NO. Possible values are:

No

Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed. Messages created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the **CANCEL PROCESS** command. If a **BACKUP DB** background process is canceled, some of the database might have already been backed up before the cancellation.

Yes

Specifies that the server processes this command in the foreground. Wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify **WAIT=YES** from the server console.

DEDUPDEVICE

Specifies that a target storage device supports data deduplication. When set to

BACKUP DB

YES, the format for backup images is optimized for data deduplication devices, making backup operations more efficient. Possible values are:

No Specifies that a target storage device does not support data deduplication. NO is the default.

Ensure that this parameter is set to NO for the following devices:

- SCSI libraries
- All devices that are defined with a FILE device class
- Virtual tape libraries (VTL) that do not support the data deduplication function

Yes

Specifies that a target device supports data deduplication and that you want to optimize backups for this function. You can set this parameter to YES if you are using a VTLs that supports data deduplication.

Example: Run an incremental backup by using a scratch volume

Run an incremental backup of the database by using a scratch volume. Use a device class of FILE for the backup.

```
backup db devclass=file type=incremental
```

Related commands

Table 23. Commands related to **BACKUP DB**

Command	Description
BACKUP DEVCONFIG	Backs up Tivoli Storage Manager device information to a file.
BACKUP VOLHISTORY	Records volume history information in external files.
CANCEL PROCESS	Cancels a background server process.
DELETE VOLHISTORY	Removes sequential volume history information from the volume history file.
EXPIRE INVENTORY	Manually starts inventory expiration processing.
MOVE DRMEDIA	Moves DRM media on-site and off-site.
PREPARE	Creates a recovery plan file.
QUERY DB	Displays allocation information about the database.
QUERY PROCESS	Displays information about background processes.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
SET DBRECOVERY	Specifies the device class to be used for automatic backups.
SET DRMDBBACKUPEXPIREDDAYS	Specifies criteria for database backup series expiration.

BACKUP DEVCONFIG (Create backup copies of device configuration information)

Use this command to back up information about device configuration for the server.

Attention: To restore a database, the server must use information from the volume history file and the device configuration file. You must make and save copies of the volume history file and the device configuration file. These files cannot be recreated.

This command backs up the following information in one or more files:

- Device class definitions
- Library definitions
- Drive definitions
- Path definitions when **SRCTYPE=SERVER**
- Server definitions
- Server name
- Server password
- Volume location information for **LIBTYPE=SCSI** libraries

You can use the DEVCONFIG server option to specify one or more files in which to store device configuration information. Tivoli Storage Manager updates the files whenever a device class, library, or drive is defined, updated, or deleted.


To ensure updates are complete before the server is halted:

- Do not halt the server for a few minutes after issuing the **BACKUP DEVCONFIG** command.
- Specify multiple DEVCONFIG options in the server options file.
- Examine the device configuration file to see if the file has been updated.

Privilege class

Any administrator can issue this command unless it includes the **FILENAMES** parameter. If the **FILENAMES** parameter is specified and the REQSYSAUTHOUTFILE server option is set to YES, the administrator must have system privilege. If the **FILENAMES** parameter is specified and the REQSYSAUTHOUTFILE server option is set to NO, the administrator must have operator, policy, storage or system privilege.

Syntax



Parameters

FileNames

Specifies the files in which to store device configuration information. You can specify multiple files by separating the names with commas and no intervening spaces. This parameter is optional.

BACKUP DEVCONFIG

If you do not specify a file name, Tivoli Storage Manager stores the information in all files specified with the DEVCONFIG option in the server options file.

Example: Backup device configuration information to a file

Back up device configuration information to a file named DEVICE.

```
backup devconfig filenames=device
```

Related commands

Table 24. Commands related to BACKUP DEVCONFIG

Command	Description
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE DEVCLASS	Defines a device class.
DEFINE DEVCLASS (z/OS media server)	Defines a device class to use storage managed by a z/OS media server.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE PATH	Defines a path from a source to a destination.
DEFINE SERVER	Defines a server for server-to-server communications.
LABEL LIBVOLUME	Labels volumes in manual or automated libraries.
QUERY LIBVOLUME	Displays information about a library volume.
SET SERVERNAME	Specifies the name by which the server is identified.
SET SERVERPASSWORD	Specifies the server password.
UPDATE DEVCLASS	Changes the attributes of a device class.
UPDATE DEVCLASS (z/OS media server)	Changes the attributes of a device class for storage managed by a z/OS media server.
UPDATE DRIVE	Changes the attributes of a drive.
UPDATE LIBRARY	Changes the attributes of a library.
UPDATE LIBVOLUME	Changes the status of a storage volume.
UPDATE PATH	Changes the attributes associated with a path.
UPDATE SERVER	Updates information about a server.

BACKUP NODE (Back up a NAS node)

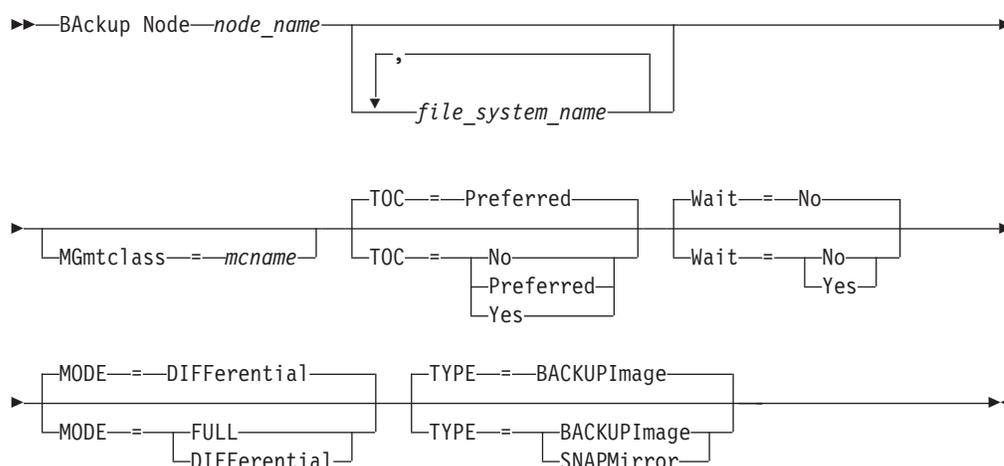
Use this command to start a backup operation for a network-attached storage (NAS) node.

Backups that are created for NAS nodes with this **BACKUP NODE** command are functionally equivalent to backups that are created by using the **BACKUP NAS** command on a Tivoli Storage Manager client. You can restore these backups with either the server's **RESTORE NODE** command or the client's **RESTORE NAS** command.

Privilege class

To issue this command, you must have system privilege, policy privilege for the domain to which the node is assigned, or client owner authority over the node.

Syntax



Parameters

node_name (Required)

Specifies the node for which the backup will be performed. You cannot use wildcard characters or specify a list of names.

file_system_name

Specifies the name of one or more file systems to back up. You can also specify names of virtual file spaces that have been defined for the NAS node. The file system name that you specify cannot contain wildcard characters. You can specify more than one file system by separating the names with commas and no intervening spaces.

If you do not specify a file system, all file systems will be backed up. Any virtual file spaces defined for the NAS node are backed up as part of the file system image, not separately.

If a file system exists on the NAS device with the same name as the virtual file space specified, Tivoli Storage Manager automatically renames the existing virtual file space in the server database, and backs up the NAS file system which matches the name specified. If the virtual file space has backup data, the file space definition associated with the virtual file space will also be renamed.

Tip: See the virtual file space name parameter in the **DEFINE VIRTUALFSMAPPING** command for more naming considerations.

BACKUP NODE

In determining the file systems to process, the server will not use any DOMAIN.NAS, INCLUDE.FS.NAS, or EXCLUDE.FS.NAS statements in any client option file or client option set. If you back up multiple file systems, the backup of each file system is a separate server process.

MGmtclass

Specifies the name of the management class to which this backup data is bound. If you do not specify a management class, the backup data is bound to the default management class of the policy domain to which the node is assigned. In determining the management class, the server will *not* use any INCLUDE.FS.NAS statements in any client option file or client option set. The destination management class may refer to a Tivoli Storage Manager native pool, in which case Network Data Management Protocol (NDMP) data is sent into the Tivoli Storage Manager native hierarchy. After this occurs, the data stays in the Tivoli Storage Manager hierarchy. Data flowing to Tivoli Storage Manager native pools goes over the LAN and data flowing to NAS pools can be directly attached or over a SAN.

When you specify a management class with the BACKUP NODE command, all versions of the backup data that belong to the NAS node are rebound to the new management class.

TOC

Specifies whether a table of contents (TOC) is saved for each file system backup. Consider the following in determining whether you want to save a table of contents.

- If a table of contents is saved, you will be able to use the **QUERY TOC** command to determine the contents of a file system backup in conjunction with the **RESTORE NODE** command to restore individual files or directory trees. You will also be able to use the Tivoli Storage Manager Web backup-archive client to examine the entire file system tree and choose files and directories to restore. Creation of a table of contents requires that you define the TOCDESTINATION attribute in the backup copy group for the management class to which this backup image is bound. Note that a table of contents creation requires additional processing, network resources, storage pool space, and possibly a mount point during the backup operation.
- A table of contents for a NAS file system cannot have a directory path greater than 1024 characters.
- If a table of contents is not saved for a file system backup, you will still be able to restore individual files or directory trees using the **RESTORE NODE** command, provided that you know the fully qualified name of each file or directory to be restored and the image in which that object was backed up. This parameter is optional. The default value is Preferred. Possible values are:

No Specifies that table of contents information is not saved for file system backups.

Preferred

Specifies that table of contents information should be saved for file system backups. However, a backup does not fail just because an error occurs during creation of the table of contents. This is the default value.

Yes

Specifies that table of contents information must be saved for each file system backup. A backup fails if an error occurs during creation of the table of contents.

Attention: If `MODE=DIFFERENTIAL` is specified and a table of contents is requested (`TOC=PREFERRED` or `TOC=YES`), but the last full image does not have a table of contents, a full backup will be performed and a table of contents will be created for that full backup.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. The default is `NO`. Possible values are:

No Specifies that the server processes this command in the background. Use the `QUERY PROCESS` command to monitor the background processing of this command.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client when the command completes. If you are backing up multiple file systems, all backup processes must complete before the command is complete.

Attention: You cannot specify `WAIT=YES` from the server console.

MODE

Specifies whether the file system backups are full or differential. The default is `DIFFERENTIAL`.

FULL

Specifies to back up the entire file system.

DIFFerential

Specifies that only the files that have changed since the most recent full backup should be backed up. If you choose a differential backup, and a full backup is not found, a full backup is performed. You cannot specify `TYPE=SNAPMIRROR` when the **MODE** parameter is set to `DIFFERENTIAL`.

TYPE

Specifies the backup method used to perform the NDMP backup operation. The default value for this parameter is `BACKUPIIMAGE` and it should be used to perform a standard NDMP base or differential backup. Other image types represent backup methods that might be specific to a particular file server. Possible values are:

BACKUPIImage

Specifies that the file system should be backed up using an NDMP dump operation. This is the default method for performing an NDMP backup. The `BACKUPIIMAGE` type operation supports full and differential backups, file-level restore processing and directory-level backup.

SNAPMirror

Specifies that the file system should be copied to a Tivoli Storage Manager storage pool using the NetApp SnapMirror to Tape function. SnapMirror images are block level full backup images of a file system. Typically, a SnapMirror backup takes significantly less time to perform than a traditional NDMP full file system backup. However there are limitations and restrictions on how SnapMirror images can be used. The SnapMirror to Tape function is intended to be used as a disaster-recovery option for copying very large NetApp file systems to secondary storage.

BACKUP NODE

For most NetApp file systems, use the standard NDMP full or differential backup method. See the *Administrator's Guide* for limitations on using SnapMirror images as a backup method.

Refer to the documentation that came with your NetApp file server for more information. When setting the TYPE parameter to SNAPMirror, note the following restrictions:

1. You cannot specify TOC=YES or TOC=PREFERRED.
2. The file_system_name cannot be a virtual filesystem name.
3. The snapshot which is created automatically by the file server during the SnapMirror copy operation will be deleted at end of the operation.
4. This parameter is valid for NetApp and IBM N-Series file servers only.

Example: Perform a full backup

Perform a full backup on the /vol/vol10 file system of NAS node NAS1.

```
backup node nas1 /vol/vol10 mode=full
```

Example: Perform a backup on a directory and create a table of contents

Back up the directory /vol/vol2/mikes on the node NAS1 and create a table of contents for the image. For the following two examples, assume Table 25 contains the virtual file space definitions exist on the server for the node NAS1.

```
backup node nas1 /mikesdir
```

Table 25. Virtual file space definitions

Virtual file space name	File system	Path
/mikesdir	/vol/vol2	/mikes
/DataDirVol2	/vol/vol2	/project1/data
/TestDirVol1	/vol/vol1	/project1/test

Example: Perform a backup on two directories

Back up the directories /vol/vol2/project1/data and /vol/vol1/project1/test of the node NAS1. Refer to Table 25 for the the virtual file space definitions that exist on the server for the node NAS1.

```
backup node nas1 /DataDirVol2,/testdirvol1 mode=full toc=yes
```

Related commands

Table 26. Commands related to BACKUP NODE

Command	Description
BACKUP NAS (Tivoli Storage Manager client command)	Creates a backup of NAS node data.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE VIRTUALFSMAPPING	Define a virtual file space mapping.
QUERY NASBACKUP	Displays information about NAS backup images.

Table 26. Commands related to BACKUP NODE (continued)

Command	Description
QUERY TOC	Displays details about the table of contents for a specified backup image.
QUERY COPYGROUP	Displays the attributes of a copy group.
RESTORE NAS (Tivoli Storage Manager client command)	Restores a backup of NAS node data.
RESTORE NODE	Restores a network-attached storage (NAS) node.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.

BACKUP STGPOOL (Back up primary storage pool to copy storage pool)

Use this command to back up primary storage pool files to a copy storage pool.

In addition to backing up primary storage pools having NATIVE or NONBLOCK data formats, this command lets you back up primary storage pools that have NDMP data formats (NETAPPDUMP, CELERRADUMP, or NDMPDUMP). The copy storage pool to which data is to be backed up must have the same data format as the primary storage pool. Tivoli Storage Manager supports backend data movement for NDMP images. For details, see the *Administrator's Guide*.

If a file already exists in the copy storage pool, the file is not backed up unless the copy of the file in the copy storage pool is marked as damaged. However, a new copy is not created if the file in the primary storage pool is also marked as damaged. In a random-access storage pool, neither cached copies of migrated files nor damaged primary files are backed up.

If migration for a storage pool starts during a storage pool backup, some files may be migrated before they are backed up. You may want to back up storage pools that are higher in the migration hierarchy before backing up storage pools that are lower. For example, when performing a storage pool backup to copy the contents of a storage pool offsite, if the process is not done according to the existing storage pool hierarchy, some files may not be copied to the copy storage pool. This could become critical for disaster recovery purposes. When performing a storage pool backup on multiple storage pools, the primary storage pool should be completed before starting the backup process on the next storage pool.


Remember: Issuing this command for a primary storage pool that is set up for data deduplication removes duplicate data, if the copy storage pool is also set up for data deduplication.


| **Restriction:** Do not run the **MOVE DRMEDIA** and **BACKUP STGPOOL** commands
 | concurrently. Ensure that the storage pool backup processes are complete before
 | you issue the **MOVE DRMEDIA** command.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the copy storage pool in which backup copies are to be produced.

Syntax





Notes:

- 1 Valid only for storage pools associated with a sequential-access device class.

Parameters

primary_pool **(Required)**

Specifies the primary storage pool.

copy_pool **(Required)**

Specifies the copy storage pool.

MAXProcess

Specifies the maximum number of parallel processes to use for backing up files. This parameter is optional. Enter a value from 1 to 999. The default is 1.

Using multiple, parallel processes may improve throughput for the backup. The expectation is that the time needed to perform the storage pool backup will be decreased by using multiple processes. However, when multiple processes are running, in some cases one or more of the processes needs to wait to use a volume that is already in use by a different backup process.

When determining this value, consider the number of logical and physical drives that can be dedicated to this operation. To access a sequential access volume, Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the backup.

Each process needs a mount point for copy storage pool volumes, and, if the device type is not FILE, each process also needs a drive. If you are backing up a sequential storage pool, each process needs an additional mount point for primary storage pool volumes and, if the device type is not FILE, an additional drive. For example, suppose you specify a maximum of 3 processes to back up a primary sequential storage pool to a copy storage pool of the same device class. Each process requires 2 mount points and 2 drives. To run all 3 processes, the device class must have a mount limit of at least 6, and at least 6 mount points and 6 drives must be available.

To preview a backup, only one process is used and no mount points or drives are needed.

Preview

Specifies if you want to preview but not perform the backup. The preview displays the number of files and bytes to be backed up and a list of the primary storage pool volumes that you must mount. This parameter is optional. The default is NO. Possible values are:

No Specifies that the backup is done.

Yes

Specifies that you want to preview the backup but not do the backup.

VOLumesonly

Specifies that you want to preview the backup only as a list of the volumes

BACKUP STGPOOL

that must be mounted. This choice requires the least processing time. The VOLUMESONLY option is valid only for storage pools associated with a sequential-access device class.

SHREDTONOshred

Specifies whether data will be backed up to a copy storage pool from a primary storage pool that enforces shredding. This parameter is optional. The default value is NO. Possible values are:

No Specifies that the server will not allow data to be backed up to a copy storage pool from a primary storage pool that enforces shredding. If the primary storage pool enforces shredding, the operation will fail.

Yes

Specifies that the server does allow data to be backed up to a copy storage pool from a primary storage pool that enforces shredding. The data in the copy storage pool will not be shredded when it is deleted.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is NO. Possible values are:

No Specifies that the server processes this command in the background.

You can continue with other tasks while the command is being processed. Messages created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the **CANCEL PROCESS** command. If you cancel this process, some files may have already been backed up prior to the cancellation.

Yes

Specifies that the server performs this operation in the foreground. You must wait for the operation to complete before continuing with other tasks. The server displays the output messages to the administrative client when the operation completes.

Note: You cannot specify WAIT=YES from the server console.

Example: Backup the primary storage pool

Back up the primary storage pool named PRIMARY_POOL to the copy storage pool named COPYSTG.

```
backup stgpool primary_pool copystg
```

Related commands

Table 27. Commands related to **BACKUP STGPOOL**

Command	Description
CANCEL PROCESS	Cancels a background server process.
MOVE DRMEDIA	Moves DRM media on-site and off-site.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY PROCESS	Displays information about background processes.

Table 27. Commands related to **BACKUP STGPOOL** (continued)

Command	Description
QUERY SHREDSTATUS	Displays information about data waiting to be shredded.
QUERY STGPOOL	Displays information about storage pools.
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.
SHRED DATA	Manually starts the process of shredding deleted data.

BACKUP VOLHISTORY (Save sequential volume history information)

Use this command to back up sequential volume history information to one or more files.

Note: You must use volume history information when you reload the database and audit affected storage pool volumes. If you cannot start the server, you can use the volume history file to query the database about these volumes.

The volume history includes information about the following types of volumes:

- Archive log volumes
- Database backup volumes
- Export volumes
- Backup set volumes
- Database snapshot volumes
- Database recovery plan file volumes
- Recovery plan file volumes
- Recovery plan file snapshot volumes
- The following sequential access storage pool volumes:
 - Volumes added to storage pools
 - Volumes reused through reclamation or MOVE DATA operations
 - Volumes removed by using the DELETE VOLUME command or during reclamation of scratch volumes

Attention: To restore a database, the server must use information from the volume history file and the device configuration file. You must make and save copies of the volume history file and the device configuration file. These files cannot be recreated.

You must use the VOLUMEHISTORY server option to specify one or more volume history files. Tivoli Storage Manager updates volume history files whenever server sequential volume history information is changed.

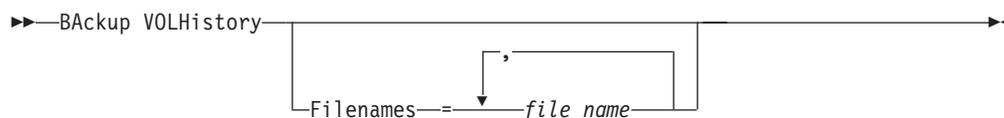
In order to ensure updates are complete before the server is halted, we recommend you:

- Not halt the server for a few minutes after issuing the BACKUP VOLHISTORY command.
- Specify multiple VOLUMEHISTORY options in the server options file.
- Examine the volume history file to see if the file has been updated.

Privilege class

Any administrator can issue this command unless it includes the FILENAMES parameter. If the FILENAMES parameter is specified and the REQSYSAUTHOUTFILE server option is set to YES, the administrator must have system privilege. If the FILENAMES parameter is specified and the REQSYSAUTHOUTFILE server option is set to NO, the administrator must have operator, policy, storage or system privilege.

Syntax



Parameters

FileNames

Specifies the names of one or more files in which to store a backup copy of volume history information. Separate multiple file names with commas and no intervening spaces. This parameter is optional.

If you do not specify a file name, Tivoli Storage Manager stores the information in all files specified with the VOLUMEHISTORY option in the server options file.

Example: Back up the volume history information to a file

Back up the volume history information to a file called VOLHIST.

```
backup volhistory filenames=volhist
```

Related commands

Table 28. Commands related to BACKUP VOLHISTORY

Command	Description
DELETE VOLHISTORY	Removes sequential volume history information from the volume history file.
DELETE VOLUME	Deletes a volume from a storage pool.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
UPDATE VOLHISTORY	Adds or changes location information for a volume in the volume history file.

BEGIN EVENTLOGGING (Begin logging events)


Use this command to begin logging events to one or more receivers. A receiver for which event logging has begun is an *active receiver*.

When the server is started, event logging automatically begins for the console and activity log and for any receivers that are started automatically based on entries in the server options file. You can use this command to begin logging events to receivers for which event logging is *not* automatically started at server startup. You can also use this command after you have disabled event logging to one or more receivers.

Privilege class

To issue this command, you must have system privilege.

Syntax



Notes:

- 1 This parameter is only available for the Windows operating system.

Parameters

Specify one or more receivers. You can specify multiple receivers by separating them with commas and no intervening spaces. If you specify ALL, logging begins for all receivers that are configured. The default is ALL.

ALL

Specifies all receivers that are configured for event logging.

CONSOLE

Specifies the server console as a receiver.

ACTLOG

Specifies the Tivoli Storage Manager activity log as a receiver.

EVENTSERVER

Specifies the event server as a receiver.

FILE

Specifies a user file as a receiver. Each logged event is a record in the file and a person cannot read each logged event easily.

FILETEXT

Specifies a user file as a receiver. Each logged event is a fixed-size, readable line.

SNMP

Specifies the simple network management protocol (SNMP) as a receiver.

TIVOLI

Specifies the Tivoli Management Environment (TME) as a receiver.

USEREXIT

Specifies a user-written routine to which Tivoli Storage Manager writes information as a receiver.

Example: Begin logging events

Begin logging events to the Tivoli Storage Manager activity log.

```
begin eventlogging actlog
```

Related commands

Table 29. Commands related to **BEGIN EVENTLOGGING**

Command	Description
DISABLE EVENTS	Disables specific events for receivers.
ENABLE EVENTS	Enables specific events for receivers.
END EVENTLOGGING	Ends event logging to a specified receiver.
QUERY ENABLED	Displays enabled or disabled events for a specific receiver.
QUERY EVENTRULES	Displays information about rules for server and client events.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

CANCEL commands

Use the **CANCEL** commands to end a task or process before it is completed.

- “CANCEL EXPIRATION (Cancel an expiration process)” on page 67
- “CANCEL EXPORT (Delete a suspended export operation)” on page 68
- “CANCEL PROCESS (Cancel an administrative process)” on page 69
- “CANCEL REPLICATION (Cancel node replication processes)” on page 71
- “CANCEL REQUEST (Cancel one or more mount requests)” on page 72
- “CANCEL RESTORE (Cancel a restartable restore session)” on page 73
- “CANCEL SESSION (Cancel one or more client sessions)” on page 74

CANCEL EXPIRATION (Cancel an expiration process)

Use this command to cancel a process that is running as a result of an inventory expiration operation.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—CANce1 EXPIration—◄◄

Example: Cancel an inventory expiration process

Cancel the process that was generated by an inventory expiration operation.

```
cancel expiration
```

Related commands

Table 30. Command related to CANCEL EXPIRATION

Command	Description
QUERY PROCESS	Displays information about background processes.
EXPIRE INVENTORY	Manually starts inventory expiration processing.

CANCEL PROCESS (Cancel an administrative process)

Use this command to cancel a background process started by an administrative command or by a process, such as storage pool migration.

The following commands generate background processes:

- AUDIT LIBRARY
- AUDIT LICENSES
- AUDIT VOLUME
- BACKUP DB
- BACKUP NODE
- BACKUP STGPOOL
- CHECKIN LIBVOLUME
- CHECKOUT LIBVOLUME
- DELETE FILESPACE
- DELETE VOLUME
- EXPIRE INVENTORY
- EXPORT ADMIN
- EXPORT NODE
- EXPORT POLICY
- EXPORT SERVER
- GENERATE BACKUPSET
- IMPORT ADMIN
- IMPORT NODE
- IMPORT POLICY
- IMPORT SERVER
- MIGRATE STGPOOL
- MOVE DATA
- MOVE DRMEDIA
- MOVE MEDIA
- PREPARE
- RECLAIM STGPOOL
- RESTORE NODE
- RESTORE STGPOOL
- RESTORE VOLUME
- VARY

The following internal server operations generate background processes:

- Inventory expiration
- Migration
- Reclamation

To cancel a process, you must have the process number, which you can obtain by issuing the **QUERY PROCESS** command.

Some processes, such as reclamation, will generate mount requests in order to complete processing. If a process has a pending mount request, the process may not respond to a **CANCEL PROCESS** command until the mount request has been answered or cancelled by using either the **REPLY** or **CANCEL REQUEST** command, or by timing out.

Issue the **QUERY REQUEST** command to list open requests, or query the activity log to determine if a given process has a pending mount request. A mount request indicates that a volume is needed for the current process, but the volume is not available in the library. It may not be available if the administrator has issued the **MOVE MEDIA** or **CHECKOUT LIBVOLUME** command, or manually removed the volume from the library.

After you issue a **CANCEL PROCESS** command for an export operation, the process cannot be restarted. To stop a server-to-server export operation but allow it to be restarted at a later time, issue the **SUSPEND EXPORT** command.

Privilege class

To issue this command, you must have system privilege.

CANCEL PROCESS

Syntax

▶▶—CANCEL PROCESS—*process_number*—▶▶

Parameters

process_number **(Required)**

Specifies the number of the background process you want to cancel.

Example: Cancel a background process using its process number

Cancel background process number 3.

```
cancel process 3
```

Related commands

Table 32. Commands related to CANCEL PROCESS

Command	Description
CANCEL EXPORT	Deletes a suspended export operation.
CANCEL REQUEST	Cancels pending volume mount requests.
QUERY EXPORT	Displays the export operations that are currently running or suspended.
QUERY PROCESS	Displays information about background processes.
REPLY	Allows a request to continue processing.
RESTART EXPORT	Restarts a suspended export operation.
SUSPEND EXPORT	Suspends a running export operation.

CANCEL REPLICATION (Cancel node replication processes)

Use this command to cancel all node replication processes.

Issue this command on the server that acts as a source for replicated data.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—CANCEL REPLICATION—◄◄

Parameters

None.

Example: Cancel node replication processes

Cancel all node replication processes.

```
cancel replication
```

Related commands

Table 33. Commands related to CANCEL REPLICATION

Command	Description
QUERY PROCESS	Displays information about background processes.
QUERY REPLICATION	Displays information about node replication processes.

CANCEL REQUEST

CANCEL REQUEST (Cancel one or more mount requests)

Use this command to cancel one or more pending media mount requests. To cancel a mount request, you need to know the request number assigned to the request. This number is included in the mount request message and can also be shown by using the **QUERY REQUEST** command.

Privilege class

To issue this command, you must have system privilege or operator privilege.

Syntax

```
►►—cANcel REQuest—request_number—ALl—PERManent—►►
```

Parameters

request_number

Specifies the request number of the mount request to cancel.

ALl

Specifies to cancel all pending mount requests.

PERManent

Specifies that you want the server to flag the volumes for which you are canceling a mount request as unavailable. This parameter is optional.

Example: Cancel a mount request

Cancel request number 2.

```
cancel request 2
```

Related commands

Table 34. Commands related to **CANCEL REQUEST**

Command	Description
QUERY REQUEST	Displays information about all pending mount requests.
UPDATE VOLUME	Updates the attributes of storage pool volumes.

CANCEL RESTORE (Cancel a restartable restore session)

Use this command to cancel a restartable restore session. You can cancel restore sessions in the active or restartable state. Any outstanding mount requests related to this session are automatically cancelled.

To display restartable restore sessions, use the **QUERY RESTORE** command.

Privilege class

To issue this command, you must have system or operator privilege.

Syntax

```
►►—CANce1—REStore—session_number—►►
                    |
                    |—AL1—
```

Parameters

session_number

Specifies the number for the restartable restore session. An active session is a positive number, and a restartable session is a negative number.

AL1

Specifies that all the restartable restore sessions are to be cancelled.

Example: Cancel restore operations

Cancel all restore operations.

```
cancel restore all
```

Related commands

Table 35. Commands related to CANCEL RESTORE

Command	Description
QUERY RESTORE	Displays information about restartable restore sessions.

CANCEL SESSION (Cancel one or more client sessions)

Use this command to cancel existing administrative or client node sessions, and to force an administrative or client node session off the server. Any outstanding mount requests related to this session are automatically cancelled. The client node must start a new session to resume activities.

If you cancel a session that is in the idle wait (IdleW) state, the client session is automatically reconnected to the server when it starts to send data again.

If this command interrupts a process, such as backup or archive, the results of any processing active at the time of interruption are rolled back and not committed to the database.

Privilege class

To issue this command, you must have system or operator privilege.

Syntax

```
▶▶ CANCEL SESSion [session_number] [AL1] ▶▶
```

Parameters

session_number
Specifies the number of the administrative, server, or client node sessions that you want to cancel.

AL1
Specifies that all client node sessions are cancelled. You cannot use this parameter to cancel administrative client or server sessions.

Example: Cancel a specific client node session

Cancel the client node session with NODEP (session 3).

```
cancel session 3
```

Related commands

Table 36. Commands related to CANCEL SESSION

Command	Description
DISABLE SESSIONS	Prevents new sessions from accessing Tivoli Storage Manager but permits existing sessions to continue.
LOCK ADMIN	Prevents an administrator from accessing Tivoli Storage Manager.
LOCK NODE	Prevents a client from accessing the server.
QUERY SESSION	Displays information about all active administrator and client sessions with Tivoli Storage Manager.

CHECKIN LIBVOLUME (Check a storage volume into a library)

Use this command to add a sequential access storage volume or a cleaning tape to the server inventory for an automated library. The server cannot use a volume that physically resides in an automated library until that volume has been checked in.

Important:

1. The **CHECKIN LIBVOLUME** command processing does not wait for a drive to become available, even if the drive is only in the IDLE state. If necessary, you can make a library drive available issuing the **DISMOUNT VOLUME** command to dismount the volume. After a library drive is available, reissue the **CHECKIN LIBVOLUME** command.
2. You do not define the drives, check in media, or label the volumes in an external library. The server provides an interface that external media management systems use to operate with the server. For more information, refer to the *Administrator's Guide*.
3. When checking in WORM tapes other than 3592, you must use **CHECKLABEL=YES** or they will be checked in as normal read-write tapes.

This command creates a background process that you can cancel with the **CANCEL PROCESS** command. To display information on background processes, use the **QUERY PROCESS** command.


For detailed and current drive and library support information, see the Supported Devices Web site for your operating system:

http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html


Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.


Syntax for SCSI libraries




CHECKIN LIBVOLUME




A (SEARCH=Yes, SEARCH=Bulk):




Syntax for 349X libraries

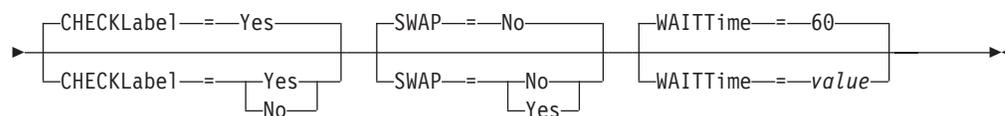
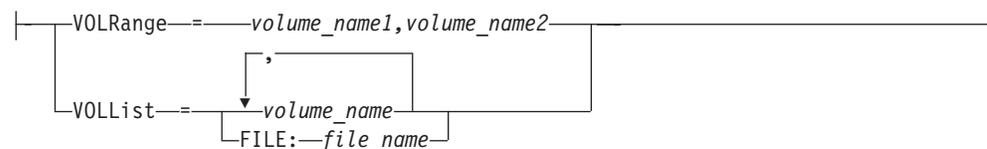


A (SEARCH=Yes):



Syntax for ACSLS libraries



**A (SEARCH=Yes):****Parameters*****library_name* (Required)**

Specifies the name of the library.

volume_name

Specifies the volume name of the storage volume being checked in. This parameter is required if SEARCH equals NO. Do not enter this parameter if the SEARCH parameter equals YES or BULK. If you are checking a volume into a SCSI library with multiple entry/exit ports, the volume in the lowest numbered slot will be checked in.

STATUS (Required)

Specifies the volume status. Possible values are:

PRivate

Specifies that the volume is a private volume that is mounted only when it is requested by name.

SCRatch

Specifies that the volume is a new scratch volume. This volume can be mounted to satisfy scratch mount requests during either data storage operations or export operations.

If a volume has an entry in volume history, you cannot check it in as a scratch volume.

CLEaner

Specifies that the volume is a cleaner cartridge and not a data cartridge. The CLEANINGS parameter is required for a cleaner cartridge and must be set to the number of cleaner uses.

CHECKLABEL=YES is not valid for checking in a cleaner cartridge. Use STATUS=CLEANER to check in a cleaner cartridge separately from a data cartridge.

OWner

Specifies which library client owns a private volume in a library that is shared across a SAN. The volume for which you specify ownership must be a private volume. You cannot specify ownership for a scratch volume. Furthermore, you cannot specify an owner when you use SEARCH=YES or SEARCH=BULK.

When you issue the **CHECKIN LIBVOLUME** command, the Tivoli Storage Manager server validates the owner. If you did not specify this parameter, then the server uses the default and delegates volume ownership to the owning library

CHECKIN LIBVOLUME

client, as recorded in the volume history file on the library manager. If the volume is not owned by any library client, then the server delegates ownership to the library manager.

SEARCH

Specifies whether the server searches the library to find volumes that were not checked in. This parameter is optional. The default is NO.

Possible values are:

No Specifies that only the named volume is checked into the library.

For SCSI libraries: The server issues a request to have the volume inserted into a cartridge slot in the library or, if available, into an entry port. The cartridge slot or entry port is identified by its element address. **For 349X libraries:** The volume could already be in the library, or you could put it into the I/O station when prompted.

Yes

Specifies that the server searches the library for volumes to be checked in. You can use the VOLRANGE or VOLLIST parameter to limit the search. When using this parameter, consider the following restrictions:

- If the library is shared between applications, the server could examine a volume required by another application. For 349X libraries, the server queries the library manager to determine all volumes that are assigned to the SCRATCH or PRIVATE category and to the INSERT category.
- For SCSI libraries, do not specify both SEARCH=YES and CHECKLABEL=NO in the same command.

Bulk

Specifies that the server searches the library's entry/exit ports for volumes that can be checked in automatically. This option only applies to SCSI libraries.

Important:

1. Do not specify both CHECKLABEL=NO and SEARCH=BULK.
2. You can use the VOLRANGE or VOLLIST parameter to limit the search.

VOLRange

Specifies a range of volume names separated by commas. You can use this parameter to limit the search for volumes to be checked in when you specify SEARCH=YES (349X, ACSLS, and SCSI libraries) or SEARCH=BULK (SCSI libraries only). If there are no volumes in the library that are within the specified range, the command completes without errors.

Specify only volume names that can be numerically incremented. In addition to the incremental area, a volume name can include an alphanumeric prefix and an alphanumeric suffix, for example:

Parameter	Description
volrange=bar110,bar130	The 21 volumes are checked in: bar110, bar111, bar112,...bar129, bar130.
volrange=bar11a,bar13a	The 3 volumes are checked in: bar11a, bar12a, bar13a.
volrange=123400,123410	The 11 volumes are checked in: 123400, 123401, ...123409, 123410.

VOLList

Specifies a list of volumes. You can use this parameter to limit the search for volumes to be checked in when you specify SEARCH=YES (349X, ACSLS, and SCSI libraries) or SEARCH=BULK (SCSI libraries only). If there are no volumes in the library that are in the list, the command completes without errors.

Possible values are:

volume_name

Specifies one or more volumes names that are separated by commas and no intervening spaces. For example: VOLLIST=TAPE01,TAPE02.

FILE: *file_name*

Specifies the name of a file that contains a list of volumes for the command. In the file, each volume name must be on a separate line. Blank lines and comment lines that begin with an asterisk are ignored. For example, to use volumes TAPE01, TAPE02 and TAPE03, create a file, TAPEVOL, that contains these lines:

```
TAPE01
TAPE02
TAPE03
```

You can specify the volumes for the command as follows:
VOLLIST=FILE:TAPEVOL.

Attention: The file name is case-sensitive.

CHECKLabel

Specifies how or whether the server should read sequential media labels of volumes. This parameter is optional. The default is YES.

Possible values are:

Yes

Specifies that an attempt is made to read the media label during check-in.

Attention:

1. For optical volumes (write-once and rewritable), you must specify YES.
2. For SCSI libraries, do not specify both SEARCH=YES and CHECKLABEL=NO in the same command.
3. For WORM media other than 3592, you must specify YES.

No Specifies that the media label is not read during check-in. However, suppressing label checking can result in future errors (for example, either a wrong label or an improperly labeled volume can cause an error). For 349X and ACSLS libraries, specify NO to avoid loading cartridges into a drive to read the media label. These libraries always return the external label information on cartridges, and Tivoli Storage Manager uses that information.

Barcode

Specifies that the server reads the bar code label if the library has a bar code reader and the volumes have external bar code labels. You can decrease the check-in time by using the bar code. This parameter applies only to SCSI libraries.

If the bar code reader cannot read the bar code label, or if the tape does not have a bar code label, the server mounts the tape and reads the internal label.

CHECKIN LIBVOLUME

DEVType

Specifies the device type for the volume being checked in. This parameter is required if none of the drives in this library have defined paths.

3590

Specifies that the device type for the volume being checked in is 3590.

3592

Specifies that the device type for the volume being checked in is 3592.

SWAP

Specifies whether the server swaps volumes if an empty library slot is not available. The volume selected for the swap operation (target swap volume) is ejected from the library and replaced with the volume being checked in. The server identifies a target swap volume by checking for an available scratch volume. If none exists, the server identifies the least frequently mounted volume.

This parameter is optional. The default is NO. This parameter only applies if there is a volume name specified in the command. Possible values are:

No Specifies that the server checks in the volume only if an empty slot is available.

Yes

Specifies that if an empty slot is not available, the server swaps cartridges to check in the volume.

WAITTime

Specifies the number of minutes that the server waits for you to reply or respond to a request. Specify a value in the range 0-9999. If you want to be prompted by the server, specify a wait time greater than zero. The default value is 60 minutes. For example, suppose the server prompts you to insert a tape into the entry/exit port of a library. If you specified a wait time of 60 minutes, the server issues a request and waits 60 minutes for you to reply. Suppose, on the other hand, you specify a wait time of 0. If you have already inserted a tape, a wait time of zero causes the operation to continue without prompting. If you have *not* inserted a tape, a wait time of zero will cause the operation to fail.

CLEAnings

Enter the recommended value for the individual cleaner cartridge (usually indicated on the cartridge). Cleanings apply only to SCSI libraries. This parameter is required if STATUS=CLEANER.

If more than one cleaner is checked into the library, only one is used until its CLEANINGS value decreases to zero. Another cleaner is then be selected, and the first cleaner can be checked out and discarded.

Example: Check a volume into a SCSI library

Check in a volume named WPDV00 into the SCSI library named AUTO.

```
checkin libvolume auto wpdv00 status=scratch
```

Example: Use a bar code reader to scan a library for a cleaner cartridge

Scan a SCSI library named AUTOLIB1 and, using the bar code reader, look for cleaner cartridge CLNV. Use SEARCH=YES, but limit the search by using the VOLLIST parameter.

```
checkin libvolume autolib1 search=yes vollist=cleanv status=cleaner
  cleanings=10 checklabel=barcode
```

Example: Scan a library to put unused volumes in a specific range in scratch status

Scan a 349X library named ABC, and limit the search to a range of unused volumes BAR110 to BAR130 and put them in scratch status.

```
checkin libvolume abc search=yes volrange=bar110,bar130
  status=scratch
```

Example: Scan a library to put a specific volume in scratch status

Use the barcode reader to scan a SCSI library named MYLIB for VOL1, and put it in scratch status.

```
checkin libvolume mylib search=yes vollist=vol1 status=scratch
  checklabel=barcode
```

Related commands

Table 37. Commands related to CHECKIN LIBVOLUME

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
CANCEL PROCESS	Cancels a background server process.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DISMOUNT VOLUME	Dismounts a sequential, removable volume by the volume name.
LABEL LIBVOLUME	Labels volumes in manual or automated libraries.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY LIBVOLUME	Displays information about a library volume.
QUERY PROCESS	Displays information about background processes.
REPLY	Allows a request to continue processing.

CHECKIN LIBVOLUME

Table 37. Commands related to CHECKIN LIBVOLUME (continued)

Command	Description
UPDATE LIBVOLUME	Changes the status of a storage volume.

CHECKOUT LIBVOLUME (Check a storage volume out of a library)

Use this command to remove a sequential access storage volume from the server inventory for an automated library. This command creates a background process that can be canceled with the **CANCEL PROCESS** command. To display information on background processes, use the **QUERY PROCESS** command.

Restrictions:

1. Check out processing does not wait for a drive to become available, even if the drive is in the IDLE state. If necessary, you can make a library drive available by dismounting the volume with the **DISMOUNT VOLUME** command. After a drive is available, the **CHECKOUT LIBVOLUME** command can be reissued.
2. Before checking out volumes from a 349X library, ensure that the 349x Cartridge Input and Output facility has enough empty slots for the volumes to be checked out. The 3494 Library Manager does not inform an application that the Cartridge Input and Output facility is full. It accepts requests to eject a cartridge and waits until the Cartridge Input and Output facility is emptied before returning to the server. Tivoli Storage Manager may appear to be hung when it is not. You should check the library and clear any intervention requests.
3. Before checking volumes out of an ACSLS library, ensure that the CAP priority in ACSLS is greater than zero. If the CAP priority is zero, then you must specify a value for the CAP parameter on the **CHECKOUT LIBVOLUME** command.


For detailed and current drive and library support information, see the Supported Devices Web site for your operating system:

http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html


Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax for SCSI library




A :

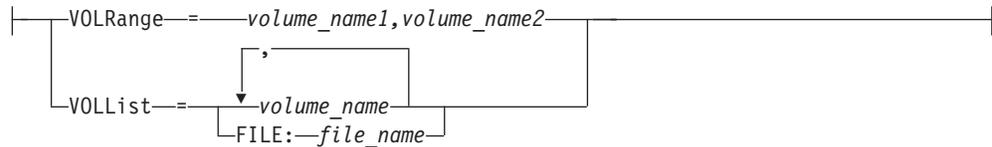


CHECKOUT LIBVOLUME

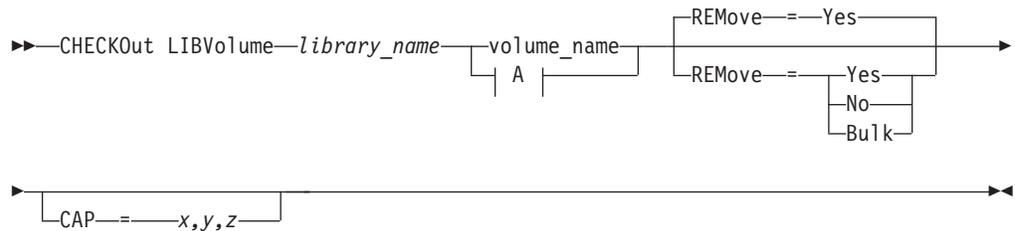
Syntax for 349X library



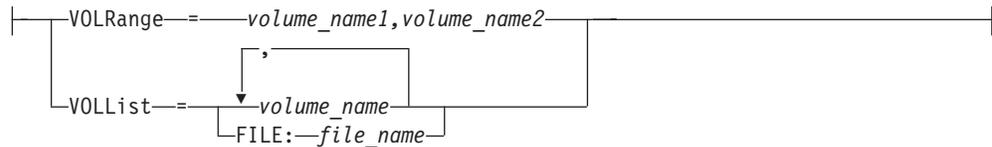
A :



Syntax for ACSLS library



A :



Parameters

library_name (Required)

Specifies the name of the library.

volume_name

Specifies the volume name.

VOLRange

Specifies two volume names separated by a comma. This parameter is a range of volumes to be checked out. If there are no volumes in the library that are within the specified range, the command completes without errors.

Specify only volume names that can be numerically incremented. In addition to the incremental area, a volume name can include an alphanumeric prefix and an alphanumeric suffix, for example:

Parameter	Description
volrange=bar110,bar130	The 21 volumes are checked out: bar110, bar111, bar112,...bar129, bar130.
volrange=bar11a,bar13a	The 3 volumes are checked out: bar11a, bar12a, bar13a.
volrange=123400,123410	The 11 volumes are checked out: 123400, 123401, ...123409, 123410.

VOLList

Specifies a list of volumes to check out. If there are no volumes in the library that are in the list, the command completes without errors.

Possible values are:

volume_name

Specifies the names of one or more values used for the command.

Example: VOLLIST=TAPE01,TAPE02.

FILE:*file_name*

Specifies the name of a file that contains a list of volumes for the command. In the file, each volume name must be on a separate line. Blank lines and comment lines that begin with an asterisk are ignored. For example, to use volumes TAPE01, TAPE02 and TAPE03, create a file, TAPEVOL, that contains these lines:

```
TAPE01
TAPE02
TAPE03
```

You can specify the volumes for the command as follows:

VOLLIST=FILE:TAPEVOL.

Attention: The file name is case-sensitive.

REMove

Specifies that the server tries to move the volume out of the library and into the convenience I/O station or entry/exit ports. This parameter is optional. Possible values, depending on the type of library, are YES, BULK, and NO. The response of the server to each of those options and the default values are described in the following sections.

349X libraries: The default is BULK. The following table shows how the server responds for 349X libraries.

Table 38. How the Tivoli Storage Manager server responds for 349X libraries

REMOVE=YES	REMOVE=BULK	REMOVE=NO
The 3494 Library Manager ejects the cartridge to the convenience I/O station.	The 3494 Library Manager ejects the cartridge to the high-capacity output facility.	The 3494 Library Manager does not eject the volume. The server leaves the cartridge in the library in the INSERT category for use by other applications.

SCSI libraries: The default is BULK. The following table shows how the server responds for a SCSI libraries.

CHECKOUT LIBVOLUME

Table 39. How the Tivoli Storage Manager server responds for SCSI libraries

	And REMOVE=YES, then...	And REMOVE=BULK, then...	And REMOVE=NO, then...
<i>Does not have entry/exit ports</i>	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server then prompts you to remove the cartridge from the slot and to issue a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.
<i>Has entry/exit ports and an entry/exit port is available</i>	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server then prompts you to remove the cartridge from the slot and to issue a REPLY command.	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server does not prompt you to remove the cartridge and does not request a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.
<i>Has entry/exit ports, but no ports are available</i>	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server then prompts you to remove the cartridge from the slot and to issue a REPLY command.	The server waits for an entry/exit port to be made available.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.

ACSLs libraries: The default is YES. If the parameter is set to YES, and the cartridge access port (CAP) has an automatic selection priority value of 0, you must specify a CAP ID. The following table shows how the server responds for ACSLS libraries.

Table 40. How the Tivoli Storage Manager server responds for ACSLS libraries

REMOVE=YES or REMOVE=BULK	REMOVE=NO
The server ejects the cartridge to the convenience I/O station, and deletes the volume entry from the server library inventory.	The server does not eject the cartridge. The server deletes the volume entry from the server library inventory and leaves the volume in the library.

CHECKLabel

Specifies how or whether the server reads sequential media labels of volumes.

Attention: This parameter does not apply to IBM 349X or ACSLS libraries.

This parameter is optional. The default is YES. Possible values are:

Yes

Specifies that the server attempts to read the media label to verify that the correct volume is being checked out.

No Specifies that during checkout the media label is not read. This improves performance because the read process does not occur.

FORCE

Specifies whether the server checks out a volume if an input/output (I/O) error occurs when reading the label.

Attention: This parameter does not apply to IBM 349X or ACSLS libraries.

This parameter is optional. The default is NO. Possible values are:

No The server does not check out a storage volume if an I/O error occurs when reading the label.

Yes

The server checks out the storage volume even if an I/O error occurs.

CAP

Specifies which cartridge access port (CAP) to use for ejecting volumes if you specify REMOVE=YES. This parameter applies to volumes in ACSLS libraries only. If the CAP priority value is set to 0 in the library, this parameter is required. If a CAP priority value greater than 0 is set in the library, this parameter is optional. By default, all CAPs initially have a priority value of 0, which means that ACSLS does not automatically select the CAP.

To display valid CAP identifiers (x,y,z), issue the **QUERY CAP** command with ALL specified from the Automated Cartridge System System Administrator (ACSSA) console on the ACSLS server host. The identifiers are as follows:

- x The Automated Cartridge System (ACS) ID. This identifier can be a number in the range 0 - 126.
- y The Library Storage Module (LSM) ID. This identifier can be a number in the range 0 - 23.
- z The CAP ID. This identifier can be a number in the range 0 - 11.

For more information, see the StorageTek documentation.

Example: Check out a volume and check the label

Check out the volume named EXB004 from the library named FOREST. Read the label to verify the volume name, but do not move the volume out of the library.

```
checkout libvolume forest exb004 checklabel=yes remove=no
```

Related commands

Table 41. Commands related to CHECKOUT LIBVOLUME

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
CANCEL PROCESS	Cancels a background server process.
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.

CHECKOUT LIBVOLUME

Table 41. Commands related to CHECKOUT LIBVOLUME (continued)

Command	Description
LABEL LIBVOLUME	Labels volumes in manual or automated libraries.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY LIBVOLUME	Displays information about a library volume.
QUERY PROCESS	Displays information about background processes.
REPLY	Allows a request to continue processing.
UPDATE LIBVOLUME	Changes the status of a storage volume.

CLEAN DRIVE (Clean a drive)

Use this command when you want Tivoli Storage Manager to immediately load a cleaner cartridge into a drive regardless of the cleaning frequency.

There are special considerations if you plan to use this command with a SCSI library that provides automatic drive cleaning through its device hardware. See the *Administrator's Guide* for details.

Restriction: You cannot run the **CLEAN DRIVE** command for a drive whose only path source is a NAS file server.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
▶▶—CLEAN DRIVE—library_name—drive_name—▶▶
```

Parameters

library_name **(Required)**

Specifies the name of the library to which the drive is assigned.

drive_name **(Required)**

Specifies the name of the drive.

Example: Clean a specific tape drive

You have already defined a library named AUTOLIB by using the **DEFINE LIBRARY** command, and you have already checked a cleaner cartridge into the library using the **CHECKIN LIBVOL** command. Inform the server that TAPEDRIVE3 in this library requires cleaning.

```
clean drive autolib tapedrive3
```

Related commands

Table 42. Commands related to **CLEAN DRIVE**

Command	Description
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE DRIVE	Deletes a drive from a library.
QUERY DRIVE	Displays information about drives.
UPDATE DRIVE	Changes the attributes of a drive.

COMMIT (Control committing of commands in a macro)

Use this command to control when a command is committed in a macro and to update the database when commands complete processing. When issued from the console mode of the administrative client, this command does not generate a message.

If an error occurs while processing the commands in a macro, the server stops processing the macro and rolls back any changes (since the last COMMIT). After a command is committed, it cannot be rolled back.

Ensure that your administrative client session is not running with the ITEMCOMMIT option if you want to control command processing. The ITEMCOMMIT option commits commands inside a script or a macro as *each* command is processed.

Privilege class

Any administrator can issue this command.

Syntax

▶—COMMIT—◀

Parameters

None.

Example: Control committing of commands in a macro

From the interactive mode of the administrative client, register and grant authority to new administrators using a macro named REG.ADM. Changes are committed after each administrator is registered and is granted authority.

Macro Contents:

```
/* REG.ADM-register policy admin & grant authority*/
REGister Admin sara hobby
GRant AUTHority sara Classes=Policy
COMMIT /* Commits changes */
REGister Admin ken plane
GRant AUTHority ken Classes=Policy
COMMIT /* Commits changes */
```

Command

```
macro reg.adm
```

Related commands

Table 43. Commands related to COMMIT

Command	Description
MACRO	Runs a specified macro file.
ROLLBACK	Discards any uncommitted changes to the database since the last COMMIT was executed.

COPY commands

Use the **COPY** commands to create a copy of Tivoli Storage Manager objects or data.

- “COPY ACTIVE DATA (Copy active backup data from a primary storage pool to an active-data pool)” on page 92
- “COPY CLOPTSET (Copy a client option set)” on page 96
- “COPY DOMAIN (Copy a policy domain)” on page 97
- “COPY MGMTCLASS (Copy a management class)” on page 99
- “COPY POLICYSET (Copy a policy set)” on page 101
- “COPY PROFILE (Copy a profile)” on page 103
- “COPY SCHEDULE (Copy a client or an administrative command schedule)” on page 105
- “COPY SCRIPT (Copy a Tivoli Storage Manager script)” on page 109
- “COPY SERVERGROUP (Copy a server group)” on page 110

COPY ACTIVE DATA (Copy active backup data from a primary storage pool to an active-data pool)

Use this command to copy active versions of backup data from a primary storage pool to an active-data pool. The primary benefit of active-data pools is fast client restores. Copy your active data regularly to ensure that the data is protected in case of a disaster.

If a file already exists in the active-data pool, the file is not copied unless the copy of the file in the active-data pool is marked damaged. However, a new copy is not created if the file in the primary storage pool is also marked damaged. In a random-access storage pool, neither cached copies of migrated files nor damaged primary files are copied.

If migration for a storage pool starts while active data is being copied, some files might be migrated before they are copied. For this reason, you should copy active data from storage pools that are higher in the migration hierarchy before copying active data from storage pools that are lower. Be sure a copy process is complete before beginning another.


Remember:

- You can only copy active data from storage pools that have a data format of NATIVE or NONBLOCK.
- Issuing this command for a primary storage pool that is set up for data deduplication removes duplicate data, if the active-data pool is also set up for data deduplication.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the active-data pool from which active versions of backup data are being copied.

Syntax



Notes:

- 1 The **VOLUMESONLY** parameter applies to sequential-access storage pools only.

Parameters

primary_pool_name (Required)

Specifies the primary storage pool.

active_data_pool_name (Required)

Specifies the active-data pool.

MAXProcess

Specifies the maximum number of parallel processes to use for copying files. This parameter is optional. Enter a value from 1 to 999. The default is 1.

Using multiple, parallel processes may improve throughput for the **COPY ACTIVE DATA** command. The expectation is that the time needed to copy active data will be decreased by using multiple processes. However, when multiple processes are running, in some cases one or more of the processes might need to wait to use a volume that is already in use by a different **COPY ACTIVE DATA** process.

When determining this value, consider the number of logical and physical drives that can be dedicated to this operation. To access a sequential-access volume, Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager and system activity and on the mount limits of the device classes for the sequential-access storage pools that are involved when copying active data.

Each process needs a mount point for active-data pool volumes, and, if the device type is not FILE, each process also needs a drive. If you are copying active data from a sequential-access storage pool, each process needs an additional mount point for primary storage pool volumes and, if the device type is not FILE, an additional drive. For example, suppose you specify a maximum of 3 processes to copy a primary sequential storage pool to an active-data pool of the same device class. Each process requires two mount points and two drives. To run all three processes, the device class must have a mount limit of at least six, and at least six mount points and six drives must be available.

To use the **PREVIEW** parameter, only one process is used, and no mount points or drives are needed.

Preview

Specifies whether you want to preview but not actually copy any active data. The preview displays the number of files and bytes to be copied and a list of the primary storage pool volumes that you must mount. This parameter is optional. The default is NO. Possible values are:

No Specifies that active data will be copied.

Yes

Specifies that you want to preview the process but not copy any data.

VOLumesonly

Specifies that you want to preview the process only as a list of the volumes that must be mounted. This choice requires the least processing time.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is NO. Possible values are:

No Specifies that the server processes this command in the background.

COPY ACTIVEDATA

You can continue with other tasks while the command is being processed. Messages created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the **CANCEL PROCESS** command. If you cancel this process, some files may have already been copied prior to the cancellation.

Yes

Specifies that the server performs this operation in the foreground. You must wait for the operation to complete before continuing with other tasks. The server displays the output messages to the administrative client when the operation completes.

You cannot specify WAIT=YES from the server console.

SHREDTONoshred

Specifies whether data should be copied from a primary storage pool that enforces shredding to an active-data pool that does not enforce shredding. This parameter is optional. The default value is NO. Possible values are:

No Specifies that the server does not allow data to be copied from a primary storage pool that enforces shredding to an active-data pool that does not enforce shredding. If the primary storage pool enforces shredding and the active-data pool does not, the operation will fail.

Yes

Specifies that the server does allow data to be copied from a primary storage pool that enforces shredding to an active-data pool that does not enforce shredding. The data in the active-data pool will not be shredded when it is deleted.

Example: Copy primary storage pool data to active-data pool

Copy the active data from a primary storage pool named PRIMARY_POOL to the active-data pool named ACTIVEPOOL. Issue the command:

```
copy activedata primary_pool activepool
```

Related commands

Table 44. Commands related to COPY ACTIVEDATA

Command	Description
DEFINE DOMAIN	Defines a policy domain that clients can be assigned to.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
IMPORT NODE	Restores client node information from external media.
IMPORT SERVER	Restores all or part of the server from external media.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.

Table 44. Commands related to COPY ACTIVE DATA (continued)

Command	Description
QUERY CONTENT	Displays information about files in a storage pool volume.
QUERY DOMAIN	Displays information about policy domains.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY STGPOOL	Displays information about storage pools.
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.
UPDATE DOMAIN	Changes the attributes of a policy domain.
UPDATE STGPOOL	Changes the attributes of a storage pool.

COPY CLOPTSET (Copy a client option set)

Use this command to copy a client option set.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax

```
►►—COpy CLOptset—current_option_set_name—new_option_set_name—◄◄
```

Parameters

current_option_set_name **(Required)**

Specifies the name of the client option set to be copied.

new_option_set_name **(Required)**

Specifies the name of the new client option set. The maximum length of the name is 64 characters.

Example: Copy a client option set

Copy a client option set named ENG to a new client option set named ENG2.

```
copy cloptset eng eng2
```

Related commands

Table 45. Commands related to COPY CLOPTSET

Command	Description
DEFINE CLIENTOPT	Adds a client option to a client option set.
DEFINE CLOPTSET	Defines a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
DELETE CLOPTSET	Deletes a client option set.
QUERY CLOPTSET	Displays information about a client option set.
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.
UPDATE CLOPTSET	Updates the description of a client option set.

COPY DOMAIN (Copy a policy domain)

Use this command to create a copy of a policy domain.

The server copies the following information to the new domain:

- Policy domain description
- Policy sets in the policy domain (including the ACTIVE policy set, if a policy set has been activated)
- Management classes in each policy set (including the default management class, if assigned)
- Copy groups in each management class

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—COPY DOMAIN—current_domain_name—new_domain_name—◄◄
```

Parameters

current_domain_name **(Required)**

Specifies the policy domain to copy.

new_domain_name **(Required)**

Specifies the name of the new policy domain. The maximum length of this name is 30 characters.

Example: Copy a policy domain to a new policy domain

Copy the policy domain PROG1 to new policy domain PROG2.

```
copy domain prog1 prog2
```

Related commands

Table 46. Commands related to COPY DOMAIN

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE DOMAIN	Defines a policy domain that clients can be assigned to.
DEFINE MGMTCLASS	Defines a management class.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE COPYGROUP	Deletes a backup or archive copy group from a policy domain and policy set.
DELETE DOMAIN	Deletes a policy domain along with any policy objects in the policy domain.

COPY DOMAIN

Table 46. Commands related to COPY DOMAIN (continued)

Command	Description
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY DOMAIN	Displays information about policy domains.
QUERY MGMTCLASS	Displays information about management classes.
QUERY POLICYSET	Displays information about policy sets.
REGISTER NODE	Defines a client node to the server and sets options for that user.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.
UPDATE DOMAIN	Changes the attributes of a policy domain.
UPDATE MGMTCLASS	Changes the attributes of a management class.
UPDATE POLICYSET	Changes the description of a policy set.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

COPY MGMTCLASS (Copy a management class)

Use this command to create a copy of a management class within the same policy set.

The server copies the following information to the new management class:

- Management class description
- Copy groups defined to the management class
- Any attributes for managing files for Tivoli Storage Manager for Space Management clients

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the new management class belongs.

Syntax

```

▶▶—COpy MGmtclass—domain_name—policy_set_name—————▶▶
▶—current_class_name—new_class_name—————▶▶

```

Parameters

domain_name **(Required)**

Specifies the policy domain to which the management class belongs.

policy_set_name **(Required)**

Specifies the policy set to which the management class belongs.

current_class_name **(Required)**

Specifies the management class to copy.

new_class_name **(Required)**

Specifies the name of the new management class. The maximum length of this name is 30 characters.

Example: Copy a management class to a new management class

Copy the management class ACTIVEFILES to a new management class, FILEHISTORY. The management class is in policy set VACATION in the EMPLOYEE_RECORDS policy domain.

```
copy mgmtclass employee_records vacation
activefiles filehistory
```

Related commands

Table 47. Commands related to COPY MGMTCLASS

Command	Description
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
QUERY COPYGROUP	Displays the attributes of a copy group.

COPY MGMTCLASS

Table 47. Commands related to COPY MGMTCLASS (continued)

Command	Description
QUERY MGMTCLASS	Displays information about management classes.
QUERY POLICYSET	Displays information about policy sets.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.
UPDATE MGMTCLASS	Changes the attributes of a management class.

COPY POLICYSET (Copy a policy set)

Use this command to copy a policy set (including the ACTIVE policy set) within the same policy domain.

The server copies the following information to the new policy set:

- Policy set description
- Management classes in the policy set (including the default management class, if assigned)
- Copy groups in each management class

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the new policy set belongs.

Syntax

```
►►—COpy POLIcYset—domain_name—current_set_name—new_set_name—►►
```

Parameters

domain_name **(Required)**

Specifies the policy domain to which the policy set belongs.

current_set_name **(Required)**

Specifies the policy set to copy.

new_set_name **(Required)**

Specifies the name of the new policy set. The maximum length of this name is 30 characters.

Example: Copy a policy set to a new policy set

Copy the policy set VACATION to the new policy set HOLIDAY in the EMPLOYEE_RECORDS policy domain.

```
copy policyset employee_records vacation holiday
```

Related commands

Table 48. Commands related to COPY POLICYSET

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE MGMTCLASS	Defines a management class.
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.
QUERY POLICYSET	Displays information about policy sets.
UPDATE POLICYSET	Changes the description of a policy set.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

COPY POLICYSET

COPY PROFILE (Copy a profile)

Use this command on a configuration manager to copy a profile and all its associated object names to a new profile.

Privilege class

To issue this command, you must have system privilege.

Syntax

►► `COPY PROFILE` *current_profile_name* *new_profile_name* ◀◀

Parameters

current_profile_name **(Required)**

Specifies the profile to copy.

new_profile_name **(Required)**

Specifies the name of the new profile. The maximum length of the profile name is 30 characters.

Example: Make a copy of a profile

Copy a profile named VAL to a new profile named VAL2.

```
copy profile val val2
```

Related commands

Table 49. Commands related to COPY PROFILE

Command	Description
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE PROFILE	Deletes a profile from a configuration manager.
DELETE SUBSCRIBER	Deletes obsolete managed server subscriptions.
DELETE SUBSCRIPTION	Deletes a specified profile subscription.
LOCK PROFILE	Prevents distribution of a configuration profile.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY PROFILE	Displays information about configuration profiles.
QUERY SUBSCRIBER	Displays information about subscribers and their subscriptions to profiles.
QUERY SUBSCRIPTION	Displays information about profile subscriptions.

COPY PROFILE

Table 49. Commands related to COPY PROFILE (continued)

Command	Description
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

COPY SCHEDULE (Copy a client or an administrative command schedule)

Use this command to create a copy of a schedule.

The COPY SCHEDULE command takes two forms, depending on whether the schedule applies to client operations or administrative commands. The syntax and parameters for each form are defined separately.

- “COPY SCHEDULE (Create a copy of a schedule for client operations)” on page 106
- “COPY SCHEDULE (Create a copy of a schedule for administrative operations)” on page 108

Table 50. Commands related to COPY SCHEDULE

Command	Description
DEFINE ASSOCIATION	Associates clients with a schedule.
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.
DELETE SCHEDULE	Deletes a schedule from the database.
QUERY SCHEDULE	Displays information about schedules.
UPDATE SCHEDULE	Changes the attributes of a schedule.

COPY SCHEDULE


COPY SCHEDULE (Create a copy of a schedule for client operations)

Use the **COPY SCHEDULE** command to create a copy of a schedule for client operations. You can copy a schedule within a policy domain or from one policy domain to another policy domain. Use the **DEFINE ASSOCIATION** command to associate the new schedule with the client nodes.

Privilege class

To copy a client schedule, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which you are copying the schedule.

Syntax



Parameters

current_domain_name (Required)

Specifies the name of the policy domain that contains the schedule you want to copy.

current_sched_name (Required)

Specifies the name of the schedule you want to copy.

new_domain_name (Required)

Specifies the name of a policy domain to which you want to copy the new schedule.

new_sched_name

Specifies the name of the new schedule. You can specify up to 30 characters for the name.

If you do not specify this name, the name of the original schedule is used.

If the schedule name is already defined in the policy domain, you must specify REPLACE=YES, or the command fails.

REPlace

Specifies whether to replace a client schedule. The default is NO. The values are:

No Specifies that a client schedule is not replaced.

Yes

Specifies that a client schedule is replaced.

Example: Copy a schedule from one policy domain to another

Copy the WEEKLY_BACKUP schedule that belongs to policy domain EMPLOYEE_RECORDS to the PROG1 policy domain and name the new schedule WEEKLY_BACK2. If there is already a schedule with this name defined in the PROG1 policy domain, do not replace it.

```
copy schedule employee_records weekly_backup  
prog1 weekly_back2
```

COPY SCHEDULE

COPY SCHEDULE (Create a copy of a schedule for administrative operations)

Use the **COPY SCHEDULE** command to create a copy of an administrative command schedule.

Privilege class

To copy an administrative command schedule, you must have system privilege.

Syntax

```
►► COPY SCHEDULE current_sched_name new_sched_name Type=Administrative►►
```



Parameters

current_schedule_name (Required)

Specifies the name of the schedule you want to copy.

new_schedule_name (Required)

Specifies the name of the new schedule. You can specify up to 30 characters for the name.

If the schedule name is already defined, you must specify REPLACE=YES, or the command fails.

Type=Administrative

Specifies that an administrative command schedule is to be copied.

REPLACE

Specifies whether to replace an administrative command schedule. The default is NO. The values are:

No Specifies that an administrative command schedule is not replaced.

Yes

Specifies that an administrative command schedule is replaced.

Example: Copy an administrative command schedule to another schedule

Copy the administrative command schedule, DATA_BACKUP and name the schedule DATA_ENG. If there is already a schedule with this name, replace it.

```
copy schedule data_backup data_eng  
type=administrative replace=yes
```

COPY SCRIPT (Copy a Tivoli Storage Manager script)

Use this command to copy a Tivoli Storage Manager script to a new script.

Privilege class

To issue this command, you must have operator, policy, storage, or system privilege.

Syntax

```
►►—COPY SCRIPT—current_script_name—new_script_name—◄◄
```

Parameters

current_script_name (Required)

Specifies the name of the script you want to copy.

new_script_name (Required)

Specifies the name of the new script. You can specify up to 30 characters for the name.

Example: Make a copy of a script

Copy script TESTDEV to a new script and name it ENGDEV.

```
copy script testdev engdev
```

Related commands

Table 51. Commands related to COPY SCRIPT

Command	Description
DEFINE SCRIPT	Defines a script to the Tivoli Storage Manager server.
DELETE SCRIPT	Deletes the script or individual lines from the script.
QUERY SCRIPT	Displays information about scripts.
RENAME SCRIPT	Renames a script to a new name.
RUN	Runs a script.
UPDATE SCRIPT	Changes or adds lines to a script.

COPY SERVERGROUP (Copy a server group)

Use this command to create a copy of a server group.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—COPY SERVERGRoup—current_group_name—new_group_name—◄◄
```

Parameters

current_group_name **(Required)**

Specifies the server group to copy.

new_group_name **(Required)**

Specifies the name of the new server group. The maximum length of this name is 64 characters.

Example: Make a copy of a server group

Copy the server group GRP_PAYROLL to the new group HQ_PAYROLL.

```
copy servergroup grp_payroll hq_payroll
```

Related commands

Table 52. Commands related to COPY SERVERGROUP

Command	Description
DEFINE GRPMEMBER	Defines a server as a member of a server group.
DEFINE SERVER	Defines a server for server-to-server communications.
DEFINE SERVERGROUP	Defines a new server group.
DELETE GRPMEMBER	Deletes a server from a server group.
DELETE SERVER	Deletes the definition of a server.
DELETE SERVERGROUP	Deletes a server group.
MOVE GRPMEMBER	Moves a server group member.
QUERY SERVER	Displays information about servers.
QUERY SERVERGROUP	Displays information about server groups.
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVER	Updates information about a server.
UPDATE SERVERGROUP	Updates a server group.

DEFINE commands

Use the **DEFINE** commands to create Tivoli Storage Manager objects.

- “DEFINE ALERTTRIGGER (Define an alert trigger)” on page 112
- “DEFINE ASSOCIATION (Associate client nodes with a schedule)” on page 115
- “DEFINE BACKUPSET (Define a backup set)” on page 117
- “DEFINE CLIENTACTION (Define a one-time client action)” on page 121
- “DEFINE CLIENTOPT (Define an option to an option set)” on page 127
- “DEFINE CLOPTSET (Define a client option set name)” on page 130
- “DEFINE COLLOGROUP (Define a collocation group)” on page 131
- “DEFINE COLLOCMEMBER (Define collocation group member)” on page 133
- “DEFINE COPYGROUP (Define a copy group)” on page 135
- “DEFINE DATAMOVER (Define a data mover)” on page 145
- “DEFINE DEVCLASS (Define a device class)” on page 148
- “DEFINE DOMAIN (Define a new policy domain)” on page 239
- “DEFINE DRIVE (Define a drive to a library)” on page 241
- “DEFINE EVENTSERVER (Define a server as the event server)” on page 245
- “DEFINE GRPMEMBER (Add a server to a server group)” on page 246
- “DEFINE LIBRARY (Define a library)” on page 247
- “DEFINE MACHINE (Define machine information for disaster recovery)” on page 267
- “DEFINE MACHNODEASSOCIATION (Associate a node with a machine)” on page 269
- “DEFINE MGMTCLASS (Define a management class)” on page 271
- “DEFINE NODEGROUP (Define a node group)” on page 274
- “DEFINE NODEGROUPMEMBER (Define node group member)” on page 275
- “DEFINE PATH (Define a path)” on page 276
- “DEFINE POLICYSET (Define a policy set)” on page 286
- “DEFINE PROFASSOCIATION (Define a profile association)” on page 288
- “DEFINE PROFILE (Define a profile)” on page 294
- “DEFINE RECMEDMACHASSOCIATION (Associate recovery media with a machine)” on page 296
- “DEFINE RECOVERYMEDIA (Define recovery media)” on page 298
- “DEFINE SCHEDULE (Define a client or an administrative command schedule)” on page 300
- “DEFINE SCRIPT (Define a Tivoli Storage Manager script)” on page 323
- “DEFINE SERVER (Define a server for server-to-server communications)” on page 326
- “DEFINE SERVERGROUP (Define a server group)” on page 333
- “DEFINE SPACETRIGGER (Define the space trigger)” on page 334
- “DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)” on page 337
- “DEFINE STGPOOL (Define a storage pool)” on page 341
- “DEFINE SUBSCRIPTION (Define a profile subscription)” on page 387
- “DEFINE VIRTUALFSMAPPING (Define a virtual file space mapping)” on page 389
- “DEFINE VOLUME (Define a volume in a storage pool)” on page 391

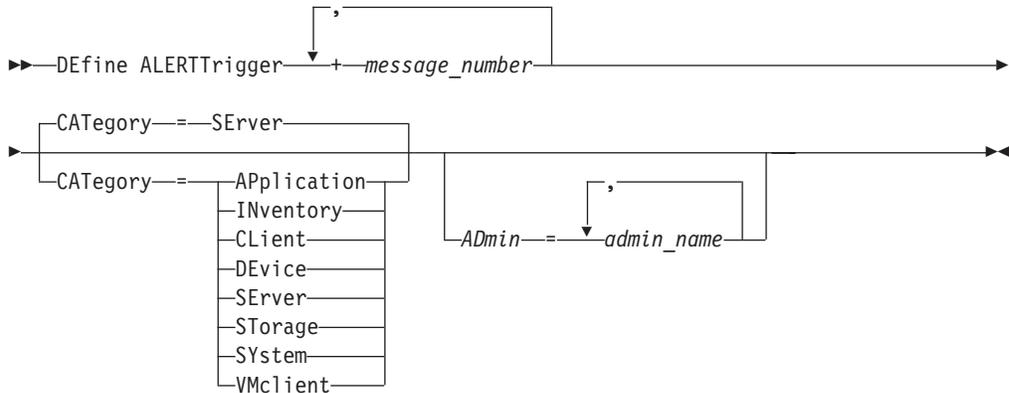
DEFINE ALERTTRIGGER (Define an alert trigger)

Use this command to trigger an alert whenever a server issues a specific error message. You can define a message number to be an alert trigger, assign it to a category, or specify administrators who can be notified of the alert by email.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

message_number (Required)

Specifies the message number that you want to associate with the alert trigger. Specify multiple message numbers, which are separated by commas, and no intervening spaces. Message numbers are a maximum of eight characters in length. Wildcard characters can be used to specify message numbers.

CATegory

Specifies the category type for the alert, which is determined by the message types. The default value is SERVER. Specify one of the following values:

APplication

Alert is classified as application category. For example, you can specify this category for messages that are associated with application (TDP) clients.

CATalog

Alert is classified as catalog category. For example, you can specify this category for messages that are associated with the database, active log file, or archive log file.

Note: Changing the category of an alert trigger does not change the category of existing alerts on the server. New alerts are categorized with the new category.

CLient

Alert is classified as client category. For example, you can specify this category for messages that are associated with general client activities.

DEvice

Alert is classified as device category. For example, you can specify this category for messages that are associated with device classes, libraries, drives, or paths.

SErver

Alert is classified as general server category. For example, you can specify this category for messages that are associated with general server activities or events.

STorage

Alert is classified as storage category. For example, you can specify this category for messages that are associated with storage pools.

SYstems

Alert is classified under system clients category. For example, you can specify this category for messages that are associated with system backup and archive or hierarchical storage management (HSM) backup-archive clients.

VMclient

Alert is classified under VMclient category. For example, you can specify this category for messages that are associated with virtual machine clients.

ADmin

This optional parameter specifies the name of the administrator who receives email notification of this alert. The alert trigger is defined successfully even if no administrator names are specified.

Assign two message numbers to an alert

Issue the following command to specify that you want two message numbers to trigger an alert:

```
define alertrigger ANR1067E,ANR1073E
```

Assign a message number to an alert and email two administrators

Issue the following command to specify the message numbers that you want to trigger an alert and have them sent by email to two administrators:

```
define alertrigger ANR1067E,ANR1073E ADmin=BILL,DJADMIN
```

Related commands

Table 53. Commands related to DEFINE ALERTRIGGER

Command	Description
“DELETE ALERTRIGGER (Remove a message from an alert trigger)” on page 400	Removes a message number that can trigger an alert.
“QUERY ALERTSTATUS (Query the status of an alert)” on page 674	Displays information about alerts that have been issued on the server.
“QUERY ALERTRIGGER (Query the list of defined alert triggers)” on page 672	Displays message numbers that trigger an alert.
“QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)” on page 796	Displays information about monitoring alerts and server status settings.
“UPDATE ALERTRIGGER (Update a defined alert trigger)” on page 1165	Updates the attributes of one or more alert triggers.

DEFINE ALERTTRIGGER

Table 53. Commands related to **DEFINE ALERTTRIGGER** (continued)

Command	Description
"UPDATE ALERTSTATUS (Update the status of an alert)" on page 1168	Updates the status of a reported alert.

DEFINE ASSOCIATION (Associate client nodes with a schedule)

Use this command to associate one or more clients with a schedule. You must assign a client node to the policy domain to which a schedule belongs. Client nodes process operations according to the schedules associated with the nodes.

Note:

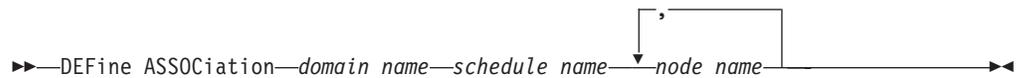
1. Tivoli Storage Manager cannot run multiple schedules concurrently for the same client node.
2. In a macro, the server may stall if some commands (such as **REGISTER NODE** and **DEFINE ASSOCIATION**) are not committed as soon as you issue them. You could follow each command in a macro with a **COMMIT** command. However, a simpler solution is to include the **-ITEMCOMMIT** option with the **DSMADMC** command.

Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege
- Restricted policy privilege for the policy domain to which the schedule belongs

Syntax



Parameters

domain_name (Required)

Specifies the name of the policy domain to which the schedule belongs.

schedule_name (Required)

Specifies the name of the schedule that you want to associate with one or more clients.

node_name (Required)

Specifies the name of a client node or a list of client nodes to associate with the specified schedule. Use commas to separate the items in the list. Do not leave spaces between the items and commas. You can use a wildcard character to specify a name. The command will not associate a listed client to the schedule if:

- The client is already associated with the specified schedule.
- The client is not assigned to the policy domain to which the schedule belongs.
- The client is a NAS node name. All NAS nodes are ignored.

Example: Associate client nodes with a schedule

Associate the client nodes SMITH or JOHN with the WEEKLY_BACKUP schedule. The associated clients are assigned to the EMPLOYEE_RECORDS policy domain.

```

define association employee_records
weekly_backup smith*,john*
    
```

DEFINE ASSOCIATION

Example: Associate client nodes with a schedule

Associate the client nodes JOE, TOM, and LARRY with the WINTER schedule. The associated clients are assigned to the EMPLOYEE_RECORDS policy domain; however, the client JOE is already associated with the WINTER schedule.

```
define association employee_records  
winter joe,tom,larry
```

Related commands

Table 54. Commands related to DEFINE ASSOCIATION

Command	Description
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.
DELETE ASSOCIATION	Deletes the association between clients and a schedule.
DELETE SCHEDULE	Deletes a schedule from the database.
QUERY ASSOCIATION	Displays the clients associated with one or more schedules.
REGISTER NODE	Defines a client node to the server and sets options for that user.

DEFINE BACKUPSET (Define a backup set)

Use this command to define a client backup set that was previously generated on one server and make it available to the server running this command. The client node has the option of restoring the backup set from the server running this command rather than the one on which the backup set was generated.

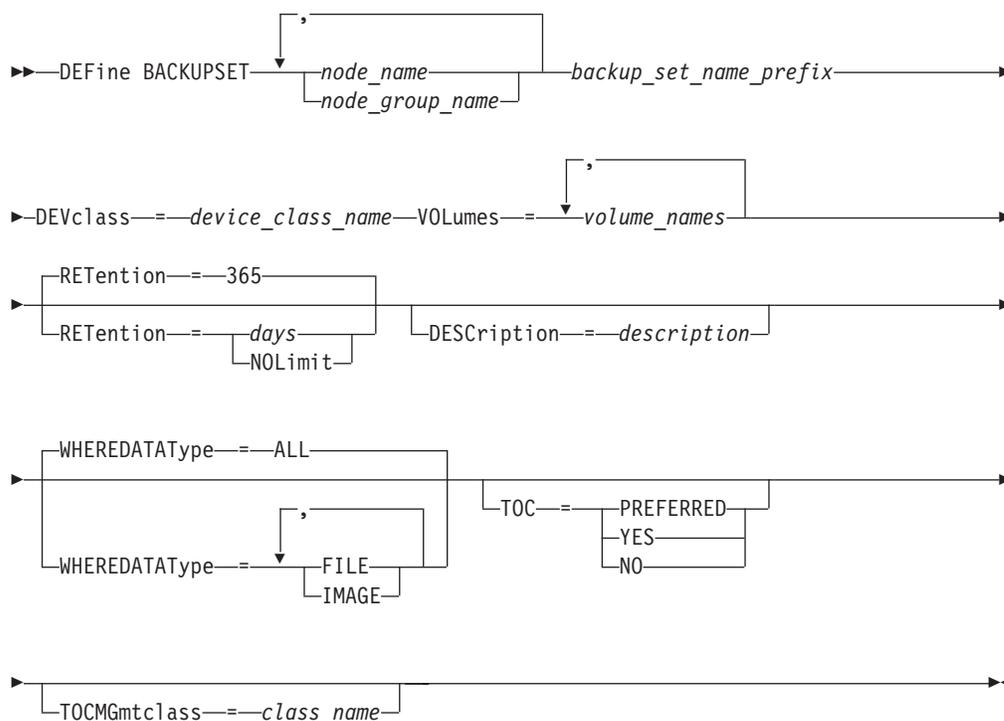
Any backup set generated on one server can be defined to another server as long as the servers share a common device type. The level of the server to which the backup set is being defined must be equal to or greater than the level of the server that generated the backup set.

You can also use the **DEFINE BACKUPSET** command to redefine a backup set that was deleted on a server.

Privilege class

If the REQSYSAUTHOUTFILE server option is set to YES (the default), the administrator must have system privilege. If the REQSYSAUTHOUTFILE server option is set to NO, the administrator must have system privilege or policy privilege for the domain to which the client node is assigned.

Syntax



Parameters

node_name or node_group_name (Required)

Specifies the name of the client nodes or node groups whose data is contained in the specified backup set volumes. To specify multiple node and node group names, separate the names with commas and no intervening spaces. Node

DEFINE BACKUPSET

names can contain wildcard characters, but node group names cannot. If the backup set volumes contain backup sets from multiple nodes, every backup set whose node name matches one of the specified node names will be defined. If the volumes contain a backup set for a node that is not currently registered, the **DEFINE BACKUPSET** command will not define the backup set for that node.

backup_set_name_prefix (Required)

Specifies the name of the backup set to define to this server. The maximum length of the name is 30 characters.

When you select a name, Tivoli Storage Manager adds a suffix to construct the backup set name. For example, if you name your backup set *mybackupset*, Tivoli Storage Manager adds a unique number such as 3099 to the name. Your backup set name is then identified as *mybackupset.3099*. To later display information about this backup set, you can include a wildcard with the name, such as *mybackupset** or you can specify the fully qualified name, such as *mybackupset.3099*.

If the backup set volumes contain backup sets for multiple nodes, then backup sets will be defined for each of the nodes using the same backup set name prefix and suffix.

DEVclass (Required)

Specifies the device class name for the volumes from which the backup set is read.

Note: The device type associated with the device class you specify must match the device class with which the backup set was originally generated.

VOLUMes (Required)

Specifies the names of the volumes used to store the backup set. You can specify multiple volumes by separating the names with commas and no intervening spaces. The volumes you specify must be available to the server that is defining the backup set.

Note: The volumes you specify must be listed in the order they were created, or the **DEFINE BACKUPSET** command will fail.

The server does not verify that every volume specified for a multiple-volume backup set actually contains part of the backup set. The first volume is always checked, and in some cases additional volumes are also checked. If these volumes are correct, the backup set is defined and all of the volumes listed in the command are protected from being overwritten. If a volume that contains part of the backup set is not listed in the command, the volume will not be protected and can potentially be overwritten during normal server operations.

Note: By default, the server attempts to create a table of contents when a backup set is defined. If an incorrect volume is specified, or if volumes are not listed in the correct order, the table of contents creation will fail. If this occurs, check the volume list in the command and consider using the **QUERY BACKUPSETCONTENTS** command to verify the contents of the backup set.

RETention

Specifies the number of days that the backup set is retained on the server. You can specify an integer from 0 to 30000. The default is 365 days. The values are:

days

Specifies the number of days to retain the backup set on the server.

NOLimit

Specifies that the backup set should be retained on the server indefinitely.

If you specify **NOLIMIT**, Tivoli Storage Manager retains the volumes containing the backup set forever, unless a user or administrator deletes the volumes from server storage.

DEscription

Specifies the description to associate with the backup set that belongs to the client node. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

WHERE DATType

Specifies the backup sets containing the specified types of data are to be defined. This parameter is optional. The default is that backup sets for all types of data (file level, image, and application) are to be defined. To specify multiple data types, separate the data types with commas and no intervening spaces. Possible values are:

ALL

Specifies that backup sets for all types of data (file level, image, and application) are to be defined. This is the default.

FILE

Specifies that a file level backup set is to be defined. File level backup sets contain files and directories backup up by the backup-archive client.

IMAGE

Specifies that an image backup set is to be defined. Image backup sets contain images created by the backup-archive client **BACKUP IMAGE** command.

TOC

Specifies whether a table of contents (TOC) should be created for the file level backup set when it is defined. The TOC parameter is ignored when defining image and application data backup sets because a table of contents is always created for these backup sets.

Consider the following in determining whether you want to create a table of contents:

- If a table of contents is created, you can use the Tivoli Storage Manager Web backup-archive client to examine the entire file system tree and choose files and directories to restore. Creation of a table of contents requires that you define the TOCDESTINATION attribute in the backup copy group for the management class specified by the **TOCMGMTCLASS** parameter. Note that a table of contents creation requires additional processing, storage pool space, and possibly a mount point during the backup set operation.
- If a table of contents is not saved for a backup set, you can still restore individual files or directory trees using the backup-archive client **RESTORE BACKUPSET** command, provided that you know the fully qualified name of each file or directory to be restored.

This parameter is optional. The default value is Preferred. Possible values are:

No Specifies that table of contents information is not saved for file level backup sets.

Preferred

Specifies that table of contents information should be saved for file level

DEFINE BACKUPSET

backup sets. However, a backup set does not fail just because an error occurs during creation of the table of contents.

Yes

Specifies that table of contents information must be saved for each file level backup set. A backup set fails if an error occurs during creation of the table of contents.

TOCMgmtclass

Specifies the name of the management class to which the table of contents should be bound. If you do not specify a management class, the table of contents is bound to the default management class for the policy domain to which the node is assigned. In this case, creation of a table of contents requires that you define the TOCDESTINATION attribute in the backup copy group for the specified management class.

Example: Define a backup set

Define the PERS_DATA backup set that belongs to client node JANE to the server running this command. Retain the backup set on the server for 50 days. Specify that volumes VOL001 and VOL002 contain the data for the backup set. The volumes are to be read by a device that is assigned to the AGADM device class. Include a description.

```
define backupset jane pers_data devclass=agadm
volumes=vol1,vol2 retention=50
description="sector 7 base image"
```

Related commands

Table 55. Commands related to **DEFINE BACKUPSET**

Command	Description
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE NODEGROUP	Deletes a node group.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
QUERY BACKUPSET	Displays backup sets.
QUERY BACKUPSETCONTENTS	Displays contents contained in backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

DEFINE CLIENTACTION (Define a one-time client action)

Use this command to schedule one or more clients to process a command for a one-time action.

The server automatically defines a schedule and associates the client node to the schedule. The server assigns the schedule priority 1, sets the PERUNITS to ONETIME, and determines the number of days to keep the schedule active. The number of days is based on the value set with the **SET CLIENTACTDURATION** command.

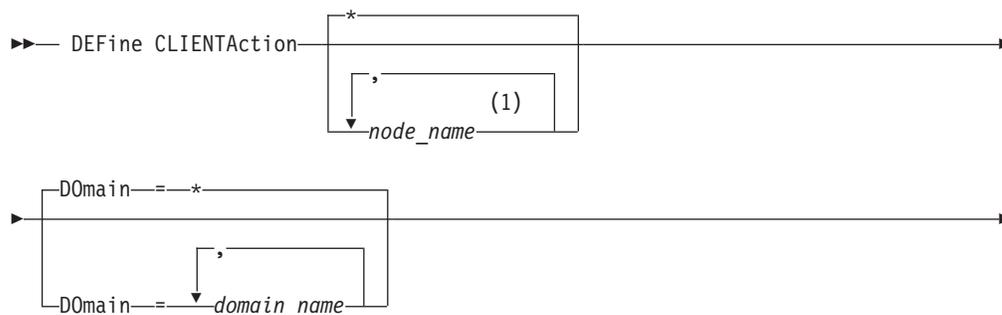
How quickly the client processes this command depends on whether the scheduling mode for the client is set to server-prompted or client-polling. The client scheduler must be started on the client workstation in order for the server to process the schedule.

Remember: The start of the Tivoli Storage Manager scheduler depends on the processing of other threads in the server and other processes on the Tivoli Storage Manager server host system. The amount of time it takes to start the scheduler also depends on network traffic and how long it takes to open a socket, to connect with the Tivoli Storage Manager client, and to receive a response from the client. In general, the greater the processing and connectivity requirements on the Tivoli Storage Manager server and client, the longer it can take to start the scheduler.

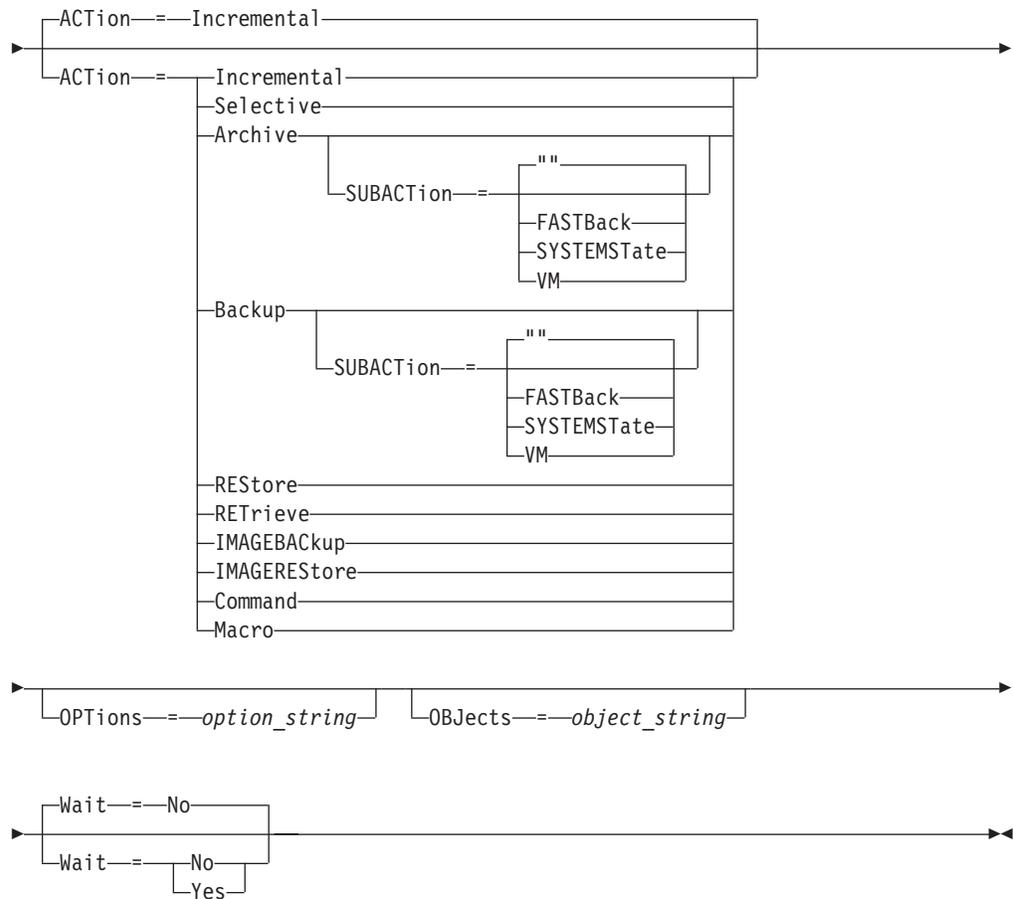
Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy for the policy domain to which the schedule belongs.

Syntax



DEFINE CLIENTACTION



Notes:

- 1 If you explicitly specify a NAS node name, the command will fail. If you provide a pattern-matching expression for the node, any NAS nodes that match the pattern will be ignored.

Parameters

node_name

Specifies the list of client nodes that will process the schedule associated with the action. You can use a wildcard character to specify a client node or a list of client nodes. Separate the client node names with commas and no intervening spaces. If you do not specify a value, all client nodes will be scheduled for the action.

Domain

Specifies the list of policy domains used to limit the list of client nodes. Only client nodes that are assigned to one of the specified policy domains will be scheduled. All clients assigned to a matching domain will be scheduled. Separate multiple domain names with commas and no intervening spaces. If you do not specify a value, all policy domains will be included in the list.

ACTION

Specifies the action that occurs when this schedule is processed. Possible values are:

Incremental

Specifies that the schedule backs up all files that are new or that have

changed since the last incremental backup. Incremental also backs up any file for which all existing backups might have expired.

Selective

Specifies that the schedule backs up only files that are specified with the OBJECTS parameter.

Archive

Specifies that the schedule archives files that are specified with the OBJECTS parameter.

Backup

Specifies that the schedule backs up files that are specified with the OBJECTS parameter.

REStore

Specifies that the schedule restores files that are specified with the OBJECTS parameter.

When you specify ACTION=RESTORE for a scheduled operation, and the REPLACE option is set to PROMPT, no prompting occurs. If you set the option to PROMPT, the files are skipped.

If you specify a second file specification, this second file specification acts as the restore destination. If you need to restore multiple groups of files, schedule one for each file specification that you need to restore.

REtrieve

Indicates that the schedule retrieves files that are specified with the OBJECTS parameter.

Remember: A second file that is specified acts as the retrieve destination. If you need to retrieve multiple groups of files, create a separate schedule for each group of files.

IMAGEBACKup

Specifies that the schedule backs up logical volumes that are specified with the OBJECTS parameter.

IMAGERESStore

Specifies that the schedule restores logical volumes that are specified with the OBJECTS parameter.

Command

Specifies that the schedule processes a client operating system command or script that is specified with the OBJECTS parameter.

Macro

Specifies that a client processes a macro whose file name is specified with the OBJECTS parameter.

SUBACTioN

Possible values are:

"" When a null string (two double quotes) is specified with ACTION=BACKUP the backup is an incremental.

FASTBBack

Specifies that a FastBack client operation that is identified by the ACTION parameter is to be scheduled for processing. The ACTION parameter must be either ARCHIVE or BACKUP.

DEFINE CLIENTACTION

SYSTEMState

Specifies that a client Systemstate backup is scheduled.

VM Specifies that a client VMware backup operation is scheduled.

OPTions

Specifies the client options that you specify to the scheduled command at the time the schedule is processed. This parameter is optional.

Only those options that are valid on the scheduled command can be specified for this parameter. Refer to the appropriate client manual for information about options that are valid from the command line. All options described there as valid only on the initial command line result in an error or are ignored when running the schedule from the server. For example, do not include the following options because they have no impact when the client processes the scheduled command:

MAXCMDRETRIES
OPTFILE
QUERYSCHEDPERIOD
RETRYPERIOD
SCHEDLOGNAME
SCHEDMODE
SERVERNAME
TCPCLIENTADDRESS
TCPCLIENTPORT

If the option string contains multiple options or options with embedded spaces, surround the entire option string with one pair of apostrophes. Enclose individual options that contain spaces in quotation marks. A leading minus sign is required in front of the option. Errors can occur if the option string contains spaces that are not quoted correctly.

The following examples show how to specify some client options:

- To specify `subdir=yes` and domain `all-local -systemobject`, enter:
`options='-subdir=yes -domain="all-local -c: -systemobject"'`
- To specify domain `all-local -c: -d:`, enter:
`options='-domain="all-local -c: -d:"'`

OBJects

Specifies the objects for which the specified action is performed. Use a single space between each object. This parameter is required except when `ACTION=INCREMENTAL`. If the action is a backup, archive, retrieve, or restore operation, the objects are file spaces, directories, or logical volumes. See the *Backup-Archive Clients Installation and User's Guide* for command syntax information. If the action is to run a command or macro, the object is the name of the command or macro to run.

When you specify `ACTION=INCREMENTAL` without specifying a value for this parameter, the scheduled command is invoked without specified objects and attempts to process the objects as defined in the client option file. To select all file spaces or directories for an action, explicitly list them in the object string. Entering only an asterisk in the object string causes the backup to occur only for the directory where the scheduler was started.

Important:

- If you specify a second file specification, and it is not a valid destination, you receive this error:
ANS1082E Invalid destination file specification <filespec> entered.
- If you specify more than two file specifications, you receive this error:
ANS1102E Excessive number of command line arguments passed to the program!

When you specify ACTION=ARCHIVE, INCREMENTAL, or SELECTIVE for this parameter, you can list a maximum of twenty (20) file specifications.

Enclose the object string in double quotes if it contains blank characters (spaces), and then surround the double quotes with single quotes. If the object string contains multiple file names, enclose each file name with its own pair of double quotes, then surround the entire string with one pair of single quotes. Errors can occur if file names contain a space that is not quoted correctly.

The following examples show how to specify some file names:

- To specify /home/file 2, /home/gif files, and /home/my test file, enter:
OBJECTS='"/home/file 2" "/home/gif files" "/home/my test file"'
- To specify /home/test file, enter:
OBJECTS='"/home/test file"'

Wait

Specifies whether to wait for a scheduled client operation to complete. This parameter is useful when defining client actions from a command script or macro. This parameter is optional. The default is No. Possible values are:

No Specifies that you do not wait for the scheduled client operation to complete. If you specify this value and the value of the ACTION parameter is COMMAND, the return code indicates whether the client action was defined.

Yes

Specifies that you wait for the scheduled client operation to complete. If you specify this value and the value of the ACTION parameter is COMMAND, the return code indicates the status of the client operation.

You cannot issue the **DEFINE CLIENTACTION** command with WAIT=YES from the server console. However, from the server console, you can:

- Specify WAIT=YES with **DEFINE CLIENTACTION** as the command line of a DEFINE SCRIPT command.
- Specify WAIT=YES with **DEFINE CLIENTACTION** as the command line of a file whose contents will be read into the script that is defined by a DEFINE SCRIPT command.

Restriction: If you specify the **DEFINE CLIENTACTION** command with WAIT=YES in a macro, the immediate schedules defined by the command will not roll back if the macro does not complete successfully.

Example: Perform a one-time incremental backup

Issue an incremental backup command for client node TOM assigned to policy domain EMPLOYEE_RECORDS. Tivoli Storage Manager defines a schedule and associates the schedule to client node TOM (assuming that the client scheduler is running).

```
define clientaction tom domain=employee_records
action=incremental
```

DEFINE CLIENTACTION

Related commands

Table 56. Commands related to DEFINE CLIENTACTION

Command	Description
DELETE SCHEDULE	Deletes a schedule from the database.
QUERY ASSOCIATION	Displays the clients associated with one or more schedules.
QUERY EVENT	Displays information about scheduled and completed events for selected clients.
QUERY SCHEDULE	Displays information about schedules.
SET CLIENTACTDURATION	Specifies the duration of a schedule defined using the DEFINE CLIENTACTION command.

DEFINE CLIENTOPT (Define an option to an option set)

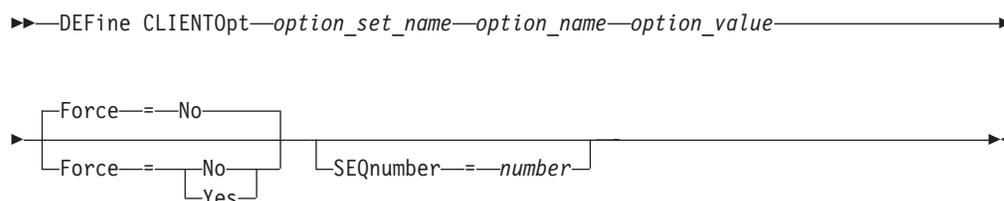
Use this command to add a client option to an option set.

For details about the options and the values you can specify, refer to *Backup-Archive Clients Installation and User's Guide*.

Privilege class

To issue this command, you must have system privilege or unrestricted policy privilege.

Syntax



Parameters

option_set_name (Required)

Specifies the name of the option set.

option_name (Required)

Specifies a client option to add to the option set.

For a list of valid client options, see "Client options that can be set by the Tivoli Storage Manager server" at the Tivoli Storage Manager information center (<http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3>).

Note: To define include-exclude values, specify the include or exclude option with *option-name*, and use *option_value* to specify any valid include or exclude statement, as you would in the client options file. For example:

```
define clientopt option_set_name incl excl "include c:\proj\text\devel.*"
```

option_value (Required)

Specifies the value for the option. If the option includes more than one value, enclose the value in quotation marks. For the values you can specify with the option refer to *Backup-Archive Clients Installation and User's Guide*.

Note:

1. The QUIET and VERBOSE options do not have an option value in the client option's file. To specify these values in a server client option set, specify a value of YES or NO.
2. To add an INCLUDE or EXCLUDE option for a file name that contains one or more spaces, put single quotation marks around the file specification, and double quotation marks around the entire option. See "Example: Add an option to a client option set" on page 128 for more information.
3. The *option_value* is limited to 1024 characters.

Force

Specifies whether the server forces the client to use the option set value. The

DEFINE CLIENTOPT

value is ignored for additive options, such as INCLEXCL and DOMAIN. The default is NO. This parameter is optional. The values are:

Yes

Specifies that the server forces the client to use the value. (The client cannot override the value.)

No Specifies that the server does not force the client to use the value. (The client can override the value.)

SEQnumber

Specifies a sequence number when an option name is specified more than once. This parameter is optional.

Example: Add an option to a client option set

Add a client option (MAXCMDRETRIES 5) to a client option set named ENG.

```
define clientopt eng maxcmdretries 5
```

Example: Add an option to exclude a file from backup

Add a client option to the option set ENGBACKUP to exclude the c:\admin\file.txt from backup services.

```
define clientopt engbackup inclexcl "exclude c:\admin\file.txt"
```

Example: Add an option to exclude a directory from backup

Add a client option to the option set WINSPEC to exclude a temporary internet directory from backup services. When you use the EXCLUDE or INCLUDE option with file names that contain spaces, put single quotation marks around the file specification, then double quotation marks around the entire option.

```
define clientopt winspec inclexcl "exclude.dir '*:\...\Temporary Internet Files'"
```

Example: Add an option to bind files in specified directories

Add client options to the option set WINSPEC to bind all files in directories C:\Data and C:\Program Files\My Apps to a management class named PRODCCLASS.

```
define clientopt winspec inclexcl "include C:\Data\...\* prodclass"  
define clientopt winspec inclexcl "include 'C:\Program  
Files\My Apps\...\*' prodclass"
```

Related commands

Table 57. Commands related to DEFINE CLIENTOPT

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLOPTSET	Defines a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
DELETE CLOPTSET	Deletes a client option set.
REGISTER NODE	Defines a client node to the server and sets options for that user.
QUERY CLOPTSET	Displays information about a client option set.

Table 57. Commands related to DEFINE CLIENTOPT (continued)

Command	Description
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.
UPDATE CLOPTSET	Updates the description of a client option set.
UPDATE NODE	Changes the attributes associated with a client node.

DEFINE CLOPTSET

DEFINE CLOPTSET (Define a client option set name)

Use this command to define a name for a set of options you can assign to clients for archive, backup, restore, and retrieve operations.

To add options to the new set, issue the **DEFINE CLIENTOPT** command.

Privilege class

To issue this command, you must have system privilege or unrestricted policy privilege.

Syntax

```
►►—DEfIne CLOptset—option_set_name—┐  
└DESCRiption—==—description—┘◄◄
```

Parameters

option_set_name (Required)

Specifies the name of the client option set. The maximum length of the name is 64 characters.

DESCRiption

Specifies a description of the client option set. The maximum length of the description is 255 characters. The description must be enclosed in quotation marks if it contains any blank characters. This parameter is optional.

Example: Define a client option set

To define a client option set named ENG issue the following command.

```
define cloptset eng
```

Related commands

Table 58. Commands related to **DEFINE CLOPTSET**

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLIENTOPT	Adds a client option to a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
DELETE CLOPTSET	Deletes a client option set.
QUERY CLOPTSET	Displays information about a client option set.
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.
UPDATE CLOPTSET	Updates the description of a client option set.

DEFINE COLLOGROUP (Define a collocation group)

Use this command to define a collocation group. A *collocation group* is a group of nodes whose data is collocated on a minimal number of sequential access volumes. Their data is collocated only if the storage pool definition is set to collocate by group (COLLOCATE=GROUP).

Privilege class

To issue this command, you must have system or unrestricted storage privilege.

Syntax

```
►►—DEfIne COLLOGrOuP—group_name—┌DEScRiption—=description—└─►
```

Parameters

group_name

Specifies the name of the collocation group name that you want to create. The maximum length of the name is 30 characters.

DEScRiption

Specifies a description of the collocation group. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Example: Define a collocation group

To define a collocation group named GROUP1 issue the following command:

```
define collogroup group1
```

Related commands

Table 59. Commands related to DEFINE COLLOGROUP

Command	Description
DEFINE COLLOCMEMBER	Adds a client node to a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node from a collocation group.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY COLLOGROUP	Displays information about collocation groups.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY STGPOOL	Displays information about storage pools.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.

DEFINE COLLOGROUP

Table 59. Commands related to DEFINE COLLOGROUP (continued)

Command	Description
UPDATE COLLOGROUP	Updates the description of a collocation group.
UPDATE STGPOOL	Changes the attributes of a storage pool.

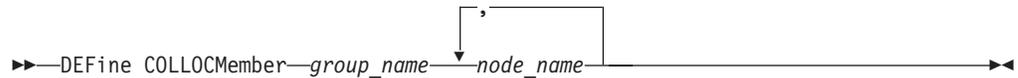
DEFINE COLLOCMEMBER (Define collocation group member)

Use this command to add a client node to a collocation group. A *collocation group* is a group of nodes whose data is collocated on a minimal number of sequential access volumes.

Privilege class

To issue this command you must have system or unrestricted storage privilege.

Syntax



Parameters

group_name

Specifies the name of the collocation group to which you want to add a client node.

node_name

Specifies the name of the client node that you want to add to the collocation group. You can specify one or more names. Separate multiple names with commas; do not use intervening spaces. You can also use wildcard characters when specifying multiple names.

Example: Define two collocation group members

Define two members, NODE1 and NODE2, to a collocation group, GROUP1.

```
define collocmember group1 node1,node2
```

Related commands

Table 60. Commands related to DEFINE COLLOCMEMBER

Command	Description
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node from a collocation group.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY COLLOGROUP	Displays information about collocation groups.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY STGPOOL	Displays information about storage pools.

DEFINE COLLOCMEMBER

Table 60. Commands related to DEFINE COLLOCMEMBER (continued)

Command	Description
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
UPDATE COLLOCGROUP	Updates the description of a collocation group.
UPDATE STGPOOL	Changes the attributes of a storage pool.

DEFINE COPYGROUP (Define a copy group)

Use this command to define a new backup or archive copy group within a specific management class, policy set, and policy domain. The server uses the backup and archive copy groups to control how clients back up and archive files, and to manage the backed-up and archived files.

To allow clients to use the new copy group, you must activate the policy set that contains the new copy group.

You can define one backup and one archive copy group for each management class. To ensure that client nodes can back up files, include a backup copy group in the default management class for a policy set.

Attention: The **DEFINE COPYGROUP** command fails if you specify a copy storage pool as a destination.

The **DEFINE COPYGROUP** command has two forms, one for defining a backup copy group and one for defining an archive copy group. The syntax and parameters for each form are defined separately.

- “DEFINE COPYGROUP (Define an archive copy group)” on page 141
- “DEFINE COPYGROUP (Define a backup copy group)” on page 136

Table 61. Commands related to DEFINE COPYGROUP

Command	Description
ASSIGN DEFMGMTCLASS	Assigns a management class as the default for a specified policy set.
BACKUP NODE	Backs up a network-attached storage (NAS) node.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE MGMTCLASS	Defines a management class.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COPYGROUP	Deletes a backup or archive copy group from a policy domain and policy set.
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
EXPIRE INVENTORY	Manually starts inventory expiration processing.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY MGMTCLASS	Displays information about management classes.
SET ARCHIVERETENTIONPROTECTION	Specifies whether data retention protection is activated.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.

DEFINE COPYGROUP–backup


DEFINE COPYGROUP (Define a backup copy group)

Use this command to define a new backup copy group within a specific management class, policy set, and policy domain.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the copy group belongs.

Syntax



Parameters

domain_name (Required)

Specifies the policy domain for which you are defining the copy group.

policy_set_name (Required)

Specifies the policy set for which you are defining the copy group.

You cannot define a copy group for a management class that belongs to the ACTIVE policy set.

class_name (Required)

Specifies the management class for which you are defining the copy group.

STANDARD

Specifies the name of the copy group, which must be STANDARD. This parameter is optional. The default value is STANDARD.

Type=Backup

Specifies that you want to define a backup copy group. The default parameter is BACKUP. This parameter is optional.

DESTination (Required)

Specifies the primary storage pool where the server initially stores backup data. You cannot specify a copy storage pool as the destination.

FREquency

Specifies how frequently Tivoli Storage Manager can back up a file. This parameter is optional. Tivoli Storage Manager backs up a file only when the specified number of days has elapsed since the last backup. The FREQUENCY value is used only during a full incremental backup operation. This value is ignored during selective backup or partial incremental backup. You can specify an integer from 0 to 9999. The default value is 0, meaning that Tivoli Storage Manager can back up a file regardless of when the file was last backed up.

VERExists

Specifies the maximum number of backup versions to retain for files that are currently on the client file system. This parameter is optional. The default value is 2.

If an incremental backup operation causes the limit to be exceeded, the server expires the oldest backup version that exists in server storage. Possible values are:

number

Specifies the number of backup versions to retain for files that are currently on the client file system. You can specify an integer from 1 to 9999.

NOLimit

Specifies that you want the server to retain all backup versions.

The number of backup versions to retain is controlled by this parameter until versions exceed the retention time specified by the RETEXTRA parameter.

VERDeleted

Specifies the maximum number of backup versions to retain for files that have been deleted from the client file system after being backed up using Tivoli Storage Manager. This parameter is optional. The default value is 1.

If a user deletes a file from the client file system, the next incremental backup causes the server to expire the oldest versions of the file in excess of this number. The expiration date for the remaining versions is determined by the retention time specified by the RETEXTRA or RETONLY parameter. Possible values are:

number

Specifies the number of backup versions to retain for files that are deleted from the client file system after being backed up. You can specify an integer from 0 to 9999.

NOLimit

Specifies that you want the server to retain all backup versions for files that are deleted from the client file system after being backed up.

RETEExtra

Specifies the number of days to retain a backup version after that version becomes inactive. A version of a file becomes inactive when the client stores a more recent backup version, or when the client deletes the file from the workstation and then runs a full incremental backup. The server deletes

DEFINE COPYGROUP–backup

inactive versions based on retention time even if the number of inactive versions does not exceed the number allowed by the VEREXISTS or VERDELETED parameters. This parameter is optional. The default value is 30 days. Possible values are:

days

Specifies the number of days to retain inactive backup versions. You can specify an integer from 0 to 9999.

NOLimit

Specifies that you want to retain inactive backup versions indefinitely.

If you specify NOLIMIT, the server deletes inactive backup versions based on the VEREXISTS parameter (when the file still exists on the client file system) VERDELETED parameter (when the file no longer exists on the client file system).

RETOnly

Specifies the number of days to retain the last backup version of a file that has been deleted from the client file system. This parameter is optional. The default value is 60. Possible values are:

days

Specifies the number of days to retain the last remaining inactive version of a file. You can specify an integer from 0 to 9999.

NOLimit

Specifies that you want to keep the last remaining inactive version of a file indefinitely.

If you specify NOLIMIT, the server retains the last remaining backup version forever, unless a user or administrator deletes the file from server storage.

MODE

Specifies whether Tivoli Storage Manager backs up a file only if the file has changed since the last backup, or whenever a client requests a backup. This parameter is optional. The default value is MODIFIED. Possible values are:

MODified

Specifies that Tivoli Storage Manager backs up the file only if it has changed since the last backup. Tivoli Storage Manager considers a file changed if any of the following is true:

- The date last modified is different
- The file size is different
- The file owner is different
- The file permissions are different

ABSolute

Specifies that Tivoli Storage Manager backs up the file regardless of whether it has been modified.

The MODE value is used only for full incremental backup. This value is ignored during partial incremental backup or selective backup.

SERialization

Specifies how Tivoli Storage Manager processes files or directories when they are modified during backup processing. This parameter is optional. The default value is SHRSTATIC. Possible values are:

SHRStatic

Specifies that Tivoli Storage Manager backs up a file or directory only if it is not being modified during backup. Tivoli Storage Manager attempts to perform a backup as many as four times, depending on the value specified for the CHANGINGRETRIES client option. If the file or directory is modified during each backup attempt, Tivoli Storage Manager does not back it up.

Static

Specifies that Tivoli Storage Manager backs up a file or directory only if it is not being modified during backup. Tivoli Storage Manager attempts to perform the backup only once.

Platforms that do not support the STATIC option default to SHRSTATIC.

SHRDynamic

Specifies that if the file or directory is being modified during a backup attempt, Tivoli Storage Manager backs up the file or directory during the last attempt even though the file or directory is being modified. Tivoli Storage Manager attempts to perform a backup as many as four times, depending on the value specified for the CHANGINGRETRIES client option.

Dynamic

Specifies that Tivoli Storage Manager backs up a file or directory on the first attempt, regardless of whether the file or directory is being modified during backup processing.

Attention: Be careful about using the SHRDYNAMIC and DYNAMIC values. Tivoli Storage Manager uses these values to determine if it backs up a file or directory while modifications are occurring. As a result, the backup version might be a fuzzy backup. A fuzzy backup does not accurately reflect what is currently in the file or directory because it contains some, but not all, modifications. If a file that contains a fuzzy backup is restored, the file may or may not be usable, depending on the application that uses the file. If a fuzzy backup is not acceptable, set **SERIALIZATION** to SHRSTATIC or STATIC so that Tivoli Storage Manager creates a backup version only if the file or directory is not being modified.

TOCDestination

Specifies the primary storage pool in which a table of contents (TOC) will initially be stored for any Network Data Management Protocol (NDMP) backup or backup set operation for which a TOC is generated. This parameter is optional. You cannot specify a copy storage pool as the destination. The storage pool specified for the destination must have NATIVE or NONBLOCK data format. To avoid mount delays, it is recommended that the storage pool have a device class of DISK or DEVTYPE=FILE. TOC generation is an option for NDMP backup operations, but is not supported for other image-backup operations.

If TOC creation is requested for a backup operation that uses NDMP and the image is bound to a management class whose backup copy group does not specify a TOC destination, the outcome will depend on the TOC parameter for the backup operation.

- If TOC=PREFERRED (the default), the backup proceeds without creation of a TOC.
- If TOC=YES, the entire backup fails because no TOC can be created.

DEFINE COPYGROUP–backup

Example: Create a backup copy group

Create a backup copy group named STANDARD for management class ACTIVEFILES in policy set VACATION in the EMPLOYEE_RECORDS policy domain. Set the backup destination to BACKUPPOOL. Set the minimum interval between backups to three days, regardless of whether the files have been modified. Retain up to five backup versions of a file while the file exists on the client file system.

```
define copygroup employee_records  
vacation activefiles standard type=backup  
destination=backuppools frequency=3  
verexists=5 mode=absolute
```


DEFINE COPYGROUP (Define an archive copy group)

Use this command to define a new archive copy group within a specific management class, policy set, and policy domain.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the copy group belongs.

Syntax



Parameters

domain_name (Required)

Specifies the name of the policy domain for which you are defining the copy group.

policy_set_name (Required)

Specifies the name of the policy set for which you are defining the copy group.

You cannot define a copy group for a management class that belongs to the ACTIVE policy set.

class_name (Required)

Specifies the name of the management class for which you are defining the copy group.

STANDARD

Specifies the name of the copy group, which must be STANDARD. This parameter is optional. The default value is STANDARD.

Type=Archive (Required)

Specifies that you want to define an archive copy group.

DEFINE COPYGROUP–archive

DESTination (Required)

Specifies the primary storage pool where the server initially stores the archive copy. You cannot specify a copy storage pool as the destination.

FREQuency=Cmd

Specifies the copy frequency, which must be CMD. This parameter is optional. The default value is CMD.

RETVer

Specifies the number of days to keep an archive copy. This parameter is optional. The default value is 365. Possible values are:

days

Specifies the length of time to keep an archive copy. You can specify an integer in the range 0 - 30000.

The **RETENTIONEXTENSION** server option can affect the volume retention if the following conditions are true:

- You specify zero for the number of days
- The destination storage pool for the archive copy group is a SnapLock storage pool (RECLAMATIONTYPE=SNAPLOCK)

If the two conditions are met, retention of the volumes is defined by the value of the **RETENTIONEXTENSION** server option. The **RETENTIONEXTENSION** server option value also applies if data is copied or moved into the SnapLock storage pool by a server process such as migration, or by using the **MOVE DATA** or **MOVE NODEDATA** commands.

NOLimit

Specifies that you want to keep an archive copy indefinitely.

If you specify **NOLIMIT**, the server retains archive copies forever, unless a user or administrator deletes the file from server storage. If you specify **NOLIMIT**, you cannot also specify **EVENT** for the **RETINIT** parameter.

The value of the **RETVER** parameter can affect the management class to which the server binds an archived directory. If the client does not use the **ARCHMC** option, the server binds directories that are archived to the default management class. If the default management class has no archive copy group, the server binds directories that are archived to the management class with the shortest retention period.

The **RETVER** parameter of the archive copy group of the management class to which an object is bound determines the retention criterion for each object. See the **SET ARCHIVERETENTIONPROTECTION** command for a description of data protection.

RETInit

Specifies when the retention time specified by the **RETVER** attribute is initiated. This parameter is optional. If you define the **RETINIT** value during copy group creation, you cannot modify it at a later time. The default value is **CREATION**. Possible values are:

CREATion

Specifies that the retention time specified by the **RETVER** attribute is initiated at the time an archive copy is stored on the Tivoli Storage Manager server.

EVent

Specifies that the retention time specified in the **RETVER** parameter is initiated at the time a client application notifies the server of a

retention-initiating event for the archive copy. If you specify RETINIT=EVENT, you cannot also specify RETVER=NOLIMIT.

Tip: You can place a deletion hold on an object that was stored with RETINIT=EVENT for which the event has not been signaled. If the event is signaled while the deletion hold is in effect, the retention period will be initiated, but the object will not be deleted while the hold is in effect.

REMin

Specifies the minimum number of days to keep an archive copy after it has been archived. This parameter is optional. The default value is 365. If you specify RETINIT=CREATION, this parameter is ignored.

MODE=ABSolute

Specifies that a file is always archived when the client requests it. The MODE must be ABSOLUTE. This parameter is optional. The default value is ABSOLUTE.

SERialization

Specifies how Tivoli Storage Manager processes files that are modified during archive. This parameter is optional. The default value is SHRSTATIC. Possible values are:

SHRStatic

Specifies that Tivoli Storage Manager archives a file only if it is not being modified. Tivoli Storage Manager attempts to perform an archive operation as many as four times, depending on the value specified for the CHANGINGRETRIES client option. If the file is modified during the archive attempt, Tivoli Storage Manager does not archive the file.

Static

Specifies that Tivoli Storage Manager archives a file only if it is not being modified. Tivoli Storage Manager attempts to perform the archive operation only once.

Platforms that do not support the STATIC option default to SHRSTATIC.

SHRDynamic

Specifies that if the file is being modified during an archive attempt, Tivoli Storage Manager archives the file during its last attempt even though the file is being modified. Tivoli Storage Manager attempts to archive the file as many as four times, depending on the value specified for the CHANGINGRETRIES client option.

Dynamic

Specifies that Tivoli Storage Manager archives a file on the first attempt, regardless of whether the file is being modified during archive processing.

Attention: Be careful about using the SHRDYNAMIC and DYNAMIC values. Tivoli Storage Manager uses them to determine if it archives a file while modifications are occurring. As a result, the archive copy might be a fuzzy backup. A fuzzy backup does not accurately reflect what is currently in the file because it contains some, but not all, modifications. If a file that contains a fuzzy backup is retrieved, the file might or might not be usable, depending on the application that uses the file. If a fuzzy backup is not acceptable, set SERIALIZATION to SHRSTATIC or STATIC so that Tivoli Storage Manager creates an archive copy only if the file is not being modified.

DEFINE COPYGROUP–archive

Example: Define an archive copy group for event-based retention

Create an archive copy group named STANDARD for management class EVENTMC in policy set SUMMER in the PROG1 policy domain. Set the archive destination to ARCHIVEPOOL, where the archive copy is kept until the server is notified of an event to initiate the retention time, after which the archive copy is kept for 30 days. The archive copy will be kept for a minimum of 90 days after being stored on the server, regardless of when the server is notified of an event to initiate the retention time.

```
define copygroup prog1 summer eventmc standard type=archive  
destination=archivepool retinit=event retver=30 retmin=90
```

DEFINE DATAMOVER (Define a data mover)

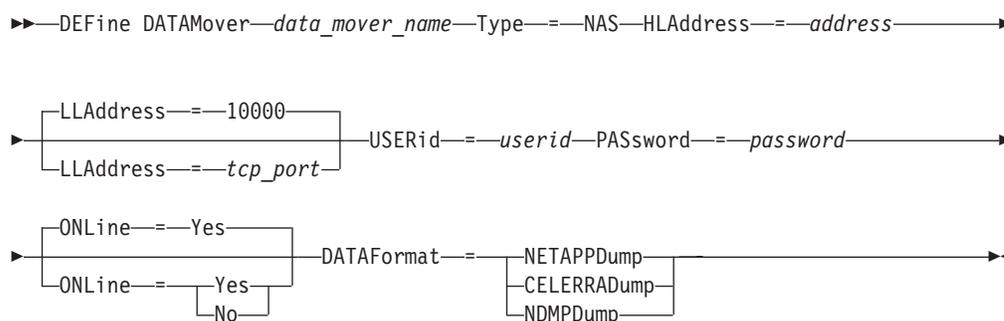
Use this command to define a data mover. A data mover is a named device that accepts a request from Tivoli Storage Manager to transfer data and can be used to perform outboard copy operations.

See the documentation for your device for guidance on specifying the parameters for this command.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

`data_mover_name` (Required)

Specifies the name of the data mover. This name must be the same as a node name that you previously registered using the `REGISTER NODE TYPE=NAS` command. The data that is backed up from this NAS data mover will be assigned to this node name in the server database. A maximum of 64 characters can be used to specify the name.

`Type=NAS` (Required)

Specifies that the data mover is a NAS file server.

`HLAddress` (Required)

Specifies either the numerical IP address or the domain name, which are used to access the NAS file server.

`LLAddress`

Specifies the TCP port number to access the NAS device for Network Data Management Protocol (NDMP) sessions. This parameter is optional. The default value is 10000.

`USERid` (Required)

Specifies the user ID for a user that is authorized to initiate an NDMP session with the NAS file server. For example, enter the administrative ID for a NetApp file server.

`PASsword` (Required)

Specifies the password for the user ID to log onto the NAS file server.

DEFINE DATAMOVER

ONLine

Specifies whether the data mover is available for use. This parameter is optional. The default is YES.

Yes

The default value. Specifies that the data mover is available for use.

No Specifies that the data mover is not available for use. When the hardware is being maintained, you can use the UPDATE DATAMOVER command to set the data mover off-line.

Important: If a library is controlled using a path from a NAS data mover to the library, and the NAS data mover is offline, the server will not be able to access the library. If the server is halted and restarted while the NAS data mover is offline, the library will not be initialized.

DATAFormat (Required)

Specifies the data format that is used by this data mover.

NETAPPDump

NETAPPDUMP must be used for NetApp NAS file servers and the IBM System Storage N Series.

CELERRADump

CELERRADUMP must be used for EMC Celerra NAS file servers.

NDMPDump

NDMPDump must be used for NAS file servers other than NetApp or EMC file servers.

Example: Define a data mover for a NetApp file server by IP address

Define a data mover for the NAS node named NAS1. The numerical IP address for the data mover is 9.67.97.103, at port 10000. The NAS file server is a NetApp device.

```
define datamover nas1 type=nas haddress=9.67.97.103 lladdress=10000 userid=root  
password=admin dataformat=netappdump
```

Example: Define a data mover by domain name

Define a data mover for the node named NAS1. The domain name for the data mover is, NETAPP2.TUCSON.IBM.COM at port 10000.

```
define datamover nas1 type=nas haddress=netapp2.tucson.ibm.com lladdress=10000  
userid=root password=admin dataformat=netappdump
```

Example: Define a data mover by IP address

Define a data mover for the node named NAS1. The numerical IP address for the data mover is 9.67.97.103, at port 10000. The NAS file server is neither a NetApp or an EMC file server.

```
define datamover nas1 type=nas haddress=9.67.97.103 lladdress=10000  
userid=root password=admin dataformat=ndmpdump
```

Related commands

Table 62. Commands related to DEFINE DATAMOVER

Command	Description
DEFINE PATH	Defines a path from a source to a destination.
DELETE DATAMOVER	Deletes a data mover.
QUERY DATAMOVER	Displays data mover definitions.
REGISTER NODE	Defines a client node to the server and sets options for that user.
UPDATE DATAMOVER	Changes the definition for a data mover.

DEFINE DEVCLASS (Define a device class)

Use this command to define a device class for a type of storage device. The server requires that a device class be defined to allow the use of a device.

For the most up-to-date list of supported devices and valid device class formats, see the Tivoli Storage Manager Supported Devices Web site:

http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html

Note: The DISK device class is defined by IBM Tivoli Storage Manager and cannot be modified with the **DEFINE DEVCLASS** command.

If you are defining a device class for devices that are to be accessed through a z/OS media server, see “DEFINE DEVCLASS - z/OS media server (Define device class for z/OS media server)” on page 212.

The following Tivoli Storage Manager device classes are ordered by device type.

- “DEFINE DEVCLASS (Define a 3590 device class)” on page 149
- “DEFINE DEVCLASS (Define a 3592 device class)” on page 153
- “DEFINE DEVCLASS (Define a 4MM device class)” on page 159
- “DEFINE DEVCLASS (Define an 8MM device class)” on page 163
- “DEFINE DEVCLASS (Define a DLT device class)” on page 169
- “DEFINE DEVCLASS (Define an ECARTRIDGE device class)” on page 175
- “DEFINE DEVCLASS (Define a FILE device class)” on page 183
- “DEFINE DEVCLASS (Define an LTO device class)” on page 186
- “DEFINE DEVCLASS (Define a NAS device class)” on page 193
- “DEFINE DEVCLASS (Define OPTICAL and WORM device classes)” on page 196
- “DEFINE DEVCLASS (Define a QIC device class)” on page 199
- “DEFINE DEVCLASS (Define a REMOVABLEFILE device class)” on page 204
- “DEFINE DEVCLASS (Define a SERVER device class)” on page 206
- “DEFINE DEVCLASS (Define a VOLSAFE device class)” on page 208

Table 63. Commands related to DEFINE DEVCLASS

Command	Description
BACKUP DEVCONFIG	Backs up Tivoli Storage Manager device information to a file.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE DEVCLASS	Deletes a device class.
QUERY DEVCLASS	Displays information about device classes.
QUERY DIRSPACE	Displays information about FILE directories.
UPDATE DEVCLASS	Changes the attributes of a device class.

DEFINE DEVCLASS (Define a 3590 device class)


Use the 3590 device class when you are using 3590 tape devices.

If you are defining a device class for devices that are to be accessed through a z/OS media server, see “DEFINE DEVCLASS (Define a 3590 device class for z/OS media server)” on page 213.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of the defined library object that contains the tape drives that can be used by this device class. For information about defining a library object, see the **DEFINE LIBRARY** command.

DEVTYPE=3590 (Required)

Specifies the 3590 device type is assigned to the device class. 3590 indicates that IBM 3590 cartridge tape devices are assigned to this device class.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional. The default value is DRIVE.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Use the specific format that the drives use.

Refer to the *Administrator's Guide* for more information.

The following tables list the recording formats, estimated capacities and recording format options for 3590 devices:

Table 64. Recording formats and default estimated capacities for 3590

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
3590B	10.0 GB	Uncompressed (basic) format
3590C	See note 20.0 GB	Compressed format
3590E-B	10.0 GB	Uncompressed (basic) format, similar to the 3590B format
3590E-C	See note 20.0 GB	Compressed format, similar to the 3590C format
3590H-B	30.0 GB (J cartridge – standard length) 60.0 GB (K cartridge - extended length)	Uncompressed (basic) format, similar to the 3590B format
3590H-C	See note 60.0 GB (J cartridge - standard length) 120.0 GB (K cartridge - extended length)	Compressed format, similar to the 3590C format

Note: If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity may be greater than the listed value.

Table 65. 3590 device recording format selections

Device	Format					
	3590B	3590C	3590E-B	3590E-C	3590H-B	3590H-C
3590	Read/Write	Read/Write	–	–	–	–
Ultra SCSI	Read/Write	Read/Write	–	–	–	–
3590E	Read	Read	Read/Write	Read/Write	–	–
3590H	Read	Read	Read	Read	Read/Write	Read/Write

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes). The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K).

For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB.

PREFIX

Specifies the high level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, use a volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the amount of time, in minutes, to retain an idle sequential access volume before dismounting it. This parameter is optional. The default value is 60. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a

simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

Note: For EXTERNAL library types, do not use the DRIVES default for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never be allowed to exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

that can be used by this device class. For information about defining a library object, see the **DEFINE LIBRARY** command.

DEVType=3592 (Required)

Specifies that the 3592 device type is assigned to the device class.

LBProtect

Specifies whether logical block protection is used to ensure the integrity of data stored on tape. When **LBPROTECT** is set to **READWRITE** or to **WRITEONLY**, the server uses this feature of the tape drive for logical block protection and generates cyclic redundancy check (CRC) protection information for each data block written on tape. The server also validates the CRC protection information when data is read from the tape.

The default is **NO**.

The following values are possible:

READWrite

Specifies that logical block protection is enabled in the server and the tape drive for both read and write operations. Data is stored with CRC information in each block. This mode affects performance because additional processor usage is required for Tivoli Storage Manager and the tape drive to calculate and compare CRC values. The **READWRITE** value does not affect backup sets and data that is generated by the **BACKUP DB** command.

When the **LBPROTECT** parameter is set to **READWRITE**, you do not have to specify the **CRCDATA** parameter in a storage pool definition because logical block protection provides better protection against data corruption.

WRITEOnly

Specifies that logical block protection is enabled in the server and the tape drive for write operations only. Data is stored containing CRC information in each block. For read operations, the server and the tape drive do not validate the CRC. This mode affects performance because additional processor usage is required for Tivoli Storage Manager to generate the CRC and for the tape drive to calculate and compare CRC values for write operations. The **WRITEONLY** value does not affect backup sets and data that are generated by the **BACKUP DB** command.

No Specifies that logical block protection is not enabled in the server and the tape drive for read and write operations. However, the server enables logical block protection on write operations for a filling volume that already has data with logical block protection.

Restriction: Logical block protection is supported only on IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.

WORM

Specifies whether the drives use **WORM** (write once, read many) media. The parameter is optional. The default is **NO**.

Yes

Specifies that the drives will use **WORM** media.

No Specifies that the drives will not use **WORM** media.

Remember:

1. To use 3592 **WORM** support in 3584 libraries, you only need to specify the **WORM** parameter. The Tivoli Storage Manager server will distinguish

between WORM and non-WORM scratch volumes. However, to use 3592 WORM support in 349X libraries, you also need to set the WORMSCRATCHCATEGORY on the **DEFINE LIBRARY** command. For details, see “DEFINE LIBRARY (Define a library)” on page 247.

2. When WORM=YES, the only valid value for the SCALECAPACITY parameter is 100.
3. Verify with your hardware vendors that your hardware is at the appropriate level of support.

For more information about WORM media, in general, refer to the *Administrator’s Guide*.

SCALECAPacity

Specifies the percentage of the media capacity that can be used to store data. This parameter is optional. The default is 100. Possible values are 20, 90, or 100.

Setting the scale capacity percentage to 100 provides maximum storage capacity. Setting it to 20 provides fastest access time.

Note: The scale capacity value will only take effect when data is first written to a volume. Any updates to the device class for scale capacity will not affect volumes that already have data written to them until the volume is returned to scratch status.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional. The default value is DRIVE.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Use the specific format that the drives use. Refer to the *Administrator’s Guide* for more information.

The following table lists the recording formats, estimated capacities and recording format options for 3592 devices:

Table 66. Recording formats and default estimated capacities for 3592

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
3592	300 GB	Uncompressed (basic) format
3592C	See note 900 GB	Compressed format
3592-2	500 GB 700 GB	Uncompressed (basic) format JA tapes Uncompressed (basic) format JB tapes
3592-2C	1.5 TB 2.1 TB	Compressed format JA tapes Compressed format JB tapes

Table 66. Recording formats and default estimated capacities for 3592 (continued)

Format	Estimated Capacity	Description
3592-3	640 GB	Uncompressed (basic) format JA tapes
	1 TB	Uncompressed (basic) format JB tapes
3592-3C	1.9 TB	Compressed format JA tapes
	3 TB	Compressed format JB tapes
3592-4	400 GB	Uncompressed (basic) format JK tapes
	1.5 TB	Uncompressed (basic) format JB tapes
	3.1 TB	Uncompressed (basic) format JC tape
3592-4C	1.2 TB	Compressed format JK tapes
	4.4 TB	Compressed format JB tapes
	9.4 TB	Compressed format JC tapes

Note: If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity might be different than the listed value.

Important: For optimal performance, avoid mixing different generations of drives in a single SCSI library. If you must mix drive generations in an SCSI library, use one of the special configurations that are described in the *Administrator's Guide* to prevent or minimize media problems.

Special configurations are also required for mixing different generations of 3592 drives in 349x and ACSLS libraries.

For details, see the *Administrator's Guide*.

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes). The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K).

For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB.

PREFIX

Specifies the high level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, use a volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the amount of time, in minutes, to retain an idle sequential access volume before dismounting it. This parameter is optional. The default value is 60. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

Tip: For EXTERNAL library types, do not use the DRIVES default for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never be allowed to exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

DRIVEEncryption

Specifies whether drive encryption will be permitted. This parameter is optional. The default is ALLOW.

ON Specifies that Tivoli Storage Manager is the key manager for drive encryption and will permit drive encryption for empty storage pool volumes only if the application method is enabled. (Other types of volumes—for example, backup sets, export volumes, and database backup volumes—will not be encrypted.) If you specify ON and you enable either the library or system method of encryption, drive encryption will not be permitted and backup operations will fail.

ALLOW

Specifies that Tivoli Storage Manager does not manage the keys for drive encryption. However, drive encryption for empty volumes is permitted if either the library or system method of encryption is enabled.

OFF

Specifies that drive encryption will not be permitted. If you enable either the library or system method of encryption, backups will fail. If you enable the application method, Tivoli Storage Manager will disable encryption and backups will be attempted.

DEFINE DEVCLASS (Define a 4MM device class)


Use the 4MM device class when you are using 4 mm tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

►►—DEFine DEVclass—*device_class_name*—LIBRary—=*library_name*—►►



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRary (Required)

Specifies the name of the defined library object that contains the 4 mm tape drives used by this device class. For information about defining a library object, see the `DEFINE LIBRARY` command.

DEVType=4MM (Required)

Specifies that the 4MM device type is assigned to the device class. The 4MM indicates that 4 mm tape devices are assigned to this device class.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional. The default value is DRIVE.

DEFINE DEVCLASS — 4MM

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Use the specific format that the drives use.

Refer to the *Administrator's Guide* for more information.

The following table lists the recording formats and estimated capacities for 4 mm devices:

Table 67. Recording formats and default estimated capacities for 4 mm tapes

Format	Estimated Capacity	Description
DRIVE	—	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
DDS1	2.6 GB (60-meter) 4.0 GB (90-meter)	Uncompressed format, only applies to 60-meter and 90-meter tapes
DDS1C	See note 1.3 GB (60-meter) 2.0 GB (90-meter)	Compressed format, only applies to 60-meter and 90-meter tapes
DDS2	4.0 GB	Uncompressed format, only applies to 120-meter tapes
DDS2C	See note 8.0 GB	Compressed format, only applies to 120-meter tapes
DDS3	12.0 GB	Uncompressed format, only applies to 125-meter tapes
DDS3C	See note 24.0 GB	Compressed format, only applies to 125-meter tapes
DDS4	20.0 GB	Uncompressed format, only applies to 150-meter tapes
DDS4C	See note 40.0 GB	Compressed format, only applies to 150-meter tapes
DDS5	36 GB	Uncompressed format, when using DAT 72 media
DDS5C	See note 72 GB	Compressed format, when using DAT 72 media
DDS6	80 GB	Uncompressed format, when using DAT 160 media
DDS6C	See note 160 GB	Compressed format, when using DAT 160 media

Note: If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity may be greater than the listed value.

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes). The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K).

For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB.

For more information on the default estimated capacity for 4 mm tapes, see Table 67 on page 160.

PREFIX

Specifies the high level qualifier of the file name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, use a volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. The default value is 60. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types (that is, a library managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be

simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

Note: For EXTERNAL library types, do not use the DRIVES default for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transaction will continue and complete but new transactions will be terminated.

DEFINE DEVCLASS — 8MM

DEVType=8MM (Required)

Specifies that the 8MM device type is assigned to the device class. 8MM indicates that 8 mm tape devices are assigned to this device class.

WORM

Specifies whether the drives will use WORM (write once, read many) media. The parameter is optional. The default is NO.

Yes

Specifies that the drives will use WORM media.

Note: If you select YES, the only options available for the FORMAT parameter are:

- DRIVE
- AIT
- AITC

No Specifies that the drives will not use WORM media.

Refer to the *Administrator's Guide* for more information.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional. The default value is DRIVE.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Use the specific format that the drives use. Refer to the *Administrator's Guide* for more information.

The following table lists the recording formats and estimated capacities for 8 mm devices:

Table 68. Recording format and default estimated capacity for 8 mm tape

Format		
Medium Type	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
8200	2.3 GB	Uncompressed (standard) format, using standard 112-meter tape cartridges
8200C	See note 3.5 GB 4.6 GB	Compressed format, using standard 112-meter tape cartridges

Table 68. Recording format and default estimated capacity for 8 mm tape (continued)

Format		
Medium Type	Estimated Capacity	Description
8500	See note	Drives (Read Write)
15m	600 MB	Eliant 820 (RW)
15m	600 MB	Exabyte 8500/8500C (RW)
15m	600 MB	Exabyte 8505 (RW)
54m	2.35 GB	Eliant 820 (RW)
54m	2.35 GB	Exabyte 8500/8500C (RW)
54m	2.35 GB	Exabyte 8505 (RW)
112m	5 GB or 10.0 GB	Eliant 820 (RW)
112m	5 GB or 10.0 GB	Exabyte 8500/8500C (RW)
112m	5 GB or 10.0 GB	Exabyte 8505 (RW)
160m XL	7 GB	Eliant 820 (RW)
8500C	See note	Drives (Read Write)
15m	1.2 GB	Eliant 820 (RW)
15m	1.2 GB	Exabyte 8500/8500C (RW)
15m	1.2 GB	Exabyte 8505 (RW)
54m	4.7 GB	Eliant 820 (RW)
54m	4.7 GB	Exabyte 8500/8500C (RW)
54m	4.7 GB	Exabyte 8505 (RW)
112m	5 GB or 10.0 GB	Eliant 820 (RW)
112m	5 GB or 10.0 GB	Exabyte 8500/8500C (RW)
112m	5 GB or 10.0 GB	Exabyte 8505 (RW)
160m XL	7 GB	Eliant 820 (RW)
8900	See note	Drive (Read Write)
15m	–	Mammoth 8900 (R)
54m	–	Mammoth 8900 (R)
112m	–	Mammoth 8900 (R)
160m XL	–	Mammoth 8900 (R)
22m	2.5 GB	Mammoth 8900 (RW)
125m	–	Mammoth 8900 (RW with upgrade)
170m	40 GB	Mammoth 8900 (RW)
AIT	See note	Drive
SDX1–25C	25 GB	AIT, AIT2 and AIT3 drives
SDX1–35C	35 GB	AIT, AIT2 and AIT3 drives
SDX2–36C	36 GB	AIT2 and AIT3 drives
SDX2–50C	50 GB	AIT2 and AIT3 drives
SDX3–100C	100 GB	AIT3, AIT4, and AIT5 drives
SDX3X–150C	150 GB	AIT3-Ex, AIT4, and AIT5 drives
SDX4–200C	200 GB	AIT4 and AIT5 drives
SDX5–400C	400 GB	AIT5 drive
AITC	See note	Drive
SDX1–25C	50 GB	AIT, AIT2 and AIT3 drives
SDX1–35C	91 GB	AIT, AIT2 and AIT3 drives
SDX2–36C	72 GB	AIT2 and AIT3 drives
SDX2–50C	130 GB	AIT2 and AIT3 drives
SDX3–100C	260 GB	AIT3, AIT4, and AIT5 drives
SDX3X–150C	390 GB	AIT3-Ex, AIT4, and AIT5 drives
SDX4–200C	520 GB	AIT4 and AIT5 drives
SDX5–400C	1040 GB	AIT5 drive

Table 68. Recording format and default estimated capacity for 8 mm tape (continued)

Format		
Medium Type	Estimated Capacity	Description
M2	See note	Drive (Read Write)
75m	20.0 GB	Mammoth II (RW)
150m	40.0 GB	Mammoth II (RW)
225m	60.0 GB	Mammoth II (RW)
M2C	See note	Drive (Read Write)
75m	50.0 GB	Mammoth II (RW)
150m	100.0 GB	Mammoth II (RW)
225m	150.0 GB	Mammoth II (RW)
SAIT	See note	Drive (Read Write)
	500 GB	Sony SAIT1-500(RW)
SAITC	See note	Drive (Read Write)
	1300 GB (1.3 TB)	Sony SAIT1-500(RW)
VXA2	See note	Drive (Read Write)
V6 (62m)	20 GB	VXA-2
V10 (124m)	40 GB	
V17 (170m)	60 GB	
VXA2C	See note	Drive (Read Write)
V6 (62m)	40 GB	VXA-2
V10 (124m)	80 GB	
V17 (170m)	120 GB	
VXA3	See note	Drive (Read Write)
X6 (62m)	40 GB	VXA-3
X10 (124m)	86 GB	
X23 (230m)	160 GB	
VXA3C	See note	Drive (Read Write)
X6 (62m)	80 GB	VXA-3
X10 (124m)	172 GB	
X23 (230m)	320 GB	

Note: The actual capacities may vary depending on which cartridges and drives are used.

- For the M2C format, the normal compression ratio is 2.5:1.
- For the AITC and SAITC formats, the normal compression ratio is 2.6:1.

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

Specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes).

For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB. The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K).

For more information on the default estimated capacity for 8 mm tapes, see Table 68 on page 164.

PREFIX

Specifies the high level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, use a volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:
 AB.CD2.E
- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. The default value is 60. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

For EXTERNAL library types (that is, a library managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

The following are possible values:

DEFINE DEVCLASS — 8MM

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

Note: For EXTERNAL library types, do not use the DRIVES default for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never be allowed to exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

Example: Define an 8 mm device class

Define a device class named 8MMTAPE for an 8 mm device in a library named AUTO. The format is DRIVE, mount limit is 2, mount retention is 10, tape volume prefix is named ADSMVOL, and the estimated capacity is 6 GB.

```
define devclass 8mmtape devtype=8mm library=auto
format=drive mountlimit=2 mountretention=10
prefix=adsmvol estcapacity=6G
```

DEFINE DEVCLASS (Define a DLT device class)


Use the DLT device class when you are using DLT tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

►►—DEFine DEVclass—*device_class_name*—LIBRARY—==—*library_name*—►►



DEFINE DEVCLASS — DLT

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of the defined library object that contains the DLT tape drives used by this device class. For information about defining a library object, see the DEFINE LIBRARY command.

DEVType=DLT (Required)

Specifies that the DLT device type is assigned to the device class. DLT indicates that DLT tape devices are assigned to this device class.

WORM

Specifies whether the drives will use WORM (write once, read many) media. The parameter is optional. The default is NO.

Yes

Specifies that the drives will use WORM media.

Note: Support for DLT WORM media is available only for SDLT-600, Quantum DLT-V4, and Quantum DLT-S4 drives in manual, SCSI, and ACSLS libraries.

No Specifies that the drives will not use WORM media.

For more information, refer to the *Administrator's Guide*.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional; the default value is DRIVE.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Use the specific format that the drives use. Refer to the *Administrator's Guide* for more information.

The following table lists the recording formats and estimated capacities for DLT devices:

Table 69. Recording format and default estimated capacity for DLT

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
DLT1	40.0 GB	Uncompressed format, using only CompacTape III cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT1C	See note 1 on page 172. 80.0 GB	Compressed format, using only CompacTape III cartridges Valid with DLT4000, DLT7000, and DLT8000 drives

Table 69. Recording format and default estimated capacity for DLT (continued)

Format	Estimated Capacity	Description
DLT10	10.0 GB	Uncompressed format, using only CompacTape III cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT10C	See note 1 on page 172. 20.0 GB	Compressed format, using only CompacTape III cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT15	15.0 GB	Uncompressed format, using only CompacTape IIIxt cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT15C	See note 1 on page 172. 30.0 GB	Compressed format, using only CompacTape IIIxt cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT20	20.0 GB	Uncompressed format, using only CompacTape IV cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT20C	See note 1 on page 172. 40.0 GB	Compressed format, using only CompacTape IV cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT35	35.0 GB	Uncompressed format, using only CompacTape IV cartridges Valid with DLT7000 and DLT8000 drives
DLT35C	See note 1 on page 172. 70.0 GB	Compressed format, using only CompacTape IV cartridges Valid with DLT7000 and DLT8000 drives
DLT40	40.0 GB	Uncompressed format, using CompacTape IV cartridges Valid with a DLT8000 drive
DLT40C	See note 1 on page 172. 80.0 GB	Compressed format, using CompacTape IV cartridges Valid with a DLT8000 drive
DLT2	80.0 GB	Uncompressed format, using Quantum DLT tape VS1 media
DLT2C	See note 1 on page 172. 160.0 GB	Compressed format, using Quantum DLT tape VS1 media
DLT4	160.0 GB	Uncompressed format, using Quantum DLTtape VS1 cartridges. Valid with Quantum DLT-V4 drive

DEFINE DEVCLASS — DLT

Table 69. Recording format and default estimated capacity for DLT (continued)

Format	Estimated Capacity	Description
DLT4C	See note 1. 320.0 GB	Compressed format, using Quantum DLTtape VS1 cartridges. Valid with Quantum DLT-V4 drive
SDLT See note 2.	100.0 GB	Uncompressed format, using Super DLT Tape 1 cartridges Valid with a Super DLT drive
SDLTC See note 2.	See note 1. 200.0 GB	Compressed format, using Super DLT Tape 1 cartridges Valid with a Super DLT drive
SDLT320 See note 2.	160.0 GB	Uncompressed format, using Quantum SDLT I media Valid with a Super DLT drive
SDLT320C See note 2.	See note 1. 320.0 GB	Compressed format, using Quantum SDLT I media Valid with a Super DLT drive
SDLT600	300.0 GB	Uncompressed format, using SuperDLTtape-II media Valid with a Super DLT drive
SDLT600C	See note 1. 600.0 GB	Compressed format, using SuperDLTtape-II media Valid with a Super DLT drive
DLTS4	800 GB	Uncompressed format, using Quantum DLT S4 media. Valid with a DLT-S4 drive
DLTS4C	See note 1. 1.6 TB	Compressed format, using Quantum DLT S4 media. Valid with a DLT-S4 drive

Note:

1. Depending on the effectiveness of compression, the actual capacity may be greater than the listed value.
2. Tivoli Storage Manager does not support a library that contains both Backward Read Compatible (BRC) SDLT and Non-Backward Read Compatible (NBRC) SDLT drives.

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes).

For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB. The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K).

For more information on estimated capacities, see Table 69 on page 170.

PREFIX

Specifies the high level qualifier of the data set name that the server writes into

the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, use a volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. The default value is 60. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types (that is, a library managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

Note: For EXTERNAL library types, do not use the DRIVES default for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

DEFINE DEVCLASS — DLT

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

DEFINE DEVCLASS (Define an ECARTRIDGE device class)


Use the ECARTRIDGE device class when you are using StorageTek drives such as the StorageTek SD-3, 9490, 9840 or 9940.

If you are defining a device class for devices that are to be accessed through a z/OS media server, see “DEFINE DEVCLASS (Define an ECARTRIDGE device class for z/OS media server)” on page 229.


Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



DEFINE DEVCLASS — ECARTRIDGE



Notes:

- 1 You cannot specify both `WORM=YES` and `DRIVEENCRYPTION=ON`.
- 2 You can use drive encryption only for Oracle StorageTek T10000B drives with a format value of `DRIVE`, `T10000B` or `T10000B-C`, and for Oracle StorageTek T10000C drives with a format value of `DRIVE`, `T10000C` or `T10000C-C`.

Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of the defined library object that contains the ECARTRIDGE tape drives that can be used by this device class. For information about defining a library object, see the **DEFINE LIBRARY** command.

DEVType=ECARTRidge (Required)

Specifies that the ECARTRIDGE device type is assigned to the device class. ECARTRIDGE indicates that a specific type of cartridge tape device (StorageTek) is assigned to this device class.

LBProtect

Specifies whether logical block protection is used to ensure the integrity of data stored on tape. When **LBPROTECT** is set to READWRITE or to WRITEONLY, the server uses this feature of the tape drive for logical block protection and generates cyclic redundancy check (CRC) protection information for each data block written on tape. The server also validates the CRC protection information when data is read from the tape.

The default is NO.

The following values are possible:

READWrite

Specifies that logical block protection is enabled in the server and the tape drive for both read and write operations. Data is stored with CRC information in each block. This mode affects performance because additional processor usage is required for Tivoli Storage Manager and the tape drive to calculate and compare CRC values. The READWRITE value does not affect backup sets and data that is generated by the **BACKUP DB** command.

When the **LBPROTECT** parameter is set to READWRITE, you do not have to specify the **CRCDATA** parameter in a storage pool definition because logical block protection provides better protection against data corruption.

WRITEOnly

Specifies that logical block protection is enabled in the server and the tape drive for write operations only. Data is stored containing CRC information in each block. For read operations, the server and the tape drive do not validate the CRC. This mode affects performance because additional processor usage is required for Tivoli Storage Manager to generate the CRC and for the tape drive to calculate and compare CRC values for write operations. The WRITEONLY value does not affect backup sets and data that are generated by the **BACKUP DB** command.

No Specifies that logical block protection is not enabled in the server and the tape drive for read and write operations. However, the server enables logical block protection on write operations for a filling volume that already has data with logical block protection.

Restriction: Logical block protection is supported only on Oracle StorageTek T10000C drives.

DEFINE DEVCLASS — ECARTRIDGE

WORM

Specifies whether the drives use WORM (write once, read many) media. The parameter is optional. The default is NO.

Restriction: If you select YES, the only options that are available for the FORMAT parameter are:

- DRIVE
- 9840
- 9840-C
- T9840D
- T9840D-C
- T10000A
- T10000A-C
- T10000B
- T10000B-C
- T10000C
- T10000C-C

Refer to the *Administrator's Guide* for more information.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional; the default value is DRIVE. If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Use the specific format that the drives use. Refer to the *Administrator's Guide* for more information.

Important: If you specify DRIVE for a device class that has non-compatible sequential access devices, then you must mount volumes on devices that are capable of reading or writing the format established when the volume was first mounted. This can cause delays if the only sequential access device that can access the volume is already in use.

The following table lists the recording formats and estimated capacities for ECARTRIDGE devices:

Table 70. Recording formats and default estimated capacities for ECARTRIDGE tapes

Format	Estimated capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
18T	360 MB	18-track uncompressed (standard) (read-only) format
18TC	1.44 GB	18-track extended (read-only), compressed format
36T	720 MB	36-track extended (read and write) format
36TC	2.88 GB	36-track extended (read and write), compressed format

Table 70. Recording formats and default estimated capacities for ECARTRIDGE tapes (continued)

Format	Estimated capacity	Description
SD3A	10 GB	Uncompressed (standard) format, using a 10 GB 'A' cartridge with 91 meters (298 feet) of tape
SD3AC	40 GB	Compressed format, using a 10 GB 'A' cartridge with 91 meters (298 feet) of tape
SD3B	25 GB	Uncompressed (standard) format, using a 25 GB 'B' cartridge with 204 meters (668 feet) of tape
SD3BC	100 GB	Compressed format, using a 25 GB 'B' cartridge with 204 meters (668 feet) of tape
SD3C	50 GB	Uncompressed (standard) format, using a 50 GB 'C' cartridge with 392 meters (1286 feet) of tape
SD3CC	200 GB	Compressed format, using a 50 GB 'C' cartridge with 392 meters (1286 feet) of tape
9840	20 GB	Uncompressed 9840 format, using a StorageTek 9840 cartridge. This format is valid for StorageTek 9840 and 9840B drives
9840-C	80 GB	LZ-1 Enhanced (4:1) compressed 9840 format, using a StorageTek 9840 cartridge
9940	60 GB	Uncompressed 9940 format, using a StorageTek 9940 cartridge
9940-C	120 GB	Compressed 9940 format, using a StorageTek 9940 cartridge
9940B	200 GB	Uncompressed 9940B format, using a StorageTek 9940 cartridge
9940B-C	400 GB	Compressed 9940B format, using a StorageTek 9940 cartridge
M8100	10 GB	Uncompressed (standard) format, using a 10 GB cartridge
M8100C	10 GB	Compressed format, using a 10 GB cartridge. Because there is no mechanism to determine the type of cartridge in an M8100 drive, the server assumes a 10 GB estimated uncompressed capacity.
T9840C	40 GB	Uncompressed T9840C format, using a StorageTek 9840 cartridge
T9840C-C	80 GB	Compressed T9840C format, using a StorageTek 9840 cartridge
T9840D	75 GB	Uncompressed T9840D format, using a StorageTek 9840 cartridge
T9840D-C	150 GB	Compressed T9840D format, using a StorageTek 9840 cartridge
T10000A	500 GB	Uncompressed T10000A format, using a StorageTek T10000 cartridge
T10000A-C	1 TB	Compressed T10000A format, using a StorageTek T10000 cartridge
T10000B	1 TB	Uncompressed T10000B format, using an Oracle StorageTek T10000 cartridge

DEFINE DEVCLASS — ECARTRIDGE

Table 70. Recording formats and default estimated capacities for ECARTRIDGE tapes (continued)

Format	Estimated capacity	Description
T10000B-C	2 TB	Compressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000C	5 TB	Uncompressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000C-C	10 TB	Compressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge

Notes:

- Some formats use a tape drive hardware compression feature. Depending on the effectiveness of compression, the actual capacity may be double or more than the listed value.
- T10000A drives can read and write the T10000A format only. T10000B drives can read, but cannot write, the T10000A format. T10000C drives can read, but cannot write, the T10000A and T10000B formats.

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes). The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K).

For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM.

If you have already established a tape label naming convention that supports your current tape management system, use a tape volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:
AB.CD2.E
- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

ADSM.BFS is an example of a tape volume filename using the default prefix and the added server qualifier.

MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. The default value is 60. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types (that is, a library managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

Tip: For EXTERNAL library types, do not use the DRIVES default for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never be allowed to exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

DRIVEEncryption

Specifies whether drive encryption is permitted. This parameter is optional. The default is ALLOW.

Restrictions:

1. You can use drive encryption only for the following drives:
 - Oracle StorageTek T10000B drives that have a format value of DRIVE, T10000B, or T10000B-C
 - Oracle StorageTek T10000C drives that have a format value of DRIVE, T10000C, or T10000C-C
2. You cannot specify Tivoli Storage Manager as the key manager for drive encryption of write once, read many (WORM) media. You cannot specify both WORM=YES and DRIVEENCRYPTION=ON.

DEFINE DEVCLASS — ECARTRIDGE

3. If encryption is enabled for a device class, and the device class is associated with a storage pool, the storage pool should not share a scratch pool with other device classes that cannot be encrypted. If a tape is encrypted, and you plan to use it on a drive that cannot be encrypted, you must manually relabel the tape before it can be used on that drive.

ON Specifies that Tivoli Storage Manager is the key manager for drive encryption and permits drive encryption for empty storage pool volumes only if the application method is enabled. (Other types of volumes are not encrypted. For example, backup sets, export volumes, and database backup volumes are not encrypted.) If you specify ON and you enable another method of encryption, drive encryption is not permitted and backup operations fail.

ALLOW

Specifies that Tivoli Storage Manager does not manage the keys for drive encryption. However, drive encryption for empty volumes is permitted if another method of encryption is enabled.

EXTERNAL

Specifies that Tivoli Storage Manager does not manage the keys for drive encryption. Use this setting with an encryption methodology that is provided by another vendor and that is used with Application Method Encryption (AME) enabled on the drive. When you specify EXTERNAL and Tivoli Storage Manager detects that AME encryption is enabled, Tivoli Storage Manager does not turn off encryption. By contrast, when you specify ALLOW and Tivoli Storage Manager detects that AME encryption is enabled, Tivoli Storage Manager turns off encryption.

OFF

Specifies that drive encryption is not permitted. If you enable another method of encryption, backups fail. If you enable the application method, Tivoli Storage Manager disables encryption and backups are not attempted.

DEFINE DEVCLASS (Define a FILE device class)

Use the FILE device class when you are using files on magnetic disk storage as volumes that store data sequentially (as on tape).

For information about disk subsystem requirements for FILE-type disk storage, see the *Administrator's Guide*.


The FILE device class does not support EXTERNAL libraries.

If you are defining a device class for devices that are to be accessed through a z/OS media server, see “DEFINE DEVCLASS (Define a FILE device class for z/OS media server)” on page 236.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

DEVType=FILE (Required)

Specifies that the FILE device type is assigned to the device class. FILE indicates that a file is assigned to this device class. When the server needs to access a volume that belongs to this device class, it opens a file and reads or writes file data.

A file is a form of sequential-access media.

MOUNTLimit

Specifies the maximum number of files that can be simultaneously open for input and output. This parameter is optional. The default value is 20. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

DEFINE DEVCLASS — FILE

MAXCAPacity

Specifies the maximum size of any data storage files defined to a storage pool in this device class.

The value of the MAXCAPACITY parameter is also used as the unit of allocation when storage pool space triggers create volumes. The default value is 2 GB (MAXCAPACITY=2G). The value specified should be less than or equal to the maximum supported size of a file on the target file system.

Specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes). The minimum size is 1 MB (MAXCAPACITY=1M). If you are defining a FILE device class for database-backup volumes, specify a value for MAXCAPACITY that is appropriate for the size of the database and that minimizes the number of database volumes.

Do not define a MAXCAPACITY value greater than 640M when this file is for REMOVABLEFILE CD support. A value less than a CD's usable space (650 MB) enables a one-to-one match between files from the FILE device class and copies that are on CD.

DIRectory

Specifies the directory location or locations of the files used in this device class. Enclose the entire list of directories within quotation marks, using commas to separate individual directory names. Special characters (for example, blank spaces) are permitted within directory names. For example, the directory list "abc def,xyz" contains two directories: abc def and xyz.

This parameter is optional.

The default is the current working directory of the server at the time the command is issued.

By specifying a directory name or names, you identify the location where the server places the files that represent storage volumes for this device class.

For NetApp SnapLock support (storage pools with RECLAMATIONTYPE=SNAPLOCK, which are going to use this device class), the directory or directories specified with DIRECTORY parameter must point to the directory or directories on the NetApp SnapLock volumes. For a detailed description of NetApp SnapLock support, refer to the *Administrator's Guide*.

While processing the command, the server expands the specified directory name or names into their fully qualified forms, starting from the root directory.

If the server needs to allocate a scratch volume, it creates a new file in one of these directories. (The server can choose any of the directories in which to create new scratch volumes.) For scratch volumes used to store client data, the file created by the server has a file name extension of .bfs. For scratch volumes used to store export data, a file name extension of .exp is used.

For example, if you define a device class with a directory of tsmstor and the server needs a scratch volume in this device class to store export data, the file that the server creates might be named /tsmstor/00566497.exp.

Important: You must ensure that storage agents can access newly created FILE volumes. Failure of the storage agent to access a FILE volume can cause operations to be retried on a LAN-only path or to fail. For more information, see the description of the DIRECTORY parameter in "DEFINE PATH (Define a path)" on page 276.

Tip: If you specify multiple directories for a device class, ensure that the directories are associated with separate file systems. Space trigger functions and storage pool space calculations take into account the space remaining in each directory. If you specify multiple directories for a device class and the directories reside in the same file system, the server will calculate space by adding values representing the space remaining in each directory. These space calculations will be inaccurate. Rather than choosing a storage pool with sufficient space for an operation, the server might choose the wrong storage pool and run out of space prematurely. For space triggers, an inaccurate calculation might result in a failure to expand the space available in a storage pool. Failure to expand space in a storage pool is one of the conditions that can cause a trigger to become disabled. If a trigger is disabled because the space in a storage pool could not be expanded, you can re-enable the trigger by issuing the following command: `update spacetrigger stg`. No further changes are required to the space trigger.

SHARED

Specifies that this FILE device class will be shared between the server and one or more storage agents. To prepare for sharing, a library will be automatically defined along with a number of drives corresponding to the MOUNTLIMIT parameter value. The drive names are the name of the library plus a number from 1 to the mount limit number. For example, if the library name is FILE and the mount limit is set to 4, the drives are named FILE11, FILE12, FILE13, FILE14.

For information about prerequisites when storage is shared by the server and storage agent, see http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Manager.

Example: Define a FILE device class with multiple directories

Define a device class that specifies multiple directories.

```
define devclass multidir devtype=file
    directory=/opt/xyz,/opt/abc,/opt/uvw
```

Example: Define a FILE device class with a 50 MB capacity

Define a device class named PLAINFILES with a FILE device type and a maximum capacity of 50 MB.

```
define devclass plainfiles devtype=file
    maxcapacity=50m
```

DEFINE DEVCLASS (Define an LTO device class)





Use the LTO device class when you are using LTO tape devices.


Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

►► DEFINE DEVCLASS *device_class_name* LIBRARY *library_name* ►►





Notes:

- 1 You cannot specify both WORM=YES and DRIVEENCRYPTION=ON.
- 2 Drive encryption is supported only for Ultrium 4, Ultrium 5, and Ultrium 6 drives and media.

Parameters

device_class_name **(Required)**

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of the defined library object that contains the LTO tape drives used by this device class. For information about defining a library object, see the **DEFINE LIBRARY** command.

DEVType=LTO (Required)

Specifies that the linear tape open (LTO) device type is assigned to the device class.

LBProtect

Specifies whether logical block protection is used to ensure the integrity of data stored on tape. When **LBPROTECT** is set to READWRITE or to WRITEONLY, the server uses this feature of the tape drive for logical block protection and generates cyclic redundancy check (CRC) protection information for each data block written on tape. The server also validates the CRC protection information when data is read from the tape.

The default is NO.

The following values are possible:

READWrite

Specifies that logical block protection is enabled in the server and the tape drive for both read and write operations. Data is stored with CRC information in each block. This mode affects performance because additional processor usage is required for Tivoli Storage Manager and the tape drive to calculate and compare CRC values. The READWRITE value does not affect backup sets and data that is generated by the **BACKUP DB** command.

When the **LBPROTECT** parameter is set to READWRITE, you do not have to specify the **CRCDATA** parameter in a storage pool definition because logical block protection provides better protection against data corruption.

WRITEOnly

Specifies that logical block protection is enabled in the server and the tape drive for write operations only. Data is stored containing CRC information in each block. For read operations, the server and the tape drive do not validate the CRC. This mode affects performance because additional processor usage is required for Tivoli Storage Manager to generate the CRC

DEFINE DEVCLASS — LTO

and for the tape drive to calculate and compare CRC values for write operations. The WRITEONLY value does not affect backup sets and data that are generated by the **BACKUP DB** command.

No Specifies that logical block protection is not enabled in the server and the tape drive for read and write operations. However, the server enables logical block protection on write operations for a filling volume that already has data with logical block protection.

Restriction: Logical block protection is supported only on IBM LTO5 and supported LTO6 drives.

WORM

Specifies whether the drives will use WORM (write once, read many) media. The parameter is optional. The default is NO.

Yes

Specifies that the drives will use WORM media.

Note:

1. To use WORM media in a library, all the drives in the library must be WORM capable.
2. You cannot specify Tivoli Storage Manager as the key manager for drive encryption of WORM (write once, read many) media. (Specifying both WORM=YES and DRIVEENCRYPTION=ON is not supported.)

No Specifies that the drives will not use WORM media.

For more information, refer to the *Administrator's Guide*.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional; the default value is DRIVE.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Use the specific format that the drives use. Refer to the *Administrator's Guide* for more information.

When migrating all drives from Ultrium to Ultrium 2 devices:

- Delete all existing Ultrium drive definitions and the paths associated with them.
- Define the new Ultrium 2 drives and paths.

If you are considering mixing different generations of LTO media and drives, be aware of the following restrictions. For more information about mixing drives and media, refer to the *Administrator's Guide*.

Table 71. Read - write capabilities for different generations of LTO drives

Drives	Generation 1 media	Generation 2 media	Generation 3 media	Generation 4 media	Generation 5 media	Generation 6 media
Generation 1	Read and write	n/a	n/a	n/a	n/a	n/a
Generation 2	Read and write	Read and write	n/a	n/a	n/a	n/a
Generation 3 ¹	Read only	Read and write	Read and write	n/a	n/a	n/a
Generation 4 ²	n/a	Read only	Read and write	Read and write	n/a	n/a

Table 71. Read - write capabilities for different generations of LTO drives (continued)

Drives	Generation 1 media	Generation 2 media	Generation 3 media	Generation 4 media	Generation 5 media	Generation 6 media
Generation 5 ³	n/a	n/a	Read only	Read and write	Read and write	n/a
Generation 6 ⁴	n/a	n/a	n/a	Read only	Read and write	Read and write

¹ In a library with a Generation 3 drive, all Generation 1 scratch volumes need to be checked out, and all Generation 1 storage pool volumes need to be updated to read-only.

² In a library with a Generation 4 drive, all Generation 2 scratch volumes need to be checked out, and all Generation 2 storage pool volumes need to be updated to read-only.

³ In a library with a Generation 5 drive, all Generation 3 scratch volumes need to be checked out, and all Generation 3 storage pool volumes need to be updated to read-only.

⁴ In a library with a Generation 6 drive, all Generation 4 scratch volumes need to be checked out, and all Generation 4 storage pool volumes need to be updated to read-only.

The following table lists the recording formats and estimated capacities for LTO devices:

Table 72. Recording format and default estimated capacity for LTO

Format	Estimated capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
ULTRIUM	100 GB	Uncompressed format, using Ultrium cartridges
ULTRIUMC	See note 200 GB	Compressed format, using Ultrium cartridges
ULTRIUM2	200 GB	Uncompressed (standard) format, using Ultrium 2 cartridges
ULTRIUM2C	See note 400 GB	Compressed format, using Ultrium 2 cartridges
ULTRIUM3	400 GB	Uncompressed (standard) format, using Ultrium 3 cartridges
ULTRIUM3C	See note 800 GB	Compressed format, using Ultrium 3 cartridges
ULTRIUM4	800 GB	Uncompressed (standard) format, using Ultrium 4 cartridges
ULTRIUM4C	See note 1.6 TB	Compressed format, using Ultrium 4 cartridges

Table 72. Recording format and default estimated capacity for LTO (continued)

Format	Estimated capacity	Description
ULTRIUM5	1.5 TB	Uncompressed (standard) format, using Ultrium 5 cartridges
ULTRIUM5C	See note 3.0 TB	Compressed format, using Ultrium 5 cartridges
ULTRIUM6	2.5 TB	Uncompressed (standard) format, using Ultrium 6 cartridges
ULTRIUM6C	See note 6.25 TB	Compressed format, using Ultrium 6 cartridges

Note: If this format uses the tape-drive hardware-compression feature, depending on the effectiveness of compression, the actual capacity may be greater than the listed value.

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), G (gigabytes), or T (terabytes). The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K).

For example, ESTCAPACITY=100G specifies that the estimated capacity for a volume in this device class is 100 GB.

For more information on estimated capacities, see Table 72 on page 189.

PREFIX

Specifies the high level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, use a volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:
AB.CD2.E
- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. The default value is 60. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types, you can enhance device sharing between applications by setting this parameter to a low value (for example, two minutes).

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

Note: For EXTERNAL library types, do not use the DRIVES default for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

DRIVEEncryption

Specifies whether drive encryption will be permitted. This parameter is optional. The default is ALLOW. Drive encryption is supported only for Ultrium 4, Ultrium 5, and Ultrium 6 drives and media.

Restriction: If encryption is enabled for a device class, and the device class is associated with a storage pool, the storage pool should not share a scratch pool with other device classes that cannot be encrypted. If a tape is encrypted, and you plan to use it on a drive that cannot be encrypted, you must manually relabel the tape before it can be used on that drive.

ON Specifies that Tivoli Storage Manager is the key manager for drive encryption and will permit drive encryption for empty storage pool volumes only if the application method is enabled. (Other types of volumes will not be encrypted. For example, backup sets, export volumes, and database backup volumes are not encrypted.) If you specify ON and

DEFINE DEVCLASS — LTO

you enable another method of encryption, drive encryption will not be permitted and backup operations will fail.

Note: You cannot specify Tivoli Storage Manager as the key manager for drive encryption of WORM (write once, read many) media. (Specifying both WORM=YES and DRIVEENCRYPTION=ON is not supported.)

ALLOW

Specifies that Tivoli Storage Manager does not manage the keys for drive encryption. However, drive encryption for empty volumes is permitted if another method of encryption is enabled.

EXTERNAL

Specifies that Tivoli Storage Manager does not manage the keys for drive encryption. Use this setting with an encryption methodology that is provided by another vendor and that is used with Application Method Encryption (AME) enabled on the drive. When you specify EXTERNAL and Tivoli Storage Manager detects that AME encryption is enabled, Tivoli Storage Manager does not turn encryption off. By contrast, when you specify ALLOW and Tivoli Storage Manager detects that AME encryption is enabled, Tivoli Storage Manager turns encryption off.

OFF

Specifies that drive encryption will not be permitted. If you enable another method of encryption, backups will fail. If you enable the application method, Tivoli Storage Manager will disable encryption and backups will be attempted.

Example: Define an LTO device class

Define a device class named LTOTAPE for an LTO drive in a library named LTOLIB. The format is ULTRIUM, mount limit is 12, mount retention is 5, tape volume prefix is named SMVOL, and the estimated capacity is 100 GB.

```
define devclass ltotape devtype=lto library=ltolib
format=ultrium mountlimit=12 mountretention=5
prefix=smvol estcapacity=100G
```

DEFINE DEVCLASS (Define a NAS device class)


Use the NAS device class when you are using NDMP (Network Data Management Protocol) operations to back up network-attached storage (NAS) file servers. The device class is for drives supported by the NAS file server for backups.

The NAS device class does not support EXTERNAL libraries.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

`DEVType=NAS` (Required)

Specifies that the network-attached storage (NAS) device type is assigned to the device class. The NAS device type is for drives attached to and used by a NAS file server for backup of NAS file systems.

`LIBRARY` (Required)

Specifies the name of the defined library object that contains the SCSI tape drives used by this device class. For information about defining a library object, see the `DEFINE LIBRARY` command.

`MOUNTRetention=0` (Required)

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. Zero (0) is the only supported value for device classes with `DEVType=NAS`.

`MOUNTWait`

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified

DEFINE DEVCLASS — NAS

amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

Tip: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

ESTCAPacity (Required)

Specifies the estimated capacity for the volumes assigned to this device class.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes). The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K).

For example, ESTCAPACITY=100G specifies that the estimated capacity for a volume in this device class is 100 GB.

PREFIX

Specifies the high level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, use a volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:
AB.CD2.E
- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is
ADSM.BFS.

Example: Define a NAS device class

Define a device class named NASTAPE for a NAS drive in a library named
NASLIB. The mount limit is DRIVES, mount retention is 0, tape volume prefix is
named SMVOL, and the estimated capacity is 200 GB.

```
define devclass nastape devtype=nas library=naslib  
mountretention=0 mountlimit=drives  
prefix=smvol estcapacity=200G
```

DEFINE DEVCLASS — OPTICAL and WORM types

DEFINE DEVCLASS (Define OPTICAL and WORM device classes)


Use the OPTICAL device class when you are using optical devices. Use the WORM device class when you are using optical devices that have WORM capability.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
▶▶—DEFine DEVclass—device_class_name—LIBRary—==—library_name————▶▶
```



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of the defined library object that contains the optical drives used by this device class. For information about defining a library object, see the DEFINE LIBRARY command.

DEVType (Required)

Specifies the device types that are assigned to the device class.

Possible values are:

OPTical

Specifies that the device class uses two-sided 5.25 inch rewritable optical media.

WORM

Specifies that the device class uses two-sided 5.25 inch write-once read-many (WORM) optical media.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional; the default value is DRIVE. If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Use the specific format that the drives use. Refer to the *Administrator's Guide* for more information.

The following table lists the recording formats and estimated capacities for OPTICAL and WORM devices:

Table 73. Estimated capacities for OPTICAL and WORM drives

Format	Estimated capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
650MB	650 MB	Using a 650 MB 5.25-inch optical drive
1300MB	1300 MB	Using a 1300 MB 5.25-inch optical drive
2600MB	2600 MB	Using a 2600 MB 5.25-inch optical drive
5200MB	5200 MB	Using a 5200 MB 5.25-inch optical drive
9100MB	9100 MB	Using a 9100 MB 5.25-inch optical drive
23GB	23 GB	Using Sony Blue Laser read-write and WORM media
30GB	30 GB	Using Plasmon UDO1 read-write and WORM media
60GB	60 GB	Using Plasmon or IBM UDO2 read-write and WORM media

Restriction: If you are considering mixing different generations of UDO media and drives, be aware of the following:

- UDO1 drives can read and write UDO1 media only.
- UDO2 drives can read from, but not write to, UDO1 media.

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes). The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K).

For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB.

For more information on estimated capacities, see Table 73.

DEFINE DEVCLASS — OPTICAL and WORM types

MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. The default value is 60. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for EXTERNAL library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true mount limit value (including online status).

Tip: For EXTERNAL library types, do not use the DRIVES default for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never be allowed to exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transaction will continue and complete but new transactions will be terminated.

DEFINE DEVCLASS (Define a QIC device class)


Use the QIC device class when you are using QIC tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

►►—DEFine DEVclass—*device_class_name*—LIBRary—=*library_name*—►►



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRary (Required)

Specifies the name of the defined library object that contains the tape drives

DEFINE DEVCLASS — QIC

that can be used by this device class. For information about defining a library object, see the **DEFINE LIBRARY** command.

DEVType=QIC (Required)

Specifies that the quarter-inch cartridge (QIC) device type is assigned to the device class.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional; the default value is DRIVE.

If the drives are in a library that includes drives of different tape technology, do not use the DRIVE value. Use the specific format that the drives use. Refer to *Administrator's Guide* for information.

The following tables list the recording formats, estimated capacities and recording format options for QIC devices:

Table 74. Recording format and default estimated capacity for quarter-inch cartridge (QIC) tape

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that can be supported by the sequential access drive on which a volume is mounted. Avoid specifying the DRIVE value when a mixture of devices is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
QIC120	26 MB – 172 MB	120 QIC format, depending on media Using DC600XTD, DC6150, DC6320, and DC6525
QIC150	31 MB – 207 MB	150 QIC format, depending on media Using DC600XTD, DC6150, DC6320, and DC6525
QIC525	65 MB – 427 MB	525 QIC format, depending on media Using DC6320 and DC6525
QIC1000	169 MB – 1.1 GB	1000 QIC format, depending on media Using DC9100 and DC9120XL
QIC2GB	2 GB	Uncompressed 2000 QIC format Using DC9100 and DC9120XL
QIC2GBC	See note 4 GB	Compressed 2000 QIC format
QIC4GB	4 GB	Uncompressed 4000 QIC format
QIC4GBC	See note 8 GB	Compressed 4000 QIC format
QIC12GB	12 GB	Uncompressed 12000 QIC format, using 343-meter tape

Table 74. Recording format and default estimated capacity for quarter-inch cartridge (QIC) tape (continued)

Format	Estimated Capacity	Description
QIC12GBC	See note 24 GB	Compressed 12000 QIC format, using 343-meter tape
QIC5010	13 GB – 16 GB	Uncompressed 5010 QIC format, depending on media
QIC5010C	See note 26 GB – 32 GB	Compressed 5010 QIC format, depending on media
QIC20GB	20 GB	Uncompressed 20000 QIC format
QIC20GBC	See note 40 GB	Compressed 20000 QIC format
QIC25GB	25 GB	Uncompressed 25000 QIC format
QIC25GBC	See note 50 GB	Compressed 25000 QIC format
QIC30GB	30 GB	Uncompressed 30000 QIC format
QIC30GBC	See note 60 GB	Compressed 30000 QIC format
QIC50GB	50 GB	Uncompressed 50GB QIC format
QIC50GBC	See note 100 GB	Compressed 50GB QIC format
QIC70GB	70 GB	Uncompressed 70GB QIC format
QIC70GBC	See note 140 GB	Compressed 70GB QIC format

Note: If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity may be greater than the listed value.

Table 75. QIC tape recording format options

Tape Format	QIC-120	QIC-150	QIC-525	QIC-1000
3M DC300XLP	–	–	–	–
3M DC600A	Read	–	–	–
3M DC600XTD	Read/Write	Read/Write	–	–
3M DC6150	Read/Write	Read/Write	–	–
3M DC6320	Read/Write	Read/Write	Read/Write	–
3M DC6525	Read/Write	Read/Write	Read/Write	–
3M DC9100	–	–	–	Read/Write
3M DC9120XL	–	–	–	Read/Write

Note: The server cannot use 3M DC300XLP and 3M DC600A tapes.

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

DEFINE DEVCLASS — QIC

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes). The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K).

For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB.

For more information on the default estimated capacity for QIC tapes, see Table 74 on page 200.

PREFIX

Specifies the high level qualifier of the file name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, use a volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:
AB.CD2.E
- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. The default value is 60. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types (that is, a library managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a

simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

Tip: For EXTERNAL library types, do not use the DRIVES default for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never be allowed to exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

DEFINE DEVCLASS — REMOVABLEFILE


DEFINE DEVCLASS (Define a REMOVABLEFILE device class)

Use the REMOVABLEFILE device class for removable media devices that are attached as local, removable file systems.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of the defined library object that contains the removable media drives used by this device class. For information about defining a library object, see the DEFINE LIBRARY command.

DEVType=REMOVABLEfile (Required)

Specifies that the REMOVABLEFILE device type is assigned to the device class. REMOVABLEFILE indicates that the volumes for this device class are files on local, removable media.

Volumes in a device class with device type REMOVABLEFILE are sequential access volumes.

Use the device manufacturer's utilities to format (if necessary) and label the media. The label on the media must meet the following restrictions:

- The label can have no more than 11 characters.
- The volume label and the name of the file on the volume must match exactly.

Refer to the *Administrator's Guide* for more information.

MAXCAPacity

Specifies the maximum size of any volumes defined to a storage pool categorized by this device class. This parameter is optional.

The MAXCAPACITY parameter must be set at less value than the capacity of the media. For CD media, the maximum capacity should be no greater than 650 MB.

space_remaining

The default maximum capacity is the space remaining on the media after it is first used.

size

You must specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes).

For example, MAXCAPACITY=5M specifies that the maximum capacity for a volume in this device class is 5 MB. The smallest value allowed is 100 KB (that is, MAXCAPACITY=100K).

MOUNTRetention

Specifies the amount of time, in minutes, to retain an idle sequential access volume before dismounting it. This parameter is optional. The default value is 60. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

Tip: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

DEFINE DEVCLASS — SERVER

DEFINE DEVCLASS (Define a SERVER device class)


Use the SERVER device class to use storage volumes or files archived in another Tivoli Storage Manager server.

If data retention protection is activated using the **SET ARCHIVERETENTIONPROTECTION** command, you cannot define a server device class. See the *Administrator's Guide* for more information.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

DEVType=SERVER (Required)

Specifies a remote connection that supports virtual volumes.

SERVERName (Required)

Specifies the name of the server. The SERVERNAME parameter must match a defined server.

MAXCAPacity

Specifies the maximum size for objects created on the target server; the default for this value is 500M. This parameter is optional.

500M

Specifies that the maximum capacity is 500M (500 MB).

size

Specify this value as an integer followed by *K* (kilobytes), *M* (megabytes), or *G* (gigabytes). The minimum value allowed is 100 KB (MAXCAPACITY=100K).

MOUNTLimit

Specifies the maximum number of simultaneous sessions between the source server and the target server. Any attempts to access more sessions than indicated by the mount limit cause the requester to wait. This parameter is optional. The default value is 1. You can specify a number from 1 to 4096.

The following are possible values:

- 1 Specifies that only one session between the source server and the target server is allowed.

number

Specifies the number of simultaneous sessions between the source server and the target server.

MOUNTRetention

Specifies the number of minutes to retain an idle connection with the target server before closing the connection. This parameter is optional. The default value is 60. You can specify a number from 0 to 9999.

PREFIX

Specifies the beginning portion of the high-level archive file name on the target server. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a high level archive file name using the default prefix is ADSM.volume1.

RETRYPeriod

Specifies the retry period in minutes. The retry period is the interval during which the server attempts to contact a target server if there is a suspected communications failure. This parameter is optional. The default value is 10 minutes.

RETRYInterval

Specifies the retry interval in seconds. The retry interval is how often retries are done within a given time period. This parameter is optional. The default value is 30 seconds.

DEFINE DEVCLASS — VOLSAFE

DEFINE DEVCLASS (Define a VOLSAFE device class)

Use the VOLSAFE device type to work with StorageTek VolSafe brand media and drives. This technology uses media that cannot be overwritten. Therefore, do not use these media for short-term backups of client files, the server database, or export tapes.

Restrictions:


1. NAS-attached libraries are not supported.
2. VolSafe media and read/write media must be in separate storage pools.
3. Check in cartridges with CHECKLABEL=YES on the CHECKIN LIBVOLUME command.
4. Label cartridges with OVERWRITE=NO on the LABEL LIBVOLUME command. If VolSafe cartridges are labeled more than one time, no additional data can be written to them.

See the *Administrator's Guide* for more information.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

`device_class_name` (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

`LIBRARY` (Required)

Specifies the name of the defined library object that contains the VolSafe drives that can be used by this device class. If any drives in a library are VolSafe-enabled, all drives in the library must be VolSafe-enabled. Consult your hardware documentation to enable VolSafe on the 9840 drives.

For information about defining a library object, see “DEFINE LIBRARY (Define a library)” on page 247.

DEVType=VOLSAFE (Required)

Specifies that the VOLSAFE device type is assigned to the device class. The label on this type of cartridge can be overwritten one time, which Tivoli Storage Manager does when it writes the first block of data. Therefore, it is important to limit the use of the LABEL LIBVOLUME command to one time per volume by using the OVERWRITE=NO parameter.

FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional; the default value is DRIVE.

Important: If you specify DRIVE for a device class that has non-compatible sequential access devices, then you must mount volumes on devices that are capable of reading or writing the format established when the volume was first mounted. This can cause delays if the only sequential access device that can access the volume is already in use.

The following table lists the recording formats and estimated capacities for VolSafe devices:

Table 76. Recording formats and default estimated capacities for Volsafe media

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives.
9840	20 GB	Uncompressed (standard) format, using a 20 GB cartridge with 270 meters (885 feet) of tape
9840-C	See note 80 GB	LZ-1 Enhanced (4:1) compressed format, using an 80 GB cartridge with 270 meters (885 feet) of tape

ESTCAPacity

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes). The smallest value that is allowed is 100 KB (that is, ESTCAPACITY=100K).

For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB.

For more information about the default estimated capacity for cartridge tapes, see Table 76.

MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. The default value is 60. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types (that is, a library managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

PREFIX

Specifies the beginning portion of the high-level archive file name on the target server. This parameter is optional. The default value is ADSM. The maximum length of this prefix is 8 characters.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a high-level archive file name that uses the default prefix is ADSM.volume1.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined to the library is used to calculate the true value (including online status).

Tip: For EXTERNAL library types, do not use the DRIVES default for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives that are used concurrently in this

device class by the server. This value must never be allowed to exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

DEFINE DEVCLASS - z/OS media server (Define device class for z/OS media server)

Use the **DEFINE DEVCLASS** command to define a device class for a type of storage device. The server requires that a device class be defined to allow the use of a device. A limited set of device class types is available for devices that are accessed through a z/OS media server.

- “DEFINE DEVCLASS (Define a 3590 device class for z/OS media server)” on page 213
- “DEFINE DEVCLASS (Define a 3592 device class for z/OS media server)” on page 218
- “DEFINE DEVCLASS (Define a cartridge device class for z/OS media server)” on page 224
- “DEFINE DEVCLASS (Define an ECARTRIDGE device class for z/OS media server)” on page 229
- “DEFINE DEVCLASS (Define a FILE device class for z/OS media server)” on page 236

Table 77. Commands related to **DEFINE DEVCLASS**

Command	Description
BACKUP DEVCONFIG	Backs up Tivoli Storage Manager device information to a file.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE DEVCLASS	Deletes a device class.
QUERY DEVCLASS	Displays information about device classes.
UPDATE DEVCLASS (z/OS media server)	Changes the attributes of a device class for storage managed by a z/OS media server.


DEFINE DEVCLASS (Define a 3590 device class for z/OS media server)

To use a z/OS media server to access 3590 devices, you must define a 3590 device class. In the device class definition, specify a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server.

DEFINE DEVCLASS — 3590, for z/OS media server

For information about defining a library, see the **DEFINE LIBRARY** command.

DEVtype=3590 (Required)

Specifies the 3590 device type is assigned to the device class. 3590 indicates that 3590 cartridge tape devices are assigned to the device class.

Restriction: The z/OS media server supports 256 KB data blocks when writing to 3590 tape drives. Verify that your hardware supports this capability.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional.

See the following table for the recording formats.

Table 78. Recording formats for 3590

Format	Description
3590B	Uncompressed (basic) format
3590C	Compressed format
3590E-B	Uncompressed (basic) format, similar to the 3590B format
3590E-C	Compressed format, similar to the 3590C format
3590H-B	Uncompressed (basic) format, similar to the 3590B format
3590H-C	Compressed format, similar to the 3590C format

Note: If the format uses the tape drive hardware compression feature the actual capacity can increase, depending on the effectiveness of compression.

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional. The default estimated capacity for 3590 tapes is 9 GB.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: K (KB), M (MB), or G (GB). For example, specify that the estimated capacity is 9 GB with the parameter **ESTCAPACITY=9G**. The smallest value that is accepted is 100 KB (**ESTCAPACITY=100K**).

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is **ADSM**. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is ADSM.BFS.

MOUNTRetention

Specifies the number of minutes that an idle tape volume is retained before it is dismounted. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. The default value is 60 minutes. Specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

MOUNTWait

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. The default value is 60. Specify a number, 1 - 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is 2.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

Possible values are:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

COMPression

Specifies whether file compression is used for this device class. This parameter is optional. The default value is **YES**.

Possible values are:

Yes

Specifies that the data for each tape volume is compressed.

No Specifies that the data for each tape volume is not compressed.

EXPIRATION

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional. There is no default value.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

Specify the expiration date using the format, *yyyymmdd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as 2014007 (the seventh day of year 2014).

If you specify the **EXPIRATION** parameter, you cannot specify the **RETENTION** parameter.

RETENTION

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

If you specify the **RETENTION** parameter, you cannot specify the **EXPIRATION** parameter.

PROTECTION

Specifies whether the RACF® program, if installed, protects volumes that are assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. The default value is **NO**. Possible values are:

No Specifies that the RACF program does not protect volumes that are assigned to this device class.

Yes

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

Tip: If sensitive data is stored on volumes that are assigned to this device class, use **PROTECTION=YES** and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Automatic

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Important: If you specify **PROTECTION=AUTOMATIC**, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify **PROTECTION=AUTOMATIC**, the z/OS media server issues **RACROUTE** commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

Do not change the setting to **PROTECTION=AUTOMATIC** for a device class that has been set to **PROTECTION=NO**. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for **PROTECTION** is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to **AUTOMATIC**, the server attempts to delete profiles when volumes are deleted.

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

UNIT

Specifies an esoteric unit name to specify a group of tape devices that support 3590 tape. This parameter is optional. The default unit name is 3590. The unit name can be up to 8 characters.

DEFINE DEVCLASS (Define a 3592 device class for z/OS media server)


To use a z/OS media server to access 3592 devices, you must define a 3592 device class. In the device class definition, specify a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

►►—DEFine DEVclass—*device_class_name*—LIBRARY—=*zos_media_library*—►►



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server.

For information about defining a library, see the **DEFINE LIBRARY** command.

DEVType=3592 (Required)

Specifies the 3592 device type is assigned to the device class.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional. The default value is **DRIVE**.

See the following table for the recording formats.

Table 79. Recording formats for 3592

Format	Description
3592	Uncompressed (basic) format
3592C	Compressed format
3592-2	Uncompressed (basic) format, similar to the 3592 format
3592-C	Compressed format, similar to the 3592C format
3592-3	Uncompressed (basic) format, similar to the 3592 format
3592-3C	Compressed format, similar to the 3592C format
3592-4	Uncompressed (basic) format, similar to the 3592 format
3592-4C	Compressed format, similar to the 3592C format
DRIVE	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
Note: If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity might be different from the listed value.	

If the drives are in a library that includes drives of different tape technology, do not use the **DRIVE** value. Use the specific format that the drives use. For optimal results, do not mix generations of drives in the same library. If a library contains mixed generations, media problems can result. For example, generation 1 and generation 2 drives cannot read generation 3 media. If possible, upgrade all drives to 3592 generation 3. If you cannot upgrade all drives to 3592 generation 3, you must use a special configuration.

WORM

Specifies whether the drives use WORM (write once, read many) media. This parameter is optional. The default is **NO**. Possible values are:

Yes

Specifies that the drives use WORM media.

No Specifies that the drives do not use WORM media.

Tip: The Tivoli Storage Manager server does not automatically delete scratch volumes in WORM storage pools after the volumes are emptied by expiration or other processes. To delete these volumes and remove them from WORM

storage pools, you must use the **DELETE VOLUME** command. Tivoli Storage Manager cannot reuse WORM volumes that were written to by the server and then deleted from a storage pool.

For more information about WORM media, see *Administrator's Guide*.

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: K (KB), M (MB), or G (GB). For example, specify that the estimated capacity is 9 GB with the parameter **ESTCAPACITY=9G**. The smallest value that is accepted is 100 KB (**ESTCAPACITY=100K**).

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is **ADSM**. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:
AB.CD2.E
- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is **ADSM.BFS**.

MOUNTRetention

Specifies the number of minutes that an idle tape volume is retained before it is dismounted. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. The default value is 60 minutes. Specify a number 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

MOUNTWait

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. The default value is 60. Specify a number, 1 - 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is 2.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

Possible values are:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

COMPression

Specifies whether file compression is used for this device class. This parameter is optional. The default value is **YES**.

Possible values are:

Yes

Specifies that the data for each tape volume is compressed.

No Specifies that the data for each tape volume is not compressed.

EXPIration

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional. There is no default value.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

Specify the expiration date using the format, *yyyymmdd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as 2014007 (the seventh day of year 2014).

If you specify the **EXPIRATION** parameter, you cannot specify the **RETENTION** parameter.

RETention

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

If you specify the **RETENTION** parameter, you cannot specify the **EXPIRATION** parameter.

PROtection

Specifies whether the RACF program, if installed, protects volumes that are assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. The default value is **NO**. Possible values are:

No Specifies that the RACF program does not protect volumes that are assigned to this device class.

Yes

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

Tip: If sensitive data is stored on volumes that are assigned to this device class, use **PROTECTION=YES** and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Automatic

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Important: If you specify **PROTECTION=AUTOMATIC**, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify **PROTECTION=AUTOMATIC**, the z/OS media server issues **RACROUTE** commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

Do not change the setting to **PROTECTION=AUTOMATIC** for a device class that has been set to **PROTECTION=NO**. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for **PROTECTION** is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to **AUTOMATIC**, the server attempts to delete profiles when volumes are deleted.

DEFINE DEVCLASS — 3592, for z/OS media server

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

UNIT

Specifies an esoteric unit name to specify a group of tape devices that support 3592 tape. This parameter is optional. The default value is 3592. The unit name can be up to 8 characters.

DEFINE DEVCLASS — CARTRIDGE, for z/OS media server

DEFINE DEVCLASS (Define a cartridge device class for z/OS media server)

To use a z/OS media server to access IBM 3480 and 3490 cartridge tape devices, you must define a **CARTRIDGE** device class. In the device class definition, specify a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server.

For information about defining a library, see the **DEFINE LIBRARY** command.

DEVtype=CARTridge (Required)

Specifies the **CARTRIDGE** device type is assigned to the device class. The **CARTRIDGE** device type is for cartridge tape devices.

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: K (KB), M (MB), or G (GB). For example, specify that the estimated capacity is 9 GB with the parameter **ESTCAPACITY=9G**. The smallest value that is accepted is 100 KB (**ESTCAPACITY=100K**).

See the following table for the recording formats.

Table 80. Recording formats and default estimated capacities for cartridge tapes

Format	Media type	Estimated capacity
3480	(Not applicable)	180 MB
3480XF	(Not applicable)	180 MB
3490 Base	CST	180 MB
3490B	CST	360 MB
3490C	CST	360 MB
3490B	ECCST	720 MB
3490C	ECCST	720 MB
Note: 3490 Base devices include models A01, A02, A10, A20, B02, B04, B20, B40, D41, D42.		

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is **ADSM**. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:
 AB.CD2.E
- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is **ADSM.BFS**.

MOUNTRetention

Specifies the number of minutes that an idle tape volume is retained before it is dismounted. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. The default value is 60 minutes. Specify a number 0 - 9999.

DEFINE DEVCLASS — CARTRIDGE, for z/OS media server

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

MOUNTWait

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. The default value is 60. Specify a number, 1 - 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is 2.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

Possible values are:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

COMPression

Specifies whether file compression is used for this device class. This parameter is optional. The default value is **YES**.

Possible values are:

Yes

Specifies that the data for each tape volume is compressed.

No Specifies that the data for each tape volume is not compressed.

EXpiration

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional. There is no default value.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

Specify the expiration date using the format, *yyyymmdd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as 2014007 (the seventh day of year 2014).

DEFINE DEVCLASS — CARTRIDGE, for z/OS media server

If you specify the **EXPIRATION** parameter, you cannot specify the **RETENTION** parameter.

RETention

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

If you specify the **RETENTION** parameter, you cannot specify the **EXPIRATION** parameter.

PROtection

Specifies whether the RACF program, if installed, protects volumes that are assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. The default value is **NO**. Possible values are:

No Specifies that the RACF program does not protect volumes that are assigned to this device class.

Yes

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

Tip: If sensitive data is stored on volumes that are assigned to this device class, use **PROTECTION=YES** and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Automatic

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Important: If you specify **PROTECTION=AUTOMATIC**, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify **PROTECTION=AUTOMATIC**, the z/OS media server issues **RACROUTE** commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

DEFINE DEVCLASS — CARTRIDGE, for z/OS media server

Do not change the setting to **PROTECTION=AUTOMATIC** for a device class that has been set to **PROTECTION=NO**. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for **PROTECTION** is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to **AUTOMATIC**, the server attempts to delete profiles when volumes are deleted.

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

UNIT

Specifies an esoteric unit name to specify a group of tape devices that support Cartridge System tape. This parameter is optional. The default value is 3480. The unit name can be up to 8 characters.

Example: Define a cartridge device class

Define a device class named SQUARE1 with a **CARTRIDGE** tape device and with RACF protection active for all tape volumes assigned to this device class. All data is compressed for this device class.

```
define devclass square1 library=zosmedia devtype=cartridge
    compression=yes protection=yes
```

DEFINE DEVCLASS (Define an ECARTRIDGE device class for z/OS media server)

To use a z/OS media server to access StorageTek drives such as the StorageTek SD-3, 9490, 9840, or 9940, you must define an **ECARTRIDGE** device class. In the device class definition, specify a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter.


StorageTek drives such as the StorageTek SD-3, 9490, 9840, or 9940 appear to be 3590 drives to the z/OS media server. If you have problems defining the devices on z/OS systems, consult your hardware vendor.

Privilege class


To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

►►—DEFine DEVclass—*device_class_name*—LIBRARY—==—*zos_media_library*—►►



DEFINE DEVCLASS — ECARTRIDGE, for z/OS media server



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

LIBRARY (Required)

Specifies the name of a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server.

For information about defining a library, see the **DEFINE LIBRARY** command.

DEVType=ECARTridge (Required)

Specifies that the **ECARTRIDGE** device type is assigned to the device class. The **ECARTRIDGE** device type is for StorageTek drives such as the StorageTek SD-3, 9490, 9840, or 9940.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional.

See the following table for the recording formats.

Table 81. Recording formats for ECARTRIDGE tapes

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. DRIVE is the default value. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
18T	360 MB	18-track uncompressed (standard) (read-only) format
18TC	1.44 GB	18-track extended (read-only), compressed format
36T	720 MB	36-track extended (read and write) format

DEFINE DEVCLASS — ECARTRIDGE, for z/OS media server

Table 81. Recording formats for ECARTRIDGE tapes (continued)

Format	Estimated Capacity	Description
36TC	2.88 GB	36-track extended (read and write), compressed format
SD3A	10 GB	Uncompressed (standard) format, using a 10 GB 'A' cartridge with 91 meters (298 feet) of tape
SD3AC	40 GB	Compressed format, using a 10 GB 'A' cartridge with 91 meters (298 feet) of tape
SD3B	25 GB	Uncompressed (standard) format, using a 25 GB 'B' cartridge with 204 meters (668 feet) of tape
SD3BC	100 GB	Compressed format, using a 25 GB 'B' cartridge with 204 meters (668 feet) of tape
SD3C	50 GB	Uncompressed (standard) format, using a 50 GB 'C' cartridge with 392 meters (1286 feet) of tape
SD3CC	200 GB	Compressed format, using a 50 GB 'C' cartridge with 392 meters (1286 feet) of tape
9840	20 GB	Uncompressed 9840 format, using a StorageTek 9840 cartridge
9840-C	80 GB	LZ-1 Enhanced (4:1) compressed 9840 format, using a StorageTek 9840 cartridge
9940	60 GB	Uncompressed 9940 format, using a StorageTek 9940 cartridge
9940-C	120 GB	Compressed 9940 format, using a StorageTek 9940 cartridge
9940B	200 GB	Uncompressed 9940B format, using a StorageTek 9940 cartridge
9940B-C	400 GB	Compressed 9940B format, using a StorageTek 9940 cartridge
M8100	10 GB	Uncompressed (standard) format, using a 10 GB cartridge
M8100C	10 GB	Compressed format, using a 10 GB cartridge. Because there is no mechanism to determine the type of cartridge in an M8100 drive, the server assumes a 10 GB estimated uncompressed capacity.
T9840C	40 GB	Uncompressed T9840C format, using a StorageTek 9840 cartridge
T9840C-C	80 GB	Compressed T9840C format, using a StorageTek 9840 cartridge
T9840D	75 GB	Uncompressed T9840D format, using a StorageTek 9840 cartridge
T9840D-C	150 GB	Compressed T9840D format, using a StorageTek 9840 cartridge
T10000A	500 GB	Uncompressed T10000A format, using a StorageTek T10000 cartridge
T10000A-C	1 TB	Compressed T10000A format, using a StorageTek T10000 cartridge
T10000B	1 TB	Uncompressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000B-C	2 TB	Compressed T10000B format, using an Oracle StorageTek T10000 cartridge

DEFINE DEVCLASS — ECARTRIDGE, for z/OS media server

Table 81. Recording formats for ECARTRIDGE tapes (continued)

Format	Estimated Capacity	Description
T10000C	5 TB	Uncompressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000C-C	10 TB	Compressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
Note: <ul style="list-style-type: none">• Some formats use a compression feature of the tape drive hardware. Depending on the effectiveness of compression, the actual capacity might be double or more than the listed value.• T10000A drives can read and write the T10000A format only. T10000B drives can read, but cannot write, the T10000A format. T10000C drives can read, but cannot write, the T10000A and T10000B formats.		

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional. The default estimated capacity is 9 GB.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: K (KB), M (MB), or G (GB). For example, specify that the estimated capacity is 9 GB with the parameter **ESTCAPACITY=9G**. The smallest value that is accepted is 100 KB (**ESTCAPACITY=100K**).

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The default value is **ADSM**. The maximum length of this prefix is 8 characters.

If you have a naming convention for media labels to support your current management system, use a volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:
AB.CD2.E
- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is **ADSM.BFS**.

MOUNTRetention

Specifies the number of minutes that an idle tape volume is retained before it is dismounted. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. The default value is 60 minutes. Specify a number 0 - 9999.

DEFINE DEVCLASS — ECARTRIDGE, for z/OS media server

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

MOUNTWait

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. The default value is 60. Specify a number, 1 - 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is 2.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

Possible values are:

DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

number

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

COMPression

Specifies whether file compression is used for this device class. This parameter is optional. The default value is **YES**.

Possible values are:

Yes

Specifies that the data for each tape volume is compressed.

No Specifies that the data for each tape volume is not compressed.

EXPIration

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional. There is no default value.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

Specify the expiration date using the format, *yyyymmdd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as 2014007 (the seventh day of year 2014).

DEFINE DEVCLASS — ECARTRIDGE, for z/OS media server

If you specify the **EXPIRATION** parameter, you cannot specify the **RETENTION** parameter.

RETention

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

If you specify the **RETENTION** parameter, you cannot specify the **EXPIRATION** parameter.

PROtection

Specifies whether the RACF program, if installed, protects volumes that are assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. The default value is **NO**. Possible values are:

No Specifies that the RACF program does not protect volumes that are assigned to this device class.

Yes

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

Tip: If sensitive data is stored on volumes that are assigned to this device class, use **PROTECTION=YES** and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Automatic

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

Important: If you specify **PROTECTION=AUTOMATIC**, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify **PROTECTION=AUTOMATIC**, the z/OS media server issues **RACROUTE** commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

DEFINE DEVCLASS — ECARTRIDGE, for z/OS media server

Do not change the setting to **PROTECTION=AUTOMATIC** for a device class that has been set to **PROTECTION=NO**. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for **PROTECTION** is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to **AUTOMATIC**, the server attempts to delete profiles when volumes are deleted.

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

UNIT

Specifies an esoteric unit name to specify a group of tape devices that support **ECARTRIDGE** tapes. Use the unit name that represents the subset of drives in the library that are attached to the z/OS system. This parameter is optional. The default value is 9840. The unit name can be up to 8 characters.

Example: Define a device class with the ECARTRIDGE device type

Define a device class named E1 with the **ECARTRIDGE** device type and with RACF protection active for all tape volumes assigned to this device class. All data is compressed for this device class. The device class is for a z/OS media server library named ZOSELIB.

```
define devclass e1 devtype=ecartridge library=zoselib compression=yes
  protection=yes
```

DEFINE DEVCLASS (Define a FILE device class for z/OS media server)

To use a z/OS media server to access storage volumes on magnetic disk devices, you must define a **FILE** device class. In the device class definition, specify a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter.

A volume in this device class is a Virtual Storage Access Method (VSAM) linear data set accessed by the z/OS media server. SCRATCH volumes can be used with device class and the z/OS media server will dynamically allocate the VSAM LDS. It is not necessary to define volumes for the server to use the device class. If you define volumes, set the high level qualifier (HLQ) so that SMS recognizes the allocation request by the z/OS media server. If you are using defined volumes, the format volume function is not supported for the server when using this device class. The z/OS media server uses a FormatWrite feature of DFSMS Media Manager when filling FILE volumes.


You can define volumes for the FILE device class by using the **DEFINE VOLUME** command. However, the z/OS media server does not allocate space for a defined volume until the volume is opened for its first use.

For information about disk subsystem requirements for **FILE** disk storage, see the *Administrator's Guide*.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

DEVType=FILE (Required)

Specifies that the **FILE** device type is assigned to the device class.

LIBRARY (Required)

Specifies the name of a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter. The disk storage that is used by this device class is accessed by the z/OS media server and managed by SMS.

For information about defining a library, see the **DEFINE LIBRARY** command.

MAXCAPacity

Specifies the maximum size of file volumes that are defined to a storage pool in this device class. This parameter is optional. The default value is 2 GB (**MAXCAPACITY=2G**).

Specify this value as an integer followed by K (KB), M (MB), or G (GB). The minimum size is 1 MB (**MAXCAPACITY=1M**). The maximum size is 16384 GB (**MAXCAPACITY=16384G**).

PRIMARYalloc

Specifies the initial amount of space that is dynamically allocated when a new volume is opened. Enough space must be available to satisfy the primary allocation amount. Storage Management Subsystem (SMS) policy determines whether multiple physical volumes can be used to satisfy the primary allocation request.

This parameter is optional. Specify this value as an integer followed by K (KB), M (MB), or G (GB). The minimum size is 100 KB (**PRIMARYALLOC=1K**). The maximum size is 16384 GB (**MAXCAPACITY=16384G**). The default size is 2600 MB (**PRIMARYALLOC=2600M**). All values are rounded to the next higher multiple of 256 KB.

To avoid wasted space, the dynamic allocation operation uses the smaller of the values specified in the two parameters, **PRIMARYALLOC** and **MAXCAPACITY**.

SMS automatic class selection (ACS) routines can affect whether the **PRIMARYALLOC** and **SECONDARYALLOC** parameter values are used.

SECONDARYalloc

Specifies the amount of space by which a file volume is extended when space that is already allocated to the file volume is used up. The data set for a file volume is extended up to the size set by the **MAXCAPACITY** parameter, then the volume is marked full.

Because secondary allocation of a linear data set cannot span a physical volume, consider the size of the physical volume when selecting a secondary allocation size. For example, physical volumes for a 3390 Model 3 are approximately 2.8 GB. To ensure that each extend request occupies nearly an entire physical volume but not more, use a secondary allocation size that is just less than 2.8 GB. A secondary allocation amount of 2600 MB allots enough space for the VSAM volume data set (VVDS), the volume label, and the volume table of contents (VTOC).

This parameter is optional. Specify this value as an integer followed by K (KB), M (MB), or G (GB). The minimum value is 0 KB (**SECONDARYALLOC=0K**). The default value is 2600 MB. The maximum value is 16384 GB. Except for 0, all values are rounded to the next higher multiple of 256 KB.

If you specify 0 (**SECONDARYALLOC=0**), the file volume cannot be extended beyond the primary allocation amount.

SMS automatic class selection (ACS) routines can affect whether the **PRIMARYALLOC** and **SECONDARYALLOC** parameter values are used.

If you specify a value for the **SECONDARYALLOCATION** parameter that is not 0, or if you allow the value to default to 2600M, the SMS DATACLAS associated with the PREFIX identifier (for example, High Level Qualifier) must have the Extended Addressability (EA) attribute specified. Without the EA attribute, the SMS DATACLAS limits the allocation of the VSAM LDS FILE volume to the primary extent. (See the description of the **PRIMARYALLOCATION** parameter). With

DEFINE DEVCLASS — FILE, for z/OS media server

the dataset limited to primary allocation size, the dataset cannot be extended by the z/OS media server, and the volume is marked FULL before the maximum capacity is reached.

Restriction: Ensure that the values that you specify for the **PRIMARYALLOC** and **SECONDARYALLOC** parameters are within practical limits for the storage device. The server cannot check whether the values exceed practical device limits, and does not check whether the two values together exceed the current **MAXCAPACITY** setting.

Tip: To fill volumes when you specify a large value for the **MAXCAPACITY** parameter, specify large values for the **PRIMARYALLOC** and **SECONDARYALLOC** parameters. Use larger MVS™ volume sizes to reduce the chance of extend failure.

PREFIX

Specifies the high-level qualifier of the data set name that is used to allocate scratch volume data sets. For all scratch file volumes created in this device class, the server uses this prefix to create the data set name. This parameter is optional. The default is **ADSM**. The maximum length of the prefix, including periods, is 32 characters.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:
AB.CD2.E
- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a file volume data set name using the default prefix is **ADSM.B0000021.BFS**.

If you have a data set naming convention, use a prefix that conforms to your naming conventions. For example, the following value is acceptable:
TSM.SERVER2.VSAMFILE.

If you are running multiple server instances for either Tivoli Storage Manager or Tivoli Storage Manager for z/OS Media you must use a unique value for the **PREFIX** parameter for each device class that you define.

MOUNTLimit

Specifies the maximum number of **FILE** volumes that can be open concurrently for this device class. This parameter is optional. The default value is 20.

If you are using IBM 3995 devices that emulate 3390 devices, set the value no higher than the number of concurrent input or output streams that are possible on the physical media.

The value that you specify in this parameter is important if there is a significant penalty switching from one volume to another. For example, switching can take place when using IBM 3995 devices to emulate 3390 devices. The value that you specify must be no higher than the number of physical drives available on the device.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

DEFINE DOMAIN (Define a new policy domain)


Use this command to define a new policy domain. A policy domain contains policy sets, management classes, and copy groups. A client is assigned to one policy domain. The ACTIVE policy set in the policy domain determines the rules for clients assigned to the domain. The rules control the archive, backup, and space management services provided for the clients.

You must activate a policy set in the domain before clients assigned to the policy domain can back up, archive, or migrate files.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

domain_name (Required)

Specifies the name of the policy domain to be defined. The maximum length of this name is 30 characters.

DESCRiption

Specifies a description of the policy domain. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

BACKRETention

Specifies the number of days (from the date the backup versions became inactive) to retain backup versions of files that are no longer on the client file system. This parameter is optional. You can specify an integer from 0 to 9999. The default value is 30. The server uses the backup retention value to manage inactive versions of files when any of the following conditions occur:

- A file is rebound to a new management class, but neither the new management class nor the default management class contains a backup copy group.
- The management class to which a file is bound no longer exists, and the default management class does not contain a backup copy group.
- The backup copy group is deleted from the management class to which a file is bound and the default management class does not contain a backup copy group.

DEFINE DOMAIN

ARCHRETention

Specifies the number of days (from the date of archive) to retain archive copies. This parameter is optional. You can specify an integer from 0 to 30000. The default value is 365. The server uses the archive retention value to manage archive copies of files when either of the following conditions occur:

- The management class to which a file is bound no longer exists, and the default management class does not contain an archive copy group.
- The archive copy group is deleted from the management class to which a file is bound and the default management class does not contain an archive copy group.

ACTIVEDESTination

This optional parameter specifies the names of active-data pools that store active versions of backup data for nodes assigned to the domain. You can specify up to ten active-data pools for a domain, separated by commas. Spaces are not permitted between the names.

Before the Tivoli Storage Manager server writes data to an active-data pool, it verifies that the node owning the data is assigned to a domain that has the active-data pool listed in the ACTIVEDESTINATION list. If the server verifies that the node meets this criteria, the data is stored in the active-data pool. If the node does not meet the criteria, then the data is not stored in the active-data pool. If the simultaneous-write function is used to write data to an active-data pool, the server performs the verification during backup operations by Tivoli Storage Manager backup-archive clients or by application clients using the Tivoli Storage Manager API. The verification is also performed when active-data is being copied using the **COPY ACTIVE** command.

Example: Define a policy domain

Define a policy domain with a name of PROG1 and the description, Programming Group Domain. Specify that archive copies are retained for 90 days when management classes or archive copy groups are deleted and the default management class does not contain an archive copy group. Also specify that backup versions are retained for 60 days when management classes or copy groups are deleted and the default management class does not contain a backup copy group.

```
define domain prog1
description="Programming Group Domain"
backretention=60 archretention=90
```

Related commands

Table 82. Commands related to DEFINE DOMAIN

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY DOMAIN	Creates a copy of a policy domain.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE DOMAIN	Deletes a policy domain along with any policy objects in the policy domain.
QUERY DOMAIN	Displays information about policy domains.
UPDATE DOMAIN	Changes the attributes of a policy domain.

DEFINE DRIVE (Define a drive to a library)

Use this command to define a drive. Each drive is assigned to a library, and so the library must be defined before you issue this command.

A path must be defined after you issue the **DEFINE DRIVE** command in order to make the drive usable by Tivoli Storage Manager software. For more information, see “DEFINE PATH (Define a path)” on page 276.

You can define more than one drive for a library by issuing the **DEFINE DRIVE** command for each drive. Stand-alone drives always require a manual library.


For detailed and current drive support information, see the Supported Devices website for your operating system:

http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Notes:

- 1 The ELEMENT parameter is valid only for drives in SCSI libraries when the drive type is not a network attached SCSI (NAS) drive. This parameter is not effective when the command is issued from a library client server (that is, when the library type is SHARED).
- 2 ACSDRVID is required for drives in ACSLS libraries. This parameter is not valid for non-ACSLs libraries.
- 3 The CLEANFREQUENCY parameter is valid only for drives in SCSI libraries.
- 4 The CLEANFREQUENCY=ASNEEDED parameter value does not work for all tape drives. See the parameter description for more information.

Parameters

library_name (Required)

Specifies the name of the library to which the drive is assigned. This parameter is required for all drives, including stand-alone drives. The specified library must have been previously defined by using the **DEFINE LIBRARY** command.

drive_name (Required)

Specifies the name that is assigned to the drive. The maximum length of this name is 30 characters.

SERial

Specifies the serial number for the drive that is being defined. This parameter is optional. The default is AUTODETECT.

If SERIAL=AUTODETECT, then the serial number reported by the drive when you define the path is used as the serial number.

If SERIAL=*serial_number*, then the serial number that is entered is used to verify that the path to the drive is correct when you define the path.

Note: Depending on the capabilities of the device, SERIAL=AUTODETECT might not be supported. In this case, the serial number is reported as blank.

ONLine

Specifies whether the drive is available for use. This parameter is optional. The default is YES.

Yes

Specifies that the drive is available for use.

No Specifies that the drive is not available for use.

ELEMent

Specifies the element address of a drive within a SCSI or virtual tape library (VTL). The server uses the element address to connect the physical location of the drive to the SCSI or VTL address of the drive. The default is AUTODETECT.

If ELEMENT=AUTODETECT, then the element number is automatically detected by Tivoli Storage Manager when the path to the drive is defined.

To find the element address for your library configuration, consult the information from the manufacturer.

Restriction:

- The ELEMENT parameter is valid only for drives in SCSI libraries or VTLs when the drive type is not a network attached SCSI (NAS) drive.
- This parameter is not effective when the command is issued from a library client server (that is, when the library type is SHARED).
- Depending on the capabilities of the library, ELEMENT=AUTODETECT might not be supported. In this case, you must supply the element address.

ACSDRVID

Specifies the ID of the drive that is being accessed in an ACSLS library. The drive ID is a set of numbers that indicates the physical location of a drive within an ACSLS library. This drive ID must be specified as *a,l,p,d*, where *a* is the ACSID, *l* is the LSM (library storage module), *p* is the panel number, and *d* is the drive ID. The server needs the drive ID to connect the physical location of the drive to the drive's SCSI address. See the StorageTek documentation for details.

CLEANFREQuency

Specifies how often the server activates drive cleaning. This parameter is optional. The default is NONE. For the most complete automation of cleaning for an automated library, you must have a cleaner cartridge checked into the library's volume inventory.

This parameter is not valid for externally managed libraries, such as 3494 libraries or StorageTek libraries that are managed under ACSLS.

Important: There are special considerations if you plan to use server-activated drive cleaning with a SCSI library that provides automatic drive cleaning support in its device hardware.

NONE

Specifies that the server does not track cleaning for this drive. This value can be used for libraries that have their own automatic cleaning.

ASNEEDED

Specifies that the server loads the drive with a checked-in cleaner cartridge only when a drive reports to the device driver that it needs cleaning.

The CLEANFREQUENCY=ASNEEDED parameter value does not work for all tape drives. Visit the Supported Devices website for your operating system to view detailed drive information. If **ASNEEDED** is not supported, you can use the *gigabytes* value for automatic cleaning.

Restriction: Tivoli Storage Manager does not control the drives that are connected to the NAS file server. If a drive is attached only to a NAS file server (no connection to a storage agent or server), do not specify ASNEEDED for the cleaning frequency.

gigabytes

Specifies, in gigabytes, how much data is processed on the drive before the server loads the drive with a cleaner cartridge. The server resets the gigabytes-processed counter each time it loads a cleaner cartridge in the drive.

Important: When CLEANFREQUENCY=gigabyte, drive cleaning can occur before the gigabyte setting is reached, if the drive notifies the device driver that a cleaning is necessary.

Consult the information from the drive manufacturer for cleaning recommendations. If the information gives recommendations for cleaning frequency in terms of hours of use, convert to a gigabytes value by doing the following:

1. Use the bytes-per-second rating for the drive to determine a gigabytes-per-hour value.
2. Multiply the gigabytes-per-hour value by the recommended hours of use between cleanings.
3. Use the result as the cleaning frequency value.

Using the cleaning frequency that is recommended by IBM for IBM drives ensures that the drives are not overcleaned.

For IBM 3590, specify a gigabyte value for the cleaning frequency to ensure that the drives receive adequate cleaning.

|
|
|

DEFINE DRIVE

Example: Define a drive to library

Define a drive in a manual library with a library name of LIB01 and a drive name of DRIVE01.

```
define drive lib01 drive01
define path server01 drive01 srctype=server desttype=drive
library=lib01 device=/dev/tmscsi/mt0
```

Example: Define a drive in an ACSLS library

Define a drive in an ACSLS library with a library name of ACSLIB and a drive name of ACSDRV1.

```
define drive acslib acsdrv1 acsdrv1=1,2,3,4
define path server01 acsdrv1 srctype=server desttype=drive
library=acslib device=/dev/tmscsi/mt0
```

Example: Define a drive in an automated library

Define a drive in an automated library with a library name of AUTO8MMLIB and a drive name of DRIVE01.

```
define drive auto8mmlib drive01 element=82
define path server01 drive01 srctype=server desttype=drive
library=auto8mmlib device=/dev/tmscsi/mt0
```

Related commands

Table 83. Commands related to DEFINE DRIVE

Command	Description
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE PATH	Defines a path from a source to a destination.
DELETE DRIVE	Deletes a drive from a library.
DELETE LIBRARY	Deletes a library.
PERFORM LIBACTION	Defines all drives and paths for a library.
QUERY DRIVE	Displays information about drives.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE DRIVE	Changes the attributes of a drive.
UPDATE PATH	Changes the attributes associated with a path.

DEFINE EVENTSERVER (Define a server as the event server)

Use this command to identify a server as the event server.

If you define an event server, one Tivoli Storage Manager server can send events to another Tivoli Storage Manager server that will log those events. See the *Administrator's Guide* for information about setting up logging events to an event server.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—DEFine EVENTSErver—*server_name*—◄◄

Parameters

server_name (Required)

Specifies the name of the event server. The server you specify must have already been defined with the **DEFINE SERVER** command.

Example: Designate the event server

Designate ASTRO to be the event server.

```
define eventserver astro
```

Related commands

Table 84. Commands related to **DEFINE EVENTSERVER**

Command	Description
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE EVENTSERVER	Deletes reference to the event server.
DISABLE EVENTS	Disables specific events for receivers.
ENABLE EVENTS	Enables specific events for receivers.
PING SERVER	Test the connections between servers.
QUERY EVENTSERVER	Displays the name of the event server.
QUERY SERVER	Displays information about servers.


DEFINE GRPMEMBER (Add a server to a server group)

Use this command to add a server as a member of a server group. You can also add one server group to another server group. A server group lets you route commands to multiple servers by specifying only the server group name.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

group_name (Required)

Specifies the name of the server group to which the member will be added.

member_name (Required)

Specifies the names of the servers or groups to be added to the group. To specify multiple servers and groups, separate the names with commas and no intervening spaces. The servers or server groups must already be defined to the server.

Example: Define a server to a server group

Define the server SANJOSE to server group CALIFORNIA.

```
define grpmember california sanjose
```

Example: Define a server and a server group to a server group

Define the server TUCSON and the server group CALIFORNIA to server group WEST_COMPLEX.

```
define grpmember west_complex tucson,california
```

Related commands

Table 85. Commands related to DEFINE GRPMEMBER

Command	Description
DEFINE SERVER	Defines a server for server-to-server communications.
DEFINE SERVERGROUP	Defines a new server group.
DELETE GRPMEMBER	Deletes a server from a server group.
DELETE SERVERGROUP	Deletes a server group.
MOVE GRPMEMBER	Moves a server group member.
QUERY SERVER	Displays information about servers.
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVERGROUP	Updates a server group.

DEFINE LIBRARY (Define a library)

Use this command to define a library. A library is a collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

A library can only be accessed by one source: a Tivoli Storage Manager server or a data mover. However, the drives in a library can be accessed by multiple sources.

The following library types can be defined to the Tivoli Storage Manager server. Syntax and parameter descriptions are available for each type.

- “DEFINE LIBRARY (Define a 349X library)” on page 249
- “DEFINE LIBRARY (Define an ACSLS library)” on page 252
- “DEFINE LIBRARY (Define an External library)” on page 254
- “DEFINE LIBRARY (Define a FILE library)” on page 256
- “DEFINE LIBRARY (Define a manual library)” on page 257
- “DEFINE LIBRARY (Define a SCSI library)” on page 259
- “DEFINE LIBRARY (Define a shared library)” on page 262
- “DEFINE LIBRARY (Define a VTL library)” on page 263
- “DEFINE LIBRARY (Define a ZOSMEDIA library type)” on page 266

For detailed and current library support information, see the Supported Devices website for your operating system:

http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html

Related commands

Table 86. Commands related to DEFINE LIBRARY

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE PATH	Defines a path from a source to a destination.
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE DRIVE	Deletes a drive from a library.
DELETE LIBRARY	Deletes a library.
DELETE PATH	Deletes a path from a source to a destination.
LABEL LIBVOLUME	Labels volumes in manual or automated libraries.
PERFORM LIBACTION	Defines all drives and paths for a library.
QUERY DRIVE	Displays information about drives.
QUERY LIBRARY	Displays information about one or more libraries.

DEFINE LIBRARY

Table 86. Commands related to **DEFINE LIBRARY** (continued)

Command	Description
QUERY LIBVOLUME	Displays information about a library volume.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE DRIVE	Changes the attributes of a drive.
UPDATE LIBRARY	Changes the attributes of a library.
UPDATE LIBVOLUME	Changes the status of a storage volume.
UPDATE PATH	Changes the attributes associated with a path.


DEFINE LIBRARY (Define a 349X library)

Use this syntax to define a 349X library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Notes:

- 1 The default value of the **RESETDRIVES** parameter is conditional. If the **SHARED** parameter is set to NO, the value of the **RESETDRIVES** parameter is NO. If the **SHARED** parameter is set to YES, the value of the **RESETDRIVES** parameter is YES.

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=349X (Required)

Specifies that the library is an IBM 3494 or 3495 Tape Library Dataserver

Restriction: IBM 3494 libraries support only one unique device type at a time.

SHARED

Specifies whether this library is shared with other Tivoli Storage Manager servers in a storage area network (SAN). This parameter is required when you define a library to the library manager.

YES

Specifies that this library can be shared with other servers. When you specify YES, the library manager server mounts volumes as requested by other servers and tracks drive and volume allocation to other servers.

DEFINE LIBRARY - 349X

NO Specifies that this library cannot be shared with other servers. **SHARED=NO** is required if the library is controlled by passing commands through a NAS file server.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional. The default is **YES**.

To use this option, you must check in the tapes with **CHECKLABEL=BARCODE** on the **CHECKIN LIBVOLUME** command.

No Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server labels only unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

SCRATCHCategory

Specifies the category number to be used for scratch volumes in the library. This parameter is optional. The default value is 301 (becomes X'12D' on the IBM 3494 since it uses hexadecimal values). You can specify a number from 1 to 65279. This number must be unique. It cannot be shared with other applications or defined libraries, and it must be different from the other category numbers in this library.

PRIVATECategory

Specifies the category number for private volumes that must be mounted by name. This parameter is optional. The default value is 300 (this value becomes X'12C' on the IBM 3494 because it uses hexadecimal values). You can specify a number from 1 to 65279. This number must be unique. It cannot be shared with other applications or defined libraries, and it must be different from the other category numbers in this library.

WORMSCRatchcategory

Specifies the category number to be used for WORM scratch volumes in the library. This parameter is required if you use WORM volumes. You can specify a number from 1 to 65279. This number must be unique. It cannot be shared with other applications or defined libraries, and it must be different from the other category numbers in this library. This parameter is only valid when 3592 WORM volumes are used.

Restriction: If the **WORMSCRATCHCATEGORY** is not defined and the **WORM** parameter is set to **YES** for the device class, the mount operation fails with an error message.

RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established. If, for example, a storage agent becomes unavailable but is still holding the path to a drive, persistent reserve allows the server to break the storage agent's reservation and access the drive.

Support for persistent reservation has the following limitations:

- If you are using the Tivoli Storage Manager device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 at <http://www.ibm.com/support/docview.wss?uid=swg21470319> for details.

- If you are using the IBM device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972> for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

Yes

Specifies that drive preemption through persistent reserve is used. YES is the default for a library that is defined with SHARED=YES. The **RESETDRIVES** parameter must be set to YES in a clustered environment or the drive is unavailable following failover.

No Specifies that drive preemption through persistent reserve is not used. NO is the default for a library that is defined with SHARED=NO.

Note: A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

Example: Define a 3494 library

Define a library named my3494 with a scratch category number of 550, a private category number of 600, and a WORM scratch category number of 400

```
define library my3494 libtype=349x scratchcategory=550
privatecategory=600 wormscratchcategory=400
```

DEFINE LIBRARY - ACSLS


DEFINE LIBRARY (Define an ACSLS library)

Use this syntax to define an ACSLS library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Notes:

- 1 The default value of the **RESETDRIVES** parameter is conditional. If the **SHARED** parameter is set to NO, the value of the **RESETDRIVES** parameter is NO. If the **SHARED** parameter is set to YES, the value of the **RESETDRIVES** parameter is YES.

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=ACSLs (Required)

Specifies that the library is a StorageTek library that is controlled by StorageTek Automated Cartridge System Library Software (ACSLs).

SHARED

Specifies whether this library is shared with other Tivoli Storage Manager servers in a storage area network (SAN). This parameter is required when you define a library to the library manager.

YES

Specifies that this library can be shared with other servers. When you specify YES, the library manager server mounts volumes as requested by other servers and tracks drive and volume allocation to other servers.

NO

Specifies that this library cannot be shared with other servers. SHARED=NO is required if the library is controlled by passing commands through a NAS file server.

RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established. If, for example, a storage agent becomes

unavailable but is still holding the path to a drive, persistent reserve allows the server to break the storage agent's reservation and access the drive.

Support for persistent reservation has the following limitations:

- If you are using the Tivoli Storage Manager device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 at <http://www.ibm.com/support/docview.wss?uid=swg21470319> for details.
- If you are using the IBM device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972> for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

Yes

Specifies that drive preemption through persistent reserve is used. YES is the default for a library that is defined with SHARED=YES. The **RESETDRIVES** parameter must be set to YES in a clustered environment or the drive is unavailable following failover.

No Specifies that drive preemption through persistent reserve is not used. NO is the default for a library that is defined with SHARED=NO.

Note: A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional. The default is YES.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the **CHECKIN LIBVOLUME** command.

No Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server labels only unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

ACSID

Specifies the number of this StorageTek library that is assigned by the ACSSA (Automatic Cartridge System System Administrator). This number can be from 0 to 126. Issue QUERY ACS on your system to get the number for your library ID. This parameter is required.

For more information, see your StorageTek documentation.

Example: Define a shared ACSLS library

Define a library named ACSLIB with the library type of ACSLS and an ACSID of 1.

```
define library acslib libtype=acsls acsid=1 shared=yes
```

DEFINE LIBRARY - EXTERNAL


DEFINE LIBRARY (Define an External library)

Use this syntax to define an External library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=EXTERNAL (Required)

Specifies that the library is managed by an external media management system. This library type does not support drive definitions with the **DEFINE DRIVE** command. Rather, the external media management system identifies the appropriate drive for media access operations.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional. The default is YES.

To use this option, you need to check in the tapes with **CHECKLABEL=BARCODE** on the **CHECKIN LIBVOLUME** command.

No Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server only labels unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

Example: Define an external library for a SAN configuration

For an IBM Tivoli Storage Manager for Storage Area Networks configuration, define a library named EXTLIB with the library type of EXTERNAL. If using Gresham Enterprise Distributable, the external library manager executable file is located in the following directory:

```
/opt/OMIdtelm/bin/elm
```

1. Define the library:

```
define library extlib libtype=external
```

2. Define the path:

```
define path server1 extlib srctype=server desttype=library
    externalmanager="/opt/OMIdtelm/bin/elm"
```

DEFINE LIBRARY - FILE


DEFINE LIBRARY (Define a FILE library)

Use this syntax to define a FILE library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=FILE (Required)

Specifies that a pseudo-library is created for sequential file volumes. When you issue the **DEFINE DEVCLASS** command with DEVTYPE=FILE and SHARED=YES parameters, this occurs automatically. FILE libraries are necessary only when sharing sequential file volumes between the server and one or more storage agents. The use of FILE libraries requires library sharing. Shared FILE libraries are supported for use in LAN-free backup configurations only. You cannot use a shared FILE library in an environment in which a library manager is used to manage library clients.

SHARED

Specifies whether this library is shared with other Tivoli Storage Manager servers in a storage area network (SAN). This parameter is required when you define a library to the library manager.

YES

Specifies that this library can be shared with other servers. When you specify YES, the library manager server mounts volumes as requested by other servers and tracks drive and volume allocation to other servers.

NO

Specifies that this library cannot be shared with other servers. SHARED=NO is required if the library is controlled by passing commands through a NAS file server.

Example: Define a shared FILE library

Define a file library with shared=yes.

```
define library file1 libtype=file shared=yes
```


DEFINE LIBRARY (Define a manual library)

Use this syntax to define a manual library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=MANUAL (Required)

Specifies that the library is not automated. When volumes must be mounted on drives in this type of library, messages are sent to operators. This type of library is used with stand-alone drives.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional. The default is YES.

To use this option, you need to check in the tapes with CHECKLABEL=BARCODE on the **CHECKIN LIBVOLUME** command.

No Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server only labels unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established. If, for example, a storage agent becomes unavailable but is still holding the path to a drive, persistent reserve allows the server to break the storage agent's reservation and access the drive.

Support for persistent reservation has the following limitations:

- If you are using the Tivoli Storage Manager device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 at <http://www.ibm.com/support/docview.wss?uid=swg21470319> for details.

DEFINE LIBRARY - MANUAL

- If you are using the IBM device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972> for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

Yes

Specifies that drive preemption through persistent reserve is used. YES is the default for a library that is defined with SHARED=YES. The **RESETDRIVES** parameter must be set to YES in a clustered environment or the drive is unavailable following failover.

No Specifies that drive preemption through persistent reserve is not used. NO is the default for a library that is defined with SHARED=NO.

Note: A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

Example: Define a manual library

Define a library named MANUALMOUNT with the library type of MANUAL.

```
define library manualmount libtype=manual
```


DEFINE LIBRARY (Define a SCSI library)

Use this syntax to define a SCSI library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Notes:

- 1 The default value of the **RESETDRIVES** parameter is conditional. If the **SHARED** parameter is set to NO, the value of the **RESETDRIVES** parameter is NO. If the **SHARED** parameter is set to YES, the value of the **RESETDRIVES** parameter is YES.

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=SCSI (Required)

Specifies that the library has a SCSI-controlled media changer device. To mount volumes on drives in this type of library, Tivoli Storage Manager uses the media changer device.

SHARED

Specifies whether this library is shared with other Tivoli Storage Manager servers in a storage area network (SAN). This parameter is required when you define a library to the library manager.

YES

Specifies that this library can be shared with other servers. When you specify YES, the library manager server mounts volumes as requested by other servers and tracks drive and volume allocation to other servers.

NO

Specifies that this library cannot be shared with other servers. SHARED=NO is required if the library is controlled by passing commands through a NAS file server.

DEFINE LIBRARY - SCSI

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional. The default is NO.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the **CHECKIN LIBVOLUME** command.

No Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server labels only unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

RELABELSCRatch

Specifies whether the server relabels volumes that were deleted and returned to scratch. When this parameter is set to YES, a LABEL LIBVOLUME operation is started and the existing volume label is overwritten. This parameter is optional and intended for use with a Virtual Tape Library (VTL).

Note: If you have both virtual and real volumes in your VTL, both types are relabeled when this parameter is enabled. If the VTL includes real volumes, specifying this option might impact performance.

No Specifies that the server does not relabel volumes that are deleted and returned to scratch.

Yes

Specifies that the server relabels volumes that are deleted and returned to scratch.

RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established. If, for example, a storage agent becomes unavailable but is still holding the path to a drive, persistent reserve allows the server to break the storage agent's reservation and access the drive.

Support for persistent reservation has the following limitations:

- If you are using the Tivoli Storage Manager device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 at <http://www.ibm.com/support/docview.wss?uid=swg21470319> for details.
- If you are using the IBM device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972> for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

Yes

Specifies that drive preemption through persistent reserve is used. YES is the default for a library that is defined with SHARED=YES. The **RESETDRIVES** parameter must be set to YES in a clustered environment or the drive is unavailable following failover.

No Specifies that drive preemption through persistent reserve is not used. NO is the default for a library that is defined with SHARED=NO.

Note: A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

SERial

Specifies the serial number for the library that is being defined. This parameter is optional. The default is AUTODETECT.

If SERIAL=AUTODETECT, then when you define the path to the library, the serial number reported by the library is used as the serial number.

If SERIAL=*serial_number*, then the number you entered is compared to the number detected by Tivoli Storage Manager.

Attention: Depending on the capabilities of the device, SERIAL=AUTODETECT might not be supported. In this case, the serial number is reported as blank.

Example: Define a SCSI library

Define a library named SCsilIB with a library type of SCSI.

```
define library scsilib libtype=scsi
```

The library requires a path. The device name for the library is:

```
/dev/tmsmcsi/lb0
```

Define the path:

```
define path server1 scsilib srctype=server desttype=library
device=/dev/tmsmcsi/lb0
```

DEFINE LIBRARY - SHARED

DEFINE LIBRARY (Define a shared library)

Use this syntax to define a shared library.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—DEFine LIBRary—library_name—LIBType=—SHARED—►►  
►—PRIMarylibmanager=—server_name—►
```

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=SHARED (Required)

Specifies that the library is shared with another Tivoli Storage Manager server over a storage area network (SAN) or a dual SCSI connection to library drives. This library type is not valid for optical devices.

Important: Specify this library type when defining the library on a library client.

PRIMarylibmanager

Specifies the name of the Tivoli Storage Manager server that is responsible for controlling access to library resources. You must define this server with the **DEFINE SERVER** command before you can use it as a library manager. This parameter is required and valid only if LIBTYPE=SHARED.

Example: Define a shared library

In a SAN, define a library named SHAREDTSM to a library client server named LIBMGR1

```
define library sharedtsm libtype=shared primarylibmanager=libmgr1
```


DEFINE LIBRARY (Define a VTL library)

Use this syntax to define a library that has a SCSI-controlled media changer device that is represented by a virtual tape library (VTL).

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Notes:

- 1 The default value of the **RESETDRIVES** parameter is conditional. If the **SHARED** parameter is set to NO, the value of the **RESETDRIVES** parameter is NO. If the **SHARED** parameter is set to YES, the value of the **RESETDRIVES** parameter is YES.

Parameters

library_name (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

LIBType=VTL (Required)

Specifies that the library has a SCSI-controlled media changer device that is represented by a virtual tape library. To mount volumes in drives in this type of library, Tivoli Storage Manager uses the media changer device.

If you are defining a VTL library, your environment must not include any mixed-media and paths must be defined between all drives in the library and all defined servers, including storage agents, that use the library. If either of these characteristics are not true, the overall performance can degrade to the same levels as the SCSI library type; especially during times of high stress.

SHARED

Specifies whether this library is shared with other Tivoli Storage Manager servers in a storage area network (SAN). This parameter is required when you define a library to the library manager.

YES

Specifies that this library can be shared with other servers. When you

DEFINE LIBRARY - VTL

specify YES, the library manager server mounts volumes as requested by other servers and tracks drive and volume allocation to other servers.

- NO** Specifies that this library cannot be shared with other servers. SHARED=NO is required if the library is controlled by passing commands through a NAS file server.

RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established. If, for example, a storage agent becomes unavailable but is still holding the path to a drive, persistent reserve allows the server to break the storage agent's reservation and access the drive.

Support for persistent reservation has the following limitations:

- If you are using the Tivoli Storage Manager device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 at <http://www.ibm.com/support/docview.wss?uid=swg21470319> for details.
- If you are using the IBM device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972> for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

Yes

Specifies that drive preemption through persistent reserve is used. YES is the default for a library that is defined with SHARED=YES. The **RESETDRIVES** parameter must be set to YES in a clustered environment or the drive is unavailable following failover.

- No** Specifies that drive preemption through persistent reserve is not used. NO is the default for a library that is defined with SHARED=NO.

Note: A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional. The default is NO.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the **CHECKIN LIBVOLUME** command.

- No** Specifies that the server does not attempt to label any volumes.

Yes

Specifies that the server labels only unlabeled volumes.

OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

RELABELScratch

Specifies whether the server relabels volumes that were deleted and returned to scratch. When this parameter is set to YES, a **LABEL LIBVOLUME** operation is started and the existing volume label is overwritten.

Note: If you have both virtual and real volumes in your VTL, both types are relabeled when this parameter is enabled. If the VTL includes real volumes, specifying this option might impact performance.

Yes

Specifies that the server relabels volumes that are deleted and returned to scratch. YES is the default.

No Specifies that the server does not relabel volumes that are deleted and returned to scratch.

SERial

Specifies the serial number for the library that is being defined. This parameter is optional. The default is AUTODETECT.

If SERIAL=AUTODETECT, then when you define the path to the library, the serial number reported by the library is used as the serial number.

If SERIAL=*serial_number*, then the number you entered is compared to the number detected by Tivoli Storage Manager.

Attention: Depending on the capabilities of the device, SERIAL=AUTODETECT might not be supported. In this case, the serial number is reported as blank.

Example: Define a VTL library

Define a library named VTLLIB with a library type of VTL.

```
define library vtllib libtype=vtl
```

The library requires a path. The device name for the library is:

```
/dev/tmscsi/lb0
```

Define the path:

```
define path server1 vtllib srctype=server desttype=library  
device=/dev/tmscsi/lb0
```

DEFINE LIBRARY - ZOSMEDIA

DEFINE LIBRARY (Define a ZOSMEDIA library type)

Use this syntax to define a library that represents a TAPE or FILE storage resource that is maintained by Tivoli Storage Manager for z/OS Media.

Define a library of type ZOSMEDIA when you want the library to be exclusively managed by Tivoli Storage Manager for z/OS Media. The library appears to the Tivoli Storage Manager server as a logical storage device that does not require DRIVE definitions. A PATH definition is required for the Tivoli Storage Manager server and any storage agents that need access to the ZOSMEDIA library resource.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—DEFine LIBRary—library_name—LIBType=—ZOSMEDIA—►►
```

Parameters

library_name (Required)

Specifies the name of the library to be defined.

LIBType=ZOSMEDIA (Required)

Specifies that the library type is the ZOSMEDIA which represents a TAPE or FILE storage resource that is maintained by Tivoli Storage Manager for z/OS Media.

Example: Configure a ZOSMEDIA library

The following example shows the steps needed to define and configure a zosmedia library. The configuration includes these components:

- A Tivoli Storage Manager server named sahara
- A library defined as type zosmedia named zebra
- A z/OS media server named oasis
- A storage agent named mirage

Define a library named ZEBRA with a library type of ZOSMEDIA:

```
define library zebra libtype=zosmedia
```

Define the z/OS media server:

```
define server oasis serverpassword=sanddune  
hladdress=9.289.19.67 lladdress=1777
```

The server requires a path to the library resource managed by Tivoli Storage Manager for z/OS Media:

```
define path sahara zebra srctype=server  
desttype=library zosmediaserver=oasis
```

The storage agent requires a path to the library resource managed by Tivoli Storage Manager for z/OS Media:

```
define path mirage zebra srctype=server  
desttype=library zosmediaserver=oasis
```


DEFINE MACHINE (Define machine information for disaster recovery)

Use this command to save disaster recovery information for a server or client node machine. This information will be included in the plan file to help you recover your machines.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

machine_name (Required)

Specifies the machine name. The name can be up to 64 characters.

DESCRIPTION

Specifies a machine description. This parameter is optional. The text can be up to 255 characters. Enclose the text in quotation marks if it contains any blank characters.

BUILDING

Specifies the building that this machine is in. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters.

FLOOR

Specifies the floor that this machine is on. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters.

ROOM

Specifies the room that this machine is in. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters.

PRIORITY

Specifies the restore priority for the machine an integer from 1 to 99. The highest priority is 1. This parameter is optional. The default is 50.

ADSMSEVER

Specifies whether the machine is a Tivoli Storage Manager server. Only one machine can be defined as a Tivoli Storage Manager server. This parameter is optional. The default is NO. Possible values are:

DEFINE MACHINE

No This machine is not a Tivoli Storage Manager server.

Yes

This machine is a Tivoli Storage Manager server.

Example: Define a machine's disaster recovery information

Define a machine named DISTRICT5, and specify a location, a floor, and a room name. This machine contains critical data and has the highest priority.

```
define machine district5 building=101 floor=27  
room=datafacilities priority=1
```

Related commands

Table 87. Commands related to DEFINE MACHINE

Command	Description
DEFINE MACHNODEASSOCIATION	Associates a Tivoli Storage Manager node with a machine.
DEFINE RECMEDMACHASSOCIATION	Associates recovery media with a machine.
DELETE MACHINE	Deletes a machine.
INSERT MACHINE	Inserts machine characteristics or recovery instructions into the Tivoli Storage Manager database.
QUERY MACHINE	Displays information about machines.
UPDATE MACHINE	Changes the information for a machine.

DEFINE MACHNODEASSOCIATION (Associate a node with a machine)

Use this command to associate client nodes with a machine. During disaster recovery, you can use this information to identify the client nodes that resided on destroyed machines.

The machine must be defined and the nodes registered to Tivoli Storage Manager.


To retrieve the information, issue the **QUERY MACHINE** command. This information will be included in the plan file to help you recover the client machines.

A node remains associated with a machine unless the node, the machine, or the association itself is deleted.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

machine_name (Required)

Specifies the machine name.

node_name (Required)

Specifies the node names. A node can only be associated with one machine. To specify multiple nodes, separate the names with commas and no intervening spaces. You can use wildcard characters to specify a name.

Example: Associate a node with a machine

Associate the node named ACCOUNTSPAYABLE with the machine named DISTRICT5.

```
define machnodeassociation district5 accountspayable
```

Related commands

Table 88. Commands related to **DEFINE MACHNODEASSOCIATION**

Command	Description
DEFINE MACHINE	Defines a machine for DRM.
DELETE MACHINE	Deletes a machine.
DELETE MACHNODEASSOCIATION	Deletes association between a machine and node.
QUERY MACHINE	Displays information about machines.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.

DEFINE MACHNODEASSOCIATION

DEFINE MGMTCLASS (Define a management class)

Use this command to define a new management class in a policy set. To allow clients to use the new management class, you must activate the policy set that contains the new class.


You can define one or more management classes for each policy set in a policy domain. A management class can contain a backup copy group, an archive copy group, or both. The user of a client node can select any management class in the active policy set or use the default management class.

Attention: The **DEFINE MGMTCLASS** command fails if a copy storage pool is specified as the destination for files that were migrated by a Tivoli Storage Manager for Space Management client.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the management class belongs.

Syntax



Parameters

domain_name (Required)

Specifies the policy domain to which the management class belongs.

policy_set_name (Required)

Specifies the policy set to which the management class belongs. You cannot define a management class to the **ACTIVE** policy set.

class_name (Required)

Specifies the name of the new management class. The maximum length of this name is 30 characters. You cannot use either *default* or *grace_period* as a class name.

SPACEMGTECHnique

Specifies whether a file using this management class is eligible for migration. This parameter is optional. The default is **NONE**. This parameter is effective

DEFINE MGMTCLASS

only for Tivoli Storage Manager for Space Management clients, not for backup-archive clients or application clients. Possible values are:

AUTOMATIC

Specifies that the file is eligible for both automatic migration and selective migration.

SELECTIVE

Specifies that the file is eligible for selective migration only.

NONE

Specifies that the file is not eligible for migration.

AUTOMIGNONUSE

Specifies the number of days that must elapse since a file was last accessed before it is eligible for automatic migration. This parameter is optional. The default value is 0. If SPACEMGTECHNIQUE is not AUTOMATIC, the server ignores this attribute. You can specify an integer from 0 to 9999.

This parameter is effective only for Tivoli Storage Manager for Space Management clients, not for backup-archive clients or application clients.

MIGREQUIRESBKUP

Specifies whether a backup version of a file must exist before a file can be migrated. This parameter is optional. The default is YES. This parameter is effective only for Tivoli Storage Manager for Space Management clients, not for backup-archive clients or application clients. Possible values are:

Yes

Specifies that a backup version must exist.

No Specifies that a backup version is optional.

MIGDESTINATION

Specifies the primary storage pool where the server initially stores files migrated by Tivoli Storage Manager for Space Management clients. This parameter is effective only for Tivoli Storage Manager for Space Management clients, and is not effective for backup-archive clients or application clients. The default is SPACEMGPOOL.

The command fails if you specify a copy storage pool as the destination.

DESCRIPTION

Specifies a description of the management class. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Example: Define a management class for a specific policy set and policy domain

Define a management class called MCLASS1 for policy set SUMMER in the PROG1 policy domain. For Tivoli Storage Manager for Space Management clients, allow both automatic and selective migration, and store migrated files in the SMPPOOL storage pool. Add the description, "Technical Support Mgmt Class."

```
define mgmtclass prog1 summer mclass1
spacemgtechnique=automatic migdestination=smpool
description="technical support mgmt class"
```

Related commands

Table 89. Commands related to DEFINE MGMTCLASS

Command	Description
ASSIGN DEFMGMTCLASS	Assigns a management class as the default for a specified policy set.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY MGMTCLASS	Displays information about management classes.
QUERY POLICYSET	Displays information about policy sets.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.
UPDATE MGMTCLASS	Changes the attributes of a management class.

DEFINE NODEGROUP (Define a node group)

Use this command to define a node group. A *node group* is a group of client nodes that are acted upon as if they were a single entity. A node can be a member of one or more node groups.

Privilege class

To issue this command, you must have system or unrestricted policy privilege.

Syntax

```
►►—DEfIne NODEGroup—group_name—[DESCRiption—==—description—]—►►
```

Parameters

group_name

Specifies the name of the node group that you want to create. The maximum length of the name is 64 characters. The specified name may not be the same as any existing client node name.

DESCRiption

Specifies a description of the node group. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Example: Define a node group

Define a node group named group1.

```
define nodegroup group1
```

Related commands

Table 90. Commands related to **DEFINE NODEGROUP**

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.


DEFINE NODEGROUPMEMBER (Define node group member)

Use this command to add a client node to a node group. A *node group* is a group of client nodes that are acted upon as if they were a single entity.

Privilege class

To issue this command you must have system or unrestricted policy privilege.

Syntax



Parameters

group_name

Specifies the name of the node group to which you want to add a client node.

node_name

Specifies the name of the client node that you want to add to the node group. You can specify one or more names. Separate multiple names with commas; do not use intervening spaces. You can also use wildcard characters when specifying multiple names.

Example: Define node group members

Define two members, node1 and node2, to a node group, group1.

```
define nodegroupmember group1 node1,node2
```

Related commands

Table 91. Commands related to **DEFINE NODEGROUPMEMBER**

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

DEFINE PATH (Define a path)

Use this command to define a path for a source to access a destination. Both the source and destination must be defined before you can define a path. For example, if a path is required between a server and a drive, you must first issue the **DEFINE DRIVE** command and then issue the **DEFINE PATH** command. A path must be defined after you issue the **DEFINE DRIVE** command in order to make the drive usable by Tivoli Storage Manager software.

Syntax and parameter descriptions are available for the following path types.

- “DEFINE PATH (Define a path when the destination is a drive)” on page 277
- “DEFINE PATH (Define a path when the destination is a library)” on page 282
- “DEFINE PATH (Define a path when the destination is a ZOSMEDIA library)” on page 285

For detailed and current device support information, see the Supported Devices website for your operating system:

http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html

Related commands

Table 92. Commands related to **DEFINE PATH**

Command	Description
DEFINE DATAMOVER	Defines a data mover to the Tivoli Storage Manager server.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE PATH	Deletes a path from a source to a destination.
PERFORM LIBACTION	Defines all drives and paths for a library.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE DATAMOVER	Changes the definition for a data mover.
UPDATE PATH	Changes the attributes associated with a path.


DEFINE PATH (Define a path when the destination is a drive)

Use this syntax when defining a path to a drive.

Privilege class

To issue this command you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

source_name (Required)

Specifies the name of source for the path. This parameter is required.

destination_name (Required)

Specifies the name of the destination. This parameter is required.

SRCType (Required)

Specifies the type of the source. This parameter is required. Possible values are:

DATAMover

Specifies that a data mover is the source.

SERVER

Specifies that a storage agent is the source.

AUTODetect

Specifies whether the serial number for a drive is automatically updated in the database at the time that the path is defined. This parameter is optional. This parameter is only valid for paths defined from the local server to a drive.

Possible values are:

No Specifies that the serial number is not automatically updated. The serial number is still compared with what is already in the database for the device. The server issues a message if there is a mismatch.

Yes

Specifies that the serial number is not automatically updated to reflect the same serial number that the drive reports to Tivoli Storage Manager.

DEFINE PATH - Destination is a drive

Important:

1. If you did not set the serial number when you defined the drive, the server always tries to detect the serial number, and AUTODETECT defaults to YES. If you have previously entered a serial number, then AUTODETECT defaults to NO.
2. The use of AUTODETECT=YES in this command means that the serial number set in the drive definition is updated with the detected serial number.
3. If you set DESTTYPE=DRIVE and AUTODETECT=YES, then the drive element number in the Tivoli Storage Manager database is automatically changed to reflect the same element number that corresponds to the serial number of that drive. This is true for drives in a SCSI library. For more information about the element number, see **DEFINE DRIVE**.
4. Depending on the capabilities of the device, the AUTODETECT parameter might not be supported.

DESTType=DRive (Required)

Specifies that a drive is the destination. When the destination is a drive, you must specify a library name.

LIBRARY

Specifies the name of the library to which the drive is assigned. The library and its drives must already be defined to the Tivoli Storage Manager server. If the path is from a NAS data mover to a library, the library must have LIBTYPE of SCSI, 349X, or ACSLS.

DEVIce

Specifies the name of the device as known to the source, or FILE if the device is a logical drive in a FILE library.

The source uses the device name to access the drive. See Table 93 for examples.

Table 93. Examples of device names

Source to destination	Example
Server to a drive (not a FILE drive)	/dev/tsm SCSI/mt3
Storage agent to a drive (not a FILE drive)	/dev/tsm SCSI/mt3
Storage agent to a drive when the drive is a logical drive in a FILE library	FILE
NAS data mover to a drive	NetApp NAS file server: rst01 EMC Celerra NAS file server: c436t011 IBM System Storage N Series: rst01

Important:

- For more complete information about device names when the source is a server, see the *Administrator's Guide*.
- For information about the device name when the source is a storage agent, see the *Storage Agent User's Guide*.
- For 349X libraries, the alias name is a symbolic name that is specified in the `/etc/ibmatl.conf` file. For more information, refer to the *IBM Tape Device Drivers Installation and User's Guide*. The guide can be downloaded from the

IBM Systems support site at the following URL: <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972>.

- For information about how to obtain names for devices that are connected to a NAS file server, consult the product information for the file server. For example, for a NetApp file server, connect to the file server using Telnet and issue the **SYSCONFIG** command. Use this command to determine device names for drives:

```
sysconfig -t
```

ONLine

Specifies whether the path is available for use. This parameter is optional. The default is YES. Possible values are:

Yes

Specifies that the path is available for use.

No Specifies that the path is not available for use.

The source and the destination must both be available to use the path.

For example, if the path from a data mover to a drive is online, but either the data mover or the drive is offline, you cannot use the path.

DIRectory

Specifies the directory location or locations where the storage agent reads and writes the files that represent storage volumes for the FILE device class that is associated with the FILE library. The DIRECTORY parameter is also used for devices of type REMOVABLEFILE. For REMOVABLEFILE devices, the DIRECTORY parameter provides information for the server (not a storage agent) along with the DRIVE parameter to describe access to the device. This parameter is optional.

For a path from a storage agent to a FILE device, this parameter is only valid when *all* of the following conditions are true:

- The source type is SERVER (meaning a storage agent that has been defined as a server to this server).
- The source name is the name of a storage agent, *not* the server.
- The destination is a logical drive that is part of a FILE library created when the device class was defined.

If you specified multiple directories for the device class associated with the FILE library, you must specify the same number of directories for each path to the FILE library. Do not change or move existing directories on the server that the storage agent is using so that the device class and the path remain synchronized. Adding directories is permitted. Specifying a mismatched number of directories can cause a run-time failure.

The default value for DIRECTORY is the directory of the server at the time the command is issued. The Windows registry is used to locate the default value.

Use a naming convention which you can use to associate the directory with a particular physical drive. This can help ensure that your configuration is valid for sharing the FILE library between the server and storage agent. If the storage agent is on a Windows system, use a universal naming convention (UNC) name. When the storage agent lacks permission to access remote storage, it experiences mount failures.

DEFINE PATH - Destination is a drive

Attention:

1. Storage agents access FILE volumes by replacing a directory name in a volume name with a directory name from a directory in the list provided with the **DEFINE PATH** command. Directories specified with this parameter are not validated on the Tivoli Storage Manager server.
2. Tivoli Storage Manager does not create shares or permissions, or mount the target file system. You must perform these actions before starting the storage agent.

Example: Define a path from a server to a drive

Define a path from a server to a drive. In this case, the server name is *NET1*, the drive name is *TAPEDRV6*, the library is *NETLIB*, and the device name is *mt4*. Set **AUTODETECT** to **NO**.

```
define path net1 tapedrv6 srctype=server autodetect=no desttype=drive
  library=netlib device=mt4
```

Example: Define a path from a data mover server to a drive for backup and restore

Define a path from the data mover that is a NAS file server to the drive that the NAS file server will use for backup and restore operations. In this example, the NAS data mover is *NAS1*, the drive name is *TAPEDRV3*, the library is *NASLIB*, and the device name for the drive is *rst0l*.

```
define path nas1 tapedrv3 srctype=datamover desttype=drive library=naslib
  device=rst0l
```

Example: Define a path from a storage agent to a drive for backup and restore

Define a path from storage agent *SA1* to the drive that the storage agent uses for backup and restore operations. In this example, the library is *TSMLIB*, the drive is *TAPEDRV4*, and the device name for the drive is */dev/tmsmcsi/mt3*.

```
define path sa1 tapedrv4 srctype=server desttype=drive library=tsmlib
  device=/dev/tmsmcsi/mt3
```

Example: Configure a storage agent to use a FILE library

The following example illustrates the importance of matching device classes and paths to ensure that storage agents can access newly created FILE volumes.

Suppose you want to use these three directories for a FILE library:

- /opt/tivoli1
- /opt/tivoli2
- /opt/tivoli3

1. Use the following command to set up a FILE library named **CLASSA** with one drive named **CLASSA1** on **SERVER1**:

```
define devclass classa devtype=file
  directory="/opt/tivoli1,/opt/tivoli2,/opt/tivoli3"
  shared=yes mountlimit=1
```

2. You want the storage agent **STA1** to be able to use the FILE library, so you define the following path for storage agent **STA1**:

```
define path sta1 classa1 srctype=server desttype=drive device=file
  directory="/opt/ibm1,/opt/ibm2,/opt/ibm3" library=classa
```

DEFINE PATH - Destination is a drive

In this scenario, the storage agent, STA1, replaces the directory name /opt/tivoli1 with the directory name /opt/ibm1/ to access FILE volumes that are in the /opt/tivoli1 directory on the server.

3. If file volume /opt/tivoli1/file1.dsm is created on SERVER1, and if the following command is issued,

```
update devclass classa directory="/opt/otherdir,/opt/tivoli2,  
/opt/tivoli3"
```

SERVER1 is still able to access file volume /opt/tivoli1/file1.dsm, but the storage agent STA1 is not able to access it because a matching directory name in the PATH directory list no longer exists. If a directory name is not available in the directory list associated with the device class, the storage agent can lose access to a FILE volume in that directory. Although the volume is still accessible from the Tivoli Storage Manager server for reading, failure of the storage agent to access the FILE volume can cause operations to be retried on a LAN-only path or to fail.

DEFINE PATH - Destination is a library


DEFINE PATH (Define a path when the destination is a library)

Use this syntax when defining a path to a library.

Privilege class

To issue this command you must have system privilege or unrestricted storage privilege.

Syntax



Notes:

- 1 DATAMOVER only applies to NAS devices.

Parameters

source_name (Required)

Specifies the name of source for the path. This parameter is required.

destination_name (Required)

Specifies the name of the destination. This parameter is required.

Attention: To define a path from a NAS data mover to a library, the library must have LIBTYPE of SCSI, 349x, or ACSLS.

SRCType (Required)

Specifies the type of the source. This parameter is required. Possible values are:

DATAMover

Specifies that a data mover is the source.

SERVer

Specifies that a storage agent is the source.

AUTODetect

Specifies whether the serial number for a drive or library will be automatically updated in the database at the time that the path is defined. This parameter is optional. This parameter is only valid for paths defined from the local server to a drive or a library. Possible values are:

No Specifies that the serial number will not be automatically updated. The serial number is still compared with what is already in the database for the device. The server issues a message if there is a mismatch.

Yes

Specifies that the serial number will be automatically updated to reflect the same serial number that the drive reports to Tivoli Storage Manager.

Important:

1. If you did not set the serial number when you defined the drive or the library, the server always tries to detect the serial number, and AUTODETECT defaults to YES. If you have previously entered a serial number, then AUTODETECT defaults to NO.
2. The use of AUTODETECT=YES in this command means that the serial number set in the drive or library definition is updated with the detected serial number.
3. Depending on the capabilities of the device, the AUTODETECT parameter may not be supported.

DESTType=LIBRARY (Required)

Specifies that a library is the destination. This parameter is required.

DEVICE

Specifies the name of the device as known to the source, or FILE if the device is a logical drive in a FILE library.

The source uses the device name to access the library. See Table 94 for examples.

Table 94. Examples of device names

Source to destination	Example
Server to a library	/dev/tmsmcsi/lb4
NAS data mover to a library	mc0

Important:

- For more complete information about device names when the source is a server, see the *Administrator's Guide*.
- For information about the device name when the source is a storage agent, see the *Storage Agent User's Guide*.
- For 349X libraries, the alias name is a symbolic name that is specified in the `/etc/ibmatl.conf` file. For more information, refer to the *IBM Tape Device Drivers Installation and User's Guide*. The guide can be downloaded from the IBM Systems support site at the following URL: <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972>.
- For information about how to obtain names for devices that are connected to a NAS file server, consult the product information for the file server. For example, for a NetApp file server, connect to the file server using Telnet and issue the **SYSCONFIG** command. Use this command to determine device names for drives:

```
sysconfig -t
```

Use this command to determine the device name for a library:

```
sysconfig -m
```

EXTERNALManager

Specifies the location of the external library manager where Tivoli Storage Manager can send media access requests. Use single quotation marks around the value of this parameter. For example, enter:

```
/opt/GESedt-acsls/bin/elmdt
```

This parameter is required when the library name is an external library.

DEFINE PATH - Destination is a library

ONLine

Specifies whether the path is available for use. This parameter is optional. The default is YES. Possible values are:

Yes

Specifies that the path is available for use.

No Specifies that the path is not available for use.

The source and the destination must both be available to use the path.

Attention: If the path to a library is offline, the server will not be able to access the library. If the server is halted and restarted while the path to the library is offline, the library will not be initialized. See the *Administrator's Guide* for additional information.

Example: Define a path from a server to a library

Define a path from the server SATURN to the SCSI type library SCSILIB:

```
define path saturn scsilib srctype=server  
desttype=library device=/dev/tmsmcsi/lb3
```


DEFINE PATH (Define a path when the destination is a ZOSMEDIA library)

Use this syntax when defining a path to a ZOSMEDIA library. You must first define the z/OS media server in your configuration with the **DEFINE SERVER** command.

Privilege class

To issue this command you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

source_name (Required)

Specifies the name of source for the path.

destination_name (Required)

Specifies the name of the ZOSMEDIA library.

SRCType=SERVer (Required)

Specifies that a storage agent or server is the source.

DESTType=LIBRary (Required)

Specifies that a library is the destination.

ZOSMEDIAServer (Required)

Specifies the name of the server that represents a Tivoli Storage Manager for z/OS Media server.

ONLine

Specifies whether the path is available for use. This parameter is optional. The default is YES. Possible values are:

Yes

Specifies that the path is available for use.

No

Specifies that the path is not available for use.

The source and the destination must both be available to use the path.

Attention: If the path to a library is offline, the server will not be able to access the library. If the server is halted and restarted while the path to the library is offline, the library will not be initialized.

If the z/OS media server cannot be accessed during initialization of the Tivoli Storage Manager server, the library path will be set offline. Use the **UPDATE PATH** command and specify **ONLINE=YES** to vary the ZOSMEDIA library back online.

DEFINE POLICYSET (Define a policy set)

Use this command to define a policy set in a policy domain. A policy set contains management classes, which contain copy groups. You can define one or more policy sets for each policy domain.


To put a policy set into effect, you must activate the policy set by using the **ACTIVATE POLICYSET** command. Only one policy set can be active in a policy domain. The copy groups and management classes within the active policy set determine the rules by which client nodes perform backup, archive, and space management operations, and how the client files stored are managed.

Use the **VALIDATE POLICYSET** command to verify that a policy set is complete and valid before activating it with the **ACTIVATE POLICYSET** command.

Privilege class

To issue this command you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax



Parameters

domain_name (Required)

Specifies the name of the policy domain to which the policy set belongs.

policy_set_name (Required)

Specifies the name of the policy set. The maximum length of this name is 30 characters. You cannot define a policy set named ACTIVE.

DESCRIPTION

Specifies a description for the new policy set. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Example: Define a policy set

Define a policy set called SUMMER for the PROG1 policy domain and include the description, "Programming Group Policies."

```

define policyset prog1 summer
description="Programming Group Policies"
  
```

Related commands

Table 95. Commands related to DEFINE POLICYSET

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY MGMTCLASS	Creates a copy of a management class.
COPY POLICYSET	Creates a copy of a policy set.

*Table 95. Commands related to **DEFINE POLICYSET** (continued)*

Command	Description
DEFINE DOMAIN	Defines a policy domain that clients can be assigned to.
DEFINE MGMTCLASS	Defines a management class.
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.
QUERY POLICYSET	Displays information about policy sets.
UPDATE POLICYSET	Changes the description of a policy set.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

DEFINE PROFASSOCIATION (Define a profile association)

Use this command on a configuration manager to associate one or more objects with a configuration profile for distribution to subscribing managed servers. After a managed server subscribes to a profile, the configuration manager sends object definitions associated with the profile to the managed server where they are stored in the database. Objects created this way in the database of a managed server become managed objects. An object can be associated with more than one profile.

You can use this command to define an initial set of profile associations and to add to existing associations.

You can associate the following types of objects with a profile:


- Administrator registrations and authorities
- Policy domains, which include the domains' policy sets, management classes, copy groups, and client schedules
- Administrative schedules
- Server command scripts
- Client option sets
- Server definitions
- Server group definitions


Tip: The configuration manager does not distribute status information for an object to managed servers. For example, information such as the number of days since an administrator last accessed the server is not distributed to managed servers. This type of information is maintained in the databases of the individual managed servers.

Privilege class

To issue this command, you must have system privilege.

Syntax





Parameters

profile_name (Required)

Specifies the name of the configuration profile.

ADMins

Specifies administrators to associate with the profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk (*) by itself, to specify all administrators that are registered with the configuration manager. If you specify the match-all definition and later add more administrators, they are automatically distributed through the profile.

The configuration manager distributes the administrator name, password, contact information, and authorities of administrators associated with the profile. The configuration manager does not distribute the following:

- The administrator named SERVER_CONSOLE, even if you use a match-all definition
- The locked or unlocked status of an administrator

When the profile already has administrators associated with it, the following apply:

- If you specify a list of administrators and a list already exists, Tivoli Storage Manager combines the new list with the existing list.
- If you specify a match-all definition and a list of administrators already exists, Tivoli Storage Manager replaces the list with the match-all definition.
- If you specify a list of administrators, and a match-all definition had previously been specified, Tivoli Storage Manager ignores the list. To remove the match-all definition, issue the DELETE PROFASSOCIATION command with the ADMINS=* parameter.

D0mains

Specifies policy domains to associate with the profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk (*) by itself, to specify all domains that are defined on the configuration manager. If you specify the match-all definition and later add more domains, they are automatically distributed through the profile.

The configuration manager distributes domain information that includes definitions of policy domains, policy sets, management classes, copy groups, and client schedules. The configuration manager does not distribute the ACTIVE policy set. Administrators on a managed server can activate any policy set within a managed domain on a managed server.

DEFINE PROFASSOCIATION

When the profile already has domains associated with it, the following apply:

- If you specify a list of domains and a list already exists, Tivoli Storage Manager combines the new list with the existing list.
- If you use a match-all definition and a list of domains already exists, Tivoli Storage Manager replaces the list with the match-all definition.
- If you specify a list of domains, and a match-all definition had previously been specified, Tivoli Storage Manager ignores the list. To remove the match-all definition, issue the **DELETE PROFASSOCIATION** command with the **DOMAINS=*** parameter.

Important: Client operations such as backup and archive fail if destination pools do not exist. Therefore, managed servers that subscribe to this profile must have definitions for any storage pools specified as destinations in the associated domains. Use the **RENAME STGPOOL** command to rename existing storage pools to match the destination names distributed.

ADSCHeds

Specifies administrative schedules to associate with the profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk (*) by itself, to specify all administrative schedules that are defined on the configuration manager. If you specify the match-all definition and later add more administrative schedules, they are automatically distributed through the profile.

Tip: Administrative schedules are not active when they are distributed by a configuration manager. An administrator on a managed server must activate any schedule to have it run on that server.

When the profile already has administrative schedules associated with it, the following apply:

- If you specify a list of administrative schedules and a list already exists, Tivoli Storage Manager combines the new list with the existing list.
- If you use a match-all definition and a list of administrative schedules already exists, Tivoli Storage Manager replaces the list with the match-all definition.
- If you specify a list of administrative schedules, and a match-all definition had previously been specified, Tivoli Storage Manager ignores the list. To remove the match-all definition, issue the **DELETE PROFASSOCIATION** command with the **ADSCHEDULES=*** parameter.

SCRipts

Specifies server command scripts to associate with the profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk (*) by itself, to specify all scripts that are defined on the configuration manager. If you specify the match-all definition and later add more scripts, they are automatically distributed through the profile.

When the profile already has scripts associated with it, the following apply:

- If you specify a list of scripts and a list already exists, Tivoli Storage Manager combines the new list with the existing list.
- If you use a match-all definition and a list of scripts already exists, Tivoli Storage Manager replaces the list with the match-all definition.

- If you specify a list of scripts, and a match-all definition had previously been specified, Tivoli Storage Manager ignores the list. To remove the match-all definition, issue the **DELETE PROFASSOCIATION** command with the **SCRIPTS=*** parameter.

CLOptsets

Specifies client option sets to associate with the profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk (*) by itself, to specify all client option sets that are defined on the configuration manager. If you specify the match-all definition and later add more client option sets, they are automatically distributed through the profile.

When the profile already has client option sets associated with it, the following apply:

- If you specify a list of client option sets and a list already exists, Tivoli Storage Manager combines the new list with the existing list.
- If you use a match-all definition and a list of client option sets already exists, Tivoli Storage Manager replaces the list with the match-all definition.
- If you specify a list of client option sets, and a match-all definition had previously been specified, Tivoli Storage Manager ignores the list. To remove the match-all definition, issue the **DELETE PROFASSOCIATION** command with the **CLOPSETS=*** parameter.

SERVERs

Specifies server definitions to associate with the profile. The definitions are distributed to managed servers that subscribe to this profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk (*) by itself, to specify all servers that are defined on the configuration manager. If you specify the match-all definition and later add more servers, they are automatically distributed through the profile.

The configuration manager distributes the following server attributes: communication method, IP address, port address, server password, URL, and the description. Distributed server definitions always have the **ALLOWREPLACE** attribute set to **YES** on the managed server, regardless of this parameter's value on the configuration manager. On the managed server, you can use the **UPDATE SERVER** command to set all other attributes.

When the profile already has servers associated with it, the following apply:

- If you specify a list of servers and a list already exists, Tivoli Storage Manager combines the new list with the existing list.
- If you use a match-all definition and a list of servers already exists, Tivoli Storage Manager replaces the list with the match-all definition.
- If you specify a list of servers, and a match-all definition had previously been specified, Tivoli Storage Manager ignores the list. To remove the match-all definition, issue the **DELETE PROFASSOCIATION** command with the **SERVERS=*** parameter.

Important:

1. A server definition on a managed server is not replaced by a definition from the configuration manager unless you have allowed replacement of the definition on the managed server. To allow replacement, on the

DEFINE PROFASSOCIATION

managed server update the server definition by using the **UPDATE SERVER** command with **ALLOWREPLACE=YES**.

2. If a configuration manager distributes a server definition to a managed server, and a server group of the same name exists on the managed server, the distributed server definition replaces the server group definition.

SERVERGroups

Specifies server groups to associate with the profile. You can use wildcard characters in the names. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all definition, an asterisk (*) by itself, to specify all server groups that are defined on the configuration manager. If you specify the match-all definition and later add more server groups, they are automatically distributed through the profile.

Tip: A configuration manager does not distribute a server group definition to a managed server if the managed server has a server defined with the same name as that of the server group.

When the profile already has server groups associated with it, the following apply:

- If you specify a list of server groups and a list already exists, Tivoli Storage Manager combines the new list with the existing list.
- If you use a match-all definition and a list of server groups already exists, Tivoli Storage Manager replaces the list with the match-all definition.
- If you specify a list of server groups, and a match-all definition had previously been specified, Tivoli Storage Manager ignores the list. To remove the match-all definition, issue the **DELETE PROFASSOCIATION** command with the **SERVERGROUPS=*** parameter.

Example: Associate a specific domain with a specific profile

Associate a domain named **MARKETING** with a profile named **DELTA**.

```
define profassociation delta domains=marketing
```

Example: Associate all domains with a specific profile

You have already associated a list of domains with a profile named **GAMMA**. Now associate all domains defined on the configuration manager with the profile.

```
define profassociation gamma domains=*
```

Related commands

Table 96. Commands related to **DEFINE PROFASSOCIATION**

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE PROFILE	Deletes a profile from a configuration manager.
LOCK PROFILE	Prevents distribution of a configuration profile.

Table 96. Commands related to DEFINE PROFASSOCIATION (continued)

Command	Description
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY PROFILE	Displays information about configuration profiles.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

DEFINE PROFILE (Define a profile)

Use this command on a configuration manager to define a profile (a set of configuration information) that can be distributed to managed servers.

After defining a profile, you can use the **DEFINE PROFASSOCIATION** command to specify objects to be distributed to managed servers subscribing to the profile.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—DEfIne PROFIle—profile_name—[DESCRiption—==—description—]—►►
```

Parameters

profile_name (Required)

Specifies the name of the profile. The maximum length of the name is 30 characters.

DESCRiption

Specifies a description of the profile. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. This parameter is optional.

Example: Define a new profile

Define a profile named ALPHA with a description of "Programming Center."

```
define profile alpha
description="Programming Center"
```

Related commands

Table 97. Commands related to DEFINE PROFILE

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE PROFILE	Deletes a profile from a configuration manager.
LOCK PROFILE	Prevents distribution of a configuration profile.
QUERY PROFILE	Displays information about configuration profiles.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.

Table 97. Commands related to DEFINE PROFILE (continued)

Command	Description
UPDATE PROFILE	Changes the description of a profile.

DEFINE RECMEDMACHASSOCIATION (Associate recovery media with a machine)


Use this command to associate recovery media with one or more machines. A machine is associated with recovery media so that the location of the boot media and its list of volume names are available to recover the machine. To retrieve the information, issue the **QUERY MACHINE** command. This information will be included in the plan file to help you recover the client machines.

To associate a machine with recovery media, both the machine and media must be defined to Tivoli Storage Manager. A machine remains associated with the media until the association, the media, or the machine is deleted.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

media_name (Required)

Specifies the name of the recovery media with which one or more machines will be associated.

machine_name (Required)

Specifies the name of the machines to be associated with the recovery media. A machine can be associated with multiple recovery media. To specify a list of machines, separate the names with commas and no intervening spaces. You can use wildcard characters to specify a name.

Example: Associate machines to recovery media

Associate machines DISTRICT1 and DISTRICT5 to the DIST5RM recovery media.

```

define recmedmachassociation dist5rm
district1,district5
  
```

Related commands

Table 98. Commands related to **DEFINE RECMEDMACHASSOCIATION**

Command	Description
DEFINE MACHINE	Defines a machine for DRM.
DEFINE RECOVERYMEDIA	Defines the media required to recover a machine.
DELETE MACHINE	Deletes a machine.
DELETE RECMEDMACHASSOCIATION	Deletes association between recovery media and a machine.
DELETE RECOVERYMEDIA	Deletes recovery media.
QUERY MACHINE	Displays information about machines.

DEFINE RECMEDMACHASSOCIATION

Table 98. Commands related to **DEFINE RECMEDMACHASSOCIATION** (continued)

Command	Description
QUERY RECOVERYMEDIA	Displays media available for machine recovery.

DEFINE RECOVERYMEDIA


DEFINE RECOVERYMEDIA (Define recovery media)

Use this command to define the media needed to recover a machine. The same media can be associated with multiple machines. To display the information, use the **QUERY MACHINE** command. This information will be included in the plan file to help you to recover the client machines.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

media_name (Required)

Specifies the name of the recovery media to be defined. The name can be up to 30 characters.

VOLumentnames

Specifies the names of volumes that contain the recoverable data (for example, operating system image copies). This parameter is required if you specify a media type of `BOOT`. Specify boot media volume names in the order in which they are to be inserted into the machine at recovery time. The maximum length of the volume names list is 255 characters. Enclose the list in quotation marks if it contains any blank characters.

DESCRiption

Specifies the description of the recovery media. This parameter is optional. The maximum length is 255 characters. Enclose the text in quotation marks if it contains any blank characters.

L0cation

Specifies the location of the recovery media. This parameter is optional. The maximum length is 255 characters. Enclose the text in quotation marks if it contains any blank characters.

Type

Specifies the type of recovery media. This parameter is optional. The default is `OTHER`.

BOot

Specifies that this is boot media. You must specify volume names if the type is BOOT.

OTHer

Specifies that this is not boot media. For example, a CD that contains operating system manuals.

PRoduct

Specifies the name of the product that wrote to this media. This parameter is optional. The maximum length is 16 characters. Enclose the text in quotation marks if it contains any blank characters.

PRoDUCTInfo

Specifies information about the product that wrote to the media. This would be information that you may need to restore the machine. This parameter is optional. The maximum length is 255 characters. Enclose the text in quotation marks if it contains any blank characters.

Example: Define the media needed to recover a machine

Define the recovery media named DIST5RM. Include a description and the location.

```
define recoverymedia dist5rm
description="district 5 base system image"
location="district 1 vault"
```

Related commands

Table 99. Commands related to DEFINE RECOVERYMEDIA

Command	Description
DEFINE RECMEDMACHASSOCIATION	Associates recovery media with a machine.
DELETE RECOVERYMEDIA	Deletes recovery media.
QUERY RECOVERYMEDIA	Displays media available for machine recovery.
UPDATE RECOVERYMEDIA	Changes the attributes of recovery media.

DEFINE SCHEDULE (Define a client or an administrative command schedule)

Use this command to create a client or administrative command schedule.

The **DEFINE SCHEDULE** command takes two forms: one if the schedule applies to client operations, one if the schedule applies to administrative commands. Within these two forms, you can select either classic or enhanced style schedules. The syntax and parameters for each form are defined separately.

- “DEFINE SCHEDULE (Define a schedule for an administrative command)” on page 314
- “DEFINE SCHEDULE (Define a client schedule)” on page 301

For each schedule, a startup window is specified. The startup window is the time period during which the schedule must be initiated. The schedule will not necessarily complete processing within this window. If the server is not running when this window starts, but is started before the end of the defined window is reached, the schedule will run when the server is restarted. Options associated with each schedule style (classic and enhanced) determine when the startup windows should begin.

Table 100. Commands related to DEFINE SCHEDULE

Command	Description
COPY SCHEDULE	Creates a copy of a schedule.
DEFINE ASSOCIATION	Associates clients with a schedule.
DELETE SCHEDULE	Deletes a schedule from the database.
QUERY EVENT	Displays information about scheduled and completed events for selected clients.
QUERY SCHEDULE	Displays information about schedules.
SET MAXCMDRETRIES	Specifies the maximum number of retries after a failed attempt to execute a scheduled command.
SET MAXSCHEDESESSIONS	Specifies the maximum number of client/server sessions available for processing scheduled work.
SET RETRYPERIOD	Specifies the time between retry attempts by the client scheduler.
UPDATE SCHEDULE	Changes the attributes of a schedule.

DEFINE SCHEDULE (Define a client schedule)

Use the **DEFINE SCHEDULE** command to define a client schedule. Tivoli Storage Manager uses this schedule to automatically perform a variety of client operations for your client workstation at specified intervals or days. After you define a schedule, use the **DEFINE ASSOCIATION** command to associate the client with the schedule.

You must start the client scheduler on the client workstation for Tivoli Storage Manager to process the schedule.

Not all clients can run all scheduled operations, even though you can define the schedule on the server and associate it with the client. For example, a Macintosh client cannot run a schedule when the action is to restore or retrieve files, or run an executable script. An executable script is also known as a command file, a batch file, or a script on different client operating systems.


Tivoli Storage Manager cannot run multiple schedules concurrently for the same client node.

Privilege class


To define a client schedule, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the schedule belongs.


Syntax

Classic client schedule



DEFINE SCHEDULE






Notes:

|
|


- 1 You can define or update a **SYSTEMSTATE** or **VM** subaction on the Administration Center but these actions are not saved.
- 2 The **OBJECTS** parameter is optional when ACTION=INCREMENTAL, but is required for other actions.

Syntax

Enhanced client schedule



DEFINE SCHEDULE



Notes:

- 1 The **OBJECTS** parameter is optional when `ACTION=INCREMENTAL`, but is required for other actions.

Parameters

domain_name (Required)

Specifies the name of the policy domain to which this schedule belongs.

schedule_name (Required)

Specifies the name of the schedule to be defined. You can specify up to 30 characters for the name.

Type=Client

Specifies that a schedule for a client is defined. This parameter is optional.

DEscription

Specifies a description of the schedule. This parameter is optional. You can specify up to 255 characters for the description. Enclose the description in quotation marks if it contains any blank characters.

ACTion

Specifies the action that occurs when this schedule is processed. Possible values are:

Incremental

Specifies that the schedule backs up all files that are new or that have changed since the last incremental backup. Incremental also backs up any file for which all existing backups might have expired.

Selective

Specifies that the schedule backs up only files that are specified with the OBJECTS parameter.

Archive

Specifies that the schedule archives files that are specified with the OBJECTS parameter.

Backup

Specifies that the schedule backs up files that are specified with the OBJECTS parameter.

REStore

Specifies that the schedule restores files that are specified with the OBJECTS parameter.

When you specify ACTION=RESTORE for a scheduled operation, and the REPLACE option is set to PROMPT, no prompting occurs. If you set the option to PROMPT, the files are skipped.

If you specify a second file specification, this second file specification acts as the restore destination. If you need to restore multiple groups of files, schedule one for each file specification that you need to restore.

REtrieve

Indicates that the schedule retrieves files that are specified with the OBJECTS parameter.

Remember: A second file that is specified acts as the retrieve destination. If you need to retrieve multiple groups of files, create a separate schedule for each group of files.

IMAGEBACKup

Specifies that the schedule backs up logical volumes that are specified with the OBJECTS parameter.

IMAGERESore

Specifies that the schedule restores logical volumes that are specified with the OBJECTS parameter.

Command

Specifies that the schedule processes a client operating system command or script that is specified with the OBJECTS parameter.

DEFINE SCHEDULE

Macro

Specifies that a client processes a macro whose file name is specified with the OBJECTS parameter.

SUBACTioN

Possible values are:

"" When a null string (two double quotes) is specified with **ACTION=BACKUP** the backup is an incremental.

FASTBacK

Specifies that a FastBack client operation that is identified by the ACTION parameter is to be scheduled for processing. The ACTION parameter must be either ARCHIVE or BACKUP.

SYSTEMState

Specifies that a client Systemstate backup is scheduled.

VM Specifies that a client VMware backup operation is scheduled.

Deploy

Specifies whether to update client workstations with deployment packages that are specified with the **OBJECTS** parameter. The **OBJECTS** parameter must contain two specifications, the package files to retrieve and the location from which to retrieve them. Ensure that the objects are in the order *files location*. For example:

```
define schedule standard deploy_1 action=DEPLOY objects=  
"\\IBM_ANR_WIN\c$\tsm\maintenance\client\v6r2\Windows\X32\v620\v6200\*  
..\IBM_ANR_WIN"
```

Values for the following options are restricted when you specify **ACTION=DEPLOY**:

PERUNITS

Specify **PERUNITS=ONETIME**. If you specify **PERUNITS=PERIOD**, the parameter is ignored.

DURUNITS

Specify **MINUTES**, **HOURS**, or **DAYS** for the **DURUNITS** parameter. Do not specify **INDEFINITE**.

SCHEDSTYLE

Specify the default style, **CLASSIC**.

The **SCHEDULE** command fails if the parameters do not conform to the required parameter values, such as the V.R.M.F.

OPTioNs

Specifies the client options that you specify to the scheduled command at the time the schedule is processed. This parameter is optional.

Only those options that are valid on the scheduled command can be specified for this parameter. Refer to the appropriate client manual for information about options that are valid from the command line. All options described there as valid only on the initial command line result in an error or are ignored when running the schedule from the server. For example, do not include the following options because they have no impact when the client processes the scheduled command:

MAXCMDRETRIES

OPTFILE

QUERYSCHEDPERIOD
 RETRYPERIOD
 SCHEDLOGNAME
 SCHEDMODE
 SERVERNAME
 TCPCLIENTADDRESS
 TCPCLIENTPORT

If the option string contains multiple options or options with embedded spaces, surround the entire option string with one pair of apostrophes. Enclose individual options that contain spaces in quotation marks. A leading minus sign is required in front of the option. Errors can occur if the option string contains spaces that are not quoted correctly.

The following examples show how to specify some client options:

- To specify `subdir=yes` and `domain all-local -systemobject`, enter:
`options='-subdir=yes -domain="all-local -c: -systemobject"'`
- To specify `domain all-local -c: -d:`, enter:
`options='-domain="all-local -c: -d:"'`

OBjects

Specifies the objects for which the specified action is performed. Use a single space between each object. This parameter is required except when `ACTION=INCREMENTAL`. If the action is a backup, archive, retrieve, or restore operation, the objects are file spaces, directories, or logical volumes. See the *Backup-Archive Clients Installation and User's Guide* for command syntax information. If the action is to run a command or macro, the object is the name of the command or macro to run.

When you specify `ACTION=INCREMENTAL` without specifying a value for this parameter, the scheduled command is invoked without specified objects and attempts to process the objects as defined in the client option file. To select all file spaces or directories for an action, explicitly list them in the object string. Entering only an asterisk in the object string causes the backup to occur only for the directory where the scheduler was started.

Important:

- If you specify a second file specification, and it is not a valid destination, you receive this error:
`ANS1082E Invalid destination file specification <filespec> entered.`
- If you specify more than two file specifications, you receive this error:
`ANS1102E Excessive number of command line arguments passed to the program!`

When you specify `ACTION=ARCHIVE`, `INCREMENTAL`, or `SELECTIVE` for this parameter, you can list a maximum of twenty (20) file specifications.

Enclose the object string in double quotes if it contains blank characters (spaces), and then surround the double quotes with single quotes. If the object string contains multiple file names, enclose each file name with its own pair of double quotes, then surround the entire string with one pair of single quotes. Errors can occur if file names contain a space that is not quoted correctly.

The following examples show how to specify some file names:

- To specify `/home/file 2`, `/home/gif files`, and `/home/my test file`, enter:
`OBJECTS='"/home/file 2" "/home/gif files" "/home/my test file"'`

DEFINE SCHEDULE

- To specify /home/test file, enter:

```
OBJECTS="/home/test file"
```

PRIority

Specifies the priority value for a schedule. This parameter is optional. You can specify an integer from 1 to 10, with 1 being the highest priority and 10 being the lowest. The default is 5.

If two or more schedules have the same window start time, the value you specify determines when Tivoli Storage Manager processes the schedule. The schedule with the highest priority starts first. For example, a schedule with PRIORITY=3 starts before a schedule with PRIORITY=5.

STARTDate

Specifies the date for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current date. Use this parameter with the **STARTTIME** parameter to specify when the initial startup window of the schedule starts.

You can specify the date using one of the values below:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+days or <i>+days</i>	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

STARTTime

Specifies the time for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current time. This parameter is used in conjunction with the **STARTDATE** parameter to specify when the initial startup window begins.

You can specify the time using one of the values below:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08
NOW	The current time	NOW

Value	Description	Example
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified	NOW+02:00 or +02:00 . If you issue this command at 5:00 with STARTTIME=NOW+02:00 or STARTTIME=+02:00 , the beginning of the startup window is at 7:00.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified	NOW-02:00 or -02:00 . If you issue this command at 5:00 with STARTTIME=NOW-02:00 or STARTTIME=-02:00 , the beginning of the startup window is at 3:00.

DURation

Specifies the number of units that define the length of the startup window of the scheduled operation. This parameter is optional. This value must be from 1 to 999. The default is 1.

Use this parameter with the **DURUNITS** parameter to specify the length of the startup window. For example, if you specify **DURATION=20** and **DURUNITS=MINUTES**, the schedule must be started within 20 minutes of the start date and start time. The default length of the startup window is 1 hour. The duration of the window must be shorter than the period between windows.

This value is ignored if you specify **DURUNITS=INDEFINITE**.

Tip: Define schedules with durations longer than 10 minutes. Doing this will give the Tivoli Storage Manager scheduler enough time to process the schedule and prompt the client.

DURUnits

Specifies the time units used to determine the duration of the window in which the schedule can start. This parameter is optional. The default is **HOURS**.

Use this parameter with the **DURATION** parameter to specify how long the startup window remains open to process the schedule. For example, if **DURATION=20** and **DURUNITS=MINUTES**, the schedule must be started within 20 minutes of the start date and start time. The schedule may not necessarily complete processing within this window. If the schedule needs to be retried for any reason, the retry attempts must begin before the startup window elapses, or the operation does not restart.

The default value for the length of the startup window is 1 hour. Possible values are:

Minutes

Specifies that the duration of the window is defined in minutes.

Hours

Specifies that the duration of the window is defined in hours.

Days

Specifies that the duration of the window is defined in days.

INDefinite

Specifies that the startup window of the scheduled operation has an indefinite duration. The schedule can run any time after the scheduled

DEFINE SCHEDULE

start time, until the schedule expires. You cannot specify `DURUNITS=INDEFINITE`, unless you specify `PERUNITS=ONETIME`. The `INDEFINITE` value is not allowed with enhanced schedules.

SCHEDStyle

This parameter is optional. `SCHEDSTYLE` defines either the interval between times when a schedule can run, or the days on which it runs. The default is the classic syntax.

Possible values are:

Classic

The parameters for the Classic syntax are: `PERIOD`, `PERUNITS`, and `DAYOFWEEK`. You cannot use these parameters: `MONTH`, `DAYOFMONTH`, and `WEEKOFMONTH`.

Enhanced

The parameters for the Enhanced syntax are: `MONTH`, `DAYOFMONTH`, `WEEKOFMONTH`, and `DAYOFWEEK`. You cannot use these parameters: `PERIOD` and `PERUNITS`.

PERiod

Specifies the length of time between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. You can specify an integer from 1 to 999. The default is 1.

Use this parameter with the **PERUNITS** parameter to specify the period between startup windows. For example, if you specify `PERIOD=5` and `PERUNITS=DAYS` (assuming that `DAYOFWEEK=ANY`), the operation is scheduled every five days after the initial start date and start time. The period between startup windows must exceed the duration of each window. The default is 1 day.

This value is ignored if you specify `PERUNITS=ONETIME`.

PERUnits

Specifies the time units used to determine the period between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. The default is `DAYS`.

Use this parameter with the **PERIOD** parameter to specify the period between startup windows. For example, if you specify `PERIOD=5` and `PERUNITS=DAYS` (assuming that `DAYOFWEEK=ANY`), the operation is scheduled every 5 days after the initial start date and start time. The default is 1 day. Possible values are:

Hours

Specifies that the time between startup windows is in hours.

Days

Specifies that the time between startup windows is in days.

Weeks

Specifies that the time between startup windows is in weeks.

Months

Specifies that the time between startup windows is in months.

When you specify `PERUNITS=MONTHS`, the scheduled operation will be processed each month on the same date. For example, if the start date for the scheduled operation is 02/04/1998, the schedule will process on the 4th of every month thereafter. However, if the date is not valid for the next month, then the scheduled operation will be processed on the last valid

date in the month. Thereafter, subsequent operations are based on this new date. For example, if the start date is 03/31/1998, the next month's operation will be scheduled for 04/30/1998. Thereafter, all subsequent operations will be on the 30th of the month until February. Because February has only 28 days, the operation will be scheduled for 02/28/1999. Subsequent operations will be processed on the 28th of the month.

Years

Specifies that the time between startup windows for the schedule is in years.

When you specify PERUNITS=YEARS, the scheduled operation will be processed on the same month and date of each year. For example, if the start date for the scheduled operation is 02/29/2004, the next year's scheduled operation will be 02/28/2005 because February only has 28 days. Thereafter, subsequent operations will be scheduled for February 28th.

Onetime

Specifies that the schedule processes once. This value overrides the value you specified for the **PERIOD** parameter.

DAYofweek

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can specify different options for the **DAYofweek** parameter, depending on whether the schedule style has been defined as Classic or Enhanced:

Classic Schedule

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can either specify one day of the week, or WEEKDAY, WEEKEND, or ANY. If the start date and start time fall on a day that does not correspond to a day you specify, the start date and start time will be shifted forward in 24-hour increments until the **DAYOFWEEK** parameter is satisfied.

If you select a value for **DAYOFWEEK** other than ANY, and depending on the values for PERIOD and PERUNITS, schedules may not be processed when you would expect. The default is ANY.

Enhanced Schedule

Specifies the days of the week on which to run the schedule. You can either specify multiple days separated by commas and no intervening blanks, or WEEKDAY, WEEKEND, or ANY. If you specify multiple days, the schedule will run on each of the specified days. If you specify WEEKDAY or WEEKEND, you must also specify either WEEKOFMONTH=FIRST or WEEKOFMONTH=LAST, and the schedule will run just once per month.

The default value is ANY, meaning the schedule will run every day of the week or on the day or days determined by other enhanced schedule parameters. **DAYOFWEEK** must have a value of ANY (either by default or specified with the command) when used with the **DAYOFMONTH** parameter.

Possible values for the **DAYofweek** parameter are:

ANY

Specifies that the startup window can begin on any day of the week.

DEFINE SCHEDULE

WEEKDay

Specifies that the startup window can begin on Monday, Tuesday, Wednesday, Thursday, or Friday.

WEEKEnd

Specifies that the startup window can begin on Saturday or Sunday.

SUnday

Specifies that the startup window begins on Sunday.

Monday

Specifies that the startup window begins on Monday.

Tuesday

Specifies that the startup window begins on Tuesday.

Wednesday

Specifies that the startup window begins on Wednesday.

Thursday

Specifies that the startup window begins on Thursday.

Friday

Specifies that the startup window begins on Friday.

SAaturday

Specifies that the startup window begins on Saturday.

MONth

Specifies the months of the year during which to run the schedule. This parameter is used only with enhanced schedules. Specify multiple values by using commas and no intervening blanks. The default value is ANY, which means that the schedule runs during every month of the year.

DAYOFMonth

Specifies the day of the month to run the schedule. This parameter is used only with enhanced schedules. You can either specify ANY or a number from -31 through 31, excluding zero. Negative values are a day from the end of the month, counting backwards. For example, the last day of the month is -1, the next-to-the-last day of the month is -2, and so on. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule runs on each of the specified days of the month. If multiple values resolve to the same day, the schedule runs only once that day.

The default value is ANY. ANY means that the schedule runs on every day of the month or on the days determined by other enhanced schedule parameters. DAYOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFWEEK or WEEKOFMONTH parameters.

WEEKofmonth

Specifies the week of the month in which to run the schedule. This parameter is used only with enhanced schedules. A week is considered any seven-day period which does not start on a particular day of the week. You can specify FIRST, SECOND, THIRD, FOURTH, LAST, or ANY. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule runs during each of the specified weeks of the month. If multiple values resolve to the same week, the schedule runs only once during that week.

The default value is ANY. ANY means that the schedule runs during every week of the month or on the day or days determined by other enhanced

schedule parameters. WEEKOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFMONTH parameter.

EXPIRATION

Specifies the date after which this schedule is no longer used. This parameter is optional. The default is NEVER. Possible values are:

Never

Specifies that the schedule never expires.

expiration_date

Specifies the date on which this schedule expires, in MM/DD/YYYY format. If you specify an expiration date, the schedule expires at 23:59:59 on the date you specify.

Example: Define a schedule for a monthly incremental backup

Define a schedule named MONTHLY_BACKUP that initiates an incremental backup of all associated nodes. Specify the start date as Tuesday, May 1, 2001. This date does not match the specified day of the week (Sunday), so the initial startup window begins on the first Sunday after May 1, 2001 (05/01/2001). The startup windows for this schedule extend from 01:00 through 03:00. This monthly schedule initiates backup of c: and d: file spaces for all associated nodes.

```
define schedule standard monthly_backup
description="Monthly Backup of c: and d: drives"
objects="c:\* d:\*"
startdate=05/01/2001 starttime=01:00
duration=2 durunits=hours period=1
perunits=months dayofweek=sunday
```

Example: Define a schedule for a weekly incremental backup

Define a schedule named WEEKLY_BACKUP that initiates an incremental backup of all associated nodes. The initial startup window for this schedule extends from 23:00 on Saturday, June 7, 1997 (06/07/1997), to 03:00 on Sunday, June 8, 1997 (06/08/1997). Subsequent windows begin at 23:00, every Saturday. No messages are returned to the client node when this schedule is run.

```
define schedule employee_records weekly_backup
startdate=06/07/1997 starttime=23:00 duration=4
durunits=hours perunits=weeks
dayofweek=saturday options=-quiet
```

Example: Define a schedule that archives a specific directory every quarter

Define a schedule that archives specific files quarterly on the last Friday of the month.

```
define schedule employee_records quarterly_archive
starttime=20:00 action=archive
object=/home/employee/records/*
duration=1 durunits=hour schedstyle=enhanced
month=mar,jun,sep,dec weekofmonth=last dayofweek=fri
```

DEFINE SCHEDULE

DEFINE SCHEDULE (Define a schedule for an administrative command)

Use the **DEFINE SCHEDULE** command to create a new schedule for processing an administrative command.

You can include scripts in an administrative command schedule so the commands are processed automatically.

Note:


1. You cannot schedule the **MACRO** command or the **QUERY ACTLOG** command.
2. If you are scheduling a command that specifies the **WAIT** parameter, the parameter must be set to **YES** in order for the process to provide a return code to the session that started it. For more information about the **WAIT** parameter, see "Server command processing" on page 15.


Privilege class

To define an administrative command schedule, you must have system privilege.

Syntax


Classic administrative schedule






Syntax

Enhanced administrative schedule



DEFINE SCHEDULE



Parameters

schedule_name (Required)

Specifies the name of the schedule to be defined. You can specify up to 30 characters for the name.

Type=Administrative

Specifies that a schedule for an administrative command is defined. This parameter is optional. An administrative command is assumed if the **CMD** parameter is specified.

CMD (Required)

Specifies the administrative command to schedule for processing. The maximum length of the command is 512 characters. Enclose the administrative command in quotation marks if it contains any blank characters.

Restriction: You cannot specify redirection characters with this parameter.

ACTIVE

Specifies whether Tivoli Storage Manager processes an administrative command schedule when the startup window occurs. This parameter is optional. The default is **NO**. The administrative command schedule must be set to the active state with the **UPDATE SCHEDULE** command so that Tivoli Storage Manager can process the schedule. Possible values are:

YES

Specifies that Tivoli Storage Manager processes an administrative command schedule when the startup window begins.

NO Specifies that Tivoli Storage Manager does not process an administrative command schedule when the startup window begins.

DESCRiption

Specifies a description of the schedule. This parameter is optional. You can specify up to 255 characters for the description. Enclose the description in quotation marks if it contains any blank characters.

PRIority

Specifies the priority value for a schedule. This parameter is optional. You can specify an integer from 1 to 10, with 1 being the highest priority and 10 being the lowest. The default is 5.

If two or more schedules have the same window start time, the value you specify determines when Tivoli Storage Manager processes the schedule. The schedule with the highest priority starts first. For example, a schedule with PRIORITY=3 starts before a schedule with PRIORITY=5.

STARTDate

Specifies the date for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current date. Use this parameter with the **STARTTIME** parameter to specify when the initial startup window of the schedule starts.

You can specify the date using one of the values below:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+days or <i>+days</i>	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
<i>EOLM-days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
<i>BOTM+days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

STARTTime

Specifies the time for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current time. This parameter is used in conjunction with the **STARTDATE** parameter to specify when the initial startup window begins.

You can specify the time using one of the values below:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or <i>+HH:MM</i>	The current time plus hours and minutes specified	NOW+02:00 or +02:00. If you issue this command at 5:00 with STARTTIME=NOW+02:00 or STARTTIME=+02:00, the beginning of the startup window is at 7:00.

DEFINE SCHEDULE

Value	Description	Example
<code>NOW-HH:MM</code> or <code>-HH:MM</code>	The current time minus hours and minutes specified	<code>NOW-02:00</code> or <code>-02:00</code> . If you issue this command at 5:00 with <code>STARTTIME=NOW-02:00</code> or <code>STARTTIME=-02:00</code> , the beginning of the startup window is at 3:00.

DURation

Specifies the number of units that define the length of the startup window of the scheduled operation. This parameter is optional. This value must be from 1 to 999. The default is 1.

Use this parameter with the **DURUNITS** parameter to specify the length of the startup window. For example, if you specify `DURATION=20` and `DURUNITS=MINUTES`, the schedule must be started within 20 minutes of the start date and start time. The default length of the startup window is 1 hour. The duration of the window must be shorter than the period between windows.

This value is ignored if you specify `DURUNITS=INDEFINITE`.

DURUnits

Specifies the time units used to determine the duration of the window in which the schedule can start. This parameter is optional. The default is `HOURS`.

Use this parameter with the **DURATION** parameter to specify how long the startup window remains open to process the schedule. For example, if `DURATION=20` and `DURUNITS=MINUTES`, the schedule must be started within 20 minutes of the start date and start time. The schedule may not necessarily complete processing within this window. If the schedule needs to be retried for any reason, the retry attempts must begin before the startup window elapses, or the operation does not restart.

The default value for the length of the startup window is 1 hour. Possible values are:

Minutes

Specifies that the duration of the window is defined in minutes.

Hours

Specifies that the duration of the window is defined in hours.

Days

Specifies that the duration of the window is defined in days.

INDefinite

Specifies that the startup window of the scheduled operation has an indefinite duration. The schedule can run any time after the scheduled start time, until the schedule expires. You cannot specify `DURUNITS=INDEFINITE`, unless you specify `PERUNITS=ONETIME`. The `INDEFINITE` value is not allowed with enhanced schedules.

SCHEDStyle

This parameter is optional. `SCHEDSTYLE` defines either the interval between times when a schedule should run, or the days on which it should run. The style can be either **classic** or **enhanced**. The default is the classic syntax.

For classic schedules, these parameters are allowed: PERIOD, PERUNITS, and DAYOFWEEK. Not allowed for classic schedules are: MONTH, DAYOFMONTH, and WEEKOFMONTH.

For enhanced schedules, these parameters are allowed: MONTH, DAYOFMONTH, WEEKOFMONTH, and DAYOFWEEK. These parameters are not allowed: PERIOD and PERUNITS.

PERiod

Specifies the length of time between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. You can specify an integer from 1 to 999. The default is 1.

Use this parameter with the **PERUNITS** parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every five days after the initial start date and start time. The period between startup windows must exceed the duration of each window. The default is 1 day.

This value is ignored if you specify PERUNITS=ONETIME.

PERUnits

Specifies the time units used to determine the period between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. The default is DAYS.

Use this parameter with the **PERIOD** parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is scheduled every 5 days after the initial start date and start time. The default is 1 day. Possible values are:

Hours

Specifies that the time between startup windows is in hours.

Days

Specifies that the time between startup windows is in days.

Weeks

Specifies that the time between startup windows is in weeks.

Months

Specifies that the time between startup windows is in months.

When you specify PERUNITS=MONTHS, the scheduled operation will be processed each month on the same date. For example, if the start date for the scheduled operation is 02/04/1998, the schedule will process on the 4th of every month thereafter. However, if the date is not valid for the next month, then the scheduled operation will be processed on the last valid date in the month. Thereafter, subsequent operations are based on this new date. For example, if the start date is 03/31/1998, the next month's operation will be scheduled for 04/30/1998. Thereafter, all subsequent operations will be on the 30th of the month until February. Because February has only 28 days, the operation will be scheduled for 02/28/1999. Subsequent operations will be processed on the 28th of the month.

Years

Specifies that the time between startup windows for the schedule is in years.

DEFINE SCHEDULE

When you specify PERUNITS=YEARS, the scheduled operation will be processed on the same month and date of each year. For example, if the start date for the scheduled operation is 02/29/2004, the next year's scheduled operation will be 02/28/2005 because February only has 28 days. Thereafter, subsequent operations will be scheduled for February 28th.

Onetime

Specifies that the schedule processes once. This value overrides the value you specified for the **PERIOD** parameter.

DAYofweek

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can specify different options for the **DAYofweek** parameter, depending on whether the schedule style has been defined as Classic or Enhanced:

Classic Schedule

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can either specify one day of the week, or WEEKDAY, WEEKEND, or ANY. If the start date and start time fall on a day that does not correspond to a day you specify, the start date and start time will be shifted forward in 24-hour increments until the **DAYOFWEEK** parameter is satisfied.

If you select a value for **DAYOFWEEK** other than ANY, and depending on the values for PERIOD and PERUNITS, schedules may not be processed when you would expect. The default is ANY.

Enhanced Schedule

Specifies the days of the week on which to run the schedule. You can either specify multiple days separated by commas and no intervening blanks, or WEEKDAY, WEEKEND, or ANY. If you specify multiple days, the schedule will run on each of the specified days. If you specify WEEKDAY or WEEKEND, you must also specify either WEEKOFMONTH=FIRST or WEEKOFMONTH=LAST, and the schedule will run just once per month.

The default value is ANY, meaning the schedule will run every day of the week or on the day or days determined by other enhanced schedule parameters. **DAYOFWEEK** must have a value of ANY (either by default or specified with the command) when used with the **DAYOFMONTH** parameter.

Possible values for the **DAYofweek** parameter are:

ANY

Specifies that the startup window can begin on any day of the week.

WEEKDay

Specifies that the startup window can begin on Monday, Tuesday, Wednesday, Thursday, or Friday.

WEEKEnd

Specifies that the startup window can begin on Saturday or Sunday.

SUunday

Specifies that the startup window begins on Sunday.

Monday

Specifies that the startup window begins on Monday.

Tuesday

Specifies that the startup window begins on Tuesday.

Wednesday

Specifies that the startup window begins on Wednesday.

Thursday

Specifies that the startup window begins on Thursday.

Friday

Specifies that the startup window begins on Friday.

Saturday

Specifies that the startup window begins on Saturday.

MONth

Specifies the months of the year during which to run the schedule. This parameter is used only with enhanced schedules. Specify multiple values by using commas and no intervening blanks. The default value is ANY. This means the schedule will run during every month of the year.

DAYOFMonth

Specifies the day of the month to run the schedule. This parameter is used only with enhanced schedules. You can either specify ANY or a number from -31 through 31, excluding zero. Negative values are a day from the end of the month, counting backwards. For example, the last day of the month is -1, the next-to-the-last day of the month is -2, etc. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule will run on each of the specified days of the month. If multiple values resolve to the same day, the schedule will run only once that day.

The default value is ANY. This means the schedule will run on every day of the month or on the days determined by other enhanced schedule parameters. DAYOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFWEEK or WEEKOFMONTH parameters.

WEEKofmonth

Specifies the week of the month in which to run the schedule. This parameter is used only with enhanced schedules. A week is considered any seven-day period which does not start on a particular day of the week. You can specify FIRST, SECOND, THIRD, FOURTH, LAST, or ANY. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule will run during each of the specified weeks of the month. If multiple values resolve to the same week, the schedule will run only once during that week.

The default value is ANY, meaning the schedule will run during every week of the month or on the day or days determined by other enhanced schedule parameters. WEEKOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFMONTH parameter.

EXPIration

Specifies the date after which this schedule is no longer used. This parameter is optional. The default is NEVER. Possible values are:

Never

Specifies that the schedule never expires.

DEFINE SCHEDULE

expiration_date

Specifies the date on which this schedule expires, in MM/DD/YYYY format. If you specify an expiration date, the schedule expires at 23:59:59 on the date you specify.

Example: Define a schedule to back up the primary storage pool every two days

Define a schedule named BACKUP_ARCHIVEPOOL that backs up the primary storage pool ARCHIVEPOOL to the copy storage pool RECOVERYPOOL. The backup runs at 8 p.m. every two days.

```
define schedule backup_archivepool type=administrative
cmd="backup stgpool archivepool recoverypool"
active=yes starttime=20:00 period=2
```

Example: Define a schedule to back up the primary storage pool twice a month

Define a schedule named BACKUP_ARCHIVEPOOL that backs up the primary storage pool ARCHIVEPOOL to the copy storage pool RECOVERYPOOL. Select an enhanced schedule and run on the first and fifteenth day of the month.

```
define schedule backup_archivepool type=administrative
cmd="backup stgpool archivepool recoverypool"
schedstyle=enhanced dayofmonth=1,15
```

DEFINE SCRIPT (Define a Tivoli Storage Manager script)

Use this command to define a Tivoli Storage Manager script or to create a new Tivoli Storage Manager script using the contents from another script.

The first line for the script may be defined with this command. To add subsequent lines to the script, use the **UPDATE SCRIPT** command.

Tips:


- The Administration Center only supports ASCII characters for input. If you need to enter characters that are not ASCII, issue the **DEFINE SCRIPT** and **UPDATE SCRIPT** commands from the server console.
- When routing commands inside scripts, enclose the server or server group in parentheses and omit the colon. Otherwise, if the syntax includes a colon, the command is not routed when the **RUN** command is issued. Instead, the command will only run on the server from which the **RUN** command is issued.
- You cannot redirect the output of a command within a Tivoli Storage Manager script. Instead, run the script and then specify command redirection. For example, to direct the output of **script1** to the `c:\temp\test.out` directory, run the script and specify command redirection as in the following example:

```
run script1 > c:\temp\test.out
```

Privilege class

To issue this command, you must have operator, policy, storage, or system privilege.

Syntax



Parameters

script_name (Required)

Specifies the name of the script to be defined. You can specify up to 30 characters for the name.

command_line

Specifies the first command to be processed in a script. You must specify either this parameter (and, optionally, the **LINE** parameter) or the **FILE** parameter.

The command you specify can include substitution variables and can be continued across multiple lines if you specify a continuation character (-) as the last character in the command. Substitution variables are specified with a '\$' character, followed by a number that indicates the value of the parameter when the script is processed. You can specify up to 1200 characters for the command line. Enclose the command in quotation marks if it contains blanks.

You can run commands serially, in parallel, or serially and in parallel by specifying the **SERIAL** or **PARALLEL** script commands for the **COMMAND_LINE**

DEFINE SCRIPT

parameter. You can run multiple commands in parallel and wait for them to complete before proceeding to the next command. Commands will run serially until the parallel command is encountered.

Conditional logic flow statements can be used. These statements include IF, EXIT, and GOTO.

Line

Specifies the line number for the command line. Because commands are specified in multiple lines, line numbers are used to determine the order for processing when the script is run. The first line, or line 001 is the default. This parameter is optional.

File

Specifies the name of the file whose contents will be read into the script to be defined. The file must reside on the server running this command. If you specify the FILE parameter, you cannot specify a command line or line number.

You can create a script by querying another script and specifying the FORMAT=RAW and OUTPUTFILE parameters. The output from querying the script is directed to a file you specify with the OUTPUTFILE parameter. To create the new script, the contents of the script to be defined are read in from the file you specified with the OUTPUTFILE parameter.

DESCRIPTION

Specifies a description for the script. You can specify up to 255 characters for the description. Enclose the description in quotation marks if it contains blank characters. This parameter is optional.

Example: Write a script to display AIX clients

Define a script that will display all AIX clients.

```
define script qaixc "select node_name from nodes where platform_name='AIX'"  
  desc='Display aix clients'
```

Example: Write and run a script to route a command to a server group

Define and run a script that will route the QUERY STGPOOL command to a server group named DEV_GROUP.

```
define script qu_stg "(dev_group) query stgpool"  
run qu_stg
```

Example: Create a script from an existing script

Define a script whose command lines are read in from a file that is named MY.SCRIPT and name the new script AGADM.

```
define script agadm file=my.script
```

Refer to the *Administrator's Guide* for more information on the following topics:

- Running commands in parallel or serially
- Using logic flow statements in scripts
- Performing tasks concurrently on multiple servers

Related commands*Table 101. Commands related to DEFINE SCRIPT*

Command	Description
COPY SCRIPT	Creates a copy of a script.
DELETE SCRIPT	Deletes the script or individual lines from the script.
QUERY SCRIPT	Displays information about scripts.
RENAME SCRIPT	Renames a script to a new name.
RUN	Runs a script.
UPDATE SCRIPT	Changes or adds lines to a script.

DEFINE SERVER

DEFINE SERVER (Define a server for server-to-server communications)

Use this command to define a server to use functions such as virtual volumes, node replication, command routing, and LAN-free data movement, among others.

Use this command to define a server for the following functions:

- Enterprise configuration
- Enterprise event logging
- Command routing
- Virtual volumes
- LAN-free data movement
- Node replication
- Data movement using z/OS media server
- Status monitoring of remote servers
- Alert monitoring of remote servers

If you use an LDAP directory server to authenticate passwords, any target servers must be configured for LDAP-authenticated passwords. Data that is replicated from a node that authenticates with an LDAP directory server is inaccessible if the target replication server is not properly configured. If your target replication server is not configured, replicated data from an LDAP node can make it to the target server. But the target replication server must be configured to use LDAP if you want to access the data.

The use of virtual volumes is not supported when the source server and the target server are on the same Tivoli Storage Manager server.

This command also is used to define a Tivoli Storage Manager storage agent as if it were a server.

Privilege class


To issue this command, you must have system privilege.

Syntax

For:

- Enterprise configuration
- Enterprise event logging
- Command routing
- Storage agent
- Node replication source and target servers
- z/OS media server


►►—DEFine—SERver—*server_name*—SERVERPAssword—=*password*—►



Notes:

- 1 The `CROSSDEFINE` parameter does not apply to storage agent definitions.
- 2 The `VALIDATEPROTOCOL` parameter applies only to storage agent definitions.

Syntax for virtual volumes



Parameters

`server_name` (Required)

Specifies the name of the server. This name must be unique on the server. The maximum length of this name is 64 characters.

For server-to-server event logging, library sharing, and node replication, you must specify a server name that matches the name that was set by issuing the **SET SERVERNAME** command at the target server.

`PAssword` (Required)

Specifies the password used to sign on to the target server for virtual volumes. If you specify the `NODENAME` parameter, you must specify the `PASSWORD` parameter. If you specify the `PASSWORD` parameter but not the `NODENAME` parameter, the node name defaults to the server name specified with the `SET SERVERNAME` command.

DEFINE SERVER

SERVERPASSword

Specifies the password of the server you are defining. This password must match the password set by the SET SERVERPASSWORD command. This parameter is required for enterprise configuration, command routing, and server-to-server event logging functions.

Tip: Command routing uses the ID and the password of the administrator who is issuing the command.

HLAddress (Required)

Specifies the IP address (in dotted decimal format) of the server.

Do not use the loopback address as the value of this parameter. Virtual volumes are not supported when the source server and the target server are the same Tivoli Storage Manager server.

LLAddress (Required)

Specifies the low-level address of the server. This address is usually the same as the address in the TCPPOINT server option of the target server. When SSL=YES, the port must already be designated for SSL communications on the target server.

COMMmethod

Specifies the communication method used to connect to the server. This parameter is optional.

URL

Specifies the URL address of this server. The parameter is optional.

DELgraceperiod

Specifies a number of days that an object remains on the target server after it was marked for deletion. You can specify a value 0 - 9999. The default is 5. This parameter is optional.

NODENAME

Specifies a node name to be used by the server to connect to the target server. This parameter is optional. If you specify the NODENAME parameter, you must also specify the PASSWORD parameter. If you specify the PASSWORD parameter but not the NODENAME parameter, the node name defaults to the server name specified with the SET SERVERNAME command.

DESCRIPTION

Specifies a description of the server. The parameter is optional. The description can be up to 255 characters. Enclose the description in quotation marks if it contains blank characters.

CROSSDEFINE

Specifies whether the server that is running this command defines itself to the server that is being specified by this command. This parameter is optional.

Important: This parameter does not apply to storage agent definitions.

If this parameter is included, you must also issue the **SET SERVERNAME**, **SET SERVERPASSWORD**, **SET SERVERHLADDRESS**, **SET CROSSDEFINE**, and **SET SERVERLLADDRESS** commands. The default is NO.

Remember:

- For replication operations, the names of the source and target replication servers must match the names that you specify in this command.

- **CROSSDEFINE** can be used with `SSL=YES` if all of the conditions specified for the `SSL=YES` parameter are in place on the source and target server.

You can specify one of the following values:

No Cross definition is not performed.

Yes

Cross definition is performed.

VALIDATEPROTOCOL

Specify a cyclic redundancy check to validate the data sent between the storage agent and Tivoli Storage Manager server. The parameter is optional. The default is `NO`. You can specify one of the following values:

No Specifies that data validation is not performed on any data sent between the storage agent and server.

All

Specifies that data validation is performed on all client file data, client file metadata, and Tivoli Storage Manager server metadata that is sent between the storage agent and server. This mode affects performance as additional resources are required to calculate and compare CRC values between the storage agent and the server.

SSL

Specifies the communication mode of the server. The default is `NO`. You can specify one of the following values:

No Specifies that unencrypted communication with the specified server occurs over a TCP/IP session without SSL.

Yes

Specifies an SSL session for communication with the specified server. The following conditions apply if you specify `YES`:

- Before starting the servers, self-signed certificates of the partner servers must be in the key database file (`cert.kdb`) of each of the servers.
- You can define multiple server names with different parameters for the same target server.
- SSL support is active if the server options file contains the `SSLTCPPORT` or `SSLTCPADMINPORT` option or if a server is defined with `SSL=YES` at startup.
- Storage agents can issue the **DSMSTA SETSTORAGESEVER** command and include the `STAKEKEYDBPW` and `SSL` parameters to create the key database.
- If third-party certificates are used, the CA (certificate authority) certificate must be kept in the key database of the partner server. If you upgrade from a release earlier than V6.3.0 and have a key database file in your instance directory, update it to allow third-party certificates. To update the key database file, issue the following command:

```
gsk8capicmd_64 -keydb -convert -populate
-db cert.kdb -pw passwordofkeydatabasefile
```

Example: Set up two servers to use SSL to communicate

You must set up two servers to use SSL to communicate. The server addresses are as follows:

- ServerA is at `bfa.tucson.ibm.com`
- ServerB is at `bfb.tucson.ibm.com`

Complete the following steps to set up the two servers for SSL:

DEFINE SERVER

1. Specify options SSLTCPPT 1542 and TCPPT 1500 for both servers in the dsmserv.opt option file.
2. Start both servers and run the following command to set up the key database file password:

```
SET SSLKEYRINGPW newpw UPDATE=Y
```
3. Shut down both servers to import the cert256 partner certificate. For ServerA, the certificate is in the /tsma instance directory. For ServerB, the certificate is in the /tsmb instance directory.
4. Start both servers. The /tsma/cert256.arm file is copied to /tsmb/cert256.bfa.arm on the bfb.tucson.ibm.com address. The /tsmb/cert256.arm file is copied to /tsmb/cert256.bfb.arm on the bfa.tucson.ibm.com address.
5. Issue the following command:
 - From ServerA:

```
gsk8capicmd_64 -cert -add -db cert.kdb -pw newpw -format ascii -label "bfb" -file /tsma/cert256.bfb.arm
```
 - From ServerB:

```
gsk8capicmd_64 -cert -add -db cert.kdb -pw newpw -format ascii -label "bfa" -file /tsmb/cert256.bfa.arm
```

From each server, you can view the certificates in the key database by issuing the following command:

```
gsk8capicmd_64 -cert -list -db cert.kdb -pw newpw
```

6. Restart the servers.
7. Issue the appropriate **DEFINE SERVER** command. For ServerA, issue the following example command:

```
DEFINE SERVER BFB h1a=bfb.tucson.ibm.com l1a=1542  
serverpa=passwordforbfb SSL=YES
```

For ServerB, issue the following example command:

```
DEFINE SERVER BFA h1a=bfa.tucson.ibm.com l1a=1542  
serverpa=passwordforbfa SSL=YES
```

If you do not use SSL, issue the following example **DEFINE SERVER** command on ServerA:

```
DEFINE SERVER BFBTCP h1a=bfb.tucson.ibm.com l1a=1500  
serverpa=passwordforbfb SSL=NO
```

If you do not use SSL, issue the following example **DEFINE SERVER** command on ServerB:

```
DEFINE SERVER BFATCP h1a=bfa.tucson.ibm.com l1a=1500  
serverpa=passwordforbfa SSL=NO
```

Example: Define a target server

A target server has a high-level address of 9.116.2.67 and a low-level address of 1570. Define that target server to the source server, name it SERVER2, and set the password to SECRET. Specify that objects remain on the target server for seven days after they are marked for deletion.

```
define server server2 password=secret  
hladdress=9.115.3.45 lladdress=1570 delgraceperiod=7
```

Example: Define a server to receive commands from other servers

Define a server that can receive commands routed from other servers. Name the server WEST_COMPLEX and set the password to CACTUS. Set the high-level address to 9.172.12.35, the low-level address to 1500, and the URL address to http://west_complex:1580/.

```
define server west_complex serverpassword=cactus
hladdress=9.172.12.35 lladdress=1500
url=http://west_complex:1580/
```

Example: Cross-define two servers

Use cross definition to define SERVER_A and SERVER_B.

1. On SERVER_B, specify the server name, password, and high- and low-level addresses of SERVER_B. Specify that cross defining is allowed.

```
set servername server_b
set serverpassword mylife
set serverhladdress 9.115.20.80
set serverlladdress 1860
set crossdefine on
```

2. On SERVER_A, specify the server name, password, and high- and low-level addresses of SERVER_A.

```
set servername server_a
set serverpassword yourlife
set serverhladdress 9.115.20.97
set serverlladdress 1500
```

3. On SERVER_A, define SERVER_B:

```
define server server_b hladdress=9.115.20.80 lladdress=1860
serverpassword=mylife crossdefine=yes
```

Related commands

See the *Tivoli Storage Manager Storage Agent User's Guide* for information about the **DSMSTA SETSTORAGESERVER** command.

Table 102. Commands related to **DEFINE SERVER**

Command	Description
DEFINE DEVCLASS	Defines a device class.
DEFINE PATH	Define a path when the destination is a z/OS media server.
DELETE DEVCLASS	Deletes a device class.
DELETE FILESPACE	Deletes data associated with clients file spaces.
DELETE SERVER	Deletes the definition of a server.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY SERVER	Displays information about servers.
RECONCILE VOLUMES	Reconciles source server virtual volume definitions and target server archive objects.
REGISTER NODE	Defines a client node to the server and sets options for that user.

DEFINE SERVER

Table 102. Commands related to **DEFINE SERVER** (continued)

Command	Description
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
SET CROSSDEFINE	Specifies whether to cross define servers.
SET SERVERNAME	Specifies the name by which the server is identified.
SET SERVERHLADDRESS	Specifies the high-level address of a server.
SET SERVERLLADDRESS	Specifies the low-level address of a server.
SET SERVERPASSWORD	Specifies the server password.
SET REPLSERVER	Specifies a target replication server.
UPDATE DEVCLASS	Changes the attributes of a device class.
UPDATE NODE	Changes the attributes associated with a client node.
UPDATE PATH	Define a path when the destination is a z/OS media server.
UPDATE SERVER	Updates information about a server.

DEFINE SERVERGROUP (Define a server group)

Use this command to define a server group. A server group lets you route commands to multiple servers by specifying only the group name. After defining the server group, add servers to the group by using the **DEFINE GRPMEMBER** command.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—DEfIne SERVERGRoup—group_name—[DESCRiption—=description—]—◄◄
```

Parameters

group_name (Required)

Specifies the name of the server group. The maximum length of the name is 64 characters.

DESCRiption

Specifies a description of the server group. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Example: Define a server group

Define a server group named WEST_COMPLEX.

```
define servergroup west_complex
```

Related commands

Table 103. Commands related to **DEFINE SERVERGROUP**

Command	Description
COPY SERVERGROUP	Creates a copy of a server group.
DEFINE GRPMEMBER	Defines a server as a member of a server group.
DELETE GRPMEMBER	Deletes a server from a server group.
DELETE SERVERGROUP	Deletes a server group.
MOVE GRPMEMBER	Moves a server group member.
QUERY SERVERGROUP	Displays information about server groups.
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVERGROUP	Updates a server group.

DEFINE SPACETRIGGER (Define the space trigger)

Use this command to define settings for triggers that determine when and how the server prepares additional space when predetermined thresholds have been exceeded in storage pools that use FILE and DISK device classes. Space triggers are not enabled for storage pools with a parameter RECLAMATIONTYPE=SNAPLOCK.

Tivoli Storage Manager allocates more space when space utilization reaches a specified value. After allocating more space, Tivoli Storage Manager either adds the space to the specified pool (random-access or sequential-access disk).

Important: Space trigger functions and storage pool space calculations take into account the space remaining in each directory. An inaccurate calculation could result in a failure to expand the space available in a storage pool. Failure to expand space in a storage pool is one of the conditions that can cause a trigger to become disabled.


For example, if you specify multiple directories for a device class and the directories reside in the same file system, the server will calculate space by adding values representing the space remaining in each directory. These space calculations will be inaccurate. Rather than choosing a storage pool with sufficient space for an operation, the server could choose the directory that is specified for the device class and run out of space prematurely.

To prevent possible problems and ensure an accurate calculation, you associate each directory with a separate file system. If a trigger becomes disabled because the space in a storage pool could not be expanded, you can re-enable the trigger by specifying the following command: update spacetrigger stg. No further changes are required to the space trigger.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

STG

Specifies a storage pool space trigger.

Fullpct

This parameter specifies the utilization percentage of the storage pool. This

parameter is optional. Specify an integer value from 0 to 99. The default is 80. A value of zero (0) disables the space trigger. When this value is exceeded, the space trigger creates new volumes. Exceeding the threshold may not cause new volumes to be created until the next space request is made.

You can determine storage pool utilization by issuing the QUERY STGPOOL command with FORMAT=DETAILED. The percentage of storage pool utilization is displayed in the field "Space Trigger Util." The calculation for this percentage does not include potential scratch volumes. The calculation for the percentage utilization used for migration and reclamation, however, does include potential scratch volumes.

SPACEexpansion

For sequential-access FILE-type storage pools, this parameter is used in determining the number of additional volumes that are created in the storage pool. Volumes are created using the MAXCAPACITY value from the storage pool's device class. For random-access DISK storage pools, the space trigger creates a single volume using the EXPANSIONPREFIX.

EXPansionprefix

For random-access DISK storage-pools, this parameter specifies the prefix that the server uses to create new storage pool files. This parameter is optional and applies only to random-access DISK device classes. The default prefix is the server installation path.

The prefix can include one or more directory separator characters, for example:
/opt/tivoli/tsm/server/bin/

You can specify up to 250 characters. If you specify an invalid prefix, automatic expansion can fail.

This parameter is not valid for space triggers for sequential-access FILE storage pools. Prefixes are obtained from the directories specified with the associated device class.

STGPOOL

Specifies the storage pool associated with this space trigger. This parameter is optional for storage pool space triggers. If you specify the STG parameter but not the STGPOOL parameter, one space trigger is created that applies to all random-access DISK and sequential-access FILE storage pools that do not have a specific space trigger.

This parameter does not apply to storage pools with the parameter RECLAMATIONTYPE=SNAPLOCK.

Example: Define a space trigger to increase storage pool space 25 percent

Set up a storage pool space trigger for increasing the amount of space in a storage pool by 25 percent when it is filled to 80 percent utilization of existing volumes. Space will be created in the directories associated with the device class.

```
define spacetrigger stg spaceexpansion=25 stgpool=file
```

Example: Define a space trigger to increase storage pool space 40 percent

Set up a space trigger for the WINPOOL1 storage pool to increase the amount of space in the storage pool by 40 percent when it is filled to 80 percent utilization of existing volumes.

DEFINE SPACETRIGGER

```
define spacetrigger stg spaceexpansion=40 stgpool=winpool1
```

Related commands

Table 104. Commands related to DEFINE SPACETRIGGER

Command	Description
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DELETE SPACETRIGGER	Deletes the storage pool space trigger.
QUERY SPACETRIGGER	Displays information about a storage pool space trigger.
UPDATE SPACETRIGGER	Changes attributes of storage pool space trigger.

DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)

Use this command to define a new status monitoring threshold.

Status monitoring thresholds compare the defined conditions to the status monitoring server queries and inserts the results in the status monitoring table.

Multiple thresholds can be defined for an activity. For example, you can create a threshold that provides a warning status if storage pool capacity utilization is greater than 80%. You can then create another threshold that provides error status if storage pool capacity utilization is greater than 90%.


Note: If a threshold is already defined for an EXISTS condition, you cannot define another threshold with one of the other condition types.

Privilege class

To issue this command, you must have system privilege.

Syntax

►► DEFINE STATUSTHRESHOLD *threshold_name* *activity* _____►



Parameters

threshold_name (Required)

Specifies the threshold name. The name cannot exceed 48 characters in length.

activity (Required)

Specifies the activity for which you want to create status indicators. Specify one of the following values:

PROCESSSUMMARY

Specifies the number of processes that are currently active.

SESSIONSUMMARY

Specifies the number of sessions that are currently active.

CLIENTSESSIONSUMMARY

Specifies the number of client sessions that are currently active.

DEFINE STATUSTHRESHOLD

SCHEDCLIENTSESSIONSUMMARY

Specifies the number of scheduled client sessions.

DBUTIL

Specifies the database utilization percentage. The default warning threshold value is 80%, and the default error threshold value is 90%.

DBFREESPACE

Specifies the free space available in the database in gigabytes.

DBUSEDSPACE

Specifies the amount of database space that is used, in gigabytes.

ARCHIVELOGFREESPACE

Specifies the free space that is available in the archive log, in gigabytes.

STGPOOLUTIL

Specifies the storage pool utilization percentage. The default warning threshold value is 80%, and the default error threshold value is 90%.

STGPOOLCAPACITY

Specifies the storage pool capacity in gigabytes.

AVGSTGPOOLUTIL

Specifies the average storage pool utilization percentage across all storage pools. The default warning threshold value is 80%, and the default error threshold value is 90%.

TOTSTGPOOLCAPACITY

Specifies the total storage pool capacity in gigabytes for all available storage pools.

TOTSTGPOLS

Specifies the number of defined storage pools.

TOTRWSTGPOLS

Specifies the number of defined storage pools that are readable or writeable.

TOTNOTRWSTGPOLS

Specifies the number of defined storage pools that are not readable or writeable.

STGPOOLINUSEANDDEFINED

Specifies the total number of defined volumes that are in use.

ACTIVELOGUTIL

Specifies the current percent utilization of the active log. The default warning threshold value is 80%, and the default error threshold value is 90%.

ARCHLOGUTIL

Specifies the current utilization of the archive log. The default warning threshold value is 80%, and the default error threshold value is 90%.

CPYSTGPOOLUTIL

Specifies the percent utilization for a copy storage pool. The default warning threshold value is 80%, and the default error threshold value is 90%.

PMRYSTGPOOLUTIL

Specifies the percent utilization for a primary storage pool. The default warning threshold value is 80%, and the default error threshold value is 90%.

DEVCLASSPCTDRVOFFLINE

Specifies the percent utilization of drives that are offline, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

DEVCLASSPCTDRVPOLLING

Specifies the drives polling, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

DEVCLASSPCTLIBPATHSOFFLINE

Specifies the library paths that are offline, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

DEVCLASSPCTPATHSOFFLINE

Specifies the percentage of device class paths that are offline, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

DEVCLASSPCTDISKSNOTRW

Specifies the percentage of disks that are not writable for the disk device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

DEVCLASSPCTDISKSUNAVAILABLE

Specifies the percentage of the disk volumes that are unavailable, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

FILEDEVCLASSPCTSCRUNALLOCATABLE

Specifies the percentage of scratch volumes that the server cannot allocate for a given non-shared file device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

Condition

Specifies the condition that is used to compare the activity output to the specified value. The default value is EXISTS. Specify one of the following values:

EXists

Creates a status monitoring indicator if the activity exists.

GT Creates a status monitoring indicator if the activity outcome is greater than the specified value.

GE Creates a status monitoring indicator if the activity outcome is greater than or equal to the specified value.

LT Creates a status monitoring indicator if the activity outcome is less than the specified value.

LE Creates a status monitoring indicator if the activity outcome is less than or equal to the specified value.

EQual

Creates a status monitoring indicator if the activity outcome is equal to the specified value.

Value (Required)

Specifies the value to compare to the activity output of the specified condition. VALUE must be specified unless CONDITION is set to EXISTS. You can specify an integer from 0 to 9223372036854775807.

DEFINE STATUSTHRESHOLD

Status

Specifies that the status indicator created in status monitoring if the condition that is being evaluated passes. This optional parameter has a default value of NORMAL. Specify one of the following values:

Normal

Specifies that the status indicator has a normal status value.

Warning

Specifies that the status indicator has a warning status value.

Error

Specifies that the status indicator has an error status value.

Define status threshold

Define a status threshold for average storage pool utilization percentage by issuing the following command:

```
define statusthreshold avgstgpl "AVGSTGPOOLUTIL" value=85
  condition=gt status=warning
```

Related commands

Table 105. Commands related to DEFINE STATUSTHRESHOLD

Command	Description
"DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)" on page 451	Deletes a status monitoring threshold.
"QUERY MONITORSTATUS (Query the monitoring status)" on page 799	Displays information about monitoring alerts and server status settings.
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"QUERY STATUSTHRESHOLD (Query status monitoring thresholds)" on page 914	Displays information about a status monitoring thresholds.
"SET STATUSMONITOR (Specifies whether to enable status monitoring)" on page 1144	Specifies whether to enable status monitoring.
"SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)" on page 1142	Specifies whether to enable client at-risk activity interval evaluation.
"SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)" on page 1146	Specifies the refresh interval for status monitoring.
"SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)" on page 1147	Specifies whether to use client at-risk skipped files as failure evaluation.
"UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)" on page 1375	Changes the attributes of an existing status monitoring threshold.

DEFINE STGPOOL (Define a storage pool)

Use this command to define a primary storage pool, copy storage pool, or an active-data pool. A primary storage pool provides a destination for backup files, archive files, or files migrated from client nodes. A copy storage pool provides a destination for backup copies of files that are in primary storage pools. An active-data pool provides a destination for active versions of backup data that are in primary storage pools.

All volumes in a storage pool belong to the same device class. Random access storage pools use the DISK device type. After you define a random access storage pool, you must define volumes for the pool to create storage space.

Sequential access storage pools use device classes that you define for tape devices, optical devices, files on disk (FILE device type), and storage on another server (SERVER device type). To create storage space in a sequential access storage pool, you must allow scratch volumes for the pool when you define or update it, or define volumes for the pool after you define the pool. You can also do both.

Restriction: If a client is using the simultaneous-write function and data deduplication, the data deduplication feature is disabled during backups to a storage pool.

The DEFINE STGPOOL command takes four forms.

- “DEFINE STGPOOL (Define a primary storage pool assigned to random access devices)” on page 343
- “DEFINE STGPOOL (Define a primary storage pool assigned to sequential access devices)” on page 352
- “DEFINE STGPOOL (Define a copy storage pool assigned to sequential access devices)” on page 369
- “DEFINE STGPOOL (Define an active-data pool assigned to sequential-access devices)” on page 378

The syntax and parameters for each form are defined separately.

Table 106. Commands related to DEFINE STGPOOL

Command	Description
BACKUP DB	Backs up the Tivoli Storage Manager database to sequential access volumes.
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
COPY ACTIVE DATA	Copies active backup data.
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node to a collocation group.
DEFINE DEVCLASS	Defines a device class.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DELETE COLLOGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node from a collocation group.
DELETE STGPOOL	Deletes a storage pool from server storage.

DEFINE STGPOOL

Table 106. Commands related to DEFINE STGPOOL (continued)

Command	Description
MOVE DATA	Moves data from a specified storage pool volume to another storage pool volume.
MOVE MEDIA	Moves storage pool volumes that are managed by an automated library.
QUERY COLLOGGROUP	Displays information about collocation groups.
QUERY DEVCLASS	Displays information about device classes.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY SHREDSTATUS	Displays information about data waiting to be shredded.
QUERY STGPOOL	Displays information about storage pools.
RENAME STGPOOL	Renames a storage pool.
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.
SET DRMPRIMSTGPOOL	Specifies that primary storage pools are managed by DRM.
SHRED DATA	Manually starts the process of shredding deleted data.
UPDATE COLLOGGROUP	Updates the description of a collocation group.
UPDATE STGPOOL	Changes the attributes of a storage pool.


DEFINE STGPOOL (Define a primary storage pool assigned to random access devices)

Use this command to define a primary storage pool assigned to random access devices.


Privilege class

To issue this command, you must have system privilege.

Syntax



DEFINE STGPOOL



Notes:

- 1 This parameter is not available for SnapLock storage pools.

Parameters

pool_name (Required)

Specifies the name of the storage pool to be defined. The name must be unique, and the maximum length is 30 characters.

DISK (Required)

Specifies that you want to define a storage pool to the DISK device class (the DISK device class is predefined during installation).

Pooltype=Primary

Specifies that you want to define a primary storage pool. This parameter is optional. The default value is PRIMARY.

DESCRIPTION

Specifies a description of the storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

ACCESS

Specifies how client nodes and server processes (such as migration and reclamation) can access files in the storage pool. This parameter is optional. The default value is READWRITE. Possible values are:

READWrite

Specifies that client nodes and server processes can read and write to files stored on volumes in the storage pool.

READOnly

Specifies that client nodes can only read files from the volumes in the storage pool.

Server processes can move files within the volumes in the storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool has been specified as a subordinate storage pool (with the **NEXTSTGPOOL** parameter) and is defined as *readonly*, the storage pool is skipped when server processes attempt to write files to the storage pool.

UNAVailable

Specifies that client nodes cannot access files stored on volumes in the storage pool.

Server processes can move files within the volumes in the storage pool and can also move or copy files from this storage pool to another storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool has been specified as a subordinate storage pool (with the **NEXTSTGPOOL** parameter) and is defined as *unavailable*, the storage pool is skipped when server processes attempt to write files to the storage pool.

MAXSize

Specifies the maximum size for a physical file that the server can store in the storage pool. This parameter is optional. The default value is **NOLIMIT**.

Possible values are:

NOLimit

Specifies that there is no maximum size limit for physical files stored in the storage pool.

maximum_file_size

Limits the maximum physical file size. Specify an integer from 1 to 999999, followed by a scale factor. For example, **MAXSIZE=5G** specifies that the maximum file size for this storage pool is 5 GB. Scale factors are:

Scale factor **Meaning**

K	kilobyte
M	megabyte
G	gigabyte
T	terabyte

If a file exceeds the maximum size and no pool is specified as the next storage pool in the hierarchy, the server does not store the file. If a file exceeds the maximum size and a pool is specified as the next storage pool, the server stores the file in the next storage pool that can accept the file size. If you specify the next storage pool parameter, at least one storage pool in your hierarchy should have no limit on the maximum size of a file. By having no limit on the size for at least one pool, you ensure that no matter what its size, the server can store the file.

For logical files that are part of an aggregate, the server considers the size of the aggregate to be the file size. Therefore, the server does not store logical files that are smaller than the maximum size limit if the files are part of an aggregate that is larger than the maximum size limit.

CRCDATA

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is optional. The default value is **NO**. By setting **CRCDATA** to **YES** and scheduling an **AUDIT VOLUME** command you can continually ensure the integrity of data stored in your storage hierarchy. Possible values are:

Yes

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because additional overhead is required to calculate and compare CRC values between the storage pool and the server.

No Specifies that data is stored without CRC information.

DEFINE STGPOOL

NEXtstgpool

Specifies a primary storage pool to which files are migrated. This parameter is optional.

If you do not specify a next storage pool, the server cannot migrate files from this storage pool and cannot store files that exceed the maximum size for this storage pool in another storage pool.

You cannot create a chain of storage pools that leads to an endless loop through the NEXTSTGPOOL parameter. At least one storage pool in the hierarchy must have no value specified for NEXTSTGPOOL.

If you specify a sequential access pool as the NEXTSTGPOOL, the pool can only be NATIVE or NONBLOCK dataformat.

HIGHmig

Specifies that the server starts migration for this storage pool when the amount of data in the pool reaches this percentage of the pool's estimated capacity. This parameter is optional. You can specify an integer from 0 to 100. The default value is 90.

When the storage pool exceeds the high migration threshold, the server can start migration of files by node, to the next storage pool, as defined with the NEXTSTGPOOL parameter. You can specify HIGHMIG=100 to prevent migration for this storage pool.

LOWmig

Specifies that the server stops migration for this storage pool when the amount of data in the pool reaches this percentage of the pool's estimated capacity. This parameter is optional. You can specify an integer from 0 to 99. The default value is 70.

When the storage pool reaches the low migration threshold, the server does not start migration of another node's files. Because all file spaces that belong to a node are migrated together, the occupancy of the storage pool can fall below the value you specified for this parameter. You can set LOWMIG=0 to permit migration to empty the storage pool.

CACHE

Specifies whether the migration process leaves a cached copy of a file in this storage pool after migrating the file to the next storage pool. This parameter is optional. The default value is NO. Possible values are:

Yes

Specifies that caching is enabled.

No Specifies that caching is disabled.

Using cache may improve the retrievability of files, but may affect the performance of other processes. See the *Administrator's Guide* for details.

MIGPRocess

Specifies the number of processes that the server uses for migrating files from this storage pool. This parameter is optional. You can specify an integer from 1 to 999. The default value is 1.

During migration, the server runs this number of processes in parallel to provide the potential for improved migration rates.

Tips:

- The number of migration processes is dependent upon the setting of the MIGPROCESS parameter and the number of nodes or the number of

collocation groups with data in the migrating storage pool. For example, if the MIGPROCESS parameter is equal to six, but there are only two nodes with data on the storage pool, migration processing only consists of two processes, not six.

- When specifying this parameter, consider whether the simultaneous-write function is enabled for server data migration. Each migration process requires a mount point and a drive for each copy storage pool and active-data pool defined to the target storage pool.

MIGDelay

Specifies the minimum number of days a file must remain in a storage pool before it becomes eligible for migration. To calculate a value to compare to the specified MIGDELAY value, the server counts the number of days that the file has been in the storage pool and the number of days, if any, since the file was retrieved by a client. The lesser of the two values is compared to the specified MIGDELAY value. For example, if all the following conditions are true, a file is not migrated:

- A file has been in a storage pool for five days.
- The file was accessed by a client within the past three days.
- The value specified for the MIGDELAY parameter is four days.

This parameter is optional. You can specify an integer from 0 to 9999. The default is 0, which means that you do not want to delay migration.

If you want the server to count the number of days based only on when a file was stored and not when it was retrieved, use the NORETRIEVEDATE server option.

MIGContinue

Specifies whether you allow the server to migrate files that do not satisfy the migration delay time. This parameter is optional. The default is YES.

Because you can require that files remain in the storage pool for a minimum number of days, the server may migrate all eligible files to the next storage pool yet not meet the low migration threshold. This parameter allows you to specify whether the server is allowed to continue the migration process by migrating files that do not satisfy the migration delay time.

Possible values are:

Yes

Specifies that, when necessary to meet the low migration threshold, the server continues to migrate files that do not satisfy the migration delay time.

If you allow more than one migration process for the storage pool, some files that do not satisfy the migration delay time may be migrated unnecessarily. As one process migrates files that satisfy the migration delay time, a second process could begin migrating files that do not satisfy the migration delay time to meet the low migration threshold. The first process that is still migrating files that satisfy the migration delay time might have, by itself, caused the low migration threshold to be met.

- No** Specifies that the server stops migration when no eligible files remain to be migrated, even before reaching the low migration threshold. The server does not migrate files unless the files satisfy the migration delay time.

DEFINE STGPOOL

AUTOCopy

Specifies when Tivoli Storage Manager performs simultaneous-write operations. The default value is CLIENT. This parameter is optional and affects the following operations:

- Client store sessions
- Server import processes
- Server data-migration processes

If an error occurs while data is being simultaneously written to a copy storage pool or active-data pool during a migration process, the server stops writing to the failing storage pools for the remainder of the process. However, the server continues to store files into the primary storage pool and any remaining copy storage pools or active-data pools. These pools remain active for the duration of the migration process. Copy storage pools are specified using the **COPYSTGPOOLS** parameter. Active-data pools are specified using the **ACTIVEDATAPOOLS** parameter.

Possible values are:

None

Specifies that the simultaneous-write function is disabled.

Client

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions or server import processes. During server import processes, data is written simultaneously to only copy storage pools. Data is not written to active-data pools during server import processes.

MIGRation

Specifies that data is written simultaneously to copy storage pools and active-data pools only during migration to this storage pool. During server data-migration processes, data is written simultaneously to copy storage pools and active-data pools only if the data does not exist in those pools. Nodes whose data is being migrated must be in a domain associated with an active-data pool. If the nodes are not in a domain associated with an active pool, the data cannot be written to the pool.

All

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions, server import processes, or server data-migration processes. Specifying this value ensures that data is written simultaneously whenever this pool is a target for any of the eligible operations.

COPYSTGpools

Specifies the names of copy storage pools where the server simultaneously writes data. The **COPYSTGPOOLS** parameter is optional. You can specify a maximum of three copy pool names separated by commas. Spaces between the names of the copy pools are not permitted. When specifying a value for the **COPYSTGPOOLS** parameter, you can also specify a value for the **COPYCONTINUE** parameter.

The combined total number of storage pools specified in the **COPYSTGPOOLS** and **ACTIVEDATAPOOLS** parameters cannot exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of copy storage pools and the **COPYCONTINUE** value from the primary storage pool. The primary storage pool is specified by the copy group of the management class that is

bound to the data. For details, refer to information about the simultaneous-write function in the *Administrator's Guide*.

The server can write data simultaneously to copy storage pools during the following operations:

- Backup and archive operations by Tivoli Storage Manager backup-archive clients or application clients using the Tivoli Storage Manager API
- Migration operations by Tivoli Storage Manager for Space Management clients
- Import operations that involve copying exported file data from external media to a primary storage pool associated with a copy storage pool list

Restriction: The simultaneous-write function is not supported for the following store operations:

- When the operation is using LAN-free data movement. Simultaneous-write operations take precedence over LAN-free data movement, causing the operations to go over the LAN. However, the simultaneous-write configuration is honored.
- NAS backup operations. If the primary storage pool specified in the DESTINATION or TOCDESTINATION in the copy group of the management class has copy storage pools defined, the copy storage pools are ignored and the data is stored into the primary storage pool only.

Attention: The function provided by the COPYSTGPOOLS parameter is not intended to replace the **BACKUP STGPOOL** command. If you use the COPYSTGPOOLS parameter, continue to use the **BACKUP STGPOOL** command to ensure that the copy storage pools are complete copies of the primary storage pool. There are cases when a copy might not be created. For more information, see the COPYCONTINUE parameter description.

COPYContinue

Specifies how the server should react to a copy storage pool write failure for any of the copy storage pools listed in the COPYSTGPOOLS parameter. This parameter is optional. The default value is YES. When specifying the COPYCONTINUE parameter, you must also specify the COPYSTGPOOLS parameter.

Possible values are:

Yes

If the **COPYCONTINUE** parameter is set to YES, the server will stop writing to the failing copy pools for the remainder of the session, but continue storing files into the primary pool and any remaining copy pools. The copy storage pool list is active only for the life of the client session and applies to all the primary storage pools in a particular storage pool hierarchy.

For additional information about the **COPYCONTINUE** parameter, refer to the information about the simultaneous-write function in the *Administrator's Guide*.

No If the **COPYCONTINUE** parameter is set to NO, the server will fail the current transaction and discontinue the store operation.

Restrictions:

- The setting of the **COPYCONTINUE** parameter does not affect active-data pools. If a write failure occurs for any of the active-data pools, the server stops writing to the failing active-data pool for the remainder of the session, but

DEFINE STGPOOL

continues storing files into the primary pool and any remaining active-data pools and copy storage pools. The active-data pool list is active only for the life of the session and applies to all the primary storage pools in a particular storage pool hierarchy.

- The setting of the **COPYCONTINUE** parameter does not affect the simultaneous-write function during server import. If data is being written simultaneously and a write failure occurs to the primary storage pool or any copy storage pool, the server import process fails.
- The setting of the **COPYCONTINUE** parameter does not affect the simultaneous-write function during server data migration. If data is being written simultaneously and a write failure occurs to any copy storage pool or active-data pool, the failing storage pool is removed and the data migration process continues. Write failures to the primary storage pool cause the migration process to fail.

ACTIVEDATApools

Specifies the names of active-data pools where the server simultaneously writes data during a client backup operation. The **ACTIVEDATAPOOLS** parameter is optional. Spaces between the names of the active-data pools are not permitted.

The combined total number of storage pools specified in the **COPYSGTPOOLS** and **ACTIVEDATAPOOLS** parameters can not exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of active-data pools from the destination storage pool specified in the copy group. The primary storage pool is specified by the copy group of the management class that is bound to the data. For details, refer to information about the simultaneous-write function in the *Administrator's Guide*.

The server can write data simultaneously to active-data pools only during backup operations by Tivoli Storage Manager backup-archive clients or application clients using the Tivoli Storage Manager API.

Restrictions:

1. This parameter is available only to primary storage pools that use **NATIVE** or **NONBLOCK** data format. This parameter is not available for storage pools that use the following data formats:
 - **NETAPPDUMP**
 - **CELERRADUMP**
 - **NDMPDUMP**
2. Writing data simultaneously to active-data pools is not supported when using LAN-free data movement. Simultaneous-write operations take precedence over LAN-free data movement, causing the operations to go over the LAN. However, the simultaneous-write configuration is honored.
3. The simultaneous-write function is not supported when a NAS backup operation is writing a TOC file. If the primary storage pool specified in the **TOCDESTINATION** in the copy group of the management class has active-data pools defined, the active-data pools are ignored and the data is stored into the primary storage pool only.
4. Data being imported will not be stored in active-data pools. After an import operation, use the **COPY ACTIVE DATA** command to store the imported data in an active-data pool.

Attention: The function provided by the `ACTIVEDATAPOOLS` parameter is not intended to replace the `COPY ACTIVE DATA` command. If you use the `ACTIVEDATAPOOLS` parameter, use the `COPY ACTIVE DATA` command to ensure that the active-data pools contain all active data of the primary storage pool.

SHRED

Specifies whether data will be physically overwritten when it is deleted. This parameter is optional. You can specify an integer from 0 to 10. The default value is 0.

If you specify a value of 0, the Tivoli Storage Manager server will delete the data from the database. However, the storage used to contain the data will not be overwritten, and the data will still exist in storage until that storage is reused for other data. It might be possible to discover and reconstruct the data after it has been deleted.

If you specify a value greater than 0, the Tivoli Storage Manager server will delete the data both logically and physically. The server will overwrite the storage used to contain the data the specified number of times. This increases the difficulty of discovering and reconstructing the data after it has been deleted.

To ensure that all copies of the data are shredded, specify a `SHRED` value greater than 0 for the storage pool specified in the `NEXTSTGPOOL` parameter, and do not specify either the `COPYSTGPOOLS` or `ACTIVEDATAPOOLS`. Specifying relatively high values for the overwrite count will generally improve the level of security, but could affect performance adversely.

Overwriting of deleted data is performed asynchronously after the delete operation is complete. Therefore, the space occupied by the deleted data will remain occupied for some period of time and will not be available as free space for new data.

A `SHRED` value greater than zero cannot be used if the value of the `CACHE` parameter is `YES`.

Important: After an export operation has finished identifying files for export, any changes to the storage pool `SHRED` value is ignored. An export operation that is suspended retains the original `SHRED` value throughout the operation. You might want to consider cancelling your export operation if changes to the storage pool `SHRED` value jeopardize the operation. You can reissue the export command after any needed cleanup.

Example: Define a primary storage pool for a DISK device class

Define a primary storage pool, `POOL1`, to use the `DISK` device class, with caching enabled. Limit the maximum file size to 5 MB. Store any files larger than 5 MB in subordinate storage pools beginning with the `PROG2` storage pool. Set the high migration threshold to 70 percent, and the low migration threshold to 30 percent.

```
define stgpool pool1 disk
description="main disk storage pool" maxsize=5m
highmig=70 lowmig=30 cache=yes
nextstgpool=prog2
```

DEFINE STGPOOL


DEFINE STGPOOL (Define a primary storage pool assigned to sequential access devices)


Use this command to define a primary storage pool assigned to sequential access devices.

Privilege class

To issue this command, you must have system privilege.

Syntax





Notes:

- 1 This parameter is not available for storage pools that use the data formats NETAPPDUMP, CELERRADUMP, or NDMPDUMP.
- 2 The RECLAMATIONTYPE=SNAPLOCK setting is valid only for storage pools defined to servers that are enabled for System Storage Archive Manager. The storage pool must be assigned to a FILE device class, and the directories specified in the device class must be NetApp SnapLock volumes.
- 3 The values NETAPPDUMP, CELERRADUMP, and NDMPDUMP are not valid for storage pools defined with a FILE-type device class.
- 4 This parameter is valid only for storage pools that are defined with a FILE-type device class.

DEFINE STGPOOL

- 5 This parameter is available only when the value of the DEDUPLICATE parameter is YES.

Parameters

pool_name (Required)

Specifies the name of the storage pool to be defined. The name must be unique, and the maximum length is 30 characters.

device_class_name (Required)

Specifies the name of the device class to which this storage pool is assigned. You can specify any device class except for the DISK device class.

Pooltype=Primary

Specifies that you want to define a primary storage pool. This parameter is optional. The default value is PRIMARY.

DESCRIPTION

Specifies a description of the storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

ACCESS

Specifies how client nodes and server processes (such as migration and reclamation) can access files in the storage pool. This parameter is optional. The default value is READWRITE. Possible values are:

READWrite

Specifies that client nodes and server processes can read and write to files stored on volumes in the storage pool.

READOnly

Specifies that client nodes can only read files from the volumes in the storage pool.

Server processes can move files within the volumes in the storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool has been specified as a subordinate storage pool (with the **NEXTSTGPOOL** parameter) and is defined as *readonly*, the storage pool is skipped when server processes attempt to write files to the storage pool.

UNAVailable

Specifies that client nodes cannot access files stored on volumes in the storage pool.

Server processes can move files within the volumes in the storage pool and can also move or copy files from this storage pool to another storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool has been specified as a subordinate storage pool (with the **NEXTSTGPOOL** parameter) and is defined as *unavailable*, the storage pool is skipped when server processes attempt to write files to the storage pool.

MAXSize

Specifies the maximum size for a physical file that the server can store in the storage pool. This parameter is optional. The default value is NOLIMIT. Possible values are:

NOLimit

Specifies that there is no maximum size limit for physical files stored in the storage pool.

maximum_file_size

Limits the maximum physical file size. Specify an integer from 1 to 999999, followed by a scale factor. For example, MAXSIZE=5G specifies that the maximum file size for this storage pool is 5 gigabytes. Scale factors are:

Scale factor **Meaning**

- K kilobyte
- M megabyte
- G gigabyte
- T terabyte

If a file exceeds the maximum size and no pool is specified as the next storage pool in the hierarchy, the server does not store the file. If a file exceeds the maximum size and a pool is specified as the next storage pool, the server stores the file in the next storage pool that can accept the file size. If you specify the next storage pool parameter, at least one storage pool in your hierarchy should have no limit on the maximum size of a file. By having no limit on the size for at least one pool, you ensure that no matter what its size, the server can store the file.

For logical files that are part of an aggregate, the server considers the size of the aggregate to be the file size. Therefore, the server does not store logical files that are smaller than the maximum size limit if the files are part of an aggregate that is larger than the maximum size limit.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

CRCDATA

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is only valid for NATIVE data format storage pools. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an **AUDIT VOLUME** command you can continually ensure the integrity of data stored in your storage hierarchy. Possible values are:

Yes

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because additional overhead is required to calculate and compare CRC values between the storage pool and the server.

No Specifies that data is stored without CRC information.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

DEFINE STGPOOL

Tip: For storage pools that are associated with the 3592, LTO, or ECARTRIDGE device type, logical block protection provides better protection against data corruption than CRC validation for a storage pool. If you specify CRC validation for a storage pool, data is validated only during volume auditing operations. Errors are identified after data is written to tape.

To enable logical block protection, specify a value of READWRITE for the **LBPROTECT** parameter on the **DEFINE DEVCLASS** and **UPDATE DEVCLASS** commands for the 3592, LTO, or ECARTRIDGE device types. Logical block protection is supported only on the following types of drives and media:

- IBM LTO5 and later.
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.
- Oracle StorageTek T10000C drives.

NEXTstgpool

Specifies a primary storage pool to which files are migrated. You cannot migrate data from a sequential access storage pool to a random access storage pool. This parameter is optional.

If this storage pool does not have a next storage pool, the server cannot migrate files from this storage pool and cannot store files that exceed the maximum size for this storage pool in another storage pool.

When there is insufficient space available in the current storage pool, the **NEXTSTGPOOL** parameter for sequential access storage pools does not allow data to be stored into the next pool. In this case the server issues a message and the transaction fails.

For next storage pools with a device type of **FILE**, the server performs a preliminary check to determine whether sufficient space is available. If space is not available, the server skips to the next storage pool in the hierarchy. If space is available, the server attempts to store data in that pool. However, it is possible that the storage operation could fail because, at the time the actual storage operation is attempted, the space is no longer available.

You cannot create a chain of storage pools that leads to an endless loop through the **NEXTSTGPOOL** parameter. At least one storage pool in the hierarchy must have no value specified for **NEXTSTGPOOL**.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

If you specify a sequential access pool as the **NEXTSTGPOOL**, the pool can only be **NATIVE** or **NONBLOCK** dataformat.

HIGHmig

Specifies that the server starts migration when storage pool utilization reaches this percentage. For sequential-access disk (**FILE**) storage pools, utilization is the ratio of data in a storage pool to the pool's total estimated data capacity, including the capacity of all scratch volumes specified for the pool. For storage pools that use tape or optical media, utilization is the ratio of volumes that contain data to the total number of volumes in the storage pool. The total

number of volumes includes the maximum number of scratch volumes. This parameter is optional. You can specify an integer from 0 to 100. The default value is 90.

When the storage pool exceeds the high migration threshold, the server can start migration of files by volume to the next storage pool defined for the pool. You can set the high migration threshold to 100 to prevent migration for the storage pool.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

LOWmig

Specifies that the server stops migration when storage pool utilization is at or below this percentage. For sequential-access disk (FILE) storage pools, utilization is the ratio of data in a storage pool to the pool's total estimated data capacity, including the capacity of all scratch volumes specified for the pool. For storage pools that use tape or optical media, utilization is the ratio of volumes that contain data to the total number of volumes in the storage pool. The total number of volumes includes the maximum number of scratch volumes. This parameter is optional. You can specify an integer from 0 to 99. The default value is 70.

When the storage pool reaches the low migration threshold, the server does not start migration of files from another volume. You can set the low migration threshold to 0 to permit migration to empty the storage pool.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

RECLAIM

Specifies when the server reclaims a volume, based on the percentage of reclaimable space on a volume. Reclaimable space is the amount of space occupied by files that are expired or deleted from the Tivoli Storage Manager database.

Reclamation makes the fragmented space on volumes usable again by moving any remaining unexpired files from one volume to another volume, thus making the original volume available for reuse. This parameter is optional. You can specify an integer from 1 to 100. The default value is 60, except for storage pools that use WORM devices.

When determining which volumes in a storage pool to reclaim, the Tivoli Storage Manager server first determines the reclamation threshold indicated by the **RECLAIM**. The server then examines the percentage of reclaimable space for each volume in the storage pool. If the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool, the volume is a candidate for reclamation.

For example, suppose storage pool FILEPOOL has a reclamation threshold of 70 percent. This value indicates that the server can reclaim any volume in the

DEFINE STGPOOL

storage pool that has a percentage of reclaimable space that is greater than 70 percent. The storage pool has three volumes:

- FILEVOL1 with 65 percent reclaimable space
- FILEVOL2 with 80 percent reclaimable space
- FILEVOL3 with 95 percent reclaimable space

When reclamation begins, the server compares the percent of reclaimable space for each volume with the reclamation threshold of 70 percent. In this example, FILEVOL2 and FILEVOL3 are candidates for reclamation because their percentages of reclaimable space are greater than 70. To determine the percentage of reclaimable space for a volume, issue the **QUERY VOLUME** command and specify **FORMAT=DETAILED**. The value in the field Pct. Reclaimable Space is the percentage of reclaimable space for the volume.

Specify a value of 50 percent or greater for this parameter so that files stored on two volumes can be combined onto a single output volume.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

RECLAIMProcess

Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool. This parameter is optional. Enter a value from 1 to 999. The default value is 1.

When calculating the value for this parameter, consider the number of sequential storage pools that will be involved with the reclamation and the number of logical and physical drives that can be dedicated to the operation. To access a sequential access volume, IBM Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the reclamation.

For example, suppose that you want to reclaim the volumes from two sequential storage pools simultaneously and that you want to specify four processes for each of the storage pools. The storage pools have the same device class. Assuming that the RECLAIMSTGPOOL parameter is not specified or that the reclaim storage pool has the same device class as the storage pool being reclaimed, each process requires two mount points and, if the device type is not FILE, two drives. (One of the drives is for the input volume, and the other drive is for the output volume.) To run eight reclamation processes simultaneously, you need a total of at least 16 mount points and 16 drives. The device class for the storage pools must have a mount limit of at least 16.

If the number of reclamation processes you specify is more than the number of available mount points or drives, the processes that do not obtain mount points or drives will wait for mount points or drives to become available. If mount points or drives do not become available within the MOUNTWAIT time, the reclamation processes will end. For information about specifying the MOUNTWAIT time, see "DEFINE DEVCLASS (Define a device class)" on page 148.

The Tivoli Storage Manager server will start the specified number of reclamation processes regardless of the number of volumes that are eligible for

reclamation. For example, if you specify ten reclamation processes and only six volumes are eligible for reclamation, the server will start ten processes and four of them will complete without processing a volume.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

RECLAIMSTGpool

Specifies another primary storage pool as a target for reclaimed data from this storage pool. This parameter is optional. When the server reclaims volumes for the storage pool, the server moves unexpired data from the volumes being reclaimed to the storage pool named with this parameter.

A reclaim storage pool is most useful for a storage pool that has only one drive in its library. When you specify this parameter, the server moves all data from reclaimed volumes to the reclaim storage pool regardless of the number of drives in the library.

To move data from the reclaim storage pool back to the original storage pool, use the storage pool hierarchy. Specify the original storage pool as the next storage pool for the reclaim storage pool.

Restriction:

- This parameter is not available for storage pools that use the following data formats:
- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

RECLAMATIONType

Specifies the method by which volumes are reclaimed and managed. This parameter is optional. The default value is THRESHOLD. Possible values are the following:

THRESHold

Specifies that volumes belonging to this storage pool are reclaimed based on the threshold value in the RECLAIM attribute for this storage pool.

SNAPlock

Specifies that FILE volumes belonging to this storage pool will be managed for retention using NetApp Data ONTAP software and NetApp SnapLock volumes. This parameter is only valid for storage pools being defined to a server that has data retention protection enabled and that is assigned to a FILE device class. Volumes in this storage pool are not reclaimed based on threshold; the RECLAIM value for the storage pool is ignored.

All volumes in this storage pool are created as FILE volumes. A retention date, derived from the retention attributes in the archive copy group for the storage pool, is set in the metadata for the FILE volume using the SnapLock feature of the NetApp Data ONTAP operating system. Until the retention date has expired, the FILE volume and any data on it cannot be deleted from the physical SnapLock volume on which it is stored.

The RECLAMATIONTYPE parameter for all storage pools being defined must be the same when defined to the same device class name. The **DEFINE**

DEFINE STGPOOL

command will fail if the RECLAMATIONTYPE parameter specified is different than what is currently defined for storage pools already defined to the device class name.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

COLlocate

Specifies whether the server attempts to keep data belonging to a single client node, group of client nodes, or client file space stored on as few volumes as possible. This parameter is optional. The default value is GROUP.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time needed to collocate files for storing and the number of volumes required. For details, see the *Administrator's Guide*.

Possible values are:

No Specifies that collocation is disabled.

GRoup

Specifies that collocation is enabled at the group level for client nodes. The server attempts to put data for nodes that belong to the same collocation group on as few volumes as possible. If the nodes in the collocation group have multiple file spaces, the server does not attempt to collocate those file spaces.

If you specify COLLOCATE=GROUP but do not define any collocation groups or if you specify COLLOCATE=GROUP but do not add nodes to a collocation group, data is collocated by node. Be sure to consider tape usage when organizing client nodes into collocation groups. For example, if a tape-based storage pool consists of data from grouped and ungrouped nodes and you specify COLLOCATE=GROUP, the server performs the following actions:

- Collocates by group the data for grouped nodes only. Whenever possible, the server collocates data that belongs to a group of nodes on a single tape or on as few tapes as possible. Data for a single node can also be spread across several tapes associated with a group.
- Collocates by node the data for ungrouped nodes. Whenever possible, the server stores the data for a single node on a single tape. All available tapes that already have data for the node are used before available space on any other tape is used.

NODe

Specifies that collocation is enabled at the client node level. The server attempts to put data for one node on as few volumes as possible. If the node has multiple file spaces, the server does not attempt to collocate those file spaces. For backward compatibility, COLLOCATE=YES is still accepted by the server to specify collocation at the client node level.

If a storage pool contains data for a node that is a member of a collocation group and you specify COLLOCATE=NODE, the data will be collocated by node not by group.

Filespace

Specifies that collocation is enabled at the file space level for client nodes. The server attempts to put data for one node and file space on as few volumes as possible. If a node has multiple file spaces, the server attempts to put data for different file spaces on different volumes.

MAXSCRatch (Required)

Specifies the maximum number of scratch volumes that the server can request for this storage pool. You can specify an integer from 0 to 100000000. By allowing the server to request scratch volumes, you avoid having to define each volume to be used.

The value specified for this parameter is used to estimate the total number of volumes available in the storage pool and the corresponding estimated capacity for the storage pool.

Scratch volumes are automatically deleted from the storage pool when they become empty. When scratch volumes with the device type of FILE are deleted, the space that the volumes occupied is freed by the server and returned to the file system.

Tip: For server-to-server operations that use virtual volumes and that store a small amount of data, consider specifying a value for the MAXSCRATCH parameter that is higher than the value you typically specify for write operations to other types of volumes. After a write operation to a virtual volume, Tivoli Storage Manager marks the volume as FULL, even if the value of the MAXCAPACITY parameter on the device-class definition has not been reached. The Tivoli Storage Manager server does not keep virtual volumes in FILLING status and does not append to them. If the value of the MAXSCRATCH parameter is too low, server-to-server operations can fail.

REUsedelay

Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer from 0 to 9999. The default value is 0, which means that a volume can be rewritten or returned to the scratch pool as soon as all the files are deleted from the volume.

Important: Use this parameter to help ensure that when you restore the database to an earlier level, database references to files in the storage pool are still valid. You must set this parameter to a value greater than the number of days you plan to retain the oldest database backup. The number of days specified for this parameter should be the same as the number specified for the **SET DRMDBBACKUPEXPIREDAYS** command. For more information, see the *Administrator's Guide*.

OVFL0cation

Specifies the overflow location for the storage pool. The server assigns this location name to a volume that is ejected from the library by the command. This parameter is optional. The location name can be a maximum length of 255 characters. Enclose the location name in quotation marks if the location name contains any blank characters.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

DEFINE STGPOOL

MIGDelay

Specifies the minimum number of days a file must remain in a storage pool before it becomes eligible for migration. All files on a volume must be eligible for migration before the server selects the volume for migration. To calculate a value to compare to the specified MIGDELAY, the server counts the number of days that the file has been in the storage pool.

This parameter is optional. You can specify an integer from 0 to 9999. The default is 0, which means that you do not want to delay migration. If you want the server to count the number of days based only on when a file was stored and not when it was retrieved, use the NORETRIEVEDATE server option.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

MIGContinue

Specifies whether you allow the server to migrate files that do not satisfy the migration delay time. This parameter is optional. The default is YES.

Because you can require that files remain in the storage pool for a minimum number of days, the server may migrate all eligible files to the next storage pool yet not meet the low migration threshold. This parameter allows you to specify whether the server is allowed to continue the migration process by migrating files that do not satisfy the migration delay time.

Possible values are:

Yes

Specifies that, when necessary to meet the low migration threshold, the server continues to migrate files that do not satisfy the migration delay time.

If you allow more than one migration process for the storage pool, some files that do not satisfy the migration delay time may be migrated unnecessarily. As one process migrates files that satisfy the migration delay time, a second process could begin migrating files that do not satisfy the migration delay time to meet the low migration threshold. The first process that is still migrating files that satisfy the migration delay time might have, by itself, caused the low migration threshold to be met.

No Specifies that the server stops migration when no eligible files remain to be migrated, even before reaching the low migration threshold. The server does not migrate files unless the files satisfy the migration delay time.

MIGProcess

Specifies the number of parallel processes to use for migrating the files from the volumes in this storage pool. This parameter is optional. Enter a value from 1 to 999. The default value is 1.

When calculating the value for this parameter, consider the number of sequential storage pools that will be involved with the migration, and the number of logical and physical drives that can be dedicated to the operation. To access a sequential-access volume, Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager

and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the migration.

For example, suppose you want to simultaneously migrate the files from volumes in two primary sequential storage pools and that you want to specify three processes for each of the storage pools. The storage pools have the same device class. Assuming that the storage pool to which files are being migrated has the same device class as the storage pool from which files are being migrated, each process requires two mount points and, if the device type is not FILE, two drives. (One drive is for the input volume, and the other drive is for the output volume.) To run six migration processes simultaneously, you need a total of at least 12 mount points and 12 drives. The device class for the storage pools must have a mount limit of at least 12.

If the number of migration processes you specify is more than the number of available mount points or drives, the processes that do not obtain mount points or drives will wait for mount points or drives to become available. If mount points or drives do not become available within the MOUNTWAIT time, the migration processes will end. For information about specifying the MOUNTWAIT time, see “DEFINE DEVCLASS (Define a device class)” on page 148.

The Tivoli Storage Manager server will start the specified number of migration processes regardless of the number of volumes that are eligible for migration. For example, if you specify ten migration processes and only six volumes are eligible for migration, the server will start ten processes and four of them will complete without processing a volume.

Tip: When specifying this parameter, consider whether the simultaneous-write function is enabled for server data migration. Each migration process requires a mount point and a drive for each copy storage pool and active-data pool defined to the target storage pool.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

DATAFormat

Specifies the data format to use to back up files to this storage pool and restore files from this storage pool. The default format is the NATIVE server format. Possible values are:

NATive

Specifies the data format is the native Tivoli Storage Manager server format and includes block headers.

NONblock

Specifies the data format is the native Tivoli Storage Manager server format and does not include block headers.

Important: The default minimum block size on a volume associated with a FILE device class is 256 KB, regardless how much data is being written to the volume. For certain tasks (for example, using content-management products, using the DIRMC client option to store directory information, or migrating very small files using Tivoli Storage Manager for Space Management or Tivoli Storage Manager HSM for Windows), you can

DEFINE STGPOOL

minimize wasted space on storage volumes by specifying the **NONBLOCK** data format. In most situations, however, the **NATIVE** format is preferred.

NETAPPDump

Specifies the data is in a NetApp dump format. This data format should be specified for file system images that are in a dump format and that have been backed up from a NetApp or an IBM System Storage N Series file server using NDMP. The server will not perform migration, reclamation, or **AUDIT VOLUME** for a storage pool with **DATAFORMAT=NETAPPDUMP**. You can use the **MOVE DATA** command to move data from one primary storage pool to another, or out of a volume if the volume needs to be reused.

CELERRADump

Specifies that the data is in an EMC Celerra dump format. This data format should be specified for file system images that are in a dump format and that have been backed up from an EMC Celerra file server using NDMP. The server will not perform migration, reclamation, or **AUDIT VOLUME** for a storage pool with **DATAFORMAT=CELERRADUMP**. You can use the **MOVE DATA** command to move data from one primary storage pool to another, or out of a volume if the volume needs to be reused.

NDMPDump

Specifies that the data is in NAS vendor-specific backup format. Use this data format for file system images that have been backed up from a NAS file server other than a NetApp or EMC Celerra file server. The server will not perform migration, reclamation, or **AUDIT VOLUME** for a storage pool with **DATAFORMAT=NDMPDUMP**. You can use the **MOVE DATA** command to move data from one primary storage pool to another, or out of a volume if the volume needs to be reused.

AUTOCopy

Specifies when Tivoli Storage Manager performs simultaneous-write operations. The default value is **CLIENT**. This parameter is optional and affects the following operations:

- Client store sessions
- Server import processes
- Server data-migration processes

If an error occurs while data is being simultaneously written to a copy storage pool or active-data pool during a migration process, the server stops writing to the failing storage pools for the remainder of the process. However, the server continues to store files into the primary storage pool and any remaining copy storage pools or active-data pools. These pools remain active for the duration of the migration process. Copy storage pools are specified using the **COPYSTGPOOLS** parameter. Active-data pools are specified using the **ACTIVEDATAPOOLS** parameter.

Possible values are:

None

Specifies that the simultaneous-write function is disabled.

CLient

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions or server import processes. During server import processes, data is written simultaneously to only copy storage pools. Data is not written to active-data pools during server import processes.

MIGRation

Specifies that data is written simultaneously to copy storage pools and active-data pools only during migration to this storage pool. During server data-migration processes, data is written simultaneously to copy storage pools and active-data pools only if the data does not exist in those pools. Nodes whose data is being migrated must be in a domain associated with an active-data pool. If the nodes are not in a domain associated with an active pool, the data cannot be written to the pool.

A11

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions, server import processes, or server data-migration processes. Specifying this value ensures that data is written simultaneously whenever this pool is a target for any of the eligible operations.

COPYSTGpools

Specifies the names of copy storage pools where the server simultaneously writes data. The COPYSTGPOOLS parameter is optional. You can specify a maximum of three copy pool names separated by commas. (In versions earlier than Version 5 Release 3, the maximum number was ten.) Spaces between the names of the copy pools are not permitted. When specifying a value for the COPYSTGPOOLS parameter, you can also specify a value for the COPYCONTINUE parameter.

The combined total number of storage pools specified in the COPYSTGPOOLS and ACTIVEDATAPOOLS parameters cannot exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of copy storage pools and the COPYCONTINUE value from the primary storage pool. The primary storage pool is specified by the copy group of the management class that is bound to the data. For details, refer to information about the simultaneous-write function in the *Administrator's Guide*.

The server can write data simultaneously to copy storage pools during the following operations:

- Backup and archive operations by Tivoli Storage Manager backup-archive clients or application clients using the Tivoli Storage Manager API
- Migration operations by Tivoli Storage Manager for Space Management clients
- Import operations that involve copying exported file data from external media to a storage pool defined with a copy storage pool list

Restrictions:

1. This parameter is available only to primary storage pools that use NATIVE or NONBLOCK data format. This parameter is not available for storage pools that use the following data formats:
 - NETAPPDUMP
 - CELERRADUMP
 - NDMPDUMP
2. Writing data simultaneously to copy storage pools is not supported when using LAN-free data movement. Simultaneous-write operations take precedence over LAN-free data movement, causing the operations to go over the LAN. However, the simultaneous-write configuration is honored.

DEFINE STGPOOL

3. The simultaneous-write function is not supported for NAS backup operations. If the primary storage pool specified in the `DESTINATION` or `TOCDESTINATION` in the copy group of the management class has copy storage pools defined, the copy storage pools are ignored and the data is stored into the primary storage pool only.

Attention: The function provided by the `COPYSTGPOOLS` parameter is not intended to replace the `BACKUP STGPOOL` command. If you use the `COPYSTGPOOLS` parameter, continue to use the `BACKUP STGPOOL` command to ensure that the copy storage pools are complete copies of the primary storage pool. There are cases when a copy may not be created. For more information, see the `COPYCONTINUE` parameter description.

COPYContinue

Specifies how the server should react to a copy storage pool write failure for any of the copy storage pools listed in the `COPYSTGPOOLS` parameter. This parameter is optional. The default value is `YES`. When specifying the `COPYCONTINUE` parameter, you must also specify the `COPYSTGPOOLS` parameter.

The `COPYCONTINUE` parameter has no effect on the simultaneous-write function during migration.

Possible values are:

Yes

If the `COPYCONTINUE` parameter is set to `YES`, the server will stop writing to the failing copy pools for the remainder of the session, but continue storing files into the primary pool and any remaining copy pools. The copy storage pool list is active only for the life of the client session and applies to all the primary storage pools in a particular storage pool hierarchy.

For additional information about the `COPYCONTINUE` parameter, refer to the information about the simultaneous-write function in the *Administrator's Guide*.

- No** If the `COPYCONTINUE` parameter is set to `NO`, the server will fail the current transaction and discontinue the store operation.

Restrictions:

- The setting of the `COPYCONTINUE` parameter does not affect active-data pools. If a write failure occurs for any of the active-data pools, the server stops writing to the failing active-data pool for the remainder of the session, but continues storing files into the primary pool and any remaining active-data pools and copy storage pools. The active-data pool list is active only for the life of the session and applies to all the primary storage pools in a particular storage pool hierarchy.
- The setting of the `COPYCONTINUE` parameter does not affect the simultaneous-write function during server import. If data is being written simultaneously and a write failure occurs to the primary storage pool or any copy storage pool, the server import process fails.
- The setting of the `COPYCONTINUE` parameter does not affect the simultaneous-write function during server data migration. If data is being written simultaneously and a write failure occurs to any copy storage pool or active-data pool, the failing storage pool is removed and the data migration process continues. Write failures to the primary storage pool cause the migration process to fail.

Restriction: This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

ACTIVEDATApools

Specifies the names of active-data pools where the server simultaneously writes data during a client backup operation. The **ACTIVEDATAPOOLS** parameter is optional. Spaces between the names of the active-data pools are not permitted.

The combined total number of storage pools specified in the **COPYSGTPOOLS** and **ACTIVEDATAPOOLS** parameters can not exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of active-data pools from the destination storage pool specified in the copy group. The primary storage pool is specified by the copy group of the management class that is bound to the data. For details, refer to information about the simultaneous-write function in the *Administrator's Guide*.

The server can write data simultaneously to active-data pools only during backup operations by Tivoli Storage Manager backup-archive clients or application clients using the Tivoli Storage Manager API.

Restrictions:

1. This parameter is available only to primary storage pools that use **NATIVE** or **NONBLOCK** data format. This parameter is not available for storage pools that use the following data formats:
 - NETAPPDUMP
 - CELERRADUMP
 - NDMPDUMP
2. Write data simultaneously to active-data pools is not supported when using LAN-free data movement. Simultaneous-write operations take precedence over LAN-free data movement, causing the operations to go over the LAN. However, the simultaneous-write configuration is honored.
3. The simultaneous-write function is not supported when a NAS backup operation is writing a TOC file. If the primary storage pool specified in the **TOCDESTINATION** in the copy group of the management class has active-data pools defined, the active-data pools are ignored, and the data is stored into the primary storage pool only.
4. Data being imported will not be stored in active-data pools. After an import operation, use the **COPY ACTIVE DATA** command to store the imported data in an active-data pool.

Attention: The function provided by the **ACTIVEDATAPOOLS** parameter is not intended to replace the **COPY ACTIVE DATA** command. If you use the **ACTIVEDATAPOOLS** parameter, use the **COPY ACTIVE DATA** command to ensure that the active-data pools contain all active data of the primary storage pool.

DEDUPLICATE

Specifies whether the data that is stored in this storage pool will be deduplicated. This parameter is optional and is valid only for storage pools that are defined with a **FILE-type** device class. The default value is **NO**.

DEFINE STGPOOL

IDENTIFYPROcess

Specifies the number of parallel processes to use for server-side duplicate identification. This parameter is optional and is valid only for storage pools that are defined with a FILE device class. Enter a value from 0 to 20. The default value is 1. If the value of the DEDUPLICATE parameter is NO, the default setting for IDENTIFYPROCESS has no effect.

When calculating the value for this parameter, consider the workload on the server and the amount of data requiring data deduplication. Server-side duplicate identification requires disk I/O and processor resources, so the more processes you allocate to data deduplication, the heavier the workload that you place on your system. In addition, consider the number of volumes that require processing. Server-side duplicate-identification processes work on volumes containing data that requires deduplication. If you update a storage pool, specifying that the data in the storage pool is to be deduplicated, all the volumes in the pool require processing. For this reason, you might have to define a high number of duplicate-identification processes initially. Over time, however, as existing volumes are processed, only the volumes containing new data have to be processed. When that happens, you can reduce the number of duplicate-identification processes.

Remember: Duplicate-identification processes can be either active or idle. Processes that are working on files are active. Processes that are waiting for files to work on are idle. Processes remain idle until volumes with data to be deduplicated become available. The output of the **QUERY PROCESS** command for a duplicate-identification process includes the total number of bytes and files that have been processed since the process first started. For example, if a duplicate-identification process processes four files, becomes idle, and then processes five more files, then the total number of files processed is nine. Processes end only when canceled or when the number of duplicate-identification processes for the storage pool is changed to a value less than the number currently specified.

Example: Define a primary storage pool with an 8MMTAPE device class

Define a primary storage pool named 8MMPOOL to the 8MMTAPE device class (with a device type of 8MM) with a maximum file size of 5 MB. Store any files larger than 5 MB in subordinate pools, beginning with POOL1. Enable collocation of files for client nodes. Allow as many as 5 scratch volumes for this storage pool.

```
define stgpool 8mmpool 8mmtape maxsize=5m
nextstgpool=pool1 collocate=node
maxscratch=5
```

DEFINE STGPOOL (Define a copy storage pool assigned to sequential access devices)


Use this command to define a copy storage pool assigned to sequential access devices.

Privilege class

To issue this command, you must have system privilege.

Syntax

►► DEFINE STGpool *pool_name* *device_class_name* POOLtype == COPY



Notes:

- 1 The `RECLAMATIONTYPE=SNAPLOCK` setting is valid only for storage

DEFINE STGPOOL

pools defined to servers that are enabled for System Storage Archive Manager. The storage pool must be assigned to a FILE device class, and the directories specified in the device class must be NetApp SnapLock volumes.

- 2 The values NETAPPDUMP, CELERRADUMP, and NDMPDUMP are not valid for storage pools that are defined with a FILE device class.
- 3 This parameter is valid only for storage pools that are defined with a FILE device class.
- 4 This parameter is available only when the value of the DEDUPLICATE parameter is YES.

Parameters

pool_name (Required)

Specifies the name of the storage pool to be defined. The name must be unique, and the maximum length is 30 characters.

device_class_name (Required)

Specifies the name of the sequential access device class to which this copy storage pool is assigned. You can specify any device class except DISK.

Pooltype=Copy (Required)

Specifies that you want to define a copy storage pool.

DESCRIPTION

Specifies a description of the copy storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

ACCESS

Specifies how client nodes and server processes (such as reclamation) can access files in the copy storage pool. This parameter is optional. The default value is READWRITE. Possible values are:

READWrite

Specifies that files can be read from and written to the volumes in the copy storage pool.

READOnly

Specifies that client nodes can only read files stored on the volumes in the copy storage pool.

Server processes can move files within the volumes in the storage pool. The server can use files in the copy storage pool to restore files to primary storage pools. However, no new writes are permitted to volumes in the copy storage pool from volumes outside the storage pool. A storage pool cannot be backed up to the copy storage pool.

UNAVailable

Specifies that client nodes cannot access files stored on volumes in the copy storage pool.

Server processes can move files within the volumes in the storage pool. The server can use files in the copy storage pool to restore files to primary storage pools. However, no new writes are permitted to volumes in the copy storage pool from volumes outside the storage pool. A storage pool cannot be backed up to the copy storage pool.

COLlocate

Specifies whether the server attempts to keep data belonging to a single client

node, group of client nodes, or client file space stored on as few volumes as possible. This parameter is optional. The default value is NO.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time needed to collocate files for storing and the number of volumes required. For details, see the *Administrator's Guide*.

Possible values are:

No Specifies that collocation is disabled.

GRoup

Specifies that collocation is enabled at the group level for client nodes. The server attempts to put data for nodes that belong to the same collocation group on as few volumes as possible. If the nodes in the collocation group have multiple file spaces, the server does not attempt to collocate those file spaces.

If you specify COLLOCATE=GROUP but do not define any collocation groups or if you specify COLLOCATE=GROUP but do not add nodes to a collocation group, data is collocated by node. Be sure to consider tape usage when organizing client nodes into collocation groups. For example, if a tape-based storage pool consists of data from grouped and ungrouped nodes and you specify COLLOCATE=GROUP, the server performs the following actions:

- Collocates by group the data for grouped nodes only. Whenever possible, the server collocates data that belongs to a group of nodes on a single tape or on as few tapes as possible. Data for a single node can also be spread across several tapes associated with a group.
- Collocates by node the data for ungrouped nodes. Whenever possible, the server stores the data for a single node on a single tape. All available tapes that already have data for the node are used before available space on any other tape is used.

NODe

Specifies that collocation is enabled at the client node level. The server attempts to put data for one node on as few volumes as possible. If the node has multiple file spaces, the server does not attempt to collocate those file spaces. For backward compatibility, COLLOCATE=YES is still accepted by the server to specify collocation at the client node level.

If a storage pool contains data for a node that is a member of a collocation group and you specify COLLOCATE=NODE, the data will be collocated by node not by group.

FIlespace

Specifies that collocation is enabled at the file space level for client nodes. The server attempts to put data for one node and file space on as few volumes as possible. If a node has multiple file spaces, the server attempts to put data for different file spaces on different volumes.

REClaim

Specifies when the server reclaims a volume, based on the percentage of reclaimable space on a volume. Reclaimable space is the amount of space occupied by files that are expired or deleted from the Tivoli Storage Manager database.

Reclamation makes the fragmented space on volumes usable again by moving any remaining unexpired files from one volume to another volume, thus

DEFINE STGPOOL

making the original volume available for reuse. This parameter is optional. You can specify an integer from 1 to 100. The default value is 100, which means that reclamation is not performed.

When determining which volumes in a storage pool to reclaim, the Tivoli Storage Manager server first determines the reclamation threshold indicated by the **RECLAIM**. The server then examines the percentage of reclaimable space for each volume in the storage pool. If the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool, the volume is a candidate for reclamation.

For example, suppose storage pool FILEPOOL has a reclamation threshold of 70 percent. This value indicates that the server can reclaim any volume in the storage pool that has a percentage of reclaimable space that is greater than 70 percent. The storage pool has three volumes:

- FILEVOL1 with 65 percent reclaimable space
- FILEVOL2 with 80 percent reclaimable space
- FILEVOL3 with 95 percent reclaimable space

When reclamation begins, the server compares the percent of reclaimable space for each volume with the reclamation threshold of 70 percent. In this example, FILEVOL2 and FILEVOL3 are candidates for reclamation because their percentages of reclaimable space are greater than 70. To determine the percentage of reclaimable space for a volume, issue the **QUERY VOLUME** command and specify **FORMAT=DETAILED**. The value in the field Pct. Reclaimable Space is the percentage of reclaimable space for the volume.

If you change the value from the default, specify a value of 50 percent or greater so that files stored on two volumes can be combined onto a single output volume.

When a copy pool volume that is offsite becomes eligible for reclamation, the reclamation process attempts to obtain the unexpired files on the reclaimable volume from a primary or copy storage pool that is onsite. The process then writes these files to an available volume in the original copy storage pool. Effectively, these files are moved back to the onsite location. However, the files could be obtained from the offsite volume after a disaster if a database backup is used that references the files on the offsite volume. Because of the way reclamation works with offsite volumes, use it carefully with copy storage pools.

RECLAIMProcess

Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool. This parameter is optional. Enter a value from 1 to 999. The default value is 1.

When calculating the value for this parameter, consider the number of sequential storage pools that will be involved with the reclamation and the number of logical and physical drives that can be dedicated to the operation. To access a sequential access volume, IBM Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the reclamation.

For example, suppose that you want to reclaim the volumes from two sequential storage pools simultaneously and that you want to specify four processes for each of the storage pools. The storage pools have the same device class. Each process requires two mount points and, if the device type is not

FILE, two drives. (One of the drives is for the input volume, and the other drive is for the output volume.) To run eight reclamation processes simultaneously, you need a total of at least 16 mount points and 16 drives. The device class for the storage pools must have a mount limit of at least 16.

If the number of reclamation processes you specify is more than the number of available mount points or drives, the processes that do not obtain mount points or drives will wait for mount points or drives to become available. If mount points or drives do not become available within the MOUNTWAIT time, the reclamation processes will end. For information about specifying the MOUNTWAIT time, see “DEFINE DEVCLASS (Define a device class)” on page 148.

The Tivoli Storage Manager server will start the specified number of reclamation processes regardless of the number of volumes that are eligible for reclamation. For example, if you specify ten reclamation processes and only six volumes are eligible for reclamation, the server will start ten processes and four of them will complete without processing a volume.

RECLAMATIONType

Specifies the method by which volumes are reclaimed and managed. This parameter is optional. The default value is THRESHOLD. Possible values are the following:

THRESHold

Specifies that volumes belonging to this storage pool are reclaimed based on the threshold value in the RECLAIM attribute for this storage pool.

SNAPlock

Specifies that FILE volumes belonging to this storage pool will be managed for retention using NetApp Data ONTAP software and NetApp SnapLock volumes. This parameter is only valid for storage pools being defined to a server that has data retention protection enabled and that is assigned to a FILE device class. Volumes in this storage pool are not reclaimed based on threshold; the RECLAIM value for the storage pool is ignored.

All volumes in this storage pool are created as FILE volumes. A retention date, derived from the retention attributes in the archive copy group for the storage pool, is set in the metadata for the FILE volume using the SnapLock feature of the NetApp Data ONTAP operating system. Until the retention date has expired, the FILE volume and any data on it cannot be deleted from the physical SnapLock volume on which it is stored.

The RECLAMATIONTYPE parameter for all storage pools being defined must be the same when defined to the same device class name. The **DEFINE** command will fail if the RECLAMATIONTYPE parameter specified is different than what is currently defined for storage pools already defined to the device class name.

OFFSITERECLAIMLimit

Specifies the number of offsite volumes to have their space reclaimed during reclamation for this storage pool. This parameter is optional. The default value is NOLIMIT. Possible values are:

NOLimit

Specifies that you want to have the space reclaimed in all of your offsite volumes.

DEFINE STGPOOL

number

Specifies the number of offsite volumes to have their space reclaimed. You can specify an integer from 0 to 99999. A value of zero means that none of the offsite volumes will be reclaimed.

Important: When determining the value for the `OFFSITERECLAIMLIMIT`, consider using the statistical information in the message issued at the end of the offsite volume reclamation operation. Alternatively, you can use the following Tivoli Storage Manager SQL select command to obtain the statistical information from the `SUMMARY` table for the offsite volume reclamation operation:

```
select * from summary where activity='OFFSITE RECLAMATION'
```

The statistical information includes the following items:

- The number of offsite volumes that were processed
- The number of parallel processes that were used
- The total amount of time required for the processing

The order in which offsite volumes are reclaimed is based on the amount of unused space in a volume. (Unused space includes both space that has never been used on the volume and space that has become empty because of file deletion.) Volumes with the largest amount of unused space are reclaimed first.

For example, suppose a copy storage pool contains three volumes: VOL1, VOL2, and VOL3. VOL1 has the largest amount of unused space, and VOL3 has the least amount of unused space. Suppose further that the percentage of unused space in each of the three volumes is greater than the value of the `RECLAIM` parameter. If you do not specify a value for the `OFFSITERECLAIMLIMIT` parameter, all three volumes will be reclaimed when the reclamation runs. If you specify a value of 2, only VOL1 and VOL2 will be reclaimed when the reclamation runs. If you specify a value of 1, only VOL1 will be reclaimed.

MAXSCRATCH (Required)

Specifies the maximum number of scratch volumes that the server can request for this storage pool. You can specify an integer from 0 to 100000000. By allowing the server to request scratch volumes as needed, you avoid having to define each volume to be used.

The value specified for this parameter is used to estimate the total number of volumes available in the copy storage pool and the corresponding estimated capacity for the copy storage pool.

Scratch volumes are automatically deleted from the storage pool when they become empty. However, if the access mode for a scratch volume is `OFFSITE`, the volume is not deleted from the copy storage pool until the access mode is changed. This allows an administrator to query the server for empty, offsite scratch volumes and return these to the onsite location.

When scratch volumes with the device type of `FILE` become empty and are deleted, the space that the volumes occupied is freed by the server and returned to the file system.

Tip: For server-to-server operations that use virtual volumes and that store a small amount of data, consider specifying a value for the `MAXSCRATCH` parameter that is higher than the value you typically specify for write operations to other types of volumes. After a write operation to a virtual volume, Tivoli Storage Manager marks the volume as `FULL`, even if the value of the `MAXCAPACITY` parameter on the device-class definition has not been

reached. The Tivoli Storage Manager server does not keep virtual volumes in FILLING status and does not append to them. If the value of the MAXSCRATCH parameter is too low, server-to-server operations can fail.

REUsedelay

Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer from 0 to 9999. The default value is 0, which means that a volume can be rewritten or returned to the scratch pool as soon as all the files are deleted from the volume.

Important: Use this parameter to help ensure that when you restore the database to an earlier level, database references to files in the copy storage pool are still valid. You must set this parameter to a value greater than the number of days you plan to retain the oldest database backup. The number of days specified for this parameter should be the same as the number specified for the **SET DRMDBBACKUPEXPIREDAYS** command. For more information, see the *Administrator's Guide*.

OVFLocation

Specifies the overflow location for the storage pool. The server assigns this location name to a volume that is ejected from the library by the command. This parameter is optional. The location name can be a maximum length of 255 characters. Enclose the location name in quotation marks if the location name contains any blank characters.

DATAFormat

Specifies the data format to use to back up files to this storage pool and restore files from this storage pool. The default format is the NATIVE server format. Possible values are:

NATive

Specifies the data format is the native IBM Tivoli Storage Manager server format and includes block headers.

NONBlock

Specifies the data format is the native IBM Tivoli Storage Manager server format and does not include block headers.

Important: The default minimum block size on a volume associated with a FILE device class is 256 KB, regardless how much data is being written to the volume. For certain tasks (for example, using content-management products, using the DIRMC client option to store directory information, or migrating very small files using Tivoli Storage Manager for Space Management or Tivoli Storage Manager HSM for Windows), you can minimize wasted space on storage volumes by specifying the NONBLOCK data format. In most situations, however, the NATIVE format is preferred.

NETAPPDump

Specifies that the data is in a NetApp dump format. Do not specify this data format for file system images that are in a dump format and that have been backed up from a NetApp file server using NDMP. The server will not perform storage pool reclamation or **AUDIT VOLUME** for a storage pool with DATAFORMAT=NETAPPDUMP. You can use the **MOVE DATA** command to move NDMP-generated data out of a volume if the volume needs to be re-used.

CELERRADump

Specifies that the data is in an EMC Celerra dump format. Do not specify

DEFINE STGPOOL

this data format for file system images that are in a dump format and that have been backed up from an EMC Celerra file server using NDMP. The server will not perform storage pool reclamation or **AUDIT VOLUME** for a storage pool with DATAFORMAT=CELERRADUMP. You can use the **MOVE DATA** command to move NDMP-generated data out of a volume if the volume needs to be re-used.

NDMPDump

Specifies that the data is in a NAS vendor-specific backup format. Do not specify this data format for file system images that are in a backup format and that have been backed up from a NAS file server other than a NetApp or EMC Celerra file server. The server will not perform storage pool reclamation or **AUDIT VOLUME** for a storage pool with DATAFORMAT=NDMPDUMP. You can use the **MOVE DATA** command to move NDMP-generated data out of a volume if the volume needs to be re-used.

CRCDATA

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is only valid for NATIVE data format storage pools. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an **AUDIT VOLUME** command you can continually ensure the integrity of data stored in your storage hierarchy. Possible values are:

Yes

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because additional overhead is required to calculate and compare CRC values between the storage pool and the server.

No Specifies that data is stored without CRC information.

Tip: For storage pools that are associated with the 3592, LTO, or ECARTRIDGE device type, logical block protection provides better protection against data corruption than CRC validation for a storage pool. If you specify CRC validation for a storage pool, data is validated only during volume auditing operations. Errors are identified after data is written to tape.

To enable logical block protection, specify a value of READWRITE for the **LBPROTECT** parameter on the **DEFINE DEVCLASS** and **UPDATE DEVCLASS** commands for the 3592, LTO, or ECARTRIDGE device types. Logical block protection is supported only on the following types of drives and media:

- IBM LTO5 and later.
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.
- Oracle StorageTek T10000C drives.

DEDUPLICATE

Specifies whether the data that is stored in this storage pool will be deduplicated. This parameter is optional and is valid only for storage pools that are defined with a FILE-type device class. The default value is NO.

IDENTIFYPROCESS

Specifies the number of parallel processes to use for server-side duplicate identification. This parameter is optional and is valid only for storage pools that are defined with a FILE device class. Enter a value from 0 to 20.

The default value for this parameter is 0. Duplicate-identification processes for a copy storage pool are not necessary if you specify duplicate-identification processes for the primary storage pool. When Tivoli Storage Manager analyzes a file in a storage pool, Tivoli Storage Manager also analyzes the file in all other storage pools.

When calculating the value for this parameter, consider the workload on the server and the amount of data requiring data deduplication. Server-side duplicate identification requires disk I/O and processor resources, so the more processes you allocate to data deduplication, the heavier the workload that you place on your system. In addition, consider the number of volumes that require processing. Server-side duplicate-identification processes work on volumes containing data that requires deduplication. If you update a storage pool, specifying that the data in the storage pool is to be deduplicated, all the volumes in the pool require processing. For this reason, you might have to define a high number of duplicate-identification processes initially. Over time, however, as existing volumes are processed, only the volumes containing new data have to be processed. When that happens, you can reduce the number of duplicate-identification processes.

Remember: Duplicate-identification processes can be either active or idle. Processes that are working on files are active. Processes that are waiting for files to work on are idle. Processes remain idle until volumes with data to be deduplicated become available. The output of the **QUERY PROCESS** command for a duplicate-identification process includes the total number of bytes and files that have been processed since the process first started. For example, if a duplicate-identification process processes four files, becomes idle, and then processes five more files, then the total number of files processed is nine. Processes end only when canceled or when the number of duplicate-identification processes for the storage pool is changed to a value less than the number currently specified.

Example: Define a copy storage pool with a DC480 device class.

Define a copy storage pool, TAPEPOOL2, to the DC480 device class. Allow up to 50 scratch volumes for this pool. Delay the reuse of volumes for 45 days.

```
define stgpool tapepool2 dc480 pooltype=copy
maxscratch=50 reusedelay=45
```

DEFINE STGPOOL

DEFINE STGPOOL (Define an active-data pool assigned to sequential-access devices)


Use this command to define an active-data pool assigned to sequential-access devices.

Privilege class

To issue this command, you must have system privilege.

Syntax

►► DEFINE STGpool *pool_name* *device_class_name* POOLtype == ACTIVEdata



Notes:

- 1 The `RECLAMATIONTYPE=SNAPLOCK` setting is valid only for storage pools defined to servers that are enabled for System Storage Archive Manager. The storage pool must be assigned to a `FILE` device class, and the directories specified in the device class must be NetApp SnapLock volumes.

- 2 This parameter is valid only for storage pools that are defined with a FILE device class.
- 3 This parameter is available only when the value of the DEDUPLICATE parameter is YES.

Parameters

pool_name (Required)

Specifies the name of the storage pool to be defined. The name must be unique, and the maximum length is 30 characters.

device_class_name (Required)

Specifies the name of the sequential access device class to which this active-data pool is assigned. You can specify any device class except DISK.

Pooltype=ACTIVEdata (Required)

Specifies that you want to define an active-data pool.

DESCRIPTION

Specifies a description of the active-data pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

ACCESS

Specifies how client nodes and server processes (such as reclamation) can access files in the active-data pool. This parameter is optional. The default value is READWRITE. Possible values are:

READWrite

Specifies that files can be read from and written to the volumes in the active-data pool.

READOnly

Specifies that client nodes can only read files stored on the volumes in the active-data pool.

Server processes can move files within the volumes in the storage pool. The server can use files in the active-data pool to restore files to primary storage pools. However, no new writes are permitted to volumes in the active-data pool from volumes outside the storage pool. A storage pool cannot be copied to the active-data pool.

UNAVailable

Specifies that client nodes cannot access files stored on volumes in the active-data pool.

Server processes can move files within the volumes in the storage pool. The server can use files in the active-data pool to restore files to primary storage pools. However, no new writes are permitted to volumes in the active-data pool from volumes outside the storage pool. A storage pool cannot be copied to the active-data pool.

COLLocate

Specifies whether the server attempts to keep data belonging to a single client node, group of client nodes, or client file space stored on as few volumes as possible. This parameter is optional. The default value is NO.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time needed to collocate files for storing and the number of volumes required. For details, see the *Administrator's Guide*.

DEFINE STGPOOL

Possible values are:

No Specifies that collocation is disabled.

GRoup

Specifies that collocation is enabled at the group level for client nodes. The server attempts to put data for nodes that belong to the same collocation group on as few volumes as possible. If the nodes in the collocation group have multiple file spaces, the server does not attempt to collocate those file spaces.

If you specify COLLOCATE=GROUP but do not define any collocation groups or if you specify COLLOCATE=GROUP but do not add nodes to a collocation group, data is collocated by node. Be sure to consider tape usage when organizing client nodes into collocation groups. For example, if a tape-based storage pool consists of data from grouped and ungrouped nodes and you specify COLLOCATE=GROUP, the server performs the following actions:

- Collocates by group the data for grouped nodes only. Whenever possible, the server collocates data that belongs to a group of nodes on a single tape or on as few tapes as possible. Data for a single node can also be spread across several tapes associated with a group.
- Collocates by node the data for ungrouped nodes. Whenever possible, the server stores the data for a single node on a single tape. All available tapes that already have data for the node are used before available space on any other tape is used.

NODe

Specifies that collocation is enabled at the client node level. The server attempts to put data for one node on as few volumes as possible. If the node has multiple file spaces, the server does not attempt to collocate those file spaces. For backward compatibility, COLLOCATE=YES is still accepted by the server to specify collocation at the client node level.

If a storage pool contains data for a node that is a member of a collocation group and you specify COLLOCATE=NODE, the data will be collocated by node not by group.

FIlespace

Specifies that collocation is enabled at the file space level for client nodes. The server attempts to put data for one node and file space on as few volumes as possible. If a node has multiple file spaces, the server attempts to put data for different file spaces on different volumes.

REClaim

Specifies when the server reclaims a volume, based on the percentage of reclaimable space on a volume. Reclaimable space is the amount of space occupied by files that are expired or deleted from the Tivoli Storage Manager database.

Reclamation makes the fragmented space and space occupied by inactive backup files on volumes usable again by moving any remaining unexpired files and active backup files from one volume to another volume. This makes the original volume available for reuse. This parameter is optional. You can specify an integer from 1 to 100. The default value is 60.

When determining which volumes in a storage pool to reclaim, the Tivoli Storage Manager server first determines the reclamation threshold indicated by the **RECLAIM**. The server then examines the percentage of reclaimable space for

each volume in the storage pool. If the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool, the volume is a candidate for reclamation.

For example, suppose storage pool FILEPOOL has a reclamation threshold of 70 percent. This value indicates that the server can reclaim any volume in the storage pool that has a percentage of reclaimable space that is greater than 70 percent. The storage pool has three volumes:

- FILEVOL1 with 65 percent reclaimable space
- FILEVOL2 with 80 percent reclaimable space
- FILEVOL3 with 95 percent reclaimable space

When reclamation begins, the server compares the percent of reclaimable space for each volume with the reclamation threshold of 70 percent. In this example, FILEVOL2 and FILEVOL3 are candidates for reclamation because their percentages of reclaimable space are greater than 70. To determine the percentage of reclaimable space for a volume, issue the **QUERY VOLUME** command and specify **FORMAT=DETAILED**. The value in the field Pct. Reclaimable Space is the percentage of reclaimable space for the volume.

If you change the value from the default, specify a value of 50 percent or greater so that files stored on two volumes can be combined onto a single output volume.

When an active-data pool volume that is offsite becomes eligible for reclamation, the reclamation process attempts to obtain the unexpired files on the reclaimable volume from a primary or active-data pool that is onsite. The process then writes these files to an available volume in the original active-data pool. Effectively, these files are moved back to the onsite location. However, the files could be obtained from the offsite volume after a disaster if a database backup is used that references the files on the offsite volume. Because of the way reclamation works with offsite volumes, use it carefully with active-data pools.

RECLAIMProcess

Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool. This parameter is optional. Enter a value from 1 to 999. The default value is 1.

When calculating the value for this parameter, consider the number of sequential storage pools that will be involved with the reclamation and the number of logical and physical drives that can be dedicated to the operation. To access a sequential access volume, IBM Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the reclamation.

For example, suppose that you want to reclaim the volumes from two sequential storage pools simultaneously and that you want to specify four processes for each of the storage pools. The storage pools have the same device class. Each process requires two mount points and, if the device type is not FILE, two drives. (One of the drives is for the input volume, and the other drive is for the output volume.) To run eight reclamation processes simultaneously, you need a total of at least 16 mount points and 16 drives. The device class for the storage pools must have a mount limit of at least 16.

If the number of reclamation processes you specify is more than the number of available mount points or drives, the processes that do not obtain mount

DEFINE STGPOOL

points or drives will wait for mount points or drives to become available. If mount points or drives do not become available within the MOUNTWAIT time, the reclamation processes will end. For information about specifying the MOUNTWAIT time, see “DEFINE DEVCLASS (Define a device class)” on page 148.

The Tivoli Storage Manager server will start the specified number of reclamation processes regardless of the number of volumes that are eligible for reclamation. For example, if you specify ten reclamation processes and only six volumes are eligible for reclamation, the server will start ten processes and four of them will complete without processing a volume.

RECLAMATIONType

Specifies the method by which volumes are reclaimed and managed. This parameter is optional. The default value is THRESHOLD. Possible values are the following:

THRESHold

Specifies that volumes belonging to this storage pool are reclaimed based on the threshold value in the RECLAIM attribute for this storage pool.

SNAPlock

Specifies that FILE volumes belonging to this storage pool will be managed for retention using NetApp Data ONTAP software and NetApp SnapLock volumes. This parameter is only valid for storage pools being defined to a server that has data retention protection enabled and that is assigned to a FILE device class. Volumes in this storage pool are not reclaimed based on threshold; the RECLAIM value for the storage pool is ignored.

All volumes in this storage pool are created as FILE volumes. A retention date, derived from the retention attributes in the archive copy group for the storage pool, is set in the metadata for the FILE volume using the SnapLock feature of the NetApp Data ONTAP operating system. Until the retention date has expired, the FILE volume and any data on it cannot be deleted from the physical SnapLock volume on which it is stored.

The RECLAMATIONTYPE parameter for all storage pools being defined must be the same when defined to the same device class name. The **DEFINE** command will fail if the RECLAMATIONTYPE parameter specified is different than what is currently defined for storage pools already defined to the device class name.

OFFSITERECLAIMLimit

Specifies the number of offsite volumes to have their space reclaimed during reclamation for this storage pool. This parameter is optional. The default value is NOLIMIT. Possible values are:

NOLimit

Specifies that you want to have the space reclaimed in all of your offsite volumes.

number

Specifies the number of offsite volumes to have their space reclaimed. You can specify an integer from 0 to 99999. A value of zero means that none of the offsite volumes will be reclaimed.

Important: When determining the value for the OFFSITERECLAIMLIMIT, consider using the statistical information in the message issued at the end of the offsite volume reclamation operation. Alternatively, you can use the

following Tivoli Storage Manager SQL select command to obtain the statistical information from the SUMMARY table for the offsite volume reclamation operation:

```
select * from summary where activity='OFFSITE RECLAMATION'
```

The statistical information includes the following items:

- The number of offsite volumes that were processed
- The number of parallel processes that were used
- The total amount of time required for the processing

The order in which offsite volumes are reclaimed is based on the amount of unused space in a volume. (Unused space includes both space that has never been used on the volume and space that has become empty because of file deletion.) Volumes with the largest amount of unused space are reclaimed first.

For example, suppose an active-data pool contains three volumes: VOL1, VOL2, and VOL3. VOL1 has the largest amount of unused space, and VOL3 has the least amount of unused space. Suppose further that the percentage of unused space in each of the three volumes is greater than the value of the RECLAIM parameter. If you do not specify a value for the OFFSITERECLAIMLIMIT parameter, all three volumes will be reclaimed when the reclamation runs. If you specify a value of 2, only VOL1 and VOL2 will be reclaimed when the reclamation runs. If you specify a value of 1, only VOL1 will be reclaimed.

MAXSCRatch (Required)

Specifies the maximum number of scratch volumes that the server can request for this storage pool. You can specify an integer from 0 to 100000000. By allowing the server to request scratch volumes as needed, you avoid having to define each volume to be used.

The value specified for this parameter is used to estimate the total number of volumes available in the active-data pool and the corresponding estimated capacity for the active-data pool.

Scratch volumes are automatically deleted from the storage pool when they become empty. However, if the access mode for a scratch volume is OFFSITE, the volume is not deleted from the active-data pool until the access mode is changed. This allows an administrator to query the server for empty, offsite scratch volumes and return these to the onsite location.

When scratch volumes with the device type of FILE become empty and are deleted, the space that the volumes occupied is freed by the server and returned to the file system.

Tip: For server-to-server operations that use virtual volumes and that store a small amount of data, consider specifying a value for the MAXSCRATCH parameter that is higher than the value you typically specify for write operations to other types of volumes. After a write operation to a virtual volume, Tivoli Storage Manager marks the volume as FULL, even if the value of the MAXCAPACITY parameter on the device-class definition has not been reached. The Tivoli Storage Manager server does not keep virtual volumes in FILLING status and does not append to them. If the value of the MAXSCRATCH parameter is too low, server-to-server operations can fail.

REUsedelay

Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool.

DEFINE STGPOOL

This parameter is optional. You can specify an integer from 0 to 9999. The default value is 0, which means that a volume can be rewritten or returned to the scratch pool as soon as all the files are deleted from the volume.

Important: Use this parameter to help ensure that when you restore the database to an earlier level, database references to files in the active-data pool are still valid. You must set this parameter to a value greater than the number of days you plan to retain the oldest database backup. The number of days specified for this parameter should be the same as the number specified for the **SET DRMDBBACKUPEXPIREDAYS** command. For more information, see the *Administrator's Guide*.

OVFLocation

Specifies the overflow location for the storage pool. The server assigns this location name to a volume that is ejected from the library by the command. This parameter is optional. The location name can be a maximum length of 255 characters. Enclose the location name in quotation marks if the location name contains any blank characters.

DATAFormat

Specifies the data format to use to copy files to this storage pool and restore files from this storage pool. The default format is the NATIVE server format. Possible values are:

NATive

Specifies the data format is the native IBM Tivoli Storage Manager server format and includes block headers.

NONblock

Specifies the data format is the native IBM Tivoli Storage Manager server format and does not include block headers.

Important: The default minimum block size on a volume associated with a FILE device class is 256 KB, regardless how much data is being written to the volume. For certain tasks (for example, using content-management products, using the DIRMC client option to store directory information, or migrating very small files using Tivoli Storage Manager for Space Management or Tivoli Storage Manager HSM for Windows), you can minimize wasted space on storage volumes by specifying the NONBLOCK data format. In most situations, however, the NATIVE format is preferred.

CRCDATA

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is only valid for NATIVE data format storage pools. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an **AUDIT VOLUME** command you can continually ensure the integrity of data stored in your storage hierarchy. Possible values are:

Yes

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because additional overhead is required to calculate and compare CRC values between the storage pool and the server.

No Specifies that data is stored without CRC information.

Tip: For storage pools that are associated with the 3592, LTO, or ECARTRIDGE device type, logical block protection provides better protection

against data corruption than CRC validation for a storage pool. If you specify CRC validation for a storage pool, data is validated only during volume auditing operations. Errors are identified after data is written to tape.

To enable logical block protection, specify a value of READWRITE for the **LBPROTECT** parameter on the **DEFINE DEVCLASS** and **UPDATE DEVCLASS** commands for the 3592, LTO, or ECARTRIDGE device types. Logical block protection is supported only on the following types of drives and media:

- IBM LTO5 and later.
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.
- Oracle StorageTek T10000C drives.

DEDUPLICATE

Specifies whether the data that is stored in this storage pool will be deduplicated. This parameter is optional and is valid only for storage pools that are defined with a FILE device class. The default value is NO.

IDENTIFYPROCESS

Specifies the number of parallel processes to use for server-side duplicate identification. This parameter is optional and is valid only for storage pools that are defined with a FILE device class. Enter a value from 0 to 20.

The default value for this parameter is 0. Duplicate-identification processes for a copy storage pool are not necessary if you specify duplicate-identification processes for the primary storage pool. When Tivoli Storage Manager analyzes a file in a storage pool, Tivoli Storage Manager also analyzes the file in all other storage pools.

When calculating the value for this parameter, consider the workload on the server and the amount of data requiring data deduplication. Server-side duplicate identification requires disk I/O and processor resources, so the more processes you allocate to data deduplication, the heavier the workload that you place on your system. In addition, consider the number of volumes that require processing. Server-side duplicate-identification processes work on volumes containing data that requires deduplication. If you update a storage pool, specifying that the data in the storage pool is to be deduplicated, all the volumes in the pool require processing. For this reason, you might have to define a high number of duplicate-identification processes initially. Over time, however, as existing volumes are processed, only the volumes containing new data have to be processed. When that happens, you can reduce the number of duplicate-identification processes.

Remember: Duplicate-identification processes can be either active or idle. Processes that are working on files are active. Processes that are waiting for files to work on are idle. Processes remain idle until volumes with data to be deduplicated become available. The output of the **QUERY PROCESS** command for a duplicate-identification process includes the total number of bytes and files that have been processed since the process first started. For example, if a duplicate-identification process processes four files, becomes idle, and then processes five more files, then the total number of files processed is nine. Processes end only when canceled or when the number of duplicate-identification processes for the storage pool is changed to a value less than the number currently specified.

DEFINE STGPOOL

Example: Define an active-data pool with a DC500 device class

Define an active-data pool, TAPEPOOL2, to the DC500 device class. Allow up to 50 scratch volumes for this pool. Delay the reuse of volumes for 45 days.

```
define stgpool tapepool3 dc500 pooltype=activedata  
maxscratch=50 reusedelay=45
```

DEFINE SUBSCRIPTION (Define a profile subscription)

Use this command on a managed server to subscribe that managed server to a profile.

When a server subscribes to its first profile, a subscription is also created to the default profile (if one exists) of the configuration manager. The server then contacts the configuration manager periodically for configuration updates.

Restrictions:

1. A server cannot subscribe to profiles from more than one configuration manager.
2. If a server subscribes to a profile with an associated object that is already defined on the server, the local definition is replaced by the definition from the configuration manager. For example, if a server has an administrative schedule named WEEKLY_BACKUP, then subscribes to a profile that also has an administrative schedule named WEEKLY_BACKUP, the local definition is replaced.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—DEfIne SUBSCRIPtion—profile_name—┬──SERVer—==—server_name—┘
```

Parameters

profile_name (Required)

Specifies the name of the profile to which the server subscribes.

SERVer

Specifies the name of the configuration manager from which the configuration information is obtained. This parameter is required, if the managed server does not have at least one subscription. If the managed server has a subscription, you can omit this parameter and it defaults to the configuration manager for that subscription.

Example: Define a profile subscription

Subscribe a profile named BETA that resides on a configuration manager named TOM.

```
define subscription beta server=tom
```

Related commands

Table 107. Commands related to DEFINE SUBSCRIPTION

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DELETE PROFILE	Deletes a profile from a configuration manager.

DEFINE SUBSCRIPTION

Table 107. Commands related to **DEFINE SUBSCRIPTION** (continued)

Command	Description
DELETE SUBSCRIBER	Deletes obsolete managed server subscriptions.
DELETE SUBSCRIPTION	Deletes a specified profile subscription.
LOCK PROFILE	Prevents distribution of a configuration profile.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY PROFILE	Displays information about configuration profiles.
QUERY SUBSCRIBER	Displays information about subscribers and their subscriptions to profiles.
QUERY SUBSCRIPTION	Displays information about profile subscriptions.
SET CONFIGREFRESH	Specifies a time interval for managed servers to contact configuration managers.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

DEFINE VIRTUALFSMAPPING (Define a virtual file space mapping)

Use this command to define a virtual file space mapping.

Virtual file space names can be used in the NAS data operations **BACKUP NODE** and **RESTORE NODE** similar to a file system name. Refer to the documentation about your NAS device for guidance on specifying the parameters for this command.


Note: The NAS node must have an associated data mover definition because when the Tivoli Storage Manager server updates a virtual file space mapping, the server attempts to contact the NAS device to validate the virtual file system and file system name.

Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege
- Restricted policy privilege for the domain to which the NAS node is assigned.

Syntax



Parameters

node_name (Required)

Specifies the NAS node on which the file system and path reside. You cannot use wildcard characters or specify a list of names.

virtual_filespace_name (Required)

Specifies the name which refers to this virtual file space definition. The virtual file space name is case sensitive and the first character must be a forward slash /. The length of the name cannot be more than 64 characters, including the required forward slash. Virtual file space names are restricted to the same character set as all other objects in the Tivoli Storage Manager server except that the forward slash / character is also allowed.

The virtual file space name cannot be identical to any file system on the NAS node. When selecting a virtual file space name, consider the following restrictions:

- If a file system is created on the NAS device with the same name as a virtual file system, a name conflict will occur on the Tivoli Storage Manager server when the new file space is backed up. Use a string for the virtual file space name that is unlikely to be used as a real file system name on your NAS device in the future.

For example: A user follows a naming convention for creating file spaces on a NAS device with names of the form /vol1, /vol2, /vol3. The user defines

DEFINE VIRTUALFSMAPPING

a virtual file space to the Tivoli Storage Manager server with the name /vol9. If the user continues to use the same naming convention, the virtual file space name is likely to conflict with a real file space name at some point in the future.

- During backup and restore operations, Tivoli Storage Manager verifies that a name conflict does not occur prior to starting the operation.
- The virtual file space name appears as a file space in the output of the QUERY FILESPACE command, and also in the backup and restore panels of the Tivoli Storage Manager Web client. Therefore, consider selecting a name that unambiguously identifies this object as a directory path on the NAS device.

file_system_name (Required)

Specifies the name of the file system in which the path is located. The file system name must exist on the specified NAS node. The file system name cannot contain wildcard characters.

path (Required)

Specifies the path from the root of the file system to the directory. The path can only reference a directory. The maximum length of the path is 1024 characters. The path name is case sensitive.

NAMEType

Specifies how the server should interpret the path name specified. This parameter is useful when a path contains characters that are not part of the code page in which the server is running. The default value is SERVER.

Possible values are:

SERVER

The server uses the server code page to interpret the path name.

HEXadecimal

The server interprets the path that you enter as the hexadecimal representation of the path. This option should be used when a path contains characters that cannot be entered. This could occur if the NAS file system is set to a language different from the one in which the server is running.

Example: Define a virtual file space mapping

Define the virtual file space mapping name /mikeshomedir for the path /home/mike on the file system /vol/vol1 on the NAS node named NAS1.

```
define virtualfsmapping nas1 /mikeshomedir /vol/vol1 /home/mike
```

Related commands

Table 108. Commands related to **DEFINE VIRTUALFSMAPPING**

Command	Description
DELETE VIRTUALFSMAPPING	Delete a virtual file space mapping.
QUERY VIRTUALFSMAPPING	Query a virtual file space mapping.
UPDATE VIRTUALFSMAPPING	Update a virtual file space mapping.

DEFINE VOLUME (Define a volume in a storage pool)

Use this command to assign a random or sequential access volume to a storage pool.

When defining a random-access (DISK) storage-pool volume or a sequential access storage pool volume that is associated with a FILE device class, you can have the server create the volume before it is assigned. Alternatively, you can use space triggers to create preassigned volumes when predetermined space-utilization thresholds are exceeded. For details about space triggers, see “DEFINE SPACETRIGGER (Define the space trigger)” on page 334. For volumes associated with device classes other than DISK or device types other than FILE, you can use the DEFINE VOLUME command to assign an already-created volume to a storage pool.

When using a **FILE** device class for storage managed by a z/OS media server, it is not necessary to format or define volumes. If you define a volume for such a **FILE** device class by using the **DEFINE VOLUME** command, the z/OS media server does not allocate space for the volume until the volume is opened for its first use.

Attention: Volumes for the z/OS media server that are created using the **DEFINE VOLUME** command remain physically full or allocated after Tivoli Storage Manager empties the volume, for example, after expiration or reclamation. For FILE volumes, the DASD space is *not* relinquished to the system when the volume is emptied. If a storage pool requires an empty or filling volume, the FILE volume can be used. In contrast, tape volumes that are logically empty are the same as physically empty. FILE and tape volumes remain defined in Tivoli Storage Manager. In contrast, SCRATCH volumes, including the physical storage allocated for SCRATCH FILE volumes, are returned to the system when emptied.

To create space in sequential access storage pools, you can define volumes or allow the server to request scratch volumes as needed, as specified by the **MAXSCRATCH** parameter for the storage pool. For storage pools associated with the FILE device class, the server can create private volumes as needed using storage-pool space triggers. For DISK storage pools, the scratch mechanism is not available. However, you can create space by creating volumes and then defining them to the server. Alternatively, you can have the server create volumes that use storage-pool space triggers.

Tivoli Storage Manager does not validate the existence of a volume name when defining a volume in a storage pool that is associated with a library. The defined volume has “0” EST capacity until data is written to the volume.

Attention: The size of a storage pool volume cannot be changed after it has been defined to the Tivoli Storage Manager server.

You cannot use this command to define volumes in storage pools with the parameter setting RECLAMATIONTYPE=SNAPLOCK. Volumes in this type of storage pool are allocated by using the **MAXSCRATCH** parameter on the storage pool definition.


Physical files that are allocated with **DEFINE VOLUME** command are not removed from a file space if you issue the **DELETE VOLUME** command.

DEFINE VOLUME

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to which the volume is assigned.

Syntax



Notes:

- 1 This value is valid only for volumes assigned to copy storage pools.
- 2 This parameter is valid only for DISK or FILE volumes.
- 3 This parameter is valid only for sequential access volumes.

Parameters

pool_name (Required)

Specifies the name of the storage pool to which the volume is assigned.

volume_name (Required)

Specifies the name of the storage pool volume to be defined. If you specify a number greater than 1 for the **NUMBEROFVOLUMES** parameter, the volume name is used as a prefix to generate multiple volume names. The volume name that you specify depends on the type of device that the storage pool uses.

Each volume used by a server for any purpose must have a unique name. This requirement applies to all volumes, whether the volumes are used for storage pools, or used for operations such as database backup or export. The requirement also applies to volumes that reside in different libraries but that are used by the same server.

Remember: Volume names cannot contain embedded blanks or equal signs.

See the following tables for volume name requirements:

- Table 109 on page 393: DISK
- Table 110 on page 393: FILE
- Table 111 on page 393: FILE for the z/OS media server

- Table 112: Tape
- Table 113 on page 394: Tape for z/OS media server
- Table 114 on page 394: REMOVABLEFILE

Table 109. Volume name requirements for DISK

Volume Name Requirements	Example
The name of the file to contain the volume data, with either the fully qualified path name or a path name relative to the current working directory.	/usr/storage/sbkup01.dsm

Table 110. Volume name requirements for FILE

Volume Name Requirements	Example
The name of the file to contain the volume data, with either the fully qualified path name or the path name relative to a directory identified in the DIRECTORY parameter for the device class.	/data/fpoo101.dsm

Place FILE volumes in one of the directories specified with the DIRECTORY parameter of the **DEFINE DEVCLASS** command. Otherwise, storage agents might not have access to the volumes. For details, see “DEFINE PATH (Define a path)” on page 276.

Table 111. z/OS media server: Volume name requirements for FILE

Volume Name Requirements	Example
For FILE volumes used with the z/OS media server server, specify a data set name. The data set name can consist of one or more qualifiers that are delimited by a period. The qualifiers can contain up to 8 characters. The maximum length of the data set name is 44 characters. The first letter of each qualifier must be alphabetic or national (@#\$), followed by alphabetic, national, hyphen, or numeric characters.	SERVER1.BFS.POOL3.VOLA

To allocate the associated VSAM Linear Dataset when the volume is tendered on the z/OS system, the High Level Qualifier (HLQ) is typically filtered by specific ACS routines within the SMS policy constraints on the system where the z/OS media server is running.

The behavior of the HLQ is similar to the behavior of the PREFIX name on a scratch request. The HLQ is typically used by DFSMS to affect allocation attributes, such as Extended Addressability for data sets that are expected to extend when space that is already allocated to the file volume is used up.

If the data set does not exist, the server creates it when the volume is used for a specific Tivoli Storage Manager storage operation. The data set is not created when the volume is defined. Data loss can result when defining volumes because the z/OS media server reuses the volume or VSAM LDS if it exists at the time of allocation time.

Important: To allow Tivoli Storage Manager to generate volume names, consider using SCRATCH volumes.

Table 112. Volume name requirements for tape

Volume Name Requirements	Example
Use 1 - 32 alphanumeric characters.	DSMT01
The volume name cannot contain any embedded blanks or equal signs.	

DEFINE VOLUME

Table 113. z/OS media server: Volume name requirements for tape

Volume Name Requirements	Example
For tape cartridges, specify a tape volume name with 1 - 6 alphanumeric characters. The server converts tape volume names to uppercase.	DSMT01

The volume name cannot contain any embedded blanks or equal signs.

Each volume used by a server for any purpose must have a unique name. This requirement applies to all volumes, whether the volumes are used for storage pools, or used for operations such as database backup or export. The requirement also applies to volumes that reside in different z/OS media libraries but that are used by the same server.

Table 114. Volume name requirements for REMOVABLEFILE

Volume Name Requirements	Example
1-6 alphanumeric characters	DSM01

The server converts volume names to uppercase.

ACcEss

Specifies how client nodes and server processes (such as migration) can access files in the storage pool volume. This parameter is optional. The default value is READWRITE. Possible values are:

READWrite

Specifies that client nodes and server processes can read from and write to files stored on the volume.

READOnly

Specifies that client nodes and server processes can only read files stored on the volume.

UNAVailable

Specifies that client nodes or server processes cannot access files stored on the volume.

If you define a random access volume as UNAVAILABLE, you cannot vary the volume online.

If you define a sequential access volume as UNAVAILABLE, the server does not attempt to access the volume.

Offsite

Specifies that the volume is at an off-site location from which it cannot be mounted. You can specify this value only for volumes in copy storage pools.

Use this value to help you track volumes at off-site locations. The server treats volumes designated as off-site differently:

- The server does not generate mount requests for volumes designated off-site.
- The server reclaims or moves data from off-site volumes by retrieving files from other storage pools.
- The server does not automatically delete empty, off-site scratch volumes from a copy storage pool.

LOcation

Specifies the location of the volume. This parameter is optional. It can be

specified only for volumes in sequential access storage pools. The location information can be a maximum length of 255 characters. Enclose the location in quotation marks if it contains any blank characters.

Format size

Specifies the size of the random access volume or FILE volume that is created and formatted in one step. The value is specified in megabytes. The maximum size is 8 000 000 MB (8 terabytes). This parameter is required if any of the following conditions are true:

- A single FILE or DISK volume is specified, which is to be created and formatted in one step.
- The value for the **NUMBEROFVOLUMES** parameter is greater than 1, and DISK volumes are being created.
- The value of the **NUMBEROFVOLUMES** parameter is greater than 1, and the value of the **FORMATSIZE** parameter is less than or equal to the **MAXCAPACITY** parameter of the **DEFINE DEVCLASS** command.

If you are allocating volumes on a z/OS media server, this parameter is not valid.

For a FILE volume, you must specify a value less than or equal to the value of the **MAXCAPACITY** parameter of the device class associated with the storage pool.

You cannot use this parameter for multiple, predefined volumes. Unless you specify **WAIT=YES** is specified, the operation is performed as a background process.

Number of volumes

Specifies the number of volumes that are created and formatted in one step. This parameter applies only to storage pools with DISK or FILE device classes. This parameter is optional. The default is 1. If you specify a value greater than 1, you must also specify a value for the **FORMATSIZE** parameter. Specify a number from 1 to 256.

If you are allocating volumes on a z/OS media server, the only value that this parameter supports is the default value of 1.

If the value for the **NUMBEROFVOLUMES** parameter is greater than 1, the volume name you specified will have a numeric suffix appended to create each name, for example, tivoli001 and tivoli002. Be sure to choose a volume name so that a valid file name for the target file system is created when the suffix is appended.

Important: You must ensure that storage agents can access newly created FILE volumes. For more information, see “DEFINE PATH (Define a path)” on page 276.

Wait

Specifies whether volume creation and formatting operation is performed in the foreground or background. This parameter is optional. It is ignored unless you also specify the **FORMATSIZE** parameter.

No Specifies that a volume creation and formatting operation is performed in the background. The NO value is the default when you also specify a format size.

Yes

Specifies that a volume creation and formatting operation is performed in the foreground.

DEFINE VOLUME

Remember: You cannot specify WAIT=YES from the server console.

Example: Use a background process to define a new 100 MB volume for a disk storage pool

Create a volume of 100 MB in the disk storage pool named BACKUPPOOL. The volume name is /var/storage/bf.dsm. Let the volume be created as a background process.

```
define volume backuppool
/var/storage/bf.dsm formatsize=100
```

Example: Define a volume to a disk storage pool with read and write access

A storage pool named POOL1 is assigned to a tape device class. Define a volume named TAPE01 to this storage pool, with READWRITE access.

```
define volume pool1 tape01 access=readwrite
```

Example: Define a volume to a file storage pool

A storage pool named FILEPOOL is assigned to a device class with a device type of FILE. Define a volume named filepool_vol01 to this storage pool.

```
define volume filepool /usr/storage/filepool_vol01
```

Example: Example: Use a background process to define 10 volumes for a file storage pool with a device class 5 GB maximum capacity

Define ten volumes in a sequential storage pool that uses a FILE device class. The storage pool is named FILEPOOL. The value of the **MAXCAPACITY** parameter for the device class associated with this storage pool is 5 GB. Creation must occur in the background.

```
define volume filepool filevol numberofvolumes=10 formatsize=5000
```

The Tivoli Storage Manager server creates volume names filevol001 through filevol010.

Volumes are created in the directory or directories specified with the DIRECTORY parameter of the device class associated with storage pool filepool. If you specified multiple directories for the device class, individual volumes can be created in any of the directories in the list.

Related commands

Table 115. Commands related to DEFINE VOLUME

Command	Description
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE VOLUME	Deletes a volume from a storage pool.
QUERY VOLUME	Displays information about storage pool volumes.
UPDATE DEVCLASS	Changes the attributes of a device class.
UPDATE LIBVOLUME	Changes the status of a storage volume.

Table 115. Commands related to DEFINE VOLUME (continued)

Command	Description
UPDATE VOLUME	Updates the attributes of storage pool volumes.

DELETE commands

Use the **DELETE** commands to delete or remove a Tivoli Storage Manager object.

- “DELETE ASSOCIATION (Delete the node association to a schedule)” on page 401
- “DELETE ALERTTRIGGER (Remove a message from an alert trigger)” on page 400
- “DELETE BACKUPSET (Delete a backup set)” on page 403
- “DELETE CLIENTOPT (Delete an option in an option set)” on page 408
- “DELETE CLOPTSET (Delete a client option set)” on page 409
- “DELETE COLLOGROUP (Delete a collocation group)” on page 410
- “DELETE COLLOCMEMBER (Delete collocation group member)” on page 411
- “DELETE COPYGROUP (Delete a backup or archive copy group)” on page 413
- “DELETE DATAMOVER (Delete a data mover)” on page 415
- “DELETE DEVCLASS (Delete a device class)” on page 416
- “DELETE DOMAIN (Delete a policy domain)” on page 417
- “DELETE DRIVE (Delete a drive from a library)” on page 418
- “DELETE EVENT (Delete event records)” on page 419
- “DELETE EVENTSERVER (Delete the definition of the event server)” on page 421
- “DELETE FILESPACE (Delete client node data from the server)” on page 422
- “DELETE GRPMEMBER (Delete a server from a server group)” on page 426
- “DELETE LIBRARY (Delete a library)” on page 428
- “DELETE MACHINE (Delete machine information)” on page 429
- “DELETE MACHNODEASSOCIATION (Delete association between a machine and a node)” on page 430
- “DELETE MGMTCLASS (Delete a management class)” on page 431
- “DELETE NODEGROUP (Delete a node group)” on page 432
- “DELETE NODEGROUPMEMBER (Delete node group member)” on page 433
- “DELETE PATH (Delete a path)” on page 434
- “DELETE POLICYSET (Delete a policy set)” on page 436
- “DELETE PROFASSOCIATION (Delete a profile association)” on page 437
- “DELETE PROFILE (Delete a profile)” on page 440
- “DELETE RECMEDMACHASSOCIATION (Delete recovery media and machine association)” on page 442
- “DELETE RECOVERYMEDIA (Delete recovery media)” on page 443
- “DELETE SCHEDULE (Delete a client or an administrative command schedule)” on page 444
- “DELETE SCRIPT (Delete command lines from a script or delete the entire script)” on page 447
- “DELETE SERVER (Delete a server definition)” on page 448
- “DELETE SERVERGROUP (Delete a server group)” on page 449
- “DELETE SPACETRIGGER (Delete the storage pool space triggers)” on page 450
- “DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)” on page 451
- “DELETE STGPOOL (Delete a storage pool)” on page 453

- “DELETE SUBSCRIBER (Delete subscriptions from a configuration manager database)” on page 454
- “DELETE SUBSCRIPTION (Delete a profile subscription)” on page 455
- “DELETE VIRTUALFSMAPPING (Delete a virtual file space mapping)” on page 456
- “DELETE VOLHISTORY (Delete sequential volume history information)” on page 457
- “DELETE VOLUME (Delete a storage pool volume)” on page 462

DELETE ALERTTRIGGER

DELETE ALERTTRIGGER (Remove a message from an alert trigger)

Use this command to remove a message from the list of alert triggers.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
DELEte ALERTTrigger +message_number
```

Parameters

message_number (Required)

Specifies the message number that you want to remove from the list of alert triggers. Specify multiple message numbers, which are separated by commas, and no intervening spaces. Message numbers are a maximum of eight characters in length. Wildcard characters can be used to specify message numbers.

Delete alert trigger

Delete two message numbers that are designated as alerts, by issuing the following command:

```
delete alerttrigger ANR1067E,ANR1073E
```

Related commands

Table 116. Commands related to DELETE ALERTTRIGGER

Command	Description
"DEFINE ALERTTRIGGER (Define an alert trigger)" on page 112	Associates specified messages to an alert trigger.
"QUERY ALERTSTATUS (Query the status of an alert)" on page 674	Displays information about alerts that have been issued on the server.
"QUERY ALERTTRIGGER (Query the list of defined alert triggers)" on page 672	Displays message numbers that trigger an alert.
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"UPDATE ALERTTRIGGER (Update a defined alert trigger)" on page 1165	Updates the attributes of one or more alert triggers.
"UPDATE ALERTSTATUS (Update the status of an alert)" on page 1168	Updates the status of a reported alert.

DELETE ASSOCIATION (Delete the node association to a schedule)

Use this command to delete the association of a client node to a client schedule. Tivoli Storage Manager no longer runs the schedule on the client node.

If you try to disassociate a client from a schedule to which it is not associated, this command has no effect for that client.

Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege
- Restricted policy privilege for the domain to which the schedule belongs

Syntax

```
▶▶—DELEte ASSOCIation—domain_name—schedule_name—node_name—▶▶
```

Parameters

domain_name (Required)

Specifies the name of the policy domain to which the schedule belongs.

schedule_name (Required)

Specifies the name of the schedule from which clients are to be disassociated.

node_name (Required)

Specifies the name of the client node that is no longer associated with the client schedule. You can specify a list of clients which are to be no longer associated with the specified schedule. Commas, with no intervening spaces, separate the items in the list. You can also use a wildcard character to specify a name. All matching clients are disassociated from the specified schedule.

Example: Delete a node association to a schedule

To delete the association of the node JEFF, assigned to the DOMAIN1 policy domain, to the WEEKLY_BACKUP schedule issue the following command:

```
delete association domain1 weekly_backup jeff
```

Example: Delete a node association to a schedule using a wildcard for node selection

Delete the association of selected clients, assigned to the DOMAIN1 policy domain, to the WEEKLY_BACKUP schedule so that this schedule is no longer run by these clients. The nodes that are disassociated from the schedule contain ABC or XYZ in the node name. Issue the command:

```
delete association domain1 weekly_backup *abc*,*xyz*
```

DELETE ASSOCIATION

Related commands

Table 117. Commands related to DELETE ASSOCIATION

Command	Description
DEFINE ASSOCIATION	Associates clients with a schedule.
QUERY ASSOCIATION	Displays the clients associated with one or more schedules.

DELETE BACKUPSET (Delete a backup set)

Use this command to manually delete a backup set before its retention period expires.


When the server creates a backup set, the retention period assigned to the backup set determines how long the backup set remains in the database. When that date passes, the server automatically deletes the backup set when expiration processing runs. However, you can also manually delete the client's backup set from the server before it is scheduled to expire by using the **DELETE BACKUPSET** command.

Attention: If the volumes contain multiple backup sets, they are not returned to scratch status until all the backup sets are expired or are deleted.

Privilege class

If the REQSYSAUTHOUTFILE server option is set to YES (the default), the administrator must have system privilege. If the REQSYSAUTHOUTFILE server option is set to NO, the administrator must have system privilege or policy privilege for the domain to which the client node is assigned.

Syntax



Parameters

`node_name` or `node_group_name` (Required)

Specifies the name of the client nodes or node groups whose data is contained in the specified backup set volumes. To specify multiple node and node group

DELETE BACKUPSET

names, separate the names with commas and no intervening spaces. Any node name you specify may contain wildcard characters, but node group names cannot contain wildcard characters. If backup set volumes contain backup sets from multiple nodes then every backup set whose node name matches one of the specified node names will be deleted.

backup_set_name (Required)

Specifies the name of the backup set to delete. The backup set name you specify can contain wildcard characters. You can specify more than one backup set name by separating the names with commas and no intervening spaces.

BEGINDate

Specifies the beginning date in which the backup set to delete was created. This parameter is optional. You can use this parameter with the **BEGINTIME** parameter to specify a range for the date and time. If you specify a begin date without a begin time, the time will be at 12:00 a.m. (midnight) on the date you specify.

You can specify the date by using one of the following values:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1999
TODAY	The current date	TODAY
<i>TODAY+days or +days</i>	The current date plus days specified.	TODAY +3 or +3.
<i>TODAY-days or -days</i>	The current date minus days specified.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
<i>EOLM-days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
<i>BOTM+days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time in which the backup set to delete was created. This parameter is optional. You can use this parameter in conjunction with the **BEGINDATE** parameter to specify a range for the date and time. If you specify a begin time without a begin date, the date will be the current date at the time you specify.

You can specify the time by using one of the following values:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08
NOW	The current time	NOW
<i>NOW+HH:MM or +HH:MM</i>	The current time plus hours and minutes specified	NOW+02:00 or +02:00.

Value	Description	Example
NOW- <i>HH:MM</i> or <i>-HH:MM</i>	The current time minus hours and minutes specified	NOW-02:00 or -02:00.

ENDDate

Specifies the ending date in which the backup set to delete was created. This parameter is optional. You can use this parameter in conjunction with the ENDTIME parameter to specify a range for the date and time. If you specify an end date without an end time, the time will be at 11:59:59 p.m. on the specified end date.

You can specify the date by using one of the following values:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1999
TODAY	The current date	TODAY
TODAY+ <i>days</i> or <i>+days</i>	The current date plus days specified.	TODAY +3 or +3.
TODAY- <i>days</i> or <i>-days</i>	The current date minus days specified.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time of the range in which the backup set to delete was created. This parameter is optional. You can use this parameter in conjunction with the ENDDATE parameter to specify a range for the date and time. If you specify an end time without an end date, the date will be the current date at the time you specify.

You can specify the time by using one of the following values:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08
NOW	The current time	NOW
NOW+ <i>HH:MM</i> or <i>+HH:MM</i>	The current time plus hours and minutes on the specified end date	NOW+02:00 or +02:00.
NOW- <i>HH:MM</i> or <i>-HH:MM</i>	The current time minus hours and minutes on the specified end date	NOW-02:00 or -02:00.

DELETE BACKUPSET

WHEREDATAType

Specifies the backup sets containing the specified types of data are to be deleted. This parameter is optional. The default is that backup sets for all types of data (file level, image, and application) are to be deleted. To specify multiple data types, separate the data types with commas and no intervening spaces. Possible values are:

ALL

Specifies that backup sets for all types of data (file level, image, and application) are to be deleted. This is the default.

FILE

Specifies that a file level backup set is to be deleted. File level backup sets contain files and directories backup up by the backup-archive client.

IMAGE

Specifies that an image backup set is to be deleted. Image backup sets contain images created by the backup-archive client **BACKUP IMAGE** command.

WHERERETention

Specifies the retention value, specified in days, that is associated with the backup sets to delete. You can specify an integer from 0 to 30000. The values are:

days

Specifies that backup sets that are retained this number of days are deleted.

NOLimit

Specifies that the backup sets that are retained indefinitely are deleted.

WHEREDEscription

Specifies the description that is associated with the backup set to delete. The description you specify can contain a wildcard character. This parameter is optional. Enclose the description in quotation marks if it contains any blank characters.

Preview

Specifies whether to preview the list of backup sets to delete, without actually deleting the backup sets. This parameter is optional. The default value is NO. The values are:

No Specifies that the backup sets are deleted.

Yes

Specifies that the server displays the list of backup sets to delete, without actually deleting the backup sets.

Example: Delete a backup set

Delete backup set named PERS_DATA.3099 that belongs to client node JANE. The backup set was generated on 11/19/1998 at 10:30:05 and the description is "Documentation Shop".

```
delete backupset pers_data.3099
begindate=11/19/1998 begintime=10:30:05
wheredescription="documentation shop"
```

Related commands*Table 118. Commands related to DELETE BACKUPSET*

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
QUERY BACKUPSET	Displays backup sets.
QUERY NODEGROUP	Displays information about node groups.
QUERY BACKUPSETCONTENTS	Displays contents contained in backup sets.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

DELETE CLIENTOPT (Delete an option in an option set)

DELETE CLIENTOPT (Delete an option in an option set)

Use this command to delete a client option in an option set.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege.

Syntax

```
►►—DELEte CLIENTOpt—option_set_name—option_name—►►
```

└─SEQnumber—=—┬─*number*—┘
└─ALL—┘

Parameters

option_set_name (Required)

Specifies the name of the client option set.

option_name (Required)

Specifies a valid client option.

SEQnumber

Specifies a sequence number when an option name is specified more than once. This parameter is optional. Valid values are:

n Specifies an integer of 0 or greater.

ALL

Specifies all sequence numbers.

Example: Delete the date format option

Delete the date format option in an option set named *ENG*.

```
delete clientopt eng dateformat
```

Related commands

Table 119. Commands related to DELETE CLIENTOPT

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLIENTOPT	Adds a client option to a client option set.
DEFINE CLOPTSET	Defines a client option set.
DELETE CLOPTSET	Deletes a client option set.
QUERY CLOPTSET	Displays information about a client option set.
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.
UPDATE CLOPTSET	Updates the description of a client option set.

DELETE CLOPTSET (Delete a client option set)

Use this command to delete a client option set.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege.

Syntax

►►—DELEte CLOptset—*option_set_name*—◄◄

Parameters

option_set_name **(Required)**

Specifies the name of the client option set to delete.

Example: Delete a client option set

Delete the client option set named ENG.

```
delete cloptset eng
```

Related commands

Table 120. Commands related to DELETE CLOPTSET

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLIENTOPT	Adds a client option to a client option set.
DEFINE CLOPTSET	Defines a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
QUERY CLOPTSET	Displays information about a client option set.
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.
UPDATE CLOPTSET	Updates the description of a client option set.

DELETE COLLOGROUP (Delete a collocation group)

Use this command to delete a collocation group. You cannot delete a collocation group if it has any members in it.

You can remove all the members in the collocation group by issuing the **DELETE COLLOCMEMBER** command with a wildcard in the *node_name* parameter.

Privilege class

To issue this command, you must have system or unrestricted storage privilege.

Syntax

```
►►—DELEte COLLOGGroup—group_name—◄◄
```

Parameters

group_name

Specifies the name of the collocation group that you want to delete.

Example: Delete a collocation group

Delete a collocation group named group1.

```
delete colloggroup group1
```

Related commands

Table 121. Commands related to **DELETE COLLOGROUP**

Command	Description
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node to a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOCMEMBER	Deletes a client node from a collocation group.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY COLLOGROUP	Displays information about collocation groups.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY STGPOOL	Displays information about storage pools.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
UPDATE COLLOGROUP	Updates the description of a collocation group.
UPDATE STGPOOL	Changes the attributes of a storage pool.

DELETE COLLOCMEMBER (Delete collocation group member)

Use this command to delete a client node from a collocation group.

Privilege class

To issue this command, you must have system or unrestricted storage privilege.

Syntax

```

DELEte COLLOCMember—group_name—node_name

```

Parameters

group_name

Specifies the name of the collocation group from which you want to delete a client node.

node_name

Specifies the name of the client node that you want to delete from the collocation group. You can specify one or more names. When specifying multiple names, separate the names with commas; do not use intervening spaces. You can also use wildcard characters to specify multiple nodes.

Example: Delete collocation group member

Delete two nodes, NODE1 and NODE2, from a collocation group, GROUP1.

```
delete collocmember group1 node1,node2
```

Related commands

Table 122. Commands related to DELETE COLLOCMEMBER

Command	Description
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node to a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOGROUP	Deletes a collocation group.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY COLLOGROUP	Displays information about collocation groups.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY STGPOOL	Displays information about storage pools.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
UPDATE COLLOGROUP	Updates the description of a collocation group.

DELETE COLLOCMEMBER

Table 122. Commands related to DELETE COLLOCMEMBER (continued)

Command	Description
UPDATE STGPOOL	Changes the attributes of a storage pool.

DELETE COPYGROUP

Example: Delete a backup copy group

Delete the backup copy group from the ACTIVEFILES management class that is in the VACATION policy set of the EMPLOYEE_RECORDS policy domain.

```
delete copygroup employee_records  
vacation activefiles
```

Example: Delete an archive copy group

Delete the archive copy group from the MCLASS1 management class that is in the SUMMER policy set of the PROG1 policy domain.

```
delete copygroup prog1 summer mclass1 type=archive
```

Related commands

Table 123. Commands related to DELETE COPYGROUP

Command	Description
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
QUERY COPYGROUP	Displays the attributes of a copy group.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.

DELETE DATAMOVER (Delete a data mover)

Use this command to delete a data mover. You cannot delete the data mover if any paths are defined for this data mover.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—DELEte DATAMover—data_mover_name—◄◄
```

Parameters

data_mover_name **(Required)**

Specifies the name of the data mover.

Note: This command deletes the data mover even if there is data for the corresponding NAS node.

Example: Delete a data mover

Delete the data mover for the node named NAS1.

```
delete datamover nas1
```

Related commands

Table 124. Commands related to DELETE DATAMOVER

Command	Description
DEFINE DATAMOVER	Defines a data mover to the Tivoli Storage Manager server.
DEFINE PATH	Defines a path from a source to a destination.
DELETE PATH	Deletes a path from a source to a destination.
QUERY DATAMOVER	Displays data mover definitions.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE DATAMOVER	Changes the definition for a data mover.

DELETE DEVCLASS

DELETE DEVCLASS (Delete a device class)

Use this command to delete a device class.

To use this command, you must first delete all storage pools that are assigned to the device class and, if necessary, cancel any database export or import processes that are using the device class.

You cannot delete the device class DISK, which is predefined at installation, but you can delete any device classes defined by the Tivoli Storage Manager administrator.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—DELEte DEVclass—device_class_name—◄◄
```

Parameters

device_class_name **(Required)**

Specifies the name of the device class to be deleted.

Example: Delete a device class

Delete the device class named MYTAPE. There are no storage pools assigned to the device class.

```
delete devclass mytape
```

Related commands

Table 125. Commands related to DELETE DEVCLASS

Command	Description
DEFINE DEVCLASS	Defines a device class.
DEFINE DEVCLASS (z/OS media server)	Defines a device class to use storage managed by a z/OS media server.
QUERY DEVCLASS	Displays information about device classes.
QUERY DIRSPACE	Displays information about FILE directories.
UPDATE DEVCLASS	Changes the attributes of a device class.
UPDATE DEVCLASS (z/OS media server)	Changes the attributes of a device class for storage managed by a z/OS media server.

DELETE DOMAIN (Delete a policy domain)

Use this command to delete a policy domain. All associated policy sets, management classes, and copy groups are deleted along with the policy domain. You cannot delete a policy domain to which client nodes are registered.

You can delete the predefined STANDARD policy domain. However, if you later reinstall the Tivoli Storage Manager server, the process restores all STANDARD policy objects.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
▶▶—DELEte D0main—domain_name—▶▶
```

Parameters

domain_name **(Required)**

Specifies the policy domain to delete.

Examples: Delete a policy domain

Delete the EMPLOYEE_RECORDS policy domain.

```
delete domain employee_records
```

Related commands

Table 126. Commands related to DELETE DOMAIN

Command	Description
COPY DOMAIN	Creates a copy of a policy domain.
DEFINE DOMAIN	Defines a policy domain that clients can be assigned to.
QUERY DOMAIN	Displays information about policy domains.
UPDATE DOMAIN	Changes the attributes of a policy domain.

DELETE DRIVE (Delete a drive from a library)

Use this command to delete a drive from a library. A drive that is in use cannot be deleted.

All paths related to a drive must be deleted before the drive itself can be deleted.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—DELEte DRive—library_name—drive_name—◄◄
```

Parameters

library_name **(Required)**

Specifies the name of the library where the drive is located.

drive_name **(Required)**

Specifies the name of the drive to be deleted.

Example: Delete a drive from a library

Delete DRIVE3 from the library named AUTO.

```
delete drive auto drive3
```

Related commands

Table 127. Commands related to DELETE DRIVE

Command	Description
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE LIBRARY	Deletes a library.
DELETE PATH	Deletes a path from a source to a destination.
PERFORM LIBACTION	Defines all drives and paths for a library.
QUERY DRIVE	Displays information about drives.
QUERY LIBRARY	Displays information about one or more libraries.
UPDATE DRIVE	Changes the attributes of a drive.

DELETE EVENT

time

Specifies the time used to determine which event records to delete. Use this parameter in conjunction with the DATE parameter to specify a date and time for deleting event records. Any record whose scheduled start occurs before the specified date and time is deleted. However, records are not deleted for events whose startup window has not yet passed. The default is 00:00.

You can specify the time by using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified	NOW+03:00 or +03:00
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified	NOW-03:00 or -03:00

Attention: If you issue this command at 9:00 using NOW+03:00 or +03:00, Tivoli Storage Manager deletes records with a time of 12:00 or later on the date you specify.

TYPE

Specifies the type of events to be deleted. This parameter is optional. The default is CLIENT. Possible values are:

Client

Specifies to delete event records for client schedules.

Administrative

Specifies to delete event records for administrative command schedules.

ALL

Specifies to delete event records for both client and administrative command schedules.

Example: Delete event records

Delete records for events with scheduled start times prior to 08:00 on May 26, 1998 (05/26/1998), and whose startup window has passed. Records for these events are deleted regardless of whether the retention period for event records, as specified with the **SET EVENTRETENTION** command, has passed.

```
delete event 05/26/1998 08:00
```

Related commands

Table 128. Commands related to DELETE EVENT

Command	Description
QUERY EVENT	Displays information about scheduled and completed events for selected clients.
SET EVENTRETENTION	Specifies the number of days to retain records for scheduled operations.

DELETE EVENTSERVER (Delete the definition of the event server)

Use this command to delete the definition of the event server. You must issue this command before you issue the **DELETE SERVER** command. If you specify the server defined as the event server on the **DELETE SERVER** command, you will receive an error message.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—DELeTe EVENTSErVer—————►

Example: Delete an event server definition

Delete the definition for the event server ASTRO.

```
delete eventserver
```

Related commands

Table 129. Commands related to DELETE EVENTSERVER

Command	Description
DEFINE EVENTSERVER	Defines a server as an event server.
QUERY EVENTSERVER	Displays the name of the event server.

DELETE FILESPACE (Delete client node data from the server)

Use this command to delete file spaces from the server. Files that belong to the file space are deleted from primary, active-data, and copy storage pools.

Tivoli Storage Manager deletes one or more file spaces as a series of batch database transactions, thus preventing a rollback or commit for an entire file space as a single action. If the process is canceled or if a system failure occurs, a partial deletion can occur. A subsequent **DELETE FILESPACE** command for the same node or owner can delete the remaining data.

If this command is applied to a WORM (write once, read many) volume, the volume will return to scratch if it has space remaining in which data can be written. (Note that data on WORM volumes, including deleted and expired data, cannot be overwritten. Therefore, data can only be written in space that does not contain current, deleted, or expired data.) If a WORM volume does not have any space available in which data can be written, it will remain private. To remove the volume from the library, you must use the **CHECKOUT LIBVOLUME** command.


Important:

- If archive retention protection is enabled, the server will only delete archive files with expired retention periods. See the **SET ARCHIVERETENTIONPROTECTION** command for more information.
- The server will not delete archive files that are on deletion hold until the hold is released.
- Reclamation will not start while the **DELETE FILESPACE** process is running.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax



**Notes:**

- 1 This parameter can only be used when TYPE=ANY or TYPE=BACKUP is specified.

Parameters***node_name* (Required)**

Specifies the name of the client node to which the file space belongs.

***file_space_name* (Required)**

Specifies the name of the file space to be deleted. This name is case-sensitive and must be entered exactly as it is known to the server. To determine how to enter the name, use the **QUERY FILESPACE** command. You can use wildcard characters to specify this name.

For a server that has clients with support for Unicode, you may need to have the server convert the file space name that you enter. For example, you may need to have the server convert the name you enter from the server's code page to Unicode. See the **NAMETYPE** parameter for details. If you do not specify a file space name, or specify only a single wildcard character for the name, you can use the **CODETYPE** parameter to limit the operation to Unicode file spaces or to non-Unicode file spaces.

Type

Specifies the type of data to be deleted. This parameter is optional. The default value is ANY. Possible values are:

ANY

Delete only backed-up versions of files and archived copies of files.

If you specify `delete filesystem node_name * type=any`, all backed-up data and archived data in all file spaces for that node are deleted. File spaces are deleted only if they do not contain files that are migrated from a Tivoli Storage Manager for Space Management client.

Backup

Delete backup data for the file space.

ARchive

Delete all archived data on the server for the file space.

SPacemanaged

Delete files migrated from a user's local file system by a Tivoli Storage Manager for Space Management client. The **OWNER** parameter is ignored when you specify TYPE=SPACEMANAGED.

SERver

Delete all archived files in all file spaces for a node that is registered as TYPE=SERVER.

DAta

Specifies objects to delete. This parameter is optional. The default value is ANY. Possible values are:

ANY

Delete files, directories, and images.

DELETE FILESPACE

Files

Delete files and directories.

IMages

Delete image objects. You can only use this parameter if you have specified TYPE=ANY or TYPE=BACKUP.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is No. Possible values are:

No Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed.

Messages created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

Yes

Specifies that the server processes this command in the foreground. Wait for the command to complete before continuing with other tasks. The server displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify WAIT=YES from the server console.

OWNer

Restricts the data that is deleted to files belonging to the owner. This parameter is optional; it is ignored when TYPE=SPACEMANAGED. This parameter only applies to multiuser client systems such as AIX, Linux, and Solaris OS.

NAMEType

Specify how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with support for Unicode. A backup-archive client with support for Unicode is currently available only for the following operating systems: Windows, Macintosh OS X, and NetWare operating systems.

Use this parameter only when you enter a partly or fully qualified file space name. The default value is SERVER. Possible values are:

SERVER

The server uses the server's code page to interpret the file space names.

UNIcode

The server converts the file space names from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server has a problem accessing system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

CODEType

Specify what type of file spaces are to be included in the operation. The default is BOTH, meaning that file spaces are included regardless of code page type. Use this parameter only when you enter a single wildcard character for the file space name. Possible values are:

UNICODE

Include only file spaces that are in Unicode.

NONUNICODE

Include only file spaces that are not in Unicode.

BOTH

Include file spaces regardless of code page type.

Example: Delete a file space

Delete the file space named C_Drive that belongs to the client node HTANG.

```
delete filesystem htang C_Drive
```

Example: Delete all space-managed files for a client node

Delete all files migrated from client node APOLLO (that is, all space-managed files).

```
delete filesystem apollo * type=spacemanaged
```

Related commands

Table 130. Commands related to DELETE FILESPACE

Command	Description
CANCEL PROCESS	Cancels a background server process.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY OCCUPANCY	Displays file space information by storage pool.
QUERY PROCESS	Displays information about background processes.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
RENAME FILESPACE	Renames a client filesystem on the server.

DELETE GRPMEMBER

DELETE GRPMEMBER (Delete a server from a server group)

Use this command to delete a server or server group from a server group.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
▶▶—DELEte GRPMEMber—group_name—member_name—▶▶
```

Parameters

group_name (Required)

Specifies the group.

member_name (Required)

Specifies the server or group to delete from the group. To specify multiple names, separate the names with commas and no intervening spaces.

Example: Delete a server from a server group

Delete member PHOENIX from group WEST_COMPLEX.

```
delete grpmember west_complex phoenix
```

Related commands

Table 131. Commands related to DELETE GRPMEMBER

Command	Description
DEFINE GRPMEMBER	Defines a server as a member of a server group.
DEFINE SERVERGROUP	Defines a new server group.
DELETE SERVER	Deletes the definition of a server.
DELETE SERVERGROUP	Deletes a server group.
MOVE GRPMEMBER	Moves a server group member.
QUERY SERVER	Displays information about servers.
QUERY SERVERGROUP	Displays information about server groups.
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVERGROUP	Updates a server group.

DELETE KEYRING (Delete password information in the key database)

Use this command to delete the password information in the Tivoli Storage Manager database for the key database (`cert.kdb`).

This command is needed when the `SSLTCPPORT` or `SSLTCPADMINPORT` options are in use and the `cert.kdb` file has been lost or is not recoverable. If the `cert.kdb` file does not exist and there is no entry in the database for its password, Tivoli Storage Manager automatically generates a new self-signed certificate in a replacement `cert.kdb` file at server startup. The administrator then distributes the new public key (that is, the corresponding `cert.arm` file) to the clients that are using Secure Sockets Layer (SSL).

If the password information is lost after it was updated outside of the server, use this command to delete the key database file information from the server database. You can also delete `cert.*` files from the server instance directory. When the server is restarted, it regenerates the `cert.kdb` file.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—DELEte KEYRING—◄◄

Parameters

None

Example: Delete password information in the key database

The Tivoli Storage Manager administrator has deleted the current `cert.kdb` file and wants Tivoli Storage Manager to generate a new one at server startup for use by SSL.

```
delete keyring
```

Related commands

Table 132. Commands related to DELETE KEYRING

Command	Description
QUERY SSLKEYRINGPW	Displays the Secure Sockets Layer (SSL) key database file password.
SET SSLKEYRINGPW	Sets or updates the key database file password.

DELETE LIBRARY (Delete a library)

Use this command to delete a library. Before you delete a library, you must delete other associated objects, such as the path.

Use this command to delete a library. Before you delete a library, delete the path and all associated drives.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
►►—DELEte LIBRary—library_name—◄◄
```

Parameters

library_name **(Required)**

Specifies the name of the library to be deleted.

Example: Delete a manual library

Delete the manual library named LIBR1.

```
delete library libr1
```

Related commands

Table 133. Commands related to DELETE LIBRARY

Command	Description
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE PATH	Defines a path from a source to a destination.
DELETE DRIVE	Deletes a drive from a library.
DELETE PATH	Deletes a path from a source to a destination.
PERFORM LIBACTION	Defines all drives and paths for a library.
QUERY DRIVE	Displays information about drives.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE DRIVE	Changes the attributes of a drive.
UPDATE LIBRARY	Changes the attributes of a library.
UPDATE PATH	Changes the attributes associated with a path.


DELETE MACHINE (Delete machine information)

Use this command to delete machine description information. To replace existing information, issue this command and then issue an **INSERT MACHINE** command.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

machine_name (Required)

Specifies the name of the machine whose information is to be deleted.

Type

Specifies the type of machine information. This parameter is optional. The default is ALL. Possible values are:

All

Specifies all information.

RECOVERYInstructions

Specifies the recovery instructions.

Characteristics

Specifies the machine characteristics.

Example: Delete a specific machine's information

Delete the machine characteristics associated with the DISTRICT5 machine.

```
delete machine district5 type=characteristics
```

Related commands

Table 134. Commands related to **DELETE MACHINE**

Command	Description
DEFINE MACHINE	Defines a machine for DRM.
INSERT MACHINE	Inserts machine characteristics or recovery instructions into the Tivoli Storage Manager database.
QUERY MACHINE	Displays information about machines.
QUERY RECOVERYMEDIA	Displays media available for machine recovery.
UPDATE MACHINE	Changes the information for a machine.

DELETE MACHNODEASSOCIATION

DELETE MACHNODEASSOCIATION (Delete association between a machine and a node)

Use this command to delete the association between a machine and one or more nodes. This command does not delete the node from Tivoli Storage Manager.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
▶▶—DELEte MACHNODEAssociation—machine_name—node_name—▶▶
```

Parameters

machine_name (Required)

Specifies the name of a machine that is associated with one or more nodes.

node_name (Required)

Specifies the name of a node associated with a machine. If you specify a list of node names, separate the names with commas and no intervening spaces. You can use wildcard characters to specify a name. If a node is not associated with the machine, that node is ignored.

Example: Delete an association between a node and a machine

Delete the association between the DISTRICT5 machine and the ACCOUNTSPAYABLE node.

```
delete machnodeassociation district5 accountspayable
```

Related commands

Table 135. Commands related to DELETE MACHNODEASSOCIATION

Command	Description
DEFINE MACHNODEASSOCIATION	Associates a Tivoli Storage Manager node with a machine.
QUERY MACHINE	Displays information about machines.

DELETE MGMTCLASS (Delete a management class)

Use this command to delete a management class. You cannot delete a management class in the ACTIVE policy set. All copy groups in the management class are deleted along with the management class.

You can delete the management class assigned as the default for a policy set, but a policy set cannot be activated unless it has a default management class.

You can delete the predefined STANDARD management class in the STANDARD policy domain. However, if you later reinstall the Tivoli Storage Manager server, the process restores all STANDARD policy objects.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the management class belongs.

Syntax

```
►►—DELEte MGMTclass—domain_name—policy_set_name—class_name—►►
```

Parameters

domain_name **(Required)**

Specifies the policy domain to which the management class belongs.

policy_set_name **(Required)**

Specifies the policy set to which the management class belongs.

class_name **(Required)**

Specifies the management class to delete.

Example: Delete a management class

Delete the ACTIVEFILES management class from the VACATION policy set of the EMPLOYEE_RECORDS policy domain.

```
delete mgmtclass employee_records
vacation activefiles
```

Related commands

Table 136. Commands related to DELETE MGMTCLASS

Command	Description
ASSIGN DEFMGMTCLASS	Assigns a management class as the default for a specified policy set.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE MGMTCLASS	Defines a management class.
QUERY MGMTCLASS	Displays information about management classes.
UPDATE MGMTCLASS	Changes the attributes of a management class.

DELETE NODEGROUP (Delete a node group)

Use this command to delete a node group. You cannot delete a node group if it has any members in it.

Attention: You can remove all the members in the node group by issuing the **DELETE NODEGROUPMEMBER** command with a wildcard in the *node_name* parameter.

Privilege class

To issue this command, you must have system or unrestricted policy privilege.

Syntax

```
►►—DELEte NODEGrouP—group_name—◄◄
```

Parameters

group_name

Specifies the name of the node group that you want to delete.

Example: Delete a node group

Delete a node group named group1.

```
delete nodegroup group1
```

Related commands

Table 137. Commands related to **DELETE NODEGROUP**

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

DELETE NODEGROUPMEMBER (Delete node group member)

Use this command to delete a client node from a node group.

Privilege class

To issue this command, you must have system or unrestricted policy privilege.

Syntax

```

  >> DELEte NODEGROUPMember—group_name—node_name >>>
  
```

Parameters

group_name

Specifies the name of the node group from which you want to delete a client node.

node_name

Specifies the name of the client node that you want to delete from the node group. You can specify one or more names. When specifying multiple names, separate the names with commas; do not use intervening spaces. You can also use wildcard characters to specify multiple nodes.

Example: Delete node group members

Delete two nodes, node1 and node2, from a node group, group1.

```
delete nodegroupmember group1 node1,node2
```

Related commands

Table 138. Commands related to **DELETE NODEGROUPMEMBER**

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

Related commands*Table 139. Commands related to DELETE PATH*

Command	Description
DEFINE DATAMOVER	Defines a data mover to the Tivoli Storage Manager server.
DEFINE PATH	Defines a path from a source to a destination.
PERFORM LIBACTION	Defines all drives and paths for a library.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE PATH	Changes the attributes associated with a path.

DELETE POLICYSET (Delete a policy set)

Use this command to delete a policy set. You cannot delete the ACTIVE policy set. When you delete a policy set, all management classes and copy groups that belong to the policy set are also deleted.

You can delete the predefined STANDARD policy set. However, if you later reinstall the Tivoli Storage Manager server, the process restores all STANDARD policy objects.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

Syntax

```
►►—DELEte POLicyset—domain_name—policy_set_name—◄◄
```

Parameters

domain_name **(Required)**

Specifies the policy domain to which the policy set belongs.

policy_set_name **(Required)**

Specifies the policy set to delete.

Example: Delete a policy set

Delete the VACATION policy set from the EMPLOYEE_RECORDS policy domain.

Command

```
delete policyset employee_records vacation
```

Related commands

Table 140. Commands related to DELETE POLICYSET

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY POLICYSET	Creates a copy of a policy set.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
QUERY POLICYSET	Displays information about policy sets.
UPDATE POLICYSET	Changes the description of a policy set.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

DELETE PROFASSOCIATION (Delete a profile association)


Use this command on a configuration manager to delete the association of one or more objects from a profile. If associations are deleted, the objects are no longer distributed to subscribing managed servers. When managed servers request updated configuration information, the configuration manager notifies them of the object deletions.

A managed server deletes the objects that were deleted from the profile, unless the objects are associated with another profile to which that server subscribes.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

profile_name (Required)

Specifies the profile from which to delete associations.

ADMinS

Specifies the administrators whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all character (*) to delete all administrators from the profile. If you specify a list of administrators and a match-all definition exists for the profile, the command fails.

DELETE PROFASSOCIATION

Administrator definitions are not changed on the configuration manager. However, they are automatically deleted from all subscribing managed servers at the next configuration refresh, with the following exceptions:

- An administrator is not deleted if that administrator has an open session on the server.
- An administrator is not deleted if, as a result, the managed server would have no administrators with system privilege class.

DOmains

Specifies the domains whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all character (*) to delete all domains from the profile. If you specify a list of domains and a match-all domain definition exists for the profile, the command fails.

The domain information is automatically deleted from all subscribing managed servers. However, a policy domain that has client nodes assigned will not be deleted. To delete the domain at the managed server, assign those client nodes to another policy domain.

ADSCHeds

Specifies a list of administrative schedules whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. If you specify a list of administrative schedules and a match-all administrative schedule definition exists for the profile, the command fails. Use the match-all character (*) to delete all administrative schedules from the profile.

The administrative schedules are automatically deleted from all subscribing managed servers. However, an administrative schedule is not deleted if the schedule is active on the managed server. To delete an active schedule, make the schedule inactive.

SCRipts

Specifies the server command scripts whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all character (*) to delete all scripts from the profile. If you specify a list of scripts and a match-all script definition exists for the profile, the command fails. The server command scripts are automatically deleted from all subscribing managed servers.

CLOptsets

Specifies the client option sets whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. Use the match-all character (*) to delete all client option sets from the profile. If you specify a list of client option sets and a match-all client option set definition exists for the profile, the command fails. The client option sets are automatically deleted from all subscribing managed servers.

SERVers

Specifies the servers whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. You can use the match-all character (*) to delete all servers from the profile. If you specify a list of servers and a match-all server definition exists for the profile, the command fails. The server definitions are automatically deleted from all subscribing managed servers with the following exceptions:

- A server definition is not deleted if the managed server has an open connection to another server.
- A server definition is not deleted if the managed server has a device class of the device type SERVER that refers to the other server.
- A server definition is not deleted if the server is the event server for the managed server.

SERVERGroups

Specifies the server groups whose association with the profile is deleted. You can specify more than one name by separating the names with commas and no intervening spaces. You can use the match-all character (*) to delete all server groups from the profile. If you specify a list of server groups and a match-all group definition exists for the profile, the command fails. The server group definitions are automatically deleted from all subscribing managed servers.

Example: Delete the domain associations for a specific profile

Delete all domain associations from a profile named MIKE.

```
delete profassociation mike domains=*
```

Related commands

Table 141. Commands related to DELETE PROFASSOCIATION

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DELETE PROFILE	Deletes a profile from a configuration manager.
LOCK PROFILE	Prevents distribution of a configuration profile.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY PROFILE	Displays information about configuration profiles.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

DELETE PROFILE (Delete a profile)

Use this command on a configuration manager to delete a profile and stop its distribution to managed servers.


You cannot delete a locked profile. You must first unlock the profile with the **UNLOCK PROFILE** command.

Deleting a profile from a configuration manager does not delete objects associated with that profile from the managed servers. You can use the **DELETE SUBSCRIPTION** command with the **DISCARDOBJECTS=YES** parameter on each subscribing managed server to delete subscriptions to the profile and associated objects. This also prevents the managed servers from requesting further updates to the profile.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

profile_name (Required)

Specifies the profile to delete.

Force

Specifies whether the profile is deleted if one or more managed servers have subscriptions to that profile. The default is NO. Possible values are:

No Specifies that the profile is not deleted if one or more managed servers have subscriptions to that profile. You can delete the subscriptions on each managed server using the **DELETE SUBSCRIPTION** command.

Yes

Specifies that the profile is deleted even if one or more managed servers have subscriptions to that profile. Each subscribing server continues to request updates for the deleted profile until the subscription is deleted.

Examples: Delete a profile

Delete a profile named BETA, even if one or more managed servers subscribe to it.
`delete profile beta force=yes`

Related commands

Table 142. Commands related to **DELETE PROFILE**

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.

Table 142. Commands related to DELETE PROFILE (continued)

Command	Description
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE SUBSCRIPTION	Deletes a specified profile subscription.
LOCK PROFILE	Prevents distribution of a configuration profile.
QUERY PROFILE	Displays information about configuration profiles.
QUERY SUBSCRIPTION	Displays information about profile subscriptions.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

DELETE RECMEDMACHASSOCIATION (Delete recovery media and machine association)

Use this command to remove the association of one or more machines with a recovery media. This command does not delete the machine from Tivoli Storage Manager.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
▶▶—DELEte RECMEDMACHAssociation—media_name—machine_name—▶▶
```

Parameters

media_name (Required)

Specifies the name of the recovery media that is associated with one or more machines.

machine_name (Required)

Specifies the name of the machine associated with the recovery media. To specify a list of machine names, separate the names with commas and no intervening spaces. You can use wildcard characters to specify a name. If a machine is not associated with the recovery media, the machine is ignored.

Example: Delete a machine's association with recovery media

Delete the association between the DIST5RM recovery media and the DISTRICT1 and DISTRICT5 machines.

```
delete recmedmachassociation
dist5rm district1,district5
```

Related commands

Table 143. Commands related to DELETE RECMEDMACHASSOCIATION

Command	Description
DEFINE RECMEDMACHASSOCIATION	Associates recovery media with a machine.
QUERY MACHINE	Displays information about machines.
QUERY RECOVERYMEDIA	Displays media available for machine recovery.

DELETE RECOVERYMEDIA (Delete recovery media)

Use this command to delete a recovery media definition from Tivoli Storage Manager.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—DELEte RECOVERYMedia—*media_name*—◄◄

Parameters

media_name (Required)

Specifies the name of the recovery media.

Example: Delete a recovery media definition

Delete the DIST5RM recovery media.

```
delete recoverymedia dist5rm
```

Related commands

Table 144. Commands related to DELETE RECOVERYMEDIA

Command	Description
DEFINE RECOVERYMEDIA	Defines the media required to recover a machine.
QUERY RECOVERYMEDIA	Displays media available for machine recovery.
UPDATE RECOVERYMEDIA	Changes the attributes of recovery media.

DELETE SCHEDULE (Delete a client or an administrative command schedule)

Use this command to delete schedules from the database.

The DELETE SCHEDULE command takes two forms: one if the schedule applies to client operations, one if the schedule applies to administrative commands. The syntax and parameters for each form are defined separately.

- “DELETE SCHEDULE (Delete an administrative schedule)” on page 446
- “DELETE SCHEDULE (Delete a client schedule)” on page 445

Table 145. Commands related to DELETE SCHEDULE

Command	Description
COPY SCHEDULE	Creates a copy of a schedule.
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.
QUERY SCHEDULE	Displays information about schedules.
UPDATE SCHEDULE	Changes the attributes of a schedule.

DELETE SCHEDULE (Delete a client schedule)

Use the DELETE SCHEDULE command to delete one or more client schedules from the database. Any client associations to a schedule are removed when the schedule is deleted.

Privilege class

To delete a client schedule, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the specified policy domain.

Syntax

```
▶▶—DELEte SCHEDULE—domain_name—schedule_name—Type=Client—▶▶
```

Parameters***domain_name* (Required)**

Specifies the name of the policy domain to which the schedule belongs.

***schedule_name* (Required)**

Specifies the name of the schedule to delete. You can use a wildcard character to specify this name.

Type=Client

Specifies to delete a client schedule. This parameter is optional. The default is CLIENT.

Example: Delete a specific schedule from a specific policy domain

Delete the WEEKLY_BACKUP schedule, which belongs to the EMPLOYEE_RECORDS policy domain.

```
delete schedule employee_records weekly_backup
```

DELETE SCHEDULE

DELETE SCHEDULE (Delete an administrative schedule)

Use this command to delete one or more administrative command schedules from the database.

Privilege class

To delete an administrative command schedule, you must have system authority.

Syntax

```
▶▶—DELEte SChedule—schedule_name—Type—=—Administrative————▶▶
```

Parameters

schedule_name (Required)

Specifies the name of the schedule to delete. You can use a wildcard character to specify this name.

Type=Administrative (Required)

Specifies to delete an administrative command schedule.

Example: Delete an administrative command schedule

Delete the administrative command scheduled named DATA_ENG.

```
delete schedule data_eng type=administrative
```

DELETE SCRIPT (Delete command lines from a script or delete the entire script)

Use this command to delete a single line from a Tivoli Storage Manager script or to delete the entire Tivoli Storage Manager script.

Privilege class

To issue this command, the administrator must have previously defined the script or must have system privilege.

Syntax

```
▶▶ DELETE SCRIPT script_name [Line=number]
```

Parameters

script_name (Required)

Specifies the name of the script to delete. The script is deleted unless you specify a line number.

Line

Specifies the line number to delete from the script. If you do not specify a line number, the entire script is deleted.

Example: Delete a specific line from a script

Using the following script named QSAMPLE and issue a command to delete line 005 from it.

```
001 /* This is a sample script */
005 QUERY STATUS
010 QUERY PROCESS

delete script qsample line=5
```

Related commands

Table 146. Commands related to DELETE SCRIPT

Command	Description
COPY SCRIPT	Creates a copy of a script.
DEFINE SCRIPT	Defines a script to the Tivoli Storage Manager server.
QUERY SCRIPT	Displays information about scripts.
RENAME SCRIPT	Renames a script to a new name.
RUN	Runs a script.
UPDATE SCRIPT	Changes or adds lines to a script.

DELETE SERVER (Delete a server definition)

Use this command to delete a server definition.

This command fails if the server:

- Is defined as the event server.
- Is named in a device class definition whose device type is SERVER.
- Has an open connection to or from another server.
- Is a target server for virtual volumes.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—DELEte—SERver—*server_name*—◄◄

Parameters

server_name **(Required)**

Specifies a server name.

Example: Delete a server's definition

Delete the definition for a server named SERVER2.

```
delete server server2
```

Related commands

Table 147. Commands related to DELETE SERVER

Command	Description
DEFINE SERVER	Defines a server for server-to-server communications.
QUERY EVENTSERVER	Displays the name of the event server.
QUERY SERVER	Displays information about servers.
RECONCILE VOLUMES	Reconciles source server virtual volume definitions and target server archive objects.
UPDATE SERVER	Updates information about a server.

DELETE SERVERGROUP (Delete a server group)

Use this command to delete a server group. If the group you delete is a member of other server groups, Tivoli Storage Manager also removes the group from the other groups.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—DELEte SERVERGRoup—group_name—◄◄
```

Parameters

group_name (Required)

Specifies the server group to delete.

Example: Delete a server group

Delete a server group named WEST_COMPLEX.

```
delete servergroup west_complex
```

Related commands

Table 148. Commands related to DELETE SERVERGROUP

Command	Description
COPY SERVERGROUP	Creates a copy of a server group.
DEFINE GRPMEMBER	Defines a server as a member of a server group.
DEFINE SERVERGROUP	Defines a new server group.
DELETE GRPMEMBER	Deletes a server from a server group.
MOVE GRPMEMBER	Moves a server group member.
QUERY SERVERGROUP	Displays information about server groups.
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVERGROUP	Updates a server group.

DELETE SPACETRIGGER (Delete the storage pool space triggers)

Use this command to delete the definition of the storage pool space trigger.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
▶▶ DELEte SPACETrigger—STG—┌STGPOOL—==—storage_pool_name└┘▶▶
```

Parameters

STG

Specifies a storage pool space trigger.

STGPOOL

Specifies the storage pool trigger to be deleted. If STG is specified without specifying STGPOOL, the default storage pool space trigger is the deletion target.

Example: Delete a space trigger definition

Delete the space trigger definition for the WINPOOL1 storage pool.

```
delete spacetrigger stg stgpool=winpool1
```

Related commands

Table 149. Commands related to DELETE SPACETRIGGER

Command	Description
DEFINE SPACETRIGGER	Defines a space trigger to expand the space for a storage pool.
QUERY SPACETRIGGER	Displays information about a storage pool space trigger.
UPDATE SPACETRIGGER	Changes attributes of storage pool space trigger.

DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)

Use this command to delete an existing status monitoring threshold.

Status monitoring thresholds compare the defined conditions to the status monitoring server queries and inserts the results in the status monitoring table.

Multiple thresholds can be defined for an activity. For example, you can create a threshold that provides a warning status if storage pool capacity utilization is greater than 80%. You can then create another threshold that provides error status if storage pool capacity utilization is greater than 90%.

Note: If a threshold is already defined for an EXISTS condition, you cannot define another threshold with one of the other condition types.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—DELEte STATusthreshold—threshold_name—►►
```

Parameters

threshold_name (**Required**)

Specifies the threshold name that you want to delete.

Delete an existing status threshold

Delete an existing status threshold by issuing the following command:

```
delete statusthreshold avgstgpl
```

Related commands

Table 150. Commands related to DELETE STATUSTHRESHOLD

Command	Description
"DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)" on page 337	Defines a status monitoring threshold.
"QUERY MONITORSTATUS (Query the monitoring status)" on page 799	Displays information about monitoring alerts and server status settings.
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"QUERY STATUSTHRESHOLD (Query status monitoring thresholds)" on page 914	Displays information about a status monitoring thresholds.
"SET STATUSMONITOR (Specifies whether to enable status monitoring)" on page 1144	Specifies whether to enable status monitoring.
"SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)" on page 1146	Specifies the refresh interval for status monitoring.

DELETE STATUSTHRESHOLD

Table 150. Commands related to **DELETE STATUSTHRESHOLD** (continued)

Command	Description
"UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)" on page 1375	Changes the attributes of an existing status monitoring threshold.

DELETE STGPOOL (Delete a storage pool)

Use this command to delete a storage pool. To delete a storage pool, you must first delete all volumes assigned to the storage pool.

You cannot delete a storage pool that is identified as the next storage pool for another storage pool. For more information on storage pool hierarchy, see the `NEXTSTGPOOL` parameter in the `DEFINE STGPOOL` command.

Important:

- Do not delete a storage pool that is specified as a destination for a management class or copy group in the `ACTIVE` policy set. Client operations might fail as a result.
- When deleting a copy storage pool that has been previously included in a primary storage-pool definition (specifically in the `COPYSTGPOOLS` list), you must remove the copy storage pool from the list prior to deletion. Otherwise, the `DELETE STGPOOL` command fails until all references to that copy pool have been removed. For each primary storage pool with a reference to the copy storage pool to be deleted, remove the reference by entering the `UPDATE STGPOOL` command with the `COPYSTGPOOLS` parameter with all previous copy storage pools except the copy storage pool to be deleted.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—DELEte STGpool—pool_name—◄◄
```

Parameters

pool_name (Required)

Specifies the storage pool to delete.

Example: Delete a storage pool

Delete the storage pool named POOLA.

```
delete stgpool poola
```

Related commands

Table 151. Commands related to `DELETE STGPOOL`

Command	Description
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
QUERY STGPOOL	Displays information about storage pools.
UPDATE STGPOOL	Changes the attributes of a storage pool.

DELETE SUBSCRIBER (Delete subscriptions from a configuration manager database)

Use this command on a configuration manager to delete managed server subscriptions from the configuration manager database. Use this command when a managed server no longer exists or cannot notify the configuration manager after deleting a subscription.

Attention: Use this command only in rare situations in which the configuration manager's database contains an entry for a subscription, but the managed server does not have such a subscription. For example, use this command if a managed server no longer exists or cannot notify the configuration manager after deleting a subscription.

Under normal circumstances, use the **DELETE SUBSCRIPTION** command to delete a subscription from the managed server. The managed server notifies the configuration manager, which then deletes the subscription from its database.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—DELEte SUBSCRIBer—server_name—►►
```

Parameters

server_name **(Required)**

Specifies the name of the managed server with subscription entries to be deleted.

Example: Delete subscription entries for a specific managed server

Delete all subscription entries for a managed server named DAN.

```
delete subscriber dan
```

Related commands

Table 152. Commands related to **DELETE SUBSCRIBER**

Command	Description
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE SUBSCRIPTION	Deletes a specified profile subscription.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY SUBSCRIBER	Displays information about subscribers and their subscriptions to profiles.
QUERY SUBSCRIPTION	Displays information about profile subscriptions.


DELETE SUBSCRIPTION (Delete a profile subscription)

Use this command on a managed server to delete a profile subscription. You can also delete from the managed server all objects associated with the profile.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

profile_name (Required)

Specifies the name of the profile for which the subscription is to be deleted.

DISCARDobjects

Specifies whether objects associated with the profile are to be deleted on the managed server. This parameter is optional. The default is NO.

No Specifies that the objects are not to be deleted.

Yes

Specifies that the objects are to be deleted, unless they are associated with another profile for which a subscription is defined.

Example: Delete a profile subscription

Delete a subscription to a profile named ALPHA and its associated objects from a managed server.

```
delete subscription alpha discardobjects=yes
```

Related commands

Table 153. Commands related to DELETE SUBSCRIPTION

Command	Description
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE SUBSCRIBER	Deletes obsolete managed server subscriptions.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY SUBSCRIBER	Displays information about subscribers and their subscriptions to profiles.
QUERY SUBSCRIPTION	Displays information about profile subscriptions.

DELETE VIRTUALFSMAPPING (Delete a virtual file space mapping)

Use this command to delete a virtual file space mapping definition. Virtual file spaces containing data cannot be deleted unless you use the **DELETE FILESPACE** command first.

Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege
- Restricted policy privilege for the domain to which the NAS node is assigned

Syntax

```
▶▶—DELEte VIRTUALFSmapping—node_name—virtual_filespace_name————▶▶
```

Parameters

node_name (Required)

Specifies the NAS node on which the file system and path reside. You cannot use wildcard characters or specify a list of names.

virtual_filespace_name (Required)

Specifies the name of the virtual file space mapping definition to be deleted. Wildcard characters are allowed.

Example: Delete a virtual file space mapping

Delete the virtual file space mapping definition /mikesomedir for the NAS node named NAS1.

```
delete virtualfsmapping nas1 /mikesomedir
```

Related commands

Table 154. Commands related to DELETE VIRTUALFSMAPPING

Command	Description
DEFINE VIRTUALFSMAPPING	Define a virtual file space mapping.
QUERY VIRTUALFSMAPPING	Query a virtual file space mapping.
UPDATE VIRTUALFSMAPPING	Update a virtual file space mapping.

DELETE VOLHISTORY (Delete sequential volume history information)

Use this command to delete volume history file records that are no longer needed (for example, records for obsolete database backup volumes).

When you delete records for volumes that are not in storage pools (for example, database backup or export volumes), the volumes return to scratch status even if Tivoli Storage Manager acquired them as private volumes. Scratch volumes of device type FILE are deleted. When you delete the records for storage pool volumes, the volumes remain in the Tivoli Storage Manager database. When you delete records for recovery plan file objects from a source server, the objects on the target server are marked for deletion.

Use the DELETE BACKUPSET command to delete specified backup set volume information in the volume history file. Do not use this DELETE VOLHISTORY command to delete backup set volume information in the volume history file.

For users of DRM, the database backup expiration should be controlled with the SET DRMDBBACKUPEXPIREDAYS command instead of this DELETE VOLHISTORY command. Using the DELETE VOLHISTORY command removes Tivoli Storage Manager's record of the volume. This can cause volumes to be lost that were managed by the MOVE DRMEDIA command. The recommended way to manage the automatic expiration of DRM database backup volumes is by using the SET DRMDBBACKUPEXPIREDAYS command.

Notes:

1. Volumes for the most recent database backup series are not deleted.
2. Existing volume history files are not automatically updated with this command.
3. You can use the DEFINE SCHEDULE command to periodically delete volume history records.

Privilege class

To issue this command, you must have system privilege.


Syntax

```

▶▶ DElete VOLHistory TDate = date [ TTime = 23:59:59 ] [ TTime = time ]

```

DELETE VOLHISTORY



Parameters

TODate (Required)

Specifies the date to use to select sequential volume history information to be deleted. Tivoli Storage Manager deletes only those records with a date on or before the date you specify. You can specify the date using one of the values below:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	01/23/1999
TODAY	The current date	TODAY
TODAY- <i>days</i> or - <i>days</i>	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY-30 or -30. To delete records that are 30 or more days old, you can specify TODAY-30 or simply -30.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

TOTime

Specifies that you want to delete records created on or before this time on the specified date. This parameter is optional. The default is the end of the day (23:59:59). You can specify the time using one of the values below:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time on the specified date	12:30:22
NOW	The current time on the specified date	NOW
NOW+ <i>HH:MM</i> or + <i>HH:MM</i>	The current time plus hours and minutes on the specified date	NOW+03:00 or +03:00. If you issue the DELETE VOLHISTORY command at 9:00 with TOTIME=NOW+03:00 or TOTIME=+03:00, Tivoli Storage Manager deletes records with a time of 12:00 or earlier on the specified date.
NOW- <i>HH:MM</i> or - <i>HH:MM</i>	The current time minus hours and minutes on the specified date	NOW-03:30 or -03:30. If you issue the DELETE VOLHISTORY command at 9:00 with TOTIME=NOW-3:30 or TOTIME=-3:30, Tivoli Storage Manager deletes records with a time of 5:30 or earlier on the specified date.

Type (Required)

Specifies the type of records, which also meet the date and time criteria, to delete from the volume history file. Possible values are:

A11

Specifies to delete all records.

Note: The **DELETE VOLHISTORY** command does not delete records of remote volumes.

DBBackup

Specifies to delete only records that contain information about volumes used for database full and incremental backups, that is with volume types of BACKUPFULL and BACKUPINCR, and that meet the specified date and time criteria. The latest database full and incremental backup series will not be deleted.

DEVclass=class_name

Specifies the device class name that was used to create the database backups. This optional parameter can be used to delete database backups created using a server-to-server virtual volume device class. The type of the device class must be SERVER. This parameter can only be used to delete volume history entries of type BACKUPFULL, BACKUPINCR, or DBSNAPSHOT.

A full, incremental, or snapshot database backup volume is eligible to be deleted if all of the following conditions are met:

- The device class used to create the database backup volume matches the specified device class
- The volume was created on or before the specified date and time
- The volume is not part of the latest full plus incremental database backup series if the specified volume type is DBBackup, or snapshot database backup series if the volume type is DBSnapshot

DELETE VOLHISTORY

DBSnapshot

Specifies to delete only records that contain information about volumes used for snapshot database backups, and that meet the specified date and time criteria. The latest snapshot database backup will not be deleted.

DEVclass=classname

Specifies the device class name that was used to create the database backups. This optional parameter can be used to delete database backups created using a server-to-server virtual volume device class. The type of the device class must be SERVER. This parameter can only be used to delete volume history entries of type BACKUPFULL, BACKUPINCR, or DBSNAPSHOT.

A full, incremental, or snapshot database backup volume is eligible to be deleted if all of the following conditions are met:

- The device class used to create the database backup volume matches the specified device class
- The volume was created on or before the specified date and time
- The volume is not part of the latest full plus incremental database backup series if the specified volume type is DBBackup, or snapshot database backup series if the volume type is DBSnapshot

DBRpf

Specifies to delete only records that contain information about full and incremental database backup volumes and recovery plan file volumes.

EXPort

Specifies to delete only records that contain information about export volumes.

RPFfile

Specifies to delete only records that contain information about recovery plan file objects that are stored on a target server and that meet the specified date and time criteria.

DELETEDatest

Specifies whether the latest recovery plan file is eligible for deletion. This optional parameter can be used to delete the latest recovery plan files created using a server-to-server virtual volume device class.

This parameter can only be used to delete volume history entries of type RPFfile (for instance, those recovery plan files that were created using the DEVCLASS parameter with the PREPARE command). If this parameter is not specified, the latest RPFfile entries are not deleted.

No Specifies the latest RPFfile file is not deleted.

Yes Specifies the latest RPFfile file is deleted if it meets the specified date and time criteria.

RPFSnapshot

Specifies to delete only records that contain information about recovery plan file objects that were created assuming snapshot database backups, that are stored on a target server and that meet the specified date and time criteria. The latest RPFSNAPSHOT file will not be deleted unless it meets the specified date and time criteria, and the DELETEDatest parameter is set to Yes.

DELETEDatest

Specifies whether the latest recovery plan file is eligible for deletion.

This optional parameter can be used to delete the latest recovery plan files created using a server-to-server virtual volume device class.

This parameter can only be used to delete volume history entries of type RPFSNAPSHOT (for instance, those recovery plan files that were created using the DEVCLASS parameter with the PREPARE command). If this parameter is not specified, the latest RPFSNAPSHOT entries are not deleted.

No Specifies the latest RPFSNAPSHOT file is not deleted.

Yes Specifies the latest RPFSNAPSHOT file is deleted if it meets the specified date and time criteria.

STGNew

Specifies to delete only records that contain information about new sequential access storage volumes.

STGReuse

Specifies to delete only records that contain information about reused sequential storage pool volumes.

STGDelete

Specifies to delete only records that contain information about deleted sequential storage pool volumes.

Example: Delete recovery plan file information

Delete all recovery plan file information created on or before 03/28/2005.

```
delete volhistory type=rpfile todate=03/28/2005
```

Related commands

Table 155. Commands related to DELETE VOLHISTORY

Command	Description
BACKUP VOLHISTORY	Records volume history information in external files.
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.
DELETE VOLUME	Deletes a volume from a storage pool.
EXPIRE INVENTORY	Manually starts inventory expiration processing.
MOVE DRMEDIA	Moves DRM media on-site and off-site.
PREPARE	Creates a recovery plan file.
QUERY RPFILE	Displays information about recovery plan files.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
SET DRMRPFEXPIREDAYS	Set criteria for recovery plan file expiration.
SET DRMDBBACKUPEXPIREDAYS	Specifies criteria for database backup series expiration.

DELETE VOLUME (Delete a storage pool volume)

Use this command to delete a storage pool volume and, optionally, the files stored in the volume.

If the volume has data, to delete the volume you must do one of the following:

- Before deleting the volume, use the **MOVE DATA** command to move all files to another volume.
- Explicitly request to discard all files in the volume when the volume is deleted (by specifying **DISCARDDATA=YES**).

If you are deleting several volumes, delete the volumes one at a time. Deleting more than one volume at a time can adversely affect server performance.

Storage pool volumes cannot be deleted if they are in use. For example, a volume cannot be deleted if a user is restoring or retrieving a file residing in the volume, if the server is writing information to the volume, or if a reclamation process is using the volume.

If you issue the **DELETE VOLUME** command, volume information is deleted from the Tivoli Storage Manager database. However, the physical files that are allocated with **DEFINE VOLUME** command are not removed from the file space.

If this command is applied to a WORM (write once, read many) volume, the volume returns to scratch if it has space remaining in which data can be written. (Note that data on WORM volumes, including deleted and expired data, cannot be overwritten. Therefore, data can only be written in space that does not contain current, deleted, or expired data.) If a WORM volume does not have any space available in which data can be written, it remains private. To remove the volume from the library, you must use the **CHECKOUT LIBVOLUME** command.

The **DELETE VOLUME** command automatically updates the server library inventory for sequential volumes if the volume is returned to scratch status when the volume becomes empty. To determine whether a volume will be returned to scratch status, issue the **QUERY VOLUME** command and look at the output. If the value for the attribute "Scratch Volume?" is "Yes," then the server library inventory is automatically updated.

If the value is "No," you can issue the **UPDATE LIBVOLUME** command to specify the status as scratch. It is recommended that you issue the **UPDATE LIBVOLUME** command after issuing the **DELETE VOLUME** command.

Attempting to use the **DELETE VOLUME** command to delete WORM FILE volumes in a storage pool with **RECLAMATIONTYPE=SNAPLOCK** fails with an error message. Deletion of empty WORM FILE volumes is performed only by the reclamation process.


If you issue the **DELETE VOLUME** command for a volume in a storage pool that has a **SHRED** parameter value greater than 0, the volume is placed in the pending state until shredding is run. Shredding is necessary to complete the deletion, even if the volume is empty.

If you issue the **DELETE VOLUME** command for a volume in a storage pool that is set up for data deduplication, the Tivoli Storage Manager destroys any object that is referencing data on that volume.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to which the volume is defined.

Syntax



Parameters

volume_name (Required)

Specifies the name of the volume to delete.

DISCARDdata

Specifies whether files stored in the volume are deleted. This parameter is optional. The default value is NO. Possible values are:

No Specifies that files stored in the volume are not deleted. If the volume contains any files, the volume is not deleted.

Yes

Specifies that all files stored in the volume are deleted. The server does not need to mount the volume for this type of deletion.

Remember:

1. The Tivoli Storage Manager server does not delete archive files that are on deletion hold.
2. If archive retention protection is enabled, the Tivoli Storage Manager server deletes only archive files whose retention period has expired.

If the volume being deleted is a primary storage pool volume, the server checks whether any copy storage pool has copies of files that are being deleted. When files stored in a primary storage pool volume are deleted, any copies of these files in copy storage pools are also deleted.

When you delete a disk volume in a primary storage pool, the command also deletes any files that are cached copies (copies of files that have been migrated to the next storage pool). Deleting cached copies of files does not delete the files that have already been migrated or backed up to copy storage pools. Only the cached copies of the files are affected.

If the volume being deleted is a copy storage pool volume, only files on the copy pool volume are deleted. The primary storage pool files are not affected.

Do not use the **DELETE VOLUME** command with DISCARDDATA=YES if a restore process (**RESTORE STGPOOL** or **RESTORE VOLUME**) is running. The **DELETE VOLUME** command could cause the restore to be incomplete.

DELETE VOLUME

If you cancel the **DELETE VOLUME** operation during processing or if a system failure occurs, some files might remain on the volume. You can delete the same volume again to have the server delete the remaining files and then the volume.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter affects processing only when you have also requested that any data on the volume be discarded. This parameter is optional. The default value is No. Possible values are:

No Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed.

The server displays messages that are created from the background process either in the activity log or the server console, depending on where messages are logged.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client when the command completes.

Remember: You cannot specify WAIT=YES from the server console.

Example: Delete a storage pool volume

Delete storage pool volume stgvol.1 from the storage pool FILEPOOL.

```
delete volume stgvol.1
```

Related commands

Table 156. Commands related to **DELETE VOLUME**

Command	Description
CANCEL PROCESS	Cancels a background server process.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
MOVE DATA	Moves data from a specified storage pool volume to another storage pool volume.
MOVE DRMEDIA	Moves DRM media on-site and off-site.
QUERY CONTENT	Displays information about files in a storage pool volume.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY PROCESS	Displays information about background processes.
QUERY VOLUME	Displays information about storage pool volumes.
UPDATE VOLUME	Updates the attributes of storage pool volumes.

DISABLE commands

Use **DISABLE** commands to prevent some types of operations by the server.

- “DISABLE EVENTS (Disable events for event logging)” on page 466
- “DISABLE REPLICATION (Prevent outbound replication processing on a server)” on page 470
- “DISABLE SESSIONS (Prevent new sessions from accessing Tivoli Storage Manager)” on page 471

DISABLE EVENTS (Disable events for event logging)

Use this command to disable the processing of one or more events. If you specify a receiver that is not supported on any platform, or if you specify an invalid event or name, Tivoli Storage Manager issues an error message. However, any valid receivers, events, or names that you specified are still enabled.

Tip: Messages in the SEVERE category and message ANR9999D can provide valuable diagnostic information if there are serious server problems. For this reason, you should not disable these messages.

Restriction:


- Certain messages are displayed on the console even if they are disabled. These include some messages issued during server startup and shutdown and responses to administrative commands.
- Server messages from the server on which this command is issued cannot be disabled for the activity log.

ANR1822I indicates that event logging is being ended for the specified receiver. When the **DISABLE EVENTS** command is issued, this message is logged to the receiver even if it is one of the events that has been disabled. This is done to confirm that event logging has ended to that receiver, but subsequent ANR1822I messages are not logged to that receiver.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

receivers (Required)

Specifies the name of the receivers for which to disable events. Specify multiple receivers by separating them with commas and no intervening spaces. Possible values are:

ALL

All receivers, except for server events on the activity log receiver (ACTLOG). Only client events can be disabled for the activity log receiver.

CONSOLE

The standard server console as a receiver.

ACTLOG

The activity log as a receiver. You can disable only client events, not server events, for the activity log.

EVENTSERVER

The event server as a receiver.

FILE

A user file as a receiver. Each logged event is a record in the file. The records are not easily readable by people.

FILETEXT

A user file as a receiver. Each logged event is a fixed-size, readable line.

NTEVENTLOG

The Windows application log as a receiver.

SNMP

The simple network management protocol (SNMP) as a receiver.

TIVOLI

The Tivoli Enterprise Console® (TEC) as a receiver.

USEREXIT

A user-written program as a receiver. The server writes information to the program.

events (Required)

Specifies the events to be disabled. You can specify multiple events by separating them with commas and no intervening spaces. Possible values are:

ALL

All events.

event_name

A four-digit message number preceded by ANR for a server event or ANE for a client event. Valid ranges are from ANR0001 to ANR9999 and from ANE4000 to ANE4999. Specify the NODENAMES parameter if client events are to be disabled for matching nodes. Specify the SERVERNAME parameter if server events are to be disabled for matching servers.

For the TIVOLI event receiver only, you can specify the following events names for the IBM Tivoli Storage Manager application clients:

Tivoli Storage Manager application client	Prefix	Range
Data Protection for Microsoft Exchange Server	ACN	3500–3649
Data Protection for Lotus® Domino®	ACD	5200–5299

DISABLE EVENTS

Tivoli Storage Manager application client	Prefix	Range
Data Protection for Oracle	ANS	500–599
Data Protection for Informix®	ANS	600–699
Data Protection for Microsoft SQL Server	ACO	3000–3999

Remember: Specifying ALL disables these messages. However, the INFO, WARNING, ERROR, and SEVERE options have no effect on the messages.

severity categories

If the event list contains a severity category, all events of that severity are disabled for the specified nodes. The message types are:

INFO

Information messages (type of I).

WARNING

Warning messages (type of W).

ERROR

Error messages (type of E).

SEVERE

Severe error messages (type of S).

NODENAME

Specifies the name of one or more node names for which events are to be disabled. You can use the wildcard character (*) to specify all nodes. You can specify NODENAME or SERVERNAME. If neither parameter is specified, the events are disabled for the server running this command.

SERVERNAME

Specifies the name of one or more server names for which events are to be disabled. You can use the wildcard character (*) to specify all servers other than the server running this command. You can specify NODENAME or SERVERNAME. If neither parameter is specified, the events are disabled for the server running this command.

Example: Disable specific categories of events

Disable all client events in the INFO and WARNING categories for the activity log and console receivers for all nodes.

```
disable events actlog,console  
info,warning nodename=*
```

Related commands

Table 157. Commands related to **DISABLE EVENTS**

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
ENABLE EVENTS	Enables specific events for receivers.
END EVENTLOGGING	Ends event logging to a specified receiver.
QUERY ENABLED	Displays enabled or disabled events for a specific receiver.
QUERY EVENTRULES	Displays information about rules for server and client events.

*Table 157. Commands related to **DISABLE EVENTS** (continued)*

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

DISABLE REPLICATION (Prevent outbound replication processing on a server)

Use this command to prevent a source replication server from starting new replication processes.

The use of this command does not stop running replication processes. Running replication processes continue until they complete or until they end without completing. Use this command and the **ENABLE REPLICATION** command to control replication processing.

Issue this command on the server that acts as a source for replicated data.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—DISAbLe REPLicAtion—◄◄

Parameters

None.

Example: Disable replication processing

Disable replication processing on a source replication server.

```
disable replication
```

Related commands

Table 158. Commands related to **DISABLE REPLICATION**

Command	Description
CANCEL REPLICATION	Cancels node replication processes.
DISABLE SESSIONS	Prevents new sessions from accessing Tivoli Storage Manager but permits existing sessions to continue.
ENABLE REPLICATION	Allows outbound replication processing on a server.
ENABLE SESSIONS	Resumes server activity following the DISABLE command or the ACCEPT DATE command.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.

DISABLE SESSIONS (Prevent new sessions from accessing Tivoli Storage Manager)


Use this command to prevent new sessions from accessing Tivoli Storage Manager. Active sessions will complete. For a particular server, you can specify whether to disable inbound sessions, outbound sessions, or both.

Server processes, such as migration and reclamation, are not affected when you issue the **DISABLE SESSIONS** command.

Privilege class

To issue this command, you must have system privilege or operator privilege.

Syntax



Parameters

Specifies the type of session to be disabled. This parameter is optional. The default value is CLIENT. You can specify one of the following values:

CLient

Disables only backup and archive client sessions.

ALL

Disables all session types.

ADMin

Disables only administrative sessions.

SERVer

Disables only server-to-server sessions. Only the following types of sessions are disabled:

- Server-to-server event logging
- Enterprise management
- Server registration
- LAN-free: storage agent - server
- Virtual volumes
- Node replication

You can also specify whether to disable inbound sessions, outbound sessions, or both for a particular server.

DISABLE SESSIONS

server_name

Specifies the name of a server whose sessions you want to disable. This parameter is optional. If you do not specify this parameter, new sessions with other servers do not start. Running sessions are not canceled.

DIRection

Specifies whether to disable inbound sessions, outbound sessions, or both. This parameter is optional. The default is BOTH. The following values are possible:

Both

Specifies that inbound sessions from the specified server and outbound sessions to the specified server are disabled.

INbound

Specifies that only inbound sessions from the specified server are disabled.

OUTbound

Specifies that only outbound sessions to the specified server are disabled.

Example: Prevent new client node backup and archive sessions on the server

Temporarily prevent new client node sessions from accessing the server.

```
disable sessions
```

Example: Prevent all new sessions on the server

Temporarily prevent any new sessions from accessing the server.

```
disable sessions all
```

Example: Disable outbound sessions to a server

Disable outbound sessions to a server named REPLSRV.

```
disable sessions server replsrv direction=outbound
```

Related commands

Table 159. Commands related to DISABLE SESSIONS

Command	Description
CANCEL SESSION	Cancels active sessions with the server.
DISABLE REPLICATION	Prevents outbound replication processing on a server.
ENABLE SESSIONS	Resumes server activity following the DISABLE command or the ACCEPT DATE command.
QUERY SESSION	Displays information about all active administrator and client sessions with Tivoli Storage Manager.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

DISMOUNT command

Use the **DISMOUNT** command to dismount a volume by the real device address or by volume name.

- “DISMOUNT VOLUME (Dismount a volume by volume name)” on page 474

DISMOUNT VOLUME

DISMOUNT VOLUME (Dismount a volume by volume name)

Use this command to dismount an idle volume by volume name. If a drive cannot dismount the volume, manual intervention is required.

Privilege class

To issue this command, you must have system privilege or operator privilege.

Syntax

►►—DISMOUNT Volume—*volume_name*—◄◄

Parameters

volume_name (Required)

Specifies the name of the volume to dismount.

Example: Dismount a specific volume

Dismount the volume BTV005.

```
dismount volume btv005
```

Related commands

Table 160. Command related to DISMOUNT VOLUME

Command	Description
QUERY MOUNT	Displays information about mounted sequential access media.

DISPLAY OBJNAME (Display a full object name)

Use this command when you want Tivoli Storage Manager to display a full object name if the name displayed in a message or query output has been abbreviated due to length. Object names that are very long can be difficult to display and use through normal operating system facilities. The Tivoli Storage Manager server will abbreviate long names and assign them a token ID which might be used if the object path name exceeds 1024 bytes. The token ID is displayed in a string that includes identifiers for the node, filesystem, and object name. The format is: [TSMOBJ:*nID.fsID.objID*]. When specified with the **DISPLAY OBJNAME** command, the token ID can be used to show the full object name.

Privilege class

Any administrator can issue this command

Syntax

►—DISplay OBJname—*token_ID*—◄

Parameters

token_ID (Required)

Specifies the ID reported in the [TSMOBJ:] tag, when an object name is too long to display.

Example: Display the full object name of a token ID in a message

Assume the you receive the following message:

```
ANR9999D file.c(1999) Error handling file [TSMOBJ:1.1.649498] because
of lack of server resources.
```

Display the full object name for the file referenced in the error message by specifying the token ID on the DISPLAY OBJNAME command.

```
display obj 1.1.649498
```

Related commands

Table 161. Commands related to DISPLAY OBJNAME

Command	Description
QUERY CONTENT	Displays information about files in a storage pool volume.

ENABLE commands

Use **ENABLE** commands to allow some types of operations by the server.

- “ENABLE EVENTS (Enable server or client events for logging)” on page 477
- “ENABLE REPLICATION (Allow outbound replication processing on a server)” on page 480
- “ENABLE SESSIONS (Resume user activity on the server)” on page 481

ENABLE EVENTS (Enable server or client events for logging)

Use this command to enable the processing of one or more events. If you specify a receiver that is not supported on any platform, or if you specify an invalid event or name, Tivoli Storage Manager issues an error message. However, any valid receivers, events, or names that you specified are still enabled.


Restriction: Certain events, such as some messages issued during server start-up and shutdown, automatically go to the console. They do not go to other receivers even if they are enabled.

Administrative commands are returned to the command issuer and are only logged as numbered events. These numbered events are not logged to the system console, but are logged to other receivers, including administrative command-line sessions running in console mode.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

receivers (Required)

Specifies one or more receivers for which to log enabled events. You can specify multiple receivers by separating them with commas and no intervening spaces. Valid values are:

ALL

All receivers.

CONSOLE

The standard server console as a receiver.

ACTLOG

The server activity log as a receiver.

ENABLE EVENTS

EVENTSERVER

The event server as a receiver.

FILE

A user file as a receiver. Each logged event is a record in the file. The records are not easily readable by people.

FILETEXT

A user file as a receiver. Each logged event is a fixed-size, readable line.

SNMP

The simple network management protocol (SNMP) as a receiver.

TIVOLI

The Tivoli Enterprise Console (TEC) as a receiver.

USEREXIT

A user-written program as a receiver. The server writes information to the program.

events (Required)

Specifies the type of events to be enabled. You can specify multiple events by separating them with commas and no intervening spaces. Possible values are:

ALL

All events.

event_name

A four-digit message number preceded by ANR for a server event or ANE for a client event. Valid ranges are from ANR0001 to ANR9999 and from ANE4000 to ANE4999. Specify the NODENAME parameter if client events are to be enabled for matching nodes. Specify the SERVERNAME parameter if server events are to be enabled for matching servers.

For the TIVOLI event receiver, you can specify the following additional ranges for the Tivoli Storage Manager application clients:

Tivoli Storage Manager application client	Prefix	Range
Data Protection for Microsoft Exchange Server	ACN	3500–3649
Data Protection for Lotus Domino	ACD	5200–5299
Data Protection for Oracle	ANS	500–599
Data Protection for Informix	ANS	600–699
Data Protection for Microsoft SQL Server	ACO	3000–3999

Restriction: The application client must have enhanced Tivoli Event Console support enabled in order to route these messages to the Tivoli Event Console.

Tip:

- Specifying the ALL option enables these messages. However, the INFO, WARNING, ERROR, and SEVERE options have no effect on the messages.
- Because of the number of messages, you should not enable all messages from a node to be logged to the Tivoli Event Console.

severity categories

If the event list contains a severity category, all events of that severity are enabled for the specified nodes. The message types are:

INFO

Information messages (type of I) are enabled.

WARNING

Warning messages (type of W) are enabled.

ERROR

Error messages (type of E) are enabled.

SEVERE

Severe error messages (type of S) are enabled.

NODENAME

Specifies one or more client nodes for which events are enabled. You can use a wildcard character to specify all client nodes. You can specify NODENAME or SERVERNAME. If neither parameter is specified, events are enabled for the server running this command.

SERVERNAME

Specifies one or more servers for which events are to be enabled. You can use a wildcard character to specify all servers other than the server from which this command is issued. You can specify SERVERNAME or NODENAME. If neither parameter is specified, the events are enabled for the server running this command.

Example: Enable specific categories of events

Enable all ERROR and SEVERE client events to the USEREXIT receiver for the node BONZO.

```
enable events userexit error,severe nodename=bonzo
```

Related commands

Table 162. Commands related to **ENABLE EVENTS**

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
DISABLE EVENTS	Disables specific events for receivers.
END EVENTLOGGING	Ends event logging to a specified receiver.
QUERY ENABLED	Displays enabled or disabled events for a specific receiver.
QUERY EVENTRULES	Displays information about rules for server and client events.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

ENABLE REPLICATION (Allow outbound replication processing on a server)

Use this command to allow a source replication server to begin normal replication processing after a database restore. You can also use this command to resume replication processing after issuing the **DISABLE REPLICATION** command.

Attention: Before enabling replication after a database restore, determine whether copies of data that are on the target server are needed. If they are, you must synchronize client node data by replicating the data from the target replication server to the source replication server. The replication process replaces the data on the source server that was lost because of the database restore.

For details about this procedure, see the *Administrator's Guide*.

Issue this command on the server that acts as a source for replicated data.

Privilege class

To issue this command, you must have system privilege.

Syntax

▶▶—ENable REPLication—▶▶

Parameters

None.

Example: Allow replication processing

Allow replication processing on a source replication server.

```
enable replication
```

Related commands

Table 163. Commands related to ENABLE REPLICATION

Command	Description
DISABLE REPLICATION	Prevents outbound replication processing on a server.
DISABLE SESSIONS	Prevents new sessions from accessing Tivoli Storage Manager but permits existing sessions to continue.
ENABLE SESSIONS	Resumes server activity following the DISABLE command or the ACCEPT DATE command.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.

ENABLE SESSIONS (Resume user activity on the server)

Use this command after issuing the **DISABLE SESSIONS** command to start new sessions that can access a server. For a particular server, you can specify whether to enable inbound sessions, outbound sessions, or both.


The processing of this command does not affect system processes, such as migration and reclamation.

Use the **QUERY STATUS** command to display the availability of the server.

Privilege class

To issue this command, you must have system privilege or operator privilege.

Syntax



Parameters

Specifies the type of session to be enabled. This parameter is optional. The default value is CLIENT. You can specify one of the following values:

CLIENT

Enables only backup and archive client sessions.

ALL

Enables all session types.

ADMIN

Enables only administrative sessions.

SERVER

Enables only server-to-server sessions. You can also specify whether to enable inbound sessions, outbound sessions, or both for a particular server.

server_name

Specifies the name of a particular server whose sessions you want to enable. This parameter is optional. If you do not specify this parameter, new sessions with all other servers are enabled.

DIRECTION

Specifies whether to enable inbound sessions, outbound sessions, or both. This parameter is optional. The default is BOTH. The following values are possible:

ENABLE SESSIONS

Both

Specifies that inbound sessions from the specified server and outbound sessions to the specified server are enabled.

INbound

Specifies that only inbound sessions to the specified server are enabled.

OUTbound

Specifies that only outbound sessions from the specified server are enabled.

Example: Resume client node activity on the server

Resume normal operation, permitting client nodes to access the server.

```
enable sessions
```

Example: Resume all activity on the server

Resume normal operation, permitting all sessions to access the server.

```
enable sessions all
```

Example: Enable outbound sessions to a server

Enable outbound sessions to a server named REPLSRV.

```
enable sessions server replsrv direction=outbound
```

Related commands

Table 164. Commands related to ENABLE SESSIONS

Command	Description
ACCEPT DATE	Accepts the current date on the server.
CANCEL SESSION	Cancels active sessions with the server.
ENABLE REPLICATION	Allows outbound replication processing on a server.
DISABLE SESSIONS	Prevents new sessions from accessing Tivoli Storage Manager but permits existing sessions to continue.
QUERY SESSION	Displays information about all active administrator and client sessions with Tivoli Storage Manager.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.


END EVENTLOGGING (Stop logging events)

Use this command to stop logging events to an active receiver.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

Specify a type of receiver. You can specify multiple receivers by separating them with commas and no intervening spaces. This is an optional parameter. The default is ALL. If you specify ALL or no receiver, logging ends for all receivers.

ALL

Specifies all receivers.

CONSOLE

Specifies the server console as a receiver.

ACTLOG

Specifies the Tivoli Storage Manager activity log as a receiver. Logging can be stopped only for client events.

EVENTSERVER

Specifies the event server as a receiver.

FILE

Specifies a user file as a receiver. Each logged event is a record in the file and a person cannot read each logged event easily.

FILETEXT

Specifies a user file as a receiver. Each logged event is a fixed-size, readable line.

SNMP

Specifies the simple network management protocol (SNMP) as a receiver.

TIVOLI

Specifies the Tivoli Management Environment (TME) as a receiver.

USEREXIT

Specifies a user-written routine to which Tivoli Storage Manager writes information as a receiver.

END EVENTLOGGING

Example: Stop logging events

End logging of events to the user exit.
end eventlogging userexit

Related commands

Table 165. Commands related to END EVENTLOGGING

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
DISABLE EVENTS	Disables specific events for receivers.
ENABLE EVENTS	Enables specific events for receivers.
QUERY ENABLED	Displays enabled or disabled events for a specific receiver.
QUERY EVENTRULES	Displays information about rules for server and client events.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

EXPIRE INVENTORY (Manually start inventory expiration processing)

Use this command to manually start inventory expiration processing. The inventory expiration process removes client backup and archive file copies from server storage. Removal is based on policy specifications in the backup and archive copy groups of the management classes to which the files are bound.

When you have the disaster recovery manager function for your Tivoli Storage Manager server, the inventory expiration process also removes eligible virtual volumes that are used by the following processes:

- Database backups of type BACKUPFULL, BACKUPINCR, and DBSNAPSHOT. The **SET DRMDBBACKUPEXPIREDAYS** command controls when these volumes are eligible for expiration.
- Recovery plan files of type RPFIL and RPFNSNAPSHOT. The **SET DRMRPFEXPIREDAYS** command controls when these volumes are eligible for expiration.

The inventory expiration process that runs during server initialization does not remove these virtual volumes.

Only one expiration process is allowed at any time, but this process can be distributed among a maximum of 40 threads. If an expiration process is currently running, you cannot start another process.

You can set up automatic expiration processing with the EXPINTERVAL server option. If you set the EXPINTERVAL option to 0, the server does not run expiration automatically, and you must issue the **EXPIRE INVENTORY** command to start expiration processing.

This command creates a background process that can be canceled with the **CANCEL PROCESS** command. To display information about background processes, use the **QUERY PROCESS** command.

If this command is applied to a WORM volume, the volume returns to being a scratch volume if it has remaining space in which data can be written. Data on WORM volumes, including deleted and expired data, cannot be overwritten. Therefore, data can be written only in space that does not contain current, deleted, or expired data. If a WORM volume does not have any space available in which data can be written, it remains private. To remove the volume from the library, you must use the **CHECKOUT LIBVOLUME** command.


Run the **EXPIRE INVENTORY** command to delete files from server storage if they were not completely deleted when you used client delete operations.

See the *Backup-Archive Clients Installation and User's Guide* for more information about client delete operations.


Privilege class

To issue this command, you must have system privilege.

Syntax



EXPIRE INVENTORY



Parameters

Quiet

Specifies whether the server suppresses detailed messages about policy changes during the expiration processing. This parameter is optional. The default is NO. Possible values are:

No Specifies that the server sends detailed informational messages.

Yes

Specifies that the server sends only summary messages. The server issues messages about policy changes only when files are deleted and either the default management class or retention grace period for the domain was used to expire the files.

You can also specify the EXPQUIET option in the server options file to automatically determine whether expiration processing is run with summary messages.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is NO. Possible values are:

No Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed.

The server displays messages that are created from the background process either in the activity log or the server console, depending on where messages are logged.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify WAIT=YES from the server console.

SKipdirs

Specifies whether the server skips directory type objects during the expiration processing. This parameter is optional. The default is NO. Possible values are:

No Specifies that the server expires files and directories based upon the appropriate policy criteria.

Yes

Specifies that the server will skip directory type backup and archive objects during expiration processing, even if the directories are eligible for expiration. By specifying YES, you prevent deletion of directories, and expiration processing can occur more quickly.

Attention: This option should not be used all of the time. With Tivoli Storage Manager version 6.0 and later, you can run multiple threads (resources) for an expiration process. Also, if you specify YES often, the database grows as the directory objects accumulate, and the time spent for expiration increases. Run SKIPDIRS=NO periodically to expire the directories and reduce the size of the database.

Node

Specifies the name of the client nodes or node groups whose data is to be processed. To specify multiple node and node group names, separate the names with commas and no intervening spaces. Node names can contain wildcard characters, but node group names cannot.

You can specify either **NODE** or **DOMAIN**. If you specify both, only those nodes that match the criteria for both specified command options are processed. If you do not specify either **NODE** or **DOMAIN** with a value, data for all nodes is processed.

Domain

Specifies that only data for client nodes assigned to the specified domain is to be processed. You can specify either **NODE** or **DOMAIN**. If you specify both, only those nodes that match the criteria for both specified command options are processed. If you do not specify either **NODE** or **DOMAIN** with a value, data for all nodes is processed.

Type

Specifies the type of data to be processed. The default value is **ALL**. Possible values are:

ALL

Process all types of data that are eligible for expiration

Archive

Process only client archive data

Backup

Process only client backup data

Other

Process only items for disaster recovery manager functions, such as recovery plan files and obsolete database backups

REsource

Specifies the number of threads that can run in parallel. Specify a number from one to 40. The default is 4.

Expiration runs as a single process, although the resources represent parallel work by the server within the single expiration process. Archive data for a node runs only on a single resource, but backup data can be spread across resources on a file space level. For example, if you specify NODE=X,Y,Z each with three file spaces and RESOURCE=5, then expiration processing for the three X, Y, and Z client nodes run in parallel. At least one resource processes each node, and at least one node uses multiple resources for processing backup data across the multiple file spaces.

EXPIRE INVENTORY

Duration

Specifies the maximum number of minutes for the expiration process to run. The process stops when the specified number of minutes pass or when all eligible expired objects are deleted, whichever comes first. You can specify a number from 1 to 999999. This parameter is optional. If this parameter is not specified, the duration of the expiration process is not limited by time.

Example: Run inventory expiration processing for a specific time period

Run the expiration process for two hours.

```
expire inventory duration=120
```

Example: Run inventory expiration processing for backup data for two client nodes

Run inventory expiration processing for the backup data for two client nodes, CHARLIE and ROBBIE. Allow the server to run expiration processing until completed.

```
expire inventory node=charlie,robbie resource=2 type=backup
```

Related commands

Table 166. Commands related to EXPIRE INVENTORY

Command	Description
AUDIT LICENSES	Verifies compliance with defined licenses.
CANCEL PROCESS	Cancels a background server process.
QUERY PROCESS	Displays information about background processes.

EXPORT commands

Use the **EXPORT** commands to copy information from a Tivoli Storage Manager server to sequential removable media.

Important: For commands that export administrators or nodes, you must consider the method of authentication. The Tivoli Storage Manager server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the **EXPORT** command, set the password by issuing the **UPDATE ADMIN** or **UPDATE NODE** command.

- “EXPORT ADMIN (Export administrator information)” on page 490
- “EXPORT NODE (Export client node information)” on page 497
- “EXPORT POLICY (Export policy information)” on page 518
- “EXPORT SERVER (Export server information)” on page 525

EXPORT ADMIN (Export administrator information)

Use this command to export administrator and authority definitions from a server. You can export the information to sequential media for later importing to another server, or you can export the information directly to another server for immediate import.

Important: For commands that export administrators or nodes, you must consider the method of authentication. The Tivoli Storage Manager server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the **EXPORT** command, set the password by issuing the **UPDATE ADMIN** or **UPDATE NODE** command.

Important: If target and source server levels are not compatible, the operation might not work. See the *Administrator's Guide* for server compatibility requirements.

Tivoli Storage Manager exports administrator information such as:

- Administrator name, password, and contact information
- Administrative privilege classes granted to the administrator
- Whether the administrator ID is locked from server access

You can use the **QUERY ACTLOG** command to view the status of the export operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the **CANCEL PROCESS** command. If you export information to sequential media and the background process is canceled, the sequential media holding the exported data are incomplete and should not be used for importing data. If a server-to-server export background process is canceled, a partial import might result. Evaluate any imported data on the target server to determine whether you want to keep or delete the imported data. Review the import messages for details. To display information about background processes, use the **QUERY PROCESS** command.

Limitation: The Tivoli Storage Manager server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue, after the import or node replication operation, update the fields with the appropriate **UPDATE** commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

The **EXPORT ADMIN** command takes two forms: Export directly to another server on the network, or export to sequential media. The syntax and parameters for each form are defined separately.

- “EXPORT ADMIN (Export administrator definitions to sequential media)” on page 492
- “EXPORT ADMIN (Export administrator information directly to another server)” on page 495

Table 167. Commands related to EXPORT ADMIN

Command	Description
CANCEL PROCESS	Cancels a background server process.
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT POLICY	Copies policy information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
IMPORT ADMIN	Restores administrative information from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY PROCESS	Displays information about background processes.

EXPORT ADMIN—to sequential media


EXPORT ADMIN (Export administrator definitions to sequential media)

You can export administrator and authority definitions from a server to sequential media for later importing to another server.

Privilege class

To issue this command, you must have system privilege.

Syntax



Notes:

- 1 If PREVIEW=NO, a device class must be specified.
- 2 If PREVIEW=NO and SCRATCH=NO, one or more volumes must be specified.

Parameters

admin_name

Specifies the administrators for which information is to be exported. This parameter is optional. The default is all administrators.

Separate the items in the list by commas, with no intervening spaces. You can use wildcard characters to specify names.

Preview

Specifies whether to preview the results of the export operation, without exporting information. You can use this parameter to preview how many bytes of data are transferred, allowing you to determine how many volumes will be required. This parameter is optional. The default value is NO. The values are:

No Specifies that the administrator information is to be exported. If you specify this value, you must specify a device class.

Yes

Specifies that the operation will be previewed but not performed. Information is reported to the server console and the activity log. If you specify this value, you do not need to specify a device class.

DEVclass

Specifies the device class to which export data is to be written. This parameter is required if you specify PREVIEW=NO.

You cannot specify the DISK device class or a device class with a device type of NAS.

If all drives for the device class are busy when the export runs, Tivoli Storage Manager cancels lower priority operations to make a drive available.

Tip: You can export data to a storage pool on another server by specifying a device class whose device type is SERVER. For details about storing data on another server, see the *Administrator's Guide*.

Scratch

Specifies whether scratch volumes can be used. The default value is YES. Possible values are:

Yes

Specifies that scratch volumes can be used for export. If you also specify a list of volumes, scratch volumes are used only if there is not enough space on the volumes specified.

No Specifies that scratch volumes cannot be used for export. To determine how many volumes you might need, you can run the command specifying PREVIEW=YES.

VOLumentnames

Specifies the volumes to be used to contain exported data. This parameter is optional, unless you specify SCRATCH=NO and PREVIEW=NO. If you do not specify a volume name, scratch volumes are used.

Possible values are:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE: *file_name*

Specifies the name of a file that contains a list of volumes. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when specifying volumes associated with the following device types:

For this device	Specify
Tape	1–6 alphanumeric characters.
FILE	Any fully qualified file name string. For example: /imdata/mt1.
REMOVABLEFILE	1–6 alphanumeric characters.
SERVER	1–250 alphanumeric characters.

EXPORT ADMIN—to sequential media

USEDVolumeList

Specifies the file where a list of volumes used in the export operation are stored. This parameter is optional.

This file can be used in the import operation. This file contains comment lines with the date and time the export was done, and the command issued to create the export.

Attention: If you specify an existing file, the file is overwritten.

ENCryptionstrength

Indicates which algorithm to use to encrypt passwords when exporting administrative and node records. This parameter is optional. The default value is AES. If you are exporting to a server that does not support AES, specify DES. Possible values are:

AES

Specifies the Advanced Encryption Standard.

DES

Specifies the Data Encryption Standard.

Example: Export administrator definitions to tape volumes

From the server, export the information for all defined administrators to tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device assigned to the MENU1 device class. The number and type of objects exported are reported to the system console and in the activity log. Issue the command:

```
export admin devclass=menu1
volumenames=tape01,tape02,tape03
```

Example: Export administrator definitions to tape volumes listed in a file

From the server, export the information for all defined administrators to tape volumes that are listed in the following file:

```
TAPEVOL
```

This file contains the following lines:

```
TAPE01
TAPE02
TAPE03
```

Specify that these tape volumes be used by a device assigned to the MENU1 device class. Issue the command:

```
export admin devclass=menu1 volumenames=file:tapevol
```

The number and type of objects exported are reported to the system console and in the activity log.

EXPORT ADMIN (Export administrator information directly to another server)


Use this command to export administrator and authority definitions directly to another server on the network. This results in an immediate import on the target server.

You can issue a **QUERY PROCESS** command from the target server to monitor the progress of the import operation.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

admin_name

Specifies the administrators for which information is to be exported. This parameter is optional. The default is all administrators.

Separate the items in the list by commas, with no intervening spaces. You can use wildcard characters to specify names.

TOserver

Specifies the name of a server to which the export data is sent directly over the network for immediate import.

Important: The target server must be defined on the originating server with the DEFINE SERVER command. The administrator that issues the export command must be defined with the same administrator name and password and have system authority on the target server.

When you specify TOSERVER, you cannot specify the DEVCLASS, VOLUMENAMES, and SCRATCH, USEDVOLUMELIST, and PREVIEW parameters.

PREVIEWImport

Specifies whether to view how much data is transferred, without actually moving any data. This information can be used to determine how much storage pool space is required on the target server. The default is NO.

Valid values are:

EXPORT ADMIN—directly to another server

Yes

Specifies that you want to preview the results of the import operation on the target server, without importing the data. Information is reported to the server console and the activity log.

No Specifies that you want the data to be imported on the target server without previewing the results.

Replacedefs

Specifies whether to replace definitions (not file data) on the server. The default is NO.

Valid values are:

Yes

Specifies that definitions are replaced on the server if definitions having the same name as those being imported exist on the target server.

No Specifies that imported definitions are skipped if their names conflict with definitions that are already defined on the target server.

ENCryptionstrength

Indicates which algorithm to use to encrypt passwords when exporting administrative and node records. This parameter is optional. The default value is AES. If you are exporting to a server that does not support AES, specify DES. Possible values are:

AES

Specifies the Advanced Encryption Standard.

DES

Specifies the Data Encryption Standard.

Example: Export administrator definitions to a target server

Export all the administrator definitions to the target server defined as OTHERSERVER. Preview the import operations on the target server. Issue the command:

```
export admin * toserver=otherserver previewimport=yes
```

From the target server, OTHERSERVER, you can view the import operations by issuing the command:

```
query process
```

EXPORT NODE (Export client node information)

Use this command to export client node definitions or file data to sequential media or directly to another server for immediate import.

Important: For commands that export administrators or nodes, you must consider the method of authentication. The Tivoli Storage Manager server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the **EXPORT** command, set the password by issuing the **UPDATE ADMIN** or **UPDATE NODE** command.

You can export data from a server with retention protection enabled, but it is not retention protected when imported on another server.

You cannot export nodes of type NAS; export processing excludes these nodes.

The following information is included in each client node definition:

- User ID, password, and contact information.
- Name of the client's assigned policy domain.
- File compression status.
- Whether the user has the authority to delete backed-up or archived files from server storage.
- Whether the client node ID is locked from server access.

Optionally, you can also export the following items:

- File space definitions.
- Backed-up, archived, and files that were migrated by a Tivoli Storage Manager for Space Management client.
- Access authorization information pertaining to the file spaces exported.
- Archive data that is in deletion hold status (the hold status is preserved). When the archive data is imported, it remains in deletion hold.

If you use an LDAP directory server to authenticate passwords, any servers that you export to must be configured for LDAP passwords. Node data that is exported from a node that authenticates with an LDAP directory server is inaccessible if the target server is not properly configured. If your target server is not configured, exported data from an LDAP node can still be exported. But the target server must be configured to use LDAP, to access the data.

Limitation: The Tivoli Storage Manager server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue, after the import or node replication operation, update the fields with the appropriate **UPDATE** commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

The **EXPORT NODE** command generates a background process that can be canceled with the **CANCEL PROCESS** command. If you are exporting node information

EXPORT NODE

to sequential media and the background process is canceled, the sequential media that is holding the exported data are incomplete and should not be used for importing data. If a server-to-server export background process is canceled, a partial import might result. Evaluate any imported data on the target server to determine whether you want to keep or delete the imported data. Review the import messages for details. To display information about background processes, issue the **QUERY PROCESS** command.

To display information about any running and suspended server-to-server export operations, issue the **QUERY EXPORT** command. The **QUERY EXPORT** command displays information only for exports that are, or can be, suspended. Export operations that can be suspended, and then restarted, are those server-to-server exports whose **FILEDATA** has a value other than **NONE**. You can issue the **QUERY ACTLOG** command to view the status of the export operation.

Restriction: If target and source server levels are not compatible, the operation might not work. See the *Administrator's Guide* for server compatibility requirements.

Because of unpredictable results, do not run expiration, migration, backup, or archive when issuing the **EXPORT NODE** command. See the *Administrator's Guide* for more information.

For a server that has clients with support for Unicode, you might need to have the server convert the file space name that you enter, or use one of the following parameters:

- **FSID**
- **UNIFILESPACE**

The **EXPORT NODE** command takes two forms: export directly to another server on the network, or export to sequential media. The syntax and parameters for each form are defined separately.

- “EXPORT NODE (Export node definitions or file data directly to another server)” on page 509
- “EXPORT NODE (Export node definitions to sequential media)” on page 500

Table 168. Commands related to EXPORT NODE

Command	Description
CANCEL EXPORT	Deletes a suspended export operation.
CANCEL PROCESS	Cancels a background server process.
COPY ACTIVATEDATA	Copies active backup data.
EXPORT ADMIN	Copies administrative information to external media or directly to another server.
EXPORT POLICY	Copies policy information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
IMPORT NODE	Restores client node information from external media.
QUERY ACTLOG	Displays messages from the server activity log.

Table 168. Commands related to EXPORT NODE (continued)

Command	Description
QUERY EXPORT	Displays the export operations that are currently running or suspended.
QUERY PROCESS	Displays information about background processes.
RESTART EXPORT	Restarts a suspended export operation.
SUSPEND EXPORT	Suspends a running export operation.

EXPORT NODE—to sequential media


EXPORT NODE (Export node definitions to sequential media)


You can export node definitions or file data from a server to sequential media for later importing to another server.

Privilege class

To issue this command, you must have system privilege.

Syntax



**Notes:**

- 1 If PREVIEW=NO, a device class must be specified.
- 2 If PREVIEW=NO and SCRATCH=NO, one or more volumes must be specified.

Parameters*node_name*

Specifies the client node names for which information is to be exported. This parameter is optional. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. For each node entered, all file spaces in the file space, FSID, and Unicode enabled lists will be searched.

Restriction: If you use wildcard characters to specify a pattern for node names, the server will not report the node names or patterns that do not match any entries in the database. Check the summary statistics in the activity log to verify that the server exported all intended nodes.

FILESpace

Specifies the file spaces for which data is to be exported. This parameter is optional. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify a name.

FSID

Specifies the file spaces by using their file space IDs (FSIDs). The server uses the FSIDs to find the file spaces to export. To find the FSID for a file space, use the **QUERY FILESPACE** command. Separate multiple file space IDs with commas and no intervening spaces. This parameter is optional.

UNIFILESpace

Specifies the file spaces that are known to the server to be Unicode enabled. The server converts the names you enter from the server code page to the UTF-8 code page to find the file spaces to export. The success of the conversion depends on the actual characters in the name and the server's code page. Separate multiple names with commas and no intervening spaces. This parameter is optional.

D0mains

Specifies the policy domains from which nodes should be exported. This parameter is optional. Separate multiple names with commas and no intervening spaces. If you specify domains, a node is exported only if it belongs to one of the specified domains. You can use wildcard characters to specify a name.

EXPORT NODE—to sequential media

FILEData

Specifies the type of files that should be exported for all nodes being exported to the server. This parameter is optional. The default value is NONE.

Note: If you are exporting a node that has group data, data that is not a part of the target objects might be exported. An example of group data is virtual machine data or system state backup data. For example, if FILEDATA=BACKUPACTIVE when the FROMDATE or TODATE parameters are specified, it is possible to include inactive backup data. The incremental backup processing for the data can cause extra files that do not meet the filtering criteria to be exported.

If you are exporting to sequential media: the device class used by the file data is determined by the device class for the storage pool. If it is the same device class specified in this command, two drives are needed to export node information. The mount limit for the device class must be at least 2.

Important: If client nodes registered as TYPE=SERVER are being exported, specify ALL, ARCHIVE, or ALLACTIVE.

The following descriptions mention *active* and *inactive* backup file versions. An active backup file version is the most recent backup version for a file that still exists on the client workstation. All other backup versions are called inactive copies. The values are:

ALl

The server exports all backup versions of files, all archived files, and all files that were migrated by a Tivoli Storage Manager for Space Management client.

None

The server does not export files, only node definitions.

ARchive

The server exports only archived files.

Backup

The server exports only backup versions, whether active or inactive.

BACKUPActive

The server exports only active backup versions. These active backup versions are the active versions in the Tivoli Storage Manager database at the time that the **EXPORT** command is issued.

ALLActive

The server exports all active backup versions of files, all archived files, and all files that were migrated by a Tivoli Storage Manager for Space Management client. The active backup versions are the active versions in the Tivoli Storage Manager database at the time that the **EXPORT** command is issued.

SPacemanaged

The server exports only files that were migrated by a Tivoli Storage Manager for Space Management client.

Preview

Specifies whether to preview the results of the export operation, without exporting information. You can use this parameter to preview how many bytes of data would be transferred, allowing you to determine how many volumes will be required. This parameter is optional. The default value is NO. The values are:

No Specifies that the node information is to be exported. If you specify this value, you must also specify a device class.

Yes

Specifies that the operation will be previewed but not performed. Information is reported to the server console and the activity log. If you specify this value, you do not need to specify a device class.

DEVclass

Specifies the device class to which export data is to be written. This parameter is required if you specify PREVIEW=NO.

You cannot specify the DISK device class or a device class with a device type of NAS.

If all drives for the device class are busy when the export runs, Tivoli Storage Manager cancels lower priority operations to make a drive available.

Tip: You can export data to a storage pool on another server by specifying a device class whose device type is SERVER. For details about storing data on another server, see the *Administrator's Guide*.

Scratch

Specifies whether scratch volumes can be used. The default value is YES.

Possible values are:

Yes

Specifies that scratch volumes can be used for export. If you also specify a list of volumes, scratch volumes are used only if there is not enough space on the volumes specified.

No Specifies that scratch volumes cannot be used for export. To determine how many volumes you might need, you can run the command specifying PREVIEW=YES.

VOLumentnames

Specifies the volumes to be used to contain exported data. This parameter is optional, unless you specify SCRATCH=NO and PREVIEW=NO. If you do not specify a volume name, scratch volumes are used.

Possible values are:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE: *file_name*

Specifies the name of a file that contains a list of volumes. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when specifying volumes associated with the following device types:

For this device	Specify
Tape	1-6 alphanumeric characters.
FILE	Any fully qualified file name string. For example: /imdata/mt1.
REMOVABLEFILE	1-6 alphanumeric characters.

EXPORT NODE—to sequential media

For this device	Specify
SERVER	1–250 alphanumeric characters.

USEDVolumelist

Specifies the file where a list of volumes used in the export operation are stored. This parameter is optional.

This file can be used in the import operation. This file contains comment lines with the date and time the export was done, and the command issued to create the export.

Attention: If you specify an existing file, the file is overwritten.

FROMDate

Specifies the earliest date for which files to be exported were stored on the server. Files that were stored on the server earlier than the specified date are not exported. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. Tivoli Storage Manager ignores the FROMDATE parameter when the FILEDATA parameter is set to NONE.

Directory processing: The FROMDATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Important: If group data is on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can be exported. Group data on the node is, for example, virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported.

Use one of the following values to specify the date:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or <i>-days</i>	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
<i>EOLM-days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
<i>BOTM+days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

If this parameter is not specified, Tivoli Storage Manager exports all objects stored before the TODATE parameter and as qualified by the FILEDATA parameter. If no TODATE parameter is specified, then all data as qualified by the FILEDATA parameter is exported.

When a server-to-server export operation uses a relative FROMDATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the FROMDATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

TODate

Specifies the latest date for files to be exported from the server. Files stored on the server on a date later than the TODATE value are not exported. TODATE only applies to client file data and does not affect other information that is being exported, such as policy.

- Tivoli Storage Manager ignores the TODATE parameter when the FILEDATA parameter is set to NONE.
- If a TODATE parameter is specified without a TOTIME parameter, the server exports all objects inserted on or before the day specified by the TODATE parameter.
- If you specified the FROMDATE parameter, the value of TODATE must be later than or equal to the FROMDATE value. If the TODATE and FROMDATE are equal, then the TOTIME parameter must be later than the FROMTIME parameter.
- The TODATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Use one of the following values to specify the date:

Value	Description	Example
MM/DD/YYYY	A specific date	10/15/2006
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

EXPORT NODE—to sequential media

When a server-to-server export operation uses a relative TODATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the TODATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

FROMTime

Specifies the earliest time for which objects to be exported were stored on the server. When you specify FROMTIME, you must also use the FROMDATE parameter. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. Objects that were stored on the server before the specified time and date are not exported. Tivoli Storage Manager ignores the FROMTIME parameter when the FILEDATA parameter is set to NONE.

Important: If group data is on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can be exported. Group data on the node is, for example, virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported.

The default value for this parameter when used with the FROMDATE parameter is midnight (00:00:00).

Use one of the following values to specify the time:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified. The FROMTIME+ can only be used with a FROMDATE before today.	NOW+02:00 or +02:00. If you issue this command at 5:00 with FROMTIME=NOW+02:00 or FROMTIME=+02:00, the export operation only contains files that were put on the server after 7:00 on the FROMDATE that you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified	NOW-02:00 or -02:00. If you issue this command at 5:00 with FROMTIME=NOW-02:00 or FROMTIME=-2:00, the export includes files that were put on the server after 3:00.

TOTime

Specifies the latest time that objects to be exported were stored on the server. You must specify the TODATE parameter in order to use the TOTIME parameter. TOTIME only applies to client file data and does not affect other information that is being exported, such as policy. Tivoli Storage Manager ignores the TOTIME parameter if the FILEDATA parameter is set to NONE.

The default value for this parameter, when used with the TODATE parameter, is midnight minus one second (23:59:59).

Important: The value of the TOTIME and TODATE parameters must be later than the FROMDATE and the FROMTIME value.

Use one of the following values to specify the time:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08
<i>NOW+HH:MM</i> <i>or+HH:MM</i>	The current time plus hours and minutes specified.	NOW+02:00 or +02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW+02:00, the export includes files that were stored from 01:00 until 07:00.
<i>NOW-HH:MM</i> <i>or-HH:MM</i>	The current time minus hours and minutes specified.	NOW-02:00 or -02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW-02:00, the export includes files that were stored from 01:00 until 03:00.

ENCryptionstrength

Indicates which algorithm to use to encrypt passwords when exporting administrative and node records. This parameter is optional. The default value is AES. If you are exporting to a server that does not support AES, specify DES. Possible values are:

AES

Specifies the Advanced Encryption Standard.

DES

Specifies the Data Encryption Standard.

ALLOWSHREddable

Specifies whether data from a storage pool that enforces shredding is exported. This parameter is optional. The default value is NO. Possible values are:

No Specifies that data is not exported from a storage pool that enforces shredding.

Yes

Specifies that data can be exported from a storage pool that enforces shredding. The data on the export media will not be shredded.

Example: Export client node information to specific tape volumes

From the server, export client node information to tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be used by a device assigned to the MENU1 device class.

```
export node devclass=menu1 volumenames=tape01,tape02,tape03
```

EXPORT NODE—to sequential media

Example: Export client node information using the FSID

From the server, use the FSID to export active backup versions of file data for client node JOE to tape volume TAPE01. To determine the FSID, first issue a **QUERY FILESPACE** command.

1. To determine the FSID, issue a **QUERY FILESPACE** command.

```
query filesystem joe
```

Node Name	Filespace Name	FSID	Platform	Filespace Type	Is Filespace Unicode?	Capacity (MB)	Pct Util
JOE	\\joe\c\$	1	WinNT	NTFS	Yes	2,502.3	75.2
JOE	\\joe\d\$	2	WinNT	NTFS	Yes	6,173.4	59.6

2. Export the active backup versions of file data and specify that the tape volume be used by a device assigned to the MENU1 device class.

```
export node joe fsid=1,2 filedata=backupactive devclass=menu1  
volumenames=tape01
```

Example: Export client node information to tape volumes listed in a file

From the server, export client node information to tape volumes that are listed in the following file:

```
TAPEVOL
```

The file contains the following lines:

```
TAPE01  
TAPE02  
TAPE03
```

Specify that the tape volumes be used by a device assigned to the MENU1 device class. Issue the following command:

```
export node devclass=menu1 volumenames=file:tapevol
```

EXPORT NODE (Export node definitions or file data directly to another server)

Use this command to export client node definitions or file data directly to another server for immediate import.

Important: You cannot export nodes of type NAS. Export processing will exclude these nodes.

You can suspend and subsequently restart a server-to-server export operation that has a FILEDATA value other than NONE. The server saves the state and status of the export operation so that it may be restarted from the point at which the operation failed or was suspended. The export operation can be restarted at a later date by issuing the **RESTART EXPORT** command.

Important: An export operation is suspended when any of the following conditions are detected:

- A **SUSPEND EXPORT** command is issued for the running export operation
- Segment preemption - the file being read for export is deleted by some other process
- Communication errors on a server-to-server export
- No available mount points
- Necessary volumes are unavailable
- I/O errors encountered

Issue the **QUERY EXPORT** command to display information on any running and suspended export operations.


The export operation cannot be restarted if the export operation fails prior to transmitting the eligible node and file space definitions to the target server. You must reenter the command to begin a new export operation.

You can issue a **QUERY PROCESS** command from the target server to monitor the progress of the import operation. Issue the **QUERY EXPORT** command to list all restartable server-to-server export operations.


Privilege class

To issue this command, you must have system privilege.

Syntax



EXPORT NODE— directly to another server



Parameters

node_name

Specifies the client node names for which information is to be exported. This parameter is optional. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. For each node entered, all file spaces in the file space, FSID, and Unicode enabled lists will be searched.

Restriction: If you specify a list of node names or node patterns, the server will not report the node names or node patterns that do not match any of the entries in the database. Check the summary statistics in the activity log to verify that the server exported all intended nodes.

FILESpace

Specifies the file spaces for which data is to be exported. This parameter is optional. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify a name.

FSID

Specifies the file spaces by using their file space IDs (FSIDs). The server uses the FSIDs to find the file spaces to export. To find the FSID for a file space, use the **QUERY FILESPACE** command. Separate multiple file space IDs with commas and no intervening spaces. This parameter is optional.

UNIFILESpace

Specifies the file spaces that are known to the server to be Unicode enabled. The server converts the names you enter from the server code page to the UTF-8 code page to find the file spaces to export. The success of the conversion depends on the actual characters in the name and the server's code page. Separate multiple names with commas and no intervening spaces. This parameter is optional.

DOmains

Specifies the policy domains from which nodes should be exported. This parameter is optional. Separate multiple names with commas and no intervening spaces. If you specify domains, Tivoli Storage Manager exports a node only if it belongs to one of the specified domains. You can use wildcard characters to specify a name.

FILEData

Specifies the type of files to export for all nodes. This parameter is optional. The default value is NONE.

Note: If you are exporting a node that has group data, data that is not a part of the target objects might be exported. An example of group data is virtual machine data or system state backup data. For example, if `FILEDATA=BACKUPACTIVE` when the `FROMDATE` or `TODATE` parameters are specified, it is possible to include inactive backup data. The incremental backup processing for the data can cause extra files that do not meet the filtering criteria to be exported.

If you are exporting to sequential media, the device class used by the file data is determined by the device class for the storage pool. If it is the same device class specified in this command, Tivoli Storage Manager requires two drives to export node information. The mount limit for the device class must be at least 2.

Important: If you export client nodes that are registered as `TYPE=SERVER`, specify `ALL`, `ARCHIVE`, or `ALLACTIVE`.

The following descriptions mention *active* and *inactive* backup file versions. An active backup file version is the most recent backup version for a file that still exists on the client workstation. All other backup file versions are called inactive copies. The values are:

ALl

The server exports all backup versions of files, all archived files, and all files migrated by a Tivoli Storage Manager for Space Management client.

None

The server does not export files, only node definitions.

ARchive

The server exports only archived files.

Backup

The server exports only backup versions, whether they are active or inactive.

EXPORT NODE— directly to another server

BACKUPActive

The server exports only active backup versions. These active backup versions are the active versions in the Tivoli Storage Manager database at the time that the **EXPORT** command is issued.

ALLActive

The server exports all active backup versions of files, all archived files, and all files that were migrated by a Tivoli Storage Manager for Space Management client. The active backup versions are the active versions in the Tivoli Storage Manager database at the time that the **EXPORT** command is issued.

Spacemanaged

The server exports only files that were migrated by a Tivoli Storage Manager for Space Management client.

FROMDate

Specifies the earliest date for which files to be exported were stored on the server. Files that were stored on the server earlier than the specified date are not exported. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. Tivoli Storage Manager ignores the FROMDATE parameter when the FILEDATA parameter is set to NONE.

Directory processing: The FROMDATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Important: If group data is on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can be exported. Group data on the node is, for example, virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported.

Use one of the following values to specify the date:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or <i>-days</i>	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
<i>EOLM-days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM

Value	Description	Example
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

If this parameter is not specified, Tivoli Storage Manager exports all objects stored before the TODATE parameter and as qualified by the FILEDATA parameter. If no TODATE parameter is specified, then all data as qualified by the FILEDATA parameter is exported.

When a server-to-server export operation uses a relative FROMDATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the FROMDATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

TODate

Specifies the latest date for files to be exported from the server. Files stored on the server on a date later than the TODATE value are not exported. TODATE only applies to client file data and does not affect other information that is being exported, such as policy.

- Tivoli Storage Manager ignores the TODATE parameter when the FILEDATA parameter is set to NONE.
- If a TODATE parameter is specified without a TOTIME parameter, the server exports all objects inserted on or before the day specified by the TODATE parameter.
- If you specified the FROMDATE parameter, the value of TODATE must be later than or equal to the FROMDATE value. If the TODATE and FROMDATE are equal, then the TOTIME parameter must be later than the FROMTIME parameter.
- The TODATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Use one of the following values to specify the date:

Value	Description	Example
MM/DD/YYYY	A specific date	10/15/2006
TODAY	The current date	TODAY
TODAY- <i>days</i> or - <i>days</i>	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.

EXPORT NODE— directly to another server

Value	Description	Example
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

When a server-to-server export operation uses a relative TODATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the TODATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

FROMTime

Specifies the earliest time for which objects to be exported were stored on the server. When you specify FROMTIME, you must also use the FROMDATE parameter. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. Objects that were stored on the server before the specified time and date are not exported. Tivoli Storage Manager ignores the FROMTIME parameter when the FILEDATA parameter is set to NONE.

Important: If group data is on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can be exported. Group data on the node is, for example, virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported.

The default value for this parameter when used with the FROMDATE parameter is midnight (00:00:00).

Use one of the following values to specify the time:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified. The FROMTIME+ can only be used with a FROMDATE before today.	NOW+02:00 or +02:00. If you issue this command at 5:00 with FROMTIME=NOW+02:00 or FROMTIME+=+02:00, the export operation only contains files that were put on the server after 7:00 on the FROMDATE that you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified	NOW-02:00 or -02:00. If you issue this command at 5:00 with FROMTIME=NOW-02:00 or FROMTIME=-2:00, the export includes files that were put on the server after 3:00.

TOTime

Specifies the latest time that objects to be exported were stored on the server. You must specify the TODATE parameter in order to use the TOTIME parameter. TOTIME only applies to client file data and does not affect other information that is being exported, such as policy. Tivoli Storage Manager ignores the TOTIME parameter if the FILEDATA parameter is set to NONE.

The default value for this parameter, when used with the TODATE parameter, is midnight minus one second (23:59:59).

Important: The value of the TOTIME and TODATE parameters must be later than the FROMDATE and the FROMTIME value.

Use one of the following values to specify the time:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08
<i>NOW+HH:MM</i> <i>or+HH:MM</i>	The current time plus hours and minutes specified.	NOW+02:00 or +02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW+02:00, the export includes files that were stored from 01:00 until 07:00.
<i>NOW-HH:MM</i> <i>or-HH:MM</i>	The current time minus hours and minutes specified.	NOW-02:00 or -02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW-02:00, the export includes files that were stored from 01:00 until 03:00.

TOServer

Specifies the name of a server to which the export data is sent directly over the network for immediate import.

Important: The target server must be defined on the originating server with the DEFINE SERVER command. The administrator that issues the export command must be defined with the same administrator name and password and have system authority on the target server.

When you specify TOSERVER, you cannot specify the DEVCLASS, VOLUMENAMES, and SCRATCH, USEDVOLUMELIST, and PREVIEW parameters.

PREVIEWImport

Specifies whether to view how much data is transferred, without actually moving any data. This information can be used to determine how much storage pool space is required on the target server. The default is NO.

Valid values are:

Yes

Specifies that you want to preview the results of the import operation on the target server, without importing the data. Information is reported to the server console and the activity log.

EXPORT NODE— directly to another server

No Specifies that you want the data to be imported on the target server without previewing the results.

MERGEfilespaces

Specifies whether Tivoli Storage Manager merges client files into existing file spaces on the target server (if they exist), or if Tivoli Storage Manager generates new file space names. The default is NO.

Valid values are:

Yes

Specifies that imported data on the target server is merged with the existing file space, if a file space with the same name exists on the target server.

No Specifies that Tivoli Storage Manager generates a new file space name for imported data on the target server if file spaces with the same name exists.

Replacedefs

Specifies whether to replace definitions (not file data) on the server. The default is NO.

Valid values are:

Yes

Specifies that definitions are replaced on the server if definitions having the same name as those being imported exist on the target server.

No Specifies that imported definitions are skipped if their names conflict with definitions that are already defined on the target server.

PROXynodeassoc

Specifies if proxy node associations are exported. This parameter is optional. The default value is NO.

ENCryptionstrength

Indicates which algorithm to use to encrypt passwords when exporting administrative and node records. This parameter is optional. The default value is AES. If you are exporting to a server that does not support AES, specify DES. Possible values are:

AES

Specifies the Advanced Encryption Standard.

DES

Specifies the Data Encryption Standard.

ALLOWSHREDdable

Specifies whether data from a storage pool that enforces shredding is exported. This parameter is optional. The default value is NO. Possible values are:

No Specifies that the server does not export data from a storage pool that enforces shredding.

Yes

Specifies that the server does export from a storage pool that enforces shredding. The data on the export media will not be shredded.

Restriction: After an export operation finishes identifying files for export, any changes to the storage pool **ALLOWSHREDABLE** value is ignored. An export operation that is suspended retains the original **ALLOWSHREDABLE** value throughout the operation. You might want to consider cancelling your export

EXPORT NODE— directly to another server

operation if changes to the storage pool **ALLOWSHREDABLE** value jeopardize the operation. You can reissue the export command after any needed cleanup.

EXPORTIDENTIFIER

This optional parameter specifies the name that you select to identify this export operation. If you do not specify an identifier name, the server generates one for you. The export identifier name cannot be more than 64 characters, cannot contain wildcard characters, and is not case sensitive. You can use the identifier name to reference export operations in the **QUERY EXPORT**, **SUSPEND EXPORT**, **RESTART EXPORT**, or **CANCEL EXPORT** commands.

Restriction: You must specify the **TOSERVER** parameter if you are specifying the **EXPORTIDENTIFIER** parameter.
EXPORTIDENTIFIER is ignored if **FILEDATA=NONE**.

Example: Export client node information and all client files

To export client node information and all client files for NODE1 directly to SERVERB, issue the following command:

```
export node node1 filedata=all toserver=serverb
```

Example: Export client node information and all client files for a specific date range

To export client node information and all client files for NODE1 directly to SERVERB between February 1, 2009 and today.

```
export node node1 filedata=all toserver=serverb  
fromdate=02/01/2009 todate=today
```

Example: Export client node information and all client files for a specific date and time range

To export client node information and all client files for NODE1 directly to SERVERB from 8:00 AM on February 1, 2009 until today at 8:00 AM, issue the following command:

```
export node node1 filedata=all toserver=serverb  
fromdate=02/01/2009 fromtime=08:00:00  
todate=today totime=08:00:00
```

Example: Export client node information and all client files for the past three days

To export client node information and all client files for NODE1 directly to SERVERB for the past three days, issue the following command:

```
export node node1 filedata=all toserver=serverb  
fromdate=today -3
```

EXPORT POLICY (Export policy information)

Use this command to export policy information from a Tivoli Storage Manager server to sequential media or directly to another server for immediate import.

Important: If target and source server levels are not compatible, the export operation might not work. See the *Administrator's Guide* for server compatibility requirements.

The server exports policy information, such as:

- Policy domain definitions
- Policy set definitions, including the active policy set
- Management class definitions, including the default management class
- Backup copy group and archive copy group definitions
- Schedule definitions for each policy domain
- Client node associations, if the client node exists on the target server

Restriction: When a policy is exported by using the **EXPORT POLICY** command, the active data pool information in the domain is not exported.

You can use the **QUERY ACTLOG** command to view the status of the export operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the **CANCEL PROCESS** command. If you export policy information to sequential media and the background process is canceled, the sequential media holding the exported data are incomplete and should not be used for importing data. If a server-to-server export background process is canceled, a partial import might result. Evaluate any imported data on the target server to determine whether you want to keep or delete the imported data. Review the import messages for details. To display information about background processes, use the **QUERY PROCESS** command.

Limitation: The Tivoli Storage Manager server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue, after the import or node replication operation, update the fields with the appropriate **UPDATE** commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

The **EXPORT POLICY** command takes two forms: Export directly to another server on the network, or export to sequential media. The syntax and parameters for each form are defined separately.

- “EXPORT POLICY (Export a policy directly to another server)” on page 523
- “EXPORT POLICY (Export policy information to sequential media)” on page 520

Table 169. Commands related to EXPORT POLICY

Command	Description
CANCEL PROCESS	Cancels a background server process.
EXPORT ADMIN	Copies administrative information to external media or directly to another server.

Table 169. Commands related to EXPORT POLICY (continued)

Command	Description
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
IMPORT POLICY	Restores policy information from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY PROCESS	Displays information about background processes.

EXPORT POLICY— to sequential media


EXPORT POLICY (Export policy information to sequential media)

Use this command to export policy information from an IBM Tivoli Storage Manager server to sequential media for later import to another server.

Privilege class

To issue this command, you must have system privilege.

Syntax



Notes:

- 1 If PREVIEW=NO, a device class must be specified.
- 2 If PREVIEW=NO and SCRATCH=NO, one or more volumes must be specified.

Parameters

domain_name

Specifies the policy domains for which information is to be exported. This parameter is optional. The default is all policy domains. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names.

Preview

Specifies whether to preview the results of the export operation, without exporting information. You can use this parameter to preview how many bytes of data are transferred, allowing you to determine how many volumes will be required. This parameter is optional. The default value is NO. The values are:

No Specifies that the policy information is to be exported. If you specify this value, you must also specify a device class.

Yes

Specifies that the operation will be previewed but not performed.

Information is reported to the server console and the activity log. If you specify this value, you do not need to specify a device class.

DEVclass

Specifies the device class to which export data is to be written. This parameter is required if you specify PREVIEW=NO.

You cannot specify the DISK device class or a device class with a device type of NAS.

If all drives for the device class are busy when the export runs, Tivoli Storage Manager cancels lower priority operations to make a drive available.

Tip: You can export data to a storage pool on another server by specifying a device class whose device type is SERVER. For details about storing data on another server, see the *Administrator's Guide*.

Scratch

Specifies whether scratch volumes can be used. The default value is YES.

Possible values are:

Yes

Specifies that scratch volumes can be used for export. If you also specify a list of volumes, scratch volumes are used only if there is not enough space on the volumes specified.

No Specifies that scratch volumes cannot be used for export. To determine how many volumes you might need, you can run the command specifying PREVIEW=YES.

VOLumentnames

Specifies the volumes to be used to contain exported data. This parameter is optional, unless you specify SCRATCH=NO and PREVIEW=NO. If you do not specify a volume name, scratch volumes are used.

Possible values are:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE:*file_name*

Specifies the name of a file that contains a list of volumes. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when specifying volumes associated with the following device types:

For this device	Specify
Tape	1–6 alphanumeric characters.
FILE	Any fully qualified file name string. For example: /imdata/mt1.
REMOVABLEFILE	1–6 alphanumeric characters.
SERVER	1–250 alphanumeric characters.

USEDVolumeList

Specifies the file where a list of volumes used in the export operation are stored. This parameter is optional.

EXPORT POLICY— to sequential media

This file can be used in the import operation. This file contains comment lines with the date and time the export was done, and the command issued to create the export.

Attention: If you specify an existing file, the file is overwritten.

Example: Export policy information to specific tape volumes

From the server, export policy information to tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device assigned to the MENU1 device class.

```
export policy devclass=menu1  
volumenames=tape01,tape02,tape03
```

Example: Export policy information to tape volumes listed in a file

From the server, export policy information to tape volumes that are listed in the following file:

```
TAPEVOL
```

This file contains the following lines:

```
TAPE01  
TAPE02  
TAPE03
```

Specify that these tape volumes be used by a device assigned to the MENU1 device class. Issue the following command:

```
export policy devclass=menu1 volumenames=file:tapevol
```


EXPORT POLICY (Export a policy directly to another server)

Use this command to export policy information directly to another server on the network. This results in an immediate import on the target server.

To monitor the progress of the import operation, you can issue a **QUERY PROCESS** command from the target server.

Privilege class

To issue this command, you must have system privilege.

Syntax**Parameters***domain_name*

Specifies the policy domains for which information is to be exported. This parameter is optional. The default is all policy domains. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names.

TOServer

Specifies the name of a server to which the export data is sent directly over the network for immediate import.

Important: The target server must be defined on the originating server with the **DEFINE SERVER** command. The administrator that issues the export command must be defined with the same administrator name and password and have system authority on the target server.

When you specify **TOSERVER**, you cannot specify the **DEVCLASS**, **VOLUMENAMES**, and **SCRATCH**, **USEDVOLUMELIST**, and **PREVIEW** parameters.

PREVIEWImport

Specifies whether to view how much data is transferred, without actually moving any data. This information can be used to determine how much storage pool space is required on the target server. The default is **NO**.

Valid values are:

Yes

Specifies that you want to preview the results of the import operation on the target server, without importing the data. Information is reported to the server console and the activity log.

No Specifies that you want the data to be imported on the target server without previewing the results.

EXPORT POLICY—directly to another server

Replacedefs

Specifies whether to replace definitions (not file data) on the server. The default is NO.

Valid values are:

Yes

Specifies that definitions are replaced on the server if definitions having the same name as those being imported exist on the target server.

No Specifies that imported definitions are skipped if their names conflict with definitions that are already defined on the target server.

Example: Export policy to another server

To export policy information directly to SERVERB, issue the following command:
`export policy replacedefs=yes toserver=othersrv`

EXPORT SERVER (Export server information)

Use this command to export all or part of the server control information and client file data (if specified) from the server to sequential media.

When you export server information to sequential media, you can later use the media to import the information to another server with a compatible device type.

Important: For commands that import administrators or nodes, you must consider the method of authentication. The Tivoli Storage Manager server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the **IMPORT** command, set the password by issuing the **UPDATE ADMIN** or **UPDATE NODE** command.

Other important information:

- If target and source server levels are not compatible, the operation might not work.

For information about server compatibility requirements, see the *Administrator's Guide*.

- You cannot use a CENTERA device class as the target medium for an export command, or as the source medium for an import command.

You also have the option of processing an export operation directly to another server on the network. This results in an immediate import process without the need for compatible sequential device types between the two servers.

You can export the following types of server information by issuing the **EXPORT SERVER** command:

- Policy domain definitions
- Policy set definitions
- Management class and copy group definitions
- Schedules defined for each policy domain
- Administrator definitions
- Client node definitions

You can optionally export the following types of data:

- File space definitions
- Access authorization information that pertains to the file spaces exported
- Backed-up, archived, and files that were migrated by a Tivoli Storage Manager for Space Management client

You cannot export nodes of type NAS. Export processing excludes these nodes.

This command generates a background process that can be canceled by the **CANCEL PROCESS** command. If you export server information to sequential media, and the background process is canceled, the sequential media holding the exported data are incomplete and should not be used for importing data. If a server-to-server export background process is canceled, a partial import might result. Evaluate any imported data on the target server to determine whether you want to keep or delete the imported data. Review the import messages for details.

EXPORT SERVER

Issue the **QUERY PROCESS** command from the target server to monitor the progress of the import operation. Issue the **QUERY EXPORT** command to list all server-to-server export operations (that have a **FILEDATA** value other than **NONE**) that are running or suspended.

You can use the **QUERY ACTLOG** command to view the actual status information which indicates the size and the success or failure of the export operation.

Limitation: The Tivoli Storage Manager server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue, after the import or node replication operation, update the fields with the appropriate **UPDATE** commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

The **EXPORT SERVER** command takes two forms: Export directly to another server on the network, or export to sequential media. The syntax and parameters for each form are defined separately.

- “EXPORT SERVER (Export a server to sequential media)” on page 527
- “EXPORT SERVER (Export server control information and client file data to another server)” on page 535

Table 170. Commands related to EXPORT SERVER

Command	Description
CANCEL EXPORT	Deletes a suspended export operation.
CANCEL PROCESS	Cancels a background server process.
COPY ACTIVEDATA	Copies active backup data.
EXPORT ADMIN	Copies administrative information to external media or directly to another server.
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT POLICY	Copies policy information to external media or directly to another server.
IMPORT SERVER	Restores all or part of the server from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY EXPORT	Displays the export operations that are currently running or suspended.
QUERY PROCESS	Displays information about background processes.
RESTART EXPORT	Restarts a suspended export operation.
SUSPEND EXPORT	Suspends a running export operation.


EXPORT SERVER (Export a server to sequential media)

You can export all or part of the server control information and client file data from a server to sequential media for later importing to another server.

Privilege class

To issue this command, you must have system privilege.

Syntax



Notes:

- 1 If PREVIEW=NO, a device class must be specified.
- 2 If PREVIEW=NO and SCRATCH=NO, one or more volumes must be specified.

EXPORT SERVER— to sequential media

Parameters

FILEData

Specifies the type of files that should be exported for all nodes defined to the server. This parameter is optional. The default value is NONE.

If you are exporting to sequential media: the device class to access the file data is determined by the device class for the storage pool. If it is the same device class specified in this command, two drives are needed to export server information. The mount limit for the device class must be set to at least 2.

The following descriptions mention *active* and *inactive* backup file versions. An active backup file version is the most recent backup version for a file that still exists on the client workstation. All other backup file versions are called inactive copies. The following values are available:

ALl

Tivoli Storage Manager exports all backup versions of files, all archived files, and all files that were migrated by a Tivoli Storage Manager for Space Management client.

None

Tivoli Storage Manager does not export files, only definitions.

ARchive

Tivoli Storage Manager exports only archived files.

Backup

Tivoli Storage Manager exports only backup versions, whether the versions are active or inactive.

BACKUPActive

Tivoli Storage Manager exports only active backup versions.

ALLActive

Tivoli Storage Manager exports all active backup versions of files, all archived files, and all files that were migrated by a Tivoli Storage Manager for Space Management client.

SPacemanaged

Tivoli Storage Manager exports only files that were migrated by a Tivoli Storage Manager for Space Management client.

Preview

Specifies whether you want to preview the results of the export operation, without exporting information. You can use this parameter to preview how many bytes of data are transferred, so that you can determine how many volumes are required. This parameter is optional. The default value is NO. The values are:

No Specifies that the server information is to be exported. If you specify this value, you must also specify a device class.

Yes

Specifies that the operation will be previewed but not performed. Information is reported to the server console and the activity log. If you specify this value, you do not need to specify a device class.

DEVclass

Specifies the device class to which export data is to be written. This parameter is required if you specify PREVIEW=NO.

You cannot specify the DISK device class or a device class with a device type of NAS.

If all drives for the device class are busy when the export runs, Tivoli Storage Manager cancels lower priority operations to make a drive available.

Tip: You can export data to a storage pool on another server by specifying a device class whose device type is SERVER. For details about storing data on another server, see the *Administrator's Guide*.

Scratch

Specifies whether scratch volumes can be used. The default value is YES.

Possible values are:

Yes

Specifies that scratch volumes can be used for export. If you also specify a list of volumes, scratch volumes are used only if there is not enough space on the volumes specified.

No Specifies that scratch volumes cannot be used for export. To determine how many volumes you might need, you can run the command specifying PREVIEW=YES.

VOLumenames

Specifies the volumes to be used to contain exported data. This parameter is optional, unless you specify SCRATCH=NO and PREVIEW=NO. If you do not specify a volume name, scratch volumes are used.

Possible values are:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE:*file_name*

Specifies the name of a file that contains a list of volumes. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when specifying volumes associated with the following device types:

For this device	Specify
Tape	1–6 alphanumeric characters.
FILE	Any fully qualified file name string. For example: /imdata/mt1.
REMOVABLEFILE	1–6 alphanumeric characters.
SERVER	1–250 alphanumeric characters.

USEDVolumeList

Specifies the file where a list of volumes used in the export operation are stored. This parameter is optional.

This file can be used in the import operation. This file contains comment lines with the date and time the export was done, and the command issued to create the export.

Attention: If you specify an existing file, the file is overwritten.

EXPORT SERVER— to sequential media

FROMDate

Specifies the earliest date for which files to be exported were stored on the server. Files that were stored on the server earlier than the specified date are not exported. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. Tivoli Storage Manager ignores the FROMDATE parameter when the FILEDATA parameter is set to NONE.

Directory processing: The FROMDATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Important: If group data is on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can be exported. Group data on the node is, for example, virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported.

Use one of the following values to specify the date:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or <i>-days</i>	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
<i>EOLM-days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
<i>BOTM+days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

If this parameter is not specified, Tivoli Storage Manager exports all objects stored before the TODATE parameter and as qualified by the FILEDATA parameter. If no TODATE parameter is specified, then all data as qualified by the FILEDATA parameter is exported.

When a server-to-server export operation uses a relative FROMDATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the FROMDATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still

07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

TODate

Specifies the latest date for files to be exported from the server. Files stored on the server on a date later than the TODATE value are not exported. TODATE only applies to client file data and does not affect other information that is being exported, such as policy.

- Tivoli Storage Manager ignores the TODATE parameter when the FILEDATA parameter is set to NONE.
- If a TODATE parameter is specified without a TOTIME parameter, the server exports all objects inserted on or before the day specified by the TODATE parameter.
- If you specified the FROMDATE parameter, the value of TODATE must be later than or equal to the FROMDATE value. If the TODATE and FROMDATE are equal, then the TOTIME parameter must be later than the FROMTIME parameter.
- The TODATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Use one of the following values to specify the date:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	10/15/2006
TODAY	The current date	TODAY
<i>TODAY-days or -days</i>	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
<i>EOLM-days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
<i>BOTM+days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

When a server-to-server export operation uses a relative TODATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the TODATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

EXPORT SERVER— to sequential media

FROMTime

Specifies the earliest time for which objects to be exported were stored on the server. When you specify FROMTIME, you must also use the FROMDATE parameter. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. Objects that were stored on the server before the specified time and date are not exported. Tivoli Storage Manager ignores the FROMTIME parameter when the FILEDATA parameter is set to NONE.

Important: If group data is on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can be exported. Group data on the node is, for example, virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported.

The default value for this parameter when used with the FROMDATE parameter is midnight (00:00:00).

Use one of the following values to specify the time:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified. The FROMTIME+ can only be used with a FROMDATE before today.	NOW+02:00 or +02:00. If you issue this command at 5:00 with FROMTIME=NOW+02:00 or FROMTIME=+02:00, the export operation only contains files that were put on the server after 7:00 on the FROMDATE that you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified	NOW-02:00 or -02:00. If you issue this command at 5:00 with FROMTIME=NOW-02:00 or FROMTIME=-2:00, the export includes files that were put on the server after 3:00.

TOTime

Specifies the latest time that objects to be exported were stored on the server. You must specify the TODATE parameter in order to use the TOTIME parameter. TOTIME only applies to client file data and does not affect other information that is being exported, such as policy. Tivoli Storage Manager ignores the TOTIME parameter if the FILEDATA parameter is set to NONE.

The default value for this parameter, when used with the TODATE parameter, is midnight minus one second (23:59:59).

Important: The value of the TOTIME and TODATE parameters must be later than the FROMDATE and the FROMTIME value.

Use one of the following values to specify the time:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08

Value	Description	Example
NOW+HH:MM or+HH:MM	The current time plus hours and minutes specified.	NOW+02:00 or +02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW+02:00, the export includes files that were stored from 01:00 until 07:00.
NOW-HH:MM or-HH:MM	The current time minus hours and minutes specified.	NOW-02:00 or -02:00. If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW-02:00, the export includes files that were stored from 01:00 until 03:00.

ENCryptionstrength

Indicates which algorithm to use to encrypt passwords when exporting administrative and node records. This parameter is optional. The default value is AES. If you are exporting to a server that does not support AES, specify DES. Possible values are:

AES

Specifies the Advanced Encryption Standard.

DES

Specifies the Data Encryption Standard.

ALLOWSHREddable

Specifies whether data from a storage pool that enforces shredding is exported. This parameter is optional. The default value is NO. Possible values are:

No Specifies that data is not exported from a storage pool that enforces shredding.

Yes

Specifies that data can be exported from a storage pool that enforces shredding. The data on the export media will not be shredded.

Example: Export a server to specific tape volumes

From the server, export server information to tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device assigned to the MENU1 device class.

```
export server devclass=menu1
volumenames=tape01,tape02,tape03
```

Example: Export a server to tape volumes listed in a file

From the server, export server information to tape volumes that are listed in the following file:

```
TAPEVOL
```

The file contains the following lines:

```
TAPE01
TAPE02
TAPE03
```

EXPORT SERVER— to sequential media

Specify that the tape volumes be used by a device assigned to the MENU1 device class. Issue the following command:

```
export server devclass=menu1 volumenames=file:tapevol
```

EXPORT SERVER (Export server control information and client file data to another server)

Use this command to export all or part of the server control information and client file data directly to another server on the network. This results in an immediate import on the target server.

Server-to-server export operations that have a FILEDATA value other than NONE can be restarted after the operation is suspended. The server saves the state and status of the export operation so that it may be restarted from the point at which the operation failed or was suspended. The export operation can be restarted at a later date by issuing the **RESTART EXPORT** command. These export operations can be manually suspended as well as restarted. Therefore, if an export fails, it is automatically suspended if it has completed the transmitting definitions phase.

An export operation is suspended when any of the following conditions is detected:

- A **SUSPEND EXPORT** command is issued for the running export operation
- Segment preemption - the file being read for export is deleted by some other process
- Communication errors on a server-to-server export
- No available mount points
- Necessary volumes are unavailable
- I/O errors encountered


The export operation cannot be restarted if the export operation fails prior to transmitting the eligible node and filespace definitions to the target server. You must reenter the command to begin a new export operation.

Issue the **QUERY PROCESS** command from the target server to monitor the progress of the import operation. Issue the **QUERY EXPORT** command to list all server-to-server export operations (that have a FILEDATA value other than NONE) that are running or suspended.


Privilege class

To issue this command, you must have system privilege.

Syntax



EXPORT SERVER— directly to another server



Parameters

FILEData

Specifies the type of files to export for all nodes defined to the server. This parameter is optional. The default value is NONE.

If you are exporting to sequential media: The device class to access the file data is determined by the device class for the storage pool. If it is the same device class specified in this command, Tivoli Storage Manager requires two drives to export server information. You must set the mount limit for the device class to at least 2.

The following descriptions mention active and inactive backup file versions. An active backup file version is the most recent backup version for a file that still exists on the client workstation. All other backup file versions are called inactive copies. The values are:

ALl

Tivoli Storage Manager exports all backup versions of files, all archived files, and all files that were migrated by a Tivoli Storage Manager for Space Management client.

None

Tivoli Storage Manager does not export files, only definitions.

ARchive

Tivoli Storage Manager exports only archived files.

Backup

Tivoli Storage Manager exports only backup versions, whether they are active or inactive.

BACKUPActive

Tivoli Storage Manager exports only active backup versions.

ALLActive

Tivoli Storage Manager exports all active backup versions of files, all archived files, and all files that were migrated by a Tivoli Storage Manager for Space Management client.

SPacemanaged

Tivoli Storage Manager exports only files that were migrated by a Tivoli Storage Manager for Space Management client.

FROMDate

Specifies the earliest date for which files to be exported were stored on the server. Files that were stored on the server earlier than the specified date are not exported. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. Tivoli Storage Manager ignores the FROMDATE parameter when the FILEDATA parameter is set to NONE.

Directory processing: The FROMDATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Important: If group data is on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can be exported. Group data on the node is, for example, virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported.

Use one of the following values to specify the date:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or <i>-days</i>	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
<i>EOLM-days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM

EXPORT SERVER— directly to another server

Value	Description	Example
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

If this parameter is not specified, Tivoli Storage Manager exports all objects stored before the TODATE parameter and as qualified by the FILEDATA parameter. If no TODATE parameter is specified, then all data as qualified by the FILEDATA parameter is exported.

When a server-to-server export operation uses a relative FROMDATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the FROMDATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

TODate

Specifies the latest date for files to be exported from the server. Files stored on the server on a date later than the TODATE value are not exported. TODATE only applies to client file data and does not affect other information that is being exported, such as policy.

- Tivoli Storage Manager ignores the TODATE parameter when the FILEDATA parameter is set to NONE.
- If a TODATE parameter is specified without a TOTIME parameter, the server exports all objects inserted on or before the day specified by the TODATE parameter.
- If you specified the FROMDATE parameter, the value of TODATE must be later than or equal to the FROMDATE value. If the TODATE and FROMDATE are equal, then the TOTIME parameter must be later than the FROMTIME parameter.
- The TODATE parameter does not apply to directories. All directories in a file space are processed even if the directories were not backed up in the specified date range.

Use one of the following values to specify the date:

Value	Description	Example
MM/DD/YYYY	A specific date	10/15/2006
TODAY	The current date	TODAY
TODAY- <i>days</i> or - <i>days</i>	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.

Value	Description	Example
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

When a server-to-server export operation uses a relative TODATE, for example, TODAY-1, and the operation is restarted at a later date, the restarted process still uses the date that was used during the original operation. For example, if a server-to-server export operation is started on 07/04/2009 and the TODATE is specified as TODAY-1, the date that is used for selecting files is 07/03/2009. If this same export operation is suspended and restarted ten days later (07/14/2009), the date that is used for selecting files is still 07/03/2009. This behavior ensures that the entire export operation uses the same cutoff date for selecting files to export.

FROMTime

Specifies the earliest time for which objects to be exported were stored on the server. When you specify FROMTIME, you must also use the FROMDATE parameter. This parameter only applies to client file data. This parameter does not affect other information that might be exported, for example, policy. Objects that were stored on the server before the specified time and date are not exported. Tivoli Storage Manager ignores the FROMTIME parameter when the FILEDATA parameter is set to NONE.

Important: If group data is on the node that you are exporting, data that was backed up before the designated FROMDATE and FROMTIME can be exported. Group data on the node is, for example, virtual machine data or system state backup data. This export is a result of incremental backup processing for the data. The incremental backup processing can cause extra files that do not meet the filtering criteria to be exported.

The default value for this parameter when used with the FROMDATE parameter is midnight (00:00:00).

Use one of the following values to specify the time:

Value	Description	Example
HH:MM:SS	A specific time	10:30:08
NOW	The current time	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes specified. The FROMTIME+ can only be used with a FROMDATE before today.	NOW+02:00 or +02:00. If you issue this command at 5:00 with FROMTIME=NOW+02:00 or FROMTIME=+02:00, the export operation only contains files that were put on the server after 7:00 on the FROMDATE that you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes specified	NOW-02:00 or -02:00. If you issue this command at 5:00 with FROMTIME=NOW-02:00 or FROMTIME=-2:00, the export includes files that were put on the server after 3:00.

EXPORT SERVER— directly to another server

TOTime

Specifies the latest time that objects to be exported were stored on the server. You must specify the TODATE parameter in order to use the TOTIME parameter. TOTIME only applies to client file data and does not affect other information that is being exported, such as policy. Tivoli Storage Manager ignores the TOTIME parameter if the FILEDATA parameter is set to NONE.

The default value for this parameter, when used with the TODATE parameter, is midnight minus one second (23:59:59).

Important: The value of the TOTIME and TODATE parameters must be later than the FROMDATE and the FROMTIME value.

Use one of the following values to specify the time:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08
<i>NOW+HH:MM</i> <i>or+HH:MM</i>	The current time plus hours and minutes specified.	<i>NOW+02:00 or +02:00.</i> If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW+02:00, the export includes files that were stored from 01:00 until 07:00.
<i>NOW-HH:MM</i> <i>or-HH:MM</i>	The current time minus hours and minutes specified.	<i>NOW-02:00 or -02:00.</i> If you issue this command at 05:00 with FROMTIME=01:00 and TOTIME=NOW-02:00, the export includes files that were stored from 01:00 until 03:00.

TOServer

Specifies the name of a server to which the export data is sent directly over the network for immediate import.

Important: The target server must be defined on the originating server with the DEFINE SERVER command. The administrator that issues the export command must be defined with the same administrator name and password and have system authority on the target server.

When you specify TOSERVER, you cannot specify the DEVCLASS, VOLUMENAMES, and SCRATCH, USEDVOLUMELIST, and PREVIEW parameters.

PREVIEWImport

Specifies whether to view how much data is transferred, without actually moving any data. This information can be used to determine how much storage pool space is required on the target server. The default is NO.

Valid values are:

Yes

Specifies that you want to preview the results of the import operation on the target server, without importing the data. Information is reported to the server console and the activity log.

EXPORT SERVER— directly to another server

No Specifies that you want the data to be imported on the target server without previewing the results.

MERGEfilespaces

Specifies whether Tivoli Storage Manager merges client files into existing file spaces on the target server (if they exist), or if Tivoli Storage Manager generates new file space names. The default is NO.

Valid values are:

Yes

Specifies that imported data on the target server is merged with the existing file space, if a file space with the same name exists on the target server.

No Specifies that Tivoli Storage Manager generates a new file space name for imported data on the target server if file spaces with the same name exists.

Replacedefs

Specifies whether to replace definitions (not file data) on the server. The default is NO.

Valid values are:

Yes

Specifies that definitions are replaced on the server if definitions having the same name as those being imported exist on the target server.

No Specifies that imported definitions are skipped if their names conflict with definitions that are already defined on the target server.

PROXynodeassoc

Specifies if proxy node associations are exported. This parameter is optional. The default value is NO.

ENCryptionstrength

Indicates which algorithm to use to encrypt passwords when exporting administrative and node records. This parameter is optional. The default value is AES. If you are exporting to a server that does not support AES, specify DES. Possible values are:

AES

Specifies the Advanced Encryption Standard.

DES

Specifies the Data Encryption Standard.

ALLOWSHREDdable

Specifies whether data from a storage pool that enforces shredding is exported. This parameter is optional. The default value is NO. Possible values are:

No Specifies that the server does not allow data to be exported from a storage pool that enforces shredding.

Yes

Specifies that the server allows data to be exported from a storage pool that enforces shredding. The data on the export media will not be shredded.

Important: After an export operation finishes identifying files for export, any changes to the storage pool ALLOWSHREDABLE value is ignored. An export operation that is suspended retains the original ALLOWSHREDABLE value throughout the operation. You might want to consider cancelling your export

EXPORT SERVER— directly to another server

operation if changes to the storage pool ALLOWSHREDABLE value jeopardize the operation. You can reissue the export command after any needed cleanup.

EXPORTIdentifier

This optional parameter specifies the name that you selected to identify this export operation. If you do not specify a command name, the server generates one for you. The export identifier name cannot be more than 64 characters, cannot contain wildcard characters, and is not case sensitive. You can use the identifier name to reference export operations in the **QUERY EXPORT**, **SUSPEND EXPORT**, **RESTART EXPORT**, or **CANCEL EXPORT** commands. EXPORTIDENTIFIER is ignored if FILEDATA=NONE or if PREVIEWIMPORT=YES.

If you are specifying the EXPORTIDENTIFIER parameter, you must specify the TOSERVER parameter.

Example: Export server information directly to another server

To export server information directly to SERVERB, issue the following command.

```
export server filedata=all toserver=serverb
```

Example: Export server information directly to another server using a date range

To export directly to SERVERB between February 1, 2009 and today, issue the following command.

```
export server filedata=all toserver=serverb  
fromdate=02/01/2009 todate=today
```

Example: Export server information and client file data directly to another server using a date and time range

To export directly to SERVERB from 8:00 a.m. on February 1, 2009 until today at 8:00 a.m., issue the following command.

```
export server filedata=all toserver=serverb  
fromdate=02/01/2009 fromtime=08:00:00  
todate=today totime=08:00:00
```

EXTEND DBSPACE (Increase space for the database)

Use this command to increase space for the database by adding directories for the database to use.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

After a new directory is added to a Tivoli Storage Manager server, the directory might not be fully used immediately. Some Tivoli Storage Manager events can cause the added database space to become increasingly used over time. For example, table reorganizations or some temporary database transactions, such as long running select statements, can cause the added database space to begin filling up. The database space redistribution among all directories can require a few days or even weeks. If the existing database directories are nearly full when the new directory is added, the server might encounter an out of space condition (reported in the `db2diag.log`).

If this condition occurs, halt and restart the server. If the restart does not correct the condition, remove the database and then restore it to the same or new directories. See the *Administrator's Guide* for details.

Syntax

```

▶▶ EXTEND DBSPACE db_directory

```

Parameters

db_directory (Required)

Specifies the directories for database storage. The directories must be empty and accessible by the user ID of the database manager. A directory name must be a fully qualified name. The maximum length of the name is 175 characters. Enclose the name in quotation marks if it contains embedded blanks, an equal sign, or other special characters.

Example: Add directories to the storage space for the database

Add two directories (`/tsm_db/stg1` and `tsm_db/stg2`) under the `/tsm_db` directory to the storage space for the database. Issue the command:

```
extend dbspace /tsm_db/stg1,/tsm_db/stg2
```

Related commands

Table 171. Commands related to EXTEND DBSPACE

Command	Description
QUERY DB	Displays allocation information about the database.
QUERY DBSPACE	Displays information about the storage space defined for the database.

GENERATE commands

Use the **GENERATE** commands for backup sets for a selected filesystem or client node.

- “GENERATE BACKUPSET (Generate a backup set of Backup-Archive Client data)” on page 545
- “GENERATE BACKUPSETTOC (Generate a table of contents for a backup set)” on page 553

GENERATE BACKUPSET (Generate a backup set of Backup-Archive Client data)

Use this command to generate a backup set for a Backup-Archive Client node. A *backup set* is a collection of a Backup-Archive Client's active backed up data, stored and managed as a single object, on specific media, in server storage. Although you can create a backup set for any client node, a backup set can be used only by a Backup-Archive Client.

Restriction: A backup set in “deduplication format” has that designation as a result of a **GENERATE BACKUPSET** command with at least one of the following specifications:

- Includes a node at Backup-Archive Client level 6.1.x (at least 6.1.0 but less than 6.2.0).
- Includes a node that has one or more nodes that are authorized to act as a proxy. At least one of those proxy nodes is at Backup-Archive Client level 6.1.x.

Backup sets in the deduplication format can be restored only by the 6.1.2 or later Backup-Archive Client. Backup-Archive Clients before version 6.1.2 cannot restore from a backup set that is in the deduplication format.

A backup set in the “distributed deduplication format” has that designation as a result of a **GENERATE BACKUPSET** command with at least one of the following specifications:

- Includes a node at Backup-Archive Client level 6.2.0 or later.
- Includes a node that has one or more nodes that are authorized to act as a proxy. At least one of those proxy nodes is at Backup-Archive Client level 6.2.0.

Backup sets in the distributed deduplication format can be restored only by the 6.2.0 or later Backup-Archive Client.

Restriction: You cannot generate a backup set with files that were backed up to Tivoli Storage Manager using NDMP. However, you can create a backup set with files that were backed up using NetApp SnapShot Difference.

The server creates copies of active versions of a client's backed up objects that are within the one-or-more file spaces specified with this command. The server then consolidates them onto sequential media. Currently, the backup object types supported for backup sets include directories and files only.

The backup-archive client node can restore its backup set from the server and from the media to which the backup set was written.

This command generates a background process that can be canceled with the **CANCEL PROCESS** command. If the background process created by this command is canceled, the media might not contain a complete backup set. You can use the **QUERY PROCESS** command to show information about the background process that is created by this command.


Tip: When Tivoli Storage Manager generates a backup set, you can improve performance if the primary storage pools containing the client data are collocated. If a primary storage pool is collocated, client node data is likely to be on fewer tape volumes than it would be if the storage pool were not collocated. With collocation, less time is spent searching database entries, and fewer mount operations are required.


GENERATE BACKUPSET

Privilege class

To issue this command, you must have system privilege or policy privilege for the domain to which the client node is assigned.

Syntax





Parameters

node_name or *node_group_name* (Required)

Specifies the name of the client node and node groups whose data is contained in the backup set. To specify multiple node names and node group names, separate the names with commas and no intervening spaces. You can use wildcard characters with node names but not with node group names. When multiple node names are specified, the server generates a backup set for each node and places all of the backup sets together on a single set of output volumes.

backup_set_name_prefix (Required)

Specifies the name of the backup set for the client node. The maximum length of the name is 30 characters.

When you select a name, Tivoli Storage Manager adds a suffix to construct your backup set name. For example, if you name your backup set *mybackupset*, Tivoli Storage Manager adds a unique number such as 3099 to the name. The backup set name is then identified to Tivoli Storage Manager as *mybackupset.3099*. To later show information about this backup set, you can include a wildcard with the name, such as *mybackupset.** or specify the fully qualified name, such as *mybackupset.3099*.

When multiple node names or node group names are specified, the server generates a backup set for each node or node group and places all the backup sets on a single set of output volumes. Each backup set is given the same fully qualified name consisting of the *backup_set_name_prefix* and a suffix determined by the server.

file_space_name

Specifies the names of one or more file spaces that contain the data to be included in the backup set. This parameter is optional. The file space name you specify can contain wildcard characters. You can specify more than one file space by separating the names with commas and no intervening spaces. If you do not specify a file space, data from all the client nodes backed-up and active file spaces is included in the backup set.

For a server that has clients with support for Unicode-enabled file spaces, you can enter *either* a file space name or a file space ID (FSID). If you enter a file space name, you may need to have the server convert the file space name that you enter. For example, you might need to have the server convert the name you enter from the server's code page to Unicode. See the NAMETYPE parameter for details. If you do not specify a file space name, or specify only a single wildcard character for the name, you can use the CODETYPE parameter to limit the operation to Unicode file spaces or to non-Unicode file spaces.

GENERATE BACKUPSET

DEVclass (Required)

Specifies the name of the device class for the volumes to which the backup set is written. The maximum length of the name is 30 characters.

Restriction: You cannot specify a device class with a device type of NAS.

SCRatch

Specifies whether to use scratch volumes for the backup set. If you include a list of volumes using the VOLUMES parameter, the server uses scratch volumes only if the data cannot be contained in the volumes you specify. The default is SCRATCH=YES. The values are:

YES

Specifies to use scratch volumes for the backup set.

NO Specifies not to use scratch volumes for the backup set.

VOLumes

Specifies the names of one or more volumes that will contain the backup set. This parameter is optional. You can specify more than one volume by separating each volume with a comma, with no intervening spaces.

If you do not specify this parameter, scratch volumes are used for the backup set.

RETention

Specifies the number of days to retain the backup set on the server. You can specify an integer from 0 to 30000. The default is 365 days. The values are:

days

Specifies the number of days to retain the backup set on the server.

NOLimit

Specifies that the backup set should be retained on the server indefinitely.

If you specify **NOLIMIT**, the server retains the volumes containing the backup set forever, unless a user or administrator deletes the volumes from server storage.

DESCription

Specifies the description to associate with the backup set. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is NO. The values are:

Yes

Specifies the command processes in the foreground. Messages created are not displayed until the command completes processing. You cannot specify WAIT=YES from the server console.

No Specifies that the command processes in the background. Use the **QUERY PROCESS** command to monitor the background processing of this command.

NAMEType

Specify how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with support for Unicode-enabled file spaces. You can use this parameter for Tivoli Storage Manager clients using Windows, NetWare, or Macintosh OS X operating systems.

Use this parameter only when you enter a partly or fully qualified file space name. The default value is SERVER. Possible values are:

SERVER

The server uses the server's code page to interpret the file space names.

UNICODE

The server converts the file space name entered from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server has a problem accessing system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

Important: Use care when specifying this parameter if multiple node names are also specified. Different nodes may use the same file space ID for different file spaces, or different file space IDs for the same file space name. Therefore, specifying a file space ID as the file space names can result in the wrong data being written to the backup set for some nodes.

CODETYPE

Specify what type of file spaces are to be included in the operation. The default is BOTH, meaning that file spaces are included regardless of code page type. Use this parameter only when you enter a single wildcard character for the file space name or when you do not specify any file space names. Possible values are:

UNICODE

Include only file spaces that are in Unicode.

NONUNICODE

Include only file spaces that are not in Unicode.

BOTH

Include file spaces regardless of code page type.

PITDATE

Specifies that files that were active on the specified date and that are still stored on the Tivoli Storage Manager server are to be included in the backup set, even if they are inactive at the time you issue the command. This parameter is optional. The default is the date on which the **GENERATE BACKUPSET** command is run. You can specify the date using one of the following values:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998
TODAY	The current date	TODAY
<i>TODAY-days</i> or <i>days</i>	The current date minus days specified	TODAY-7 or -7. To include files that were active a week ago, specify PITDATE=TODAY-7 or PITDATE=-7
EOLM (End Of Last Month)	The last day of the previous month.	EOLM

GENERATE BACKUPSET

Value	Description	Example
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

PITTime

Specifies that files that were active on the specified time and that are still stored on the Tivoli Storage Manager server are to be included in the backup set, even if they are inactive at the time you issue the command. This parameter is optional. If a PITDate was specified, the default is midnight (00:00:00); otherwise the default is the time at which the **GENERATE BACKUPSET** command is started. You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified PIT date	12:33:28
NOW	The current date on the specified PIT date	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified PIT date	NOW+03:00 or +03:00 If you issue this command at 9:00 with PITTIME=NOW+03:00 or PITTIME=+03:00. Tivoli Storage Manager includes files that were active at 12:00 on the PIT date.

DATAType

Specifies that backup sets containing the specified types of data that are to be generated. This parameter is optional. The default is that file level backup sets are to be generated. To specify multiple data types, separate data types with commas and no intervening spaces.

The server generates a backup set for each data type and places all the backup sets on a single set of output volumes. Each backup set is given the same fully qualified name consisting of the *backup_set_name_prefix* and a suffix determined by the server. However, each backup set will have a different data type, as shown by the **QUERY BACKUPSET** command. Possible values are:

ALL

Specifies that backup sets for all types of data (file level, image, and application) that have been backed up on the server are to be generated.

FILE

Specifies that a file level backup set is to be generated. File level backup sets contain files and directories backed up by the backup client. If no files or directories have been backed up by the backup client, a file level backup set is not generated. This is the default.

IMAGE

Specifies that an image backup set is to be generated. Image backup sets contain images created by the backup client **BACKUP IMAGE** command. Image backup sets are generated only if an image has been backed up by the backup client.

TOC

Specifies whether a table of contents (TOC) is saved for each file level backup set. Tables of contents are always saved for backup sets containing image or application data. The TOC parameter is ignored when generating image and application backup sets. A table of contents will always be generated for image and application backup sets.

Consider the following in determining whether you want to save a table of contents:

- If a table of contents is saved for a backup set, you can use the Tivoli Storage Manager Web backup-archive client to examine the entire file system tree and choose files and directories to restore. To create a table of contents, you must define the TOCDESTINATION attribute in the backup copy group for the management class that is specified by the **TOCMGMTCLASS** parameter. Creating a table of contents requires additional processing, storage pool space, and possibly a mount point during the backup set operation.
- If a table of contents is not saved for a backup set, you can still restore individual files or directory trees using the backup-archive client **RESTORE BACKUPSET** command, if you know the fully qualified name of each file or directory to be restored.

To display the contents of backup sets, you can also use the **QUERY BACKUPSETCONTENTS** command.

This parameter is optional. Possible values are:

No Specifies that table of contents information is not saved for file level backup sets.

Preferred

Specifies that table of contents information should be saved for file level backup sets. This is the default. However, a backup set does not fail just because an error occurs during creation of the table of contents.

Yes

Specifies that table of contents information must be saved for each file level backup set. A backup set fails if an error occurs during creation of the table of contents.

TOCMgmtclass

Specifies the name of the management class to which the table of contents should be bound. If you do not specify a management class, the table of contents is bound to the default management class for the policy domain to which the node is assigned. In this case, creation of a table of contents requires that you define the TOCDESTINATION attribute in the backup copy group for the specified management class.

ALLOWSHREddable

Specifies whether data from a storage pool that enforces shredding is included in the backup set. This parameter is optional. Possible values are:

No Specifies that data from a storage pool that enforces shredding is not included in the backup set. This is the default.

GENERATE BACKUPSET

Yes

Specifies that data from a storage pool that enforces shredding can be included in the backup set. The data on the backup set media will not be shredded.

Example: Generate a backup set for a file space

Generate a backup set of a file space called /srvr that belongs to client node JANE. Name the backup set PERS_DATA and retain it for 75 days. Specify that volumes VOL1 and VOL2 contain the data for the backup set. The volumes are to be read by a device that is assigned to the AGADM device class. Include a description.

```
generate backupset jane pers_data /srvr devclass=agadm
retention=75 volumes=vol1,vol2
description="area 51 base image"
```

Example: Generate a backup set of a Unicode-enabled file space

Generate a backup set of the Unicode-enabled file space, \\joe\c\$, that belongs to client node JOE. Name the backup set JOES_DATA. Specify that volume VOL1 contain the data for the backup set. The volume is to be read by a device that is assigned to the AGADM device class. Have the server convert the \\joe\c\$ file space name from the server code page to the UTF-8 code page.

```
generate backupset joe joes_data \\joe\c$ devclass=agadm
volumes=vol1 nametype=unicode
```

Related commands

Table 172. Commands related to GENERATE BACKUPSET

Command	Description
COPY ACTIVATEDATA	Copies active backup data.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
QUERY BACKUPSET	Displays backup sets.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
QUERY NODEGROUP	Displays information about node groups.
QUERY BACKUPSETCONTENTS	Displays contents contained in backup sets.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.
UPDATE NODEGROUP	Updates the description of a node group.

GENERATE BACKUPSETTOC (Generate a table of contents for a backup set)


Use this command to generate a table of contents for a backup set that does not already have one. The backup-archive client uses the table of contents to display the backup set, which allows users to select individual files to be restored from the backup set.

Creating a table of contents for a backup set requires storage pool space and possibly one or more mount points during the creation operation.

Privilege class

To issue this command, you must have system privilege or policy privilege for the domain to which the client node is assigned.

Syntax



Parameters

node_name (Required)

Specifies the name of the client node whose data is contained in the backup set. You cannot use wildcard characters to specify a name, nor can you specify a list of client node names.

backup_set_name (Required)

Specifies the name of the backup set for the client node. You cannot use wildcard characters to specify a name, nor can you specify a list of backup set names.

DATATYPE

Specifies the type of data to be included in the table of contents. This parameter is optional. By default, all data is included. To specify multiple data types, separate the data types with commas and no intervening spaces.

Possible values are:

ALL

Specifies that the table of contents includes all types of data (file-level, image, and application) stored in the backup set. This is the default.

FILE

Specifies that the table of contents includes only file-level data. File-level data consists of files and directories backed up by the backup-archive client. If the backup set contains no files or directories, the table of contents is not generated.

IMAGE

Specifies that the table of contents will include only image backups. Image

GENERATE BACKUPSETTOC

backups consist of file system images created by the backup client **BACKUP IMAGE** command. If the backup set contains no image backups, the table of contents will not be generated.

TOCMgmtclass

Specifies the name of the management class to which the table of contents should be bound. If you do not specify a management class, the table of contents is bound to the default management class for the policy domain to which the node is assigned. If you create a table of contents you must define the TOCDESTINATION attribute in the backup copy group for the specified management class.

Example: Generate a table of contents

Generate a table of contents for a backup set named PROJX_DATA that contains the data for client node GARY. The table of contents is to be bound to the default management class.

```
generate backupsettoc gary projx_data
```

Related commands

Table 173. Commands related to GENERATE BACKUPSETTOC

Command	Description
COPY ACTIVE DATA	Copies active backup data.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays backup sets.
QUERY NODEGROUP	Displays information about node groups.
QUERY BACKUPSETCONTENTS	Displays contents contained in backup sets.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.
UPDATE NODEGROUP	Updates the description of a node group.

GRANT commands

Use the **GRANT** command to grant appropriate privileges or access.

- “GRANT AUTHORITY (Add administrator authority)” on page 556
- “GRANT PROXYNODE (Grant proxy authority to a client node)” on page 560

GRANT AUTHORITY

GRANT AUTHORITY (Add administrator authority)


Use this command to grant an administrator one or more administrative privilege classes, and authority to access client nodes.

You cannot grant restricted privilege to an unrestricted policy or unrestricted storage administrator. You must use the **REVOKE AUTHORITY** command to remove the administrator's unrestricted privilege, then use this command to grant restricted privilege to the administrator.


Privilege class

To issue this command, you must have system privilege.

Syntax



A:



Notes:

- 1 You must specify one or more of these parameters.

Parameters

admin_name (Required)

Specifies the name of the administrator being granted an administrative privilege class.

Classes

Specifies one or more privilege classes to grant to an administrator. This parameter is required, except when you specify the `STGPOLS` parameter. You can specify more than one privilege class by separating each with a comma. Possible classes are:

System

Specifies that you want to grant system privilege to an administrator. A system administrator has the highest level of authority in Tivoli Storage Manager. A system administrator can issue any administrative command

and has authority to manage all policy domains and all storage pools. Do not specify additional privilege classes or the `DOMAINS` or `STGPOOLS` parameters when granting system privilege to an administrator. Only a system administrator can grant authority to other administrators.

Policy

Specifies that you want to grant policy privilege to an administrator. If you do not specify the `DOMAINS` parameter, unrestricted policy privilege is granted. An unrestricted policy administrator can issue commands that affect all existing policy domains as well as any policy domains that are defined in the future. An unrestricted policy administrator cannot define, delete, or copy policy domains. Use the **GRANT AUTHORITY** command with `CLASSES=POLICY` and no `DOMAINS` parameter to upgrade a restricted policy administrator to an unrestricted policy administrator.

Storage

Specifies that you want to grant storage privilege to an administrator. If the `STGPOOLS` parameter is not specified, unrestricted storage privilege is granted. An unrestricted storage administrator can issue all commands that allocate and control storage resources for the server. An unrestricted storage administrator can issue commands that affect all existing storage pools as well as any storage pools that are defined in the future. An unrestricted storage administrator cannot define or delete storage pools. Using the **GRANT AUTHORITY** command with `CLASSES=STORAGE` and no `STGPOOLS` parameter upgrades a restricted storage administrator to an unrestricted storage administrator.

Operator

Specifies that you want to grant operator privilege to an administrator. An administrator with operator privilege can issue commands that control the immediate operation of the server and the availability of storage media.

Node

Specifies that you want to grant a node privilege to a user. A user with client node privilege can remotely access a web backup-archive client with an administrative user ID and password if they have been given owner authority or access authority. Access authority is the default for a node privilege class.

Attention: When you specify the node privilege class, you must also specify either the `DOMAIN` parameter or the `NODE` parameter, but not both.

AUTHORITY

Specifies the authority level of a user with node privilege. This parameter is optional.

If an administrator already has system or policy privilege to the policy domain to which the node belongs, this command will not change the administrator's privilege. Possible authority levels are:

Access

Specifies that you want to grant client access authority to a user with the node privilege class. This is the default when `CLASSES=NODE` is specified. A user with client access authority can access a web backup-archive client and perform backup and restore actions on that client.

GRANT AUTHORITY

Attention: A user with client access authority cannot access that client from another system by using the `-NODENAME` or `-VIRTUALNODENAME` parameter.

A client node can set the `REVOKEREMOTEACCESS` option to restrict a user that has node privilege with client access authority from accessing a client workstation that is running a web client. This option does not apply to administrators with client owner authority, system privilege, or policy privilege to the policy domain to which the node belongs.

Owner

Specifies that you want to grant client owner authority to a user with the node privilege class. A user with client owner authority can access a web backup-archive client through the web client interface and also access their data from another client using the `-NODENAME` or `-VIRTUALNODENAME` parameter.

DOmains

Specifies that you want to grant to the administrator client access or client owner authority to all clients in the specified policy domain. You cannot use this parameter together with the `NODE` parameter.

NOde

Specifies that you want to grant the administrator client access or client owner authority to the node. You cannot use this parameter together with the `DOMAIN` parameter.

DOmains

When used with `CLASSES=POLICY`, specifies that you want to grant restricted policy privilege to an administrator.

Restricted policy privilege permits an administrator to issue a subset of the policy commands for the domains to which the administrator is authorized. You can use this parameter to grant additional policy domain authority to a restricted policy administrator. This parameter is optional. You can specify more than one policy domain by delimiting each policy domain name with a comma.

You can use wildcard characters to specify a name. Authority for all matching policy domains is granted.

STGpools

Specifies that you want to grant restricted storage privilege to an administrator. If the `STGPOOLS` parameter is specified, then `CLASSES=STORAGE` is optional.

Restricted storage privilege permits you to issue a subset of the storage commands for the storage pools to which the administrator is authorized. You can use this parameter to grant additional storage pool authority to a restricted storage administrator. This parameter is optional. You can specify more than one storage pool by delimiting each storage pool name with a comma.

You can use wildcard characters to specify a name. Authority for all matching storage pools is granted.

Example: Grant system privilege to an administrator

Grant system privilege to administrator Larry.

```
grant authority larry classes=system
```

Example: Grant access to additional policy domains

Specify additional policy domains that the restricted policy administrator CLAUDIA can manage.

```
grant authority claudia domains=employee_records,prog1
```

Example: Provide an administrator with unrestricted storage privilege and restricted policy privilege

Provide administrator TOM with unrestricted storage privilege and restricted policy privilege for the domains whose names start with EMP.

```
grant authority tom classes=storage domains=emp*
```

Example: Grant an administrator authority restricted to a specific node

Grant node privilege to user HELP so that help desk personnel can assist the client node LABCLIENT in backing up or restoring data without having other higher-level Tivoli Storage Manager privileges.

```
grant authority help classes=node node=labclient
```

Related commands

Table 174. Commands related to GRANT AUTHORITY

Command	Description
QUERY ADMIN	Displays information about one or more Tivoli Storage Manager administrators.
REVOKE AUTHORITY	Revokes one or more privilege classes or restricts access to policy domains and storage pools.

GRANT PROXYNODE (Grant proxy authority to a client node)

Use this command to grant proxy authority to a client node on the Tivoli Storage Manager server.

Target client nodes own the data and agent nodes act on behalf of the target nodes. When granted proxy authority to a target client node, an agent node can perform backup and restore operations for the target node. Data that the agent node stores on behalf of the target node is stored under the target node's name in server storage.

Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege

Syntax

```
▶▶ GRANT PROXYnode TArget==target_node_name—AGent==agent_node_name▶▶
```

Parameters

TArget (Required)

Specifies the name of the node that owns the data. Wildcard names cannot be used to specify the target node name.

AGent (Required)

Specifies the name of the node performing operations for the target node. The agent node does not have to be in the same domain as the target node. Wildcard characters and comma-separated lists of node names are allowed.

Example: Grant proxy authority to a client node

Assume that MOE and JOE are agent nodes in a NAS cluster and are used to backup and restore shared NAS data. To create a proxy authority relationship for target node NASCLUSTER, issue the following command:

```
grant proxynode target=nascluster agent=moe,joe
```

Issue the following command on agent node MOE to back up NAS cluster data stored on the E: drive. The name of the target node is NASCLUSTER.

```
dsmc -asnode=nascluster incremental e:
```

Related commands

Table 175. Commands related to GRANT PROXYNODE

Command	Description
QUERY PROXYNODE	Display nodes with authority to act as proxy nodes.
REVOKE PROXYNODE	Revoke proxy authority from an agent node.

HALT (Shut down the server)

Use this command to shut down the server. The **HALT** command forces an abrupt shutdown, which cancels all the administrative and client node sessions even if they are not completed.

Any transactions in progress interrupted by the **HALT** command are rolled back when you restart the server. Use the **HALT** command only after the administrative and client node sessions are completed or canceled. To shut down the server without severely impacting administrative and client node sessions, perform the following steps:

1. Use the **DISABLE SESSIONS** command to prevent starting new client node sessions.
2. Use the **QUERY SESSIONS** command to identify any existing administrative and client node sessions.
3. Notify any existing administrative and client node sessions that you plan to shut down the server (you must do this outside of Tivoli Storage Manager).
4. Use the **CANCEL SESSIONS** command to cancel any existing administrative or client node sessions.
5. Issue the **HALT** command to shut down the server and stop any administrative and client node sessions.

Tip:

The **HALT** command can be replicated using the **ALIASHALT** server option. Use the server option to define a term other than **HALT** that performs the same function. The **HALT** command retains its normal function however, the server option provides an additional method for issuing the **HALT** command. See “**ALIASHALT**” on page 1457 for additional information.

Privilege class

To issue this command, you must have system or operator privilege.

Syntax

▶▶—HALT—◀◀

Parameters

None.

Example: Shut down the server

Shut down the server, either from the server console or from an administrative client. All user activity stops immediately and no new activity can start.

```
halt
```

Related commands

Table 176. Commands related to **HALT**

Command	Description
CANCEL PROCESS	Cancels a background server process.

HALT

Table 176. Commands related to HALT (continued)

Command	Description
CANCEL SESSION	Cancels active sessions with the server.
DISABLE SESSIONS	Prevents new sessions from accessing Tivoli Storage Manager but permits existing sessions to continue.
ENABLE SESSIONS	Resumes server activity following the DISABLE command or the ACCEPT DATE command.
QUERY PROCESS	Displays information about background processes.
QUERY SESSION	Displays information about all active administrator and client sessions with Tivoli Storage Manager.

utility_name

Specifies the name of the server utility for which you want to display information. This parameter is optional.

Example: Display the help topics

Display the help topics for the command-line interface.

```
help
```

Partial output:

```
1.0 Administering the server from the command line
  1.1 Issuing commands from the administrative client
    1.1.1 Starting and stopping the administrative client
    1.1.2 Monitoring server activities from the administrative client
```

Example: Display a help topic by using the help topic number

Display help information by using the help topic number. The topic number for the command **DEFINE DEVCLASS** for a 3590 device class is 3.13.10.2.

```
help 3.13.10.2
```

Example: Display help for one command

Display help information on the **REMOVE** commands.

```
help remove
```

```
3.44 REMOVE commands
Use the REMOVE commands to remove an object.
The following is a list of REMOVE commands:
* 3.44.1, "REMOVE ADMIN (Delete an administrator)"
* 3.44.2, "REMOVE NODE (Delete a node or an associated machine node)"
```

Example: Display help for a specific error message

Display help information on the error message ANR2535E.

```
help 2535
```

```
ANR2535E Command: The node node name cannot be removed or renamed
because it has an associated data mover.
Explanation: You attempted to remove or rename a node that has an
associated data mover.
System action: The server does not remove or rename the node.
User response: To remove or rename the node, delete the associated data
mover and reissue the command.
```

Example: Display help for a specific option

Display the description, syntax, and an example for the **COMMMETHOD** server option.

```
help commethod
```

Example: Display help for a specific utility

Display the description, syntax, and an example for the **DSMSERV** utility.

```
help dsmserv
```

IDENTIFY DUPLICATES (Identify duplicate data in a storage pool)

Use this command to start or stop processes that identify duplicate data in a storage pool. You can specify the number of duplicate-identification processes and their duration.

When you create a new storage pool for data deduplication, you can specify 0 - 20 duplicate-identification processes. Tivoli Storage Manager starts the specified number of duplicate-identification processes automatically when the server is started. If you do not stop them, they run indefinitely.

This command affects only server-side deduplication processing. In client-side data deduplication processing, duplicates are identified on the backup-archive client.

With the **IDENTIFY DUPLICATES** command, you can start additional processes, stop some or all of the processes, and specify an amount of time that the change remains in effect. If you increased or decreased the number of duplicate-identification processes, you can use the **IDENTIFY DUPLICATES** command to reset the number of processes to the number that is specified in the storage pool definition.

If you did not specify any duplicate-identification processes in the storage pool definition, you can use the **IDENTIFY DUPLICATES** command to start and stop all processes manually.

This command starts or stops a background process or processes that you can cancel with the **CANCEL PROCESS** command. To display information about background processes, use the **QUERY PROCESS** command.

Important:

- You can also change the number of duplicate-identification processes by updating the storage pool definition using the **UPDATE STGPOOL** command. However, when you update a storage pool definition, you cannot specify a duration. The processes that you specify in the storage pool definition run indefinitely, or until you issue the **IDENTIFY DUPLICATES** command, update the storage pool definition again, or cancel a process.

Issuing the **IDENTIFY DUPLICATES** does not change the setting for the number of duplicate-identification processes in the storage pool definition.

- Duplicate-identification processes can be either active or idle. Processes that are deduplicating files are active. Processes that are waiting for files to deduplicate are idle. Processes remain idle until volumes with data to be deduplicated become available. Processes stop only when cancelled or when you change the number of duplicate-identification processes for the storage pool to a value less than that currently specified. Before a duplicate-identification process stops, it must finish the file that it is deduplicating.

The output of the **QUERY PROCESS** command for a duplicate-identification process includes the total number of bytes and files that have been processed since the process first started. For example, if a duplicate-identification process processes four files, becomes idle, and then processes five more files, then the total number of files processed is nine.

Privilege class

To issue this command, you must have system privilege.

This parameter is optional. If you do not specify a value, the processes that are running after the command is issued will run indefinitely. They will end only if you reissue the **IDENTIFY DUPLICATES** command, update the storage pool definition, or cancel a process.

For example, if you define a storage pool with two duplicate-identification processes and you issue the **IDENTIFY DUPLICATES** command with `DURATION=60` and `NUMPROCESS=4`, the server starts two additional duplicate-identification processes that run for 60 minutes. At the end of that time, two processes finish the files they are working on and stop. The two processes that stop might not be the same two processes that started as a result of issuing this command.

The Tivoli Storage Manager server stops idle processes first. If, after stopping all idle processes, more processes need to be stopped, the Tivoli Storage Manager server notifies active processes to stop.

When the server stops a duplicate-identification process, the process completes the current physical file and then stops. As a result, it might take several minutes to reach the amount of time that you specified as a value for this parameter.

Example: Controlling the number and duration of duplicate-identification processes

In this example, you specified three duplicate-identification processes in the storage pool definition. You use the **IDENTIFY DUPLICATES** command to change the number of processes and to specify the amount of time the change is to remain in effect.

Table 177. Controlling duplicate-identification processes manually

The storage pool definition specifies three duplicate-identification processes. Using the IDENTIFY DUPLICATES command, you specify...	...and a duration of...	The result is...
2 duplicate-identification processes	None specified	One duplicate-identification processes finishes the file it is working on, if any, and then stops. Two processes run indefinitely, or until you reissue the IDENTIFY DUPLICATES command, update the storage pool definition, or cancel a process.
	60 minutes	One duplicate-identification process finishes the file it is working on, if any, and then stops. After 60 minutes, the server starts one process so that three are running.
4 duplicate-identification processes	None specified	The server starts one duplicate-identification process. Four processes run indefinitely, or until you reissue the IDENTIFY DUPLICATES command, update the storage pool definition, or cancel a process.
	60 minutes	The server starts one duplicate-identification process. At the end of 60 minutes, one process finishes the file it is working on, if any, and then stops. The additional process started by this command might not be the one that stops when the duration has expired.

IDENTIFY DUPLICATES

Table 177. Controlling duplicate-identification processes manually (continued)

The storage pool definition specifies three duplicate-identification processes. Using the IDENTIFY DUPLICATES command, you specify...	...and a duration of...	The result is...
0 duplicate-identification processes	None specified	All duplicate-identification processes finish the files that they are working on, if any, and stop. This change lasts indefinitely, or until you reissue the IDENTIFY DUPLICATES command, update the storage pool definition, or cancel a process.
	60 minutes	All duplicate-identification processes finish the files that they are working on, if any, and stop. At the end of 60 minutes, the server starts three processes.
None specified	Not available	The number of duplicate-identification processes resets to the number of processes specified in the storage pool definition. This change lasts indefinitely, or until you reissue the IDENTIFY DUPLICATES command, update the storage pool definition, or cancel a process.

Example: Identify duplicates in a storage pool

Identify duplicates in a storage pool, STGPOOLA, using three duplicate-identification processes. Specify that this change is to remain in effect for 60 minutes.

```
identify duplicates stgpoola duration=60 numprocess=3
```

Related commands

Table 178. Commands related to **IDENTIFY DUPLICATES**

Command	Description
CANCEL PROCESS	Cancels a background server process.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
QUERY CONTENT	Displays information about files in a storage pool volume.
QUERY PROCESS	Displays information about background processes.
QUERY STGPOOL	Displays information about storage pools.
UPDATE STGPOOL	Changes the attributes of a storage pool.

IMPORT commands

Use the **IMPORT** commands to import information from export media to a Tivoli Storage Manager server.

Important: For commands that import administrators or nodes, you must consider the method of authentication. The Tivoli Storage Manager server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the **IMPORT** command, set the password by issuing the **UPDATE ADMIN** or **UPDATE NODE** command.

- “IMPORT ADMIN (Import administrator information)” on page 570
- “IMPORT NODE (Import client node information)” on page 574
- “IMPORT POLICY (Import policy information)” on page 581
- “IMPORT SERVER (Import server information)” on page 584

IMPORT ADMIN (Import administrator information)

Use this command to import administrator and authority definitions for one or more administrators from export media to the Tivoli Storage Manager server.

Important: For commands that import administrators or nodes, you must consider the method of authentication. The Tivoli Storage Manager server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the **IMPORT** command, set the password by issuing the **UPDATE ADMIN** or **UPDATE NODE** command.

You can use the **QUERY ACTLOG** command to view the status of the import operation.

You can also view this information from the server console.

Limitation: The Tivoli Storage Manager server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue, after the import or node replication operation, update the fields with the appropriate **UPDATE** commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

This command generates a background process that can be canceled with the **CANCEL PROCESS** command. If an **IMPORT ADMIN** background process is canceled, some of the data is already imported. To display information about background processes, use the **QUERY PROCESS** command.


Restriction:


- If target and source server levels are not compatible, the operation might not work. See the *Administrator's Guide* for server compatibility requirements.
- If the administrator definition that is being imported includes analyst authority, the administrator definition is imported but not the analyst authority. Analyst authority is not valid for servers at V6.1 or later.

Privilege class

To issue this command, you must have system privilege.

Syntax





Parameters

admin_name

Specifies the administrators for which you want to import information. This parameter is optional. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names.

Preview

Specifies whether you want to preview the results of the import operation, without importing administrator information. If you specify YES for the value, you must mount the export volumes. This parameter is optional. The default value is NO. The values are:

No Specifies that the information is to be imported.

Yes

Specifies that the operation will be previewed but not performed. Information about the number and types of objects imported, together with the number of bytes transferred, are reported to the server console and the activity log.

DEVclass (Required)

Specifies the device class from which import data is to be read.

You cannot specify the DISK device class or a device class with a device type of NAS. If all drives for the device class are busy when the import runs, Tivoli Storage Manager cancels lower priority operations, such as reclamation, to make a drive available.

VOLumentname (Required)

Specifies the volumes to be used for the import operation. Volumes must be imported in the same order as they were exported. The values are:

volume_name

Specifies the volume name. To specify multiple volumes, separate names with commas and no intervening spaces.

FILE:*file_name*

Specifies the name of a file that contains a list of volumes used for the imported data. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when specifying volumes associated with the following device types:

For this device	Specify
Tape	1 – 6 alphanumeric characters.

For this device	Specify
FILE	Any fully qualified file name string. For example: /imdata/mt1.
REMOVABLEFILE	1-6 alphanumeric characters.
SERVER	1-250 alphanumeric characters.

Replacedefs

Specifies whether to replace administrator definitions on the target server. The default value is NO. The values are:

No Specifies that definitions are not to be replaced.

Yes

Specifies that definitions are to be replaced.

Example: Import administrator information from specific tape volumes

From the server, import the information for all defined administrators from tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device assigned to the MENU1 device class. Issue the command:

```
import admin devclass=menu1
volumenames=tape01,tape02,tape03
```

Example: Import administrator information from tape volumes listed in a file

From the server, import the information for all defined administrators from tape volumes that are listed in the following file:

```
TAPEVOL
```

This file contains these lines:

```
TAPE01
TAPE02
TAPE03
```

Specify that these tape volumes be read by a device assigned to the MENU1 device class. Issue the command:

```
import admin devclass=menu1 volumenames=file:tapevol
```

Related commands

Table 179. Commands related to IMPORT ADMIN

Command	Description
CANCEL PROCESS	Cancels a background server process.
EXPORT ADMIN	Copies administrative information to external media or directly to another server.
IMPORT NODE	Restores client node information from external media.
IMPORT POLICY	Restores policy information from external media.
IMPORT SERVER	Restores all or part of the server from external media.

Table 179. Commands related to IMPORT ADMIN (continued)

Command	Description
QUERY ACTLOG	Displays messages from the server activity log.
QUERY PROCESS	Displays information about background processes.

IMPORT NODE (Import client node information)

Use this command to import client node definitions from a server or sequential media to a target IBM Tivoli Storage Manager server.

Important: For commands that import administrators or nodes, you must consider the method of authentication. The Tivoli Storage Manager server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the **IMPORT** command, set the password by issuing the **UPDATE ADMIN** or **UPDATE NODE** command.

The imported nodes are associated with the same policy domain as they were on the source server, if you specify the domain and if that policy domain exists on the target server. Otherwise, imported nodes are associated with the **STANDARD** policy domain on the target server.

IBM Tivoli Storage Manager servers with retention protection enabled do not allow import operations.

Restriction:

1. If target and source server levels are not compatible, the operation might not work.
For server compatibility requirements, see the *Administrator's Guide*.
2. Importing data from a **CENTERA** device class is not supported. However, files that are being imported can be stored on a **CENTERA** storage device.
3. If you use an LDAP directory server to authenticate passwords, any target servers must be configured for LDAP passwords. Data that is imported from a node that authenticates with an LDAP directory server is inaccessible if the target server is not properly configured. If your target server is not configured, imported data from an LDAP node can still go there. But the target server must be configured to use LDAP in order for you to access the imported data.

You can use the **QUERY ACTLOG** command to view the status of the import operation. You can also view this information from the server console.

This command generates a background process that can be canceled with the **CANCEL PROCESS** command. If an **IMPORT NODE** background process is canceled, some of the data might already be imported. To display information about background processes, use the **QUERY PROCESS** command.

For a server that has clients with support for Unicode, you might need to have the server convert the file space name that you enter, or use the following parameter:

- **HEXFILESPEC** parameter
- **UNIFILESPEC** parameter


Limitation: The Tivoli Storage Manager server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue, after the import or node replication operation, update the fields with the appropriate **UPDATE** commands. This server

limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

node_name

Specifies the client nodes for which you want to import information. This parameter is optional.

Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. All matching nodes are included in the list.

FILESpace

Specifies file space names for which you want to import information. This parameter is optional. The default is all file spaces.

Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names.

Important:

1. Existing file spaces are not replaced. New file spaces are created when identical names are encountered. However, this new name might match an existing name on the client node, which can have file spaces not yet backed up to the server.
2. This parameter is only specified for non-Unicode file spaces. To import all file spaces, that are both Unicode and non-Unicode, use the **FILEDATA=ALL** parameter without the **FILESPACE** and **UNIFILESPACE** parameters.

DOmains

Specifies the policy domains from which to import node information. These domains must have been included in the data that was exported. This parameter is optional. The default is all domains that were exported.

Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify a name.

FILEData

Specifies the type of files that can be imported for all nodes specified and found on the export media. This parameter is optional. The default value is **NONE**.

If you are importing from sequential media, the device class used by the file data is determined by the device class for the storage pool. If it is the same device class specified in this command, two drives are needed to import the node information. The mount limit for the device class must be at least 2.

The following descriptions mention *active* and *inactive* backup file copies. An active backup file copy is the most recent backup copy for a file that still exists on the client workstation. All other backup file copies are called inactive copies. The values are:

ALl

The server imports all backup versions of files, all archived files, and all files that were migrated by a Tivoli Storage Manager for Space Management client. File spaces included are both Unicode and non-Unicode.

None

Only node definitions are imported. The server does not import any files.

ARchive

The server imports only archived files.

Backup

The server imports only backup versions, whether active or inactive.

BACKUPActive

The server imports only active backup versions. These active backup versions are the active versions in the Tivoli Storage Manager database at the time that the **IMPORT** command is issued.

ALLActive

The server imports all active backup versions of files, all archived files, and all files that were migrated by a Tivoli Storage Manager for Space Management client. The active backup versions are the active versions in the Tivoli Storage Manager database at the time that the **IMPORT** command is issued.

Spacemanaged

The server imports only files that were migrated by a Tivoli Storage Manager for Space Management client.

Preview

Specifies whether to preview the results of the import operation, without importing information. The **PREVIEW=YES** option requires that you mount the export volumes. This parameter is optional. The default value is **NO**. Possible values are:

No Specifies that the node information is to be imported.

Yes

Specifies that you want to preview the results of the import operation, without importing files. Information is reported to the server console and the activity log.

DEVclass (Required)

Specifies the device class from which import data is to be read. You cannot specify the **DISK** device class or a device class with a device type of **NAS**.

If all drives for the device class are busy when the import runs, the server cancels lower priority operations, such as reclamation, to make a drive available.

Dates

Specifies whether the date of the file copies are as specified when the files were exported (**ABSOLUTE**), or are adjusted to the date of import (**RELATIVE**). The default value is **ABSOLUTE**.

If the export media was idle for some time after export (sitting on a shelf for six months, for example), the original backup, or archive dates might be old enough to trigger immediate expiration of file copies when the data is imported into a server. The **RELATIVE** specification for this value adjusts for time elapsed since export so that the file copies are not immediately expired.

For example, assume that an export tape contains an archive file copy that was archived five days before the export operation. If the media are saved for six months and then imported, the archive file appears to have been inserted six months and five days ago by default (**DATES=ABSOLUTE**) and might be subject to immediate expiration depending on the retention value that is specified in the file's management class. Specifying **DATES=RELATIVE** results in resetting the archive date for the file to five days ago during import. The **DATES=RELATIVE** parameter thus adjusts file backup and archive dates for the time that elapsed since the export operation occurred.

IMPORT NODE

Possible values are:

Absolute

The dates for file copies are set to the values specified when the files were exported.

Relative

The dates for file copies are adjusted to the date of import.

VOLUMenames (Required)

Specifies the volumes to be used for the import operation. Volumes must be imported in the same order as they were exported. The values are:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE:*file_name*

Specifies the name of a file that contains a list of volumes used for the imported data. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when specifying volumes associated with the following device types:

For this device	Specify
Tape	1 – 6 alphanumeric characters.
FILE	Any fully qualified file name string. An example is /imdata/mt1.
REMOVABLEFILE	1–6 alphanumeric characters.
SERVER	1–250 alphanumeric characters.

Replacedefs

Specifies whether to replace definitions on the target server. The default value is NO. The values are:

No Objects are not to be replaced.

Yes

Objects are to be replaced.

HEXFILESpace

Specifies the hexadecimal representation of the file space names in UTF-8 format. Separate multiple names with commas and no intervening spaces. This parameter is optional.

To view the hexadecimal representation of a file space name, you can use the **QUERY FILESPACE** command with **FORMAT=DETAILED**.

UNIFILESpace

Specifies the file spaces that are known to the server to be Unicode enabled. The server converts the names you enter from the server code page to the UTF-8 code page to find the file spaces to import. The success of the conversion depends on the actual characters in the name and the server's code page. Separate multiple names with commas and no intervening spaces. This parameter is optional.

MERGEfilespaces

Specifies whether Tivoli Storage Manager merges client files into existing file

spaces on the target server (if they exist), or if Tivoli Storage Manager generates new file space names. The default is NO.

Valid values are:

Yes

Specifies that imported data on the target server is merged with the existing file space, if a file space with the same name exists on the target server.

No Specifies that Tivoli Storage Manager generates a new file space name for imported data on the target server if file spaces with the same name exists.

PROXynodeassoc

Specifies if proxy node associations are imported. This parameter is optional. The default value is NO.

Example: Import client node information from tapes

From the server, import client node information from tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device assigned to the MENU1 device class.

```
import node devclass=menu1 volumenames=tape01,tape02,tape03
```

Example: Import client node information from tapes listed in a file

From the server, import client node information from tape volumes listed in a file named TAPEVOL.

This file contains these lines:

```
TAPE01
TAPE02
TAPE03
```

Specify that these tape volumes be read by a device assigned to the MENU1 device class.

```
import node devclass=menu1 volumenames=file:tapevol
```

Example: Import the active backup for a client node

From the server, import the active backup versions of file data for client node JOE from tape volume TAPE01. The file space is Unicode.

```
import node joe unificspace=\\joe\c$ filedata=backupactive devclass=menu1
volumenames=tape01
```

Related commands

Table 180. Commands related to IMPORT NODE

Command	Description
CANCEL PROCESS	Cancels a background server process.
COPY ACTIVE DATA	Copies active backup data.
EXPORT NODE	Copies client node information to external media or directly to another server.
IMPORT ADMIN	Restores administrative information from external media.

IMPORT NODE

Table 180. Commands related to **IMPORT NODE** (continued)

Command	Description
IMPORT POLICY	Restores policy information from external media.
IMPORT SERVER	Restores all or part of the server from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY PROCESS	Displays information about background processes.

IMPORT POLICY (Import policy information)

Use this command to import policy domain information from sequential export media to the Tivoli Storage Manager server. IBM Tivoli Storage Manager servers with retention protection enabled do not allow import operations.

Tivoli Storage Manager client data can be moved between servers with export and import processing, if the same removable media type is supported on both platforms.

Restriction: If target and source server levels are not compatible, the import operation might not work. See the *Administrator's Guide* for server compatibility requirements.

You can use the **QUERY ACTLOG** command to view the status of the import operation. You can also view this information from the server console.


This command generates a background process that can be canceled with the **CANCEL PROCESS** command. If an **IMPORT POLICY** background process is canceled, some of the data is already imported. To display information about background processes, use the **QUERY PROCESS** command.

Limitation: The Tivoli Storage Manager server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue, after the import or node replication operation, update the fields with the appropriate **UPDATE** commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

domain_name

Specifies the policy domains for which information is to be imported. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. The default (*) is all policy.

Preview

Specifies whether you want to preview the results of the import operation without importing information. The PREVIEW=YES option requires that you mount the export volumes. This parameter is optional. The default value is NO. The values are:

No Specifies that the information is to be imported.

Yes

Specifies that the operation will be previewed but not performed. Information is reported to the server console and the activity log.

DEVclass (Required)

Specifies the device class from which import data is to be read. You cannot specify the DISK device class.

You also cannot specify a device class with a device type of NAS.

If all drives for the device class are busy when the import runs, Tivoli Storage Manager cancels lower priority operations, such as reclamation, to make a drive available.

VOLumentnames (Required)

Specifies the volumes to be used for the import operation. Volumes must be imported in the same order as they were exported. The values are:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE:*file_name*

Specifies the name of a file that contains a list of volumes. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when specifying volumes associated with the following device types:

For this device	Specify
Tape	1 – 6 alphanumeric characters.
FILE	Any fully qualified file name string. For example: /imdata/mt1
REMOVABLEFILE	1–6 alphanumeric characters.
SERVER	1–250 alphanumeric characters.

Replacedefs

Specifies whether to replace policy definitions on the target server. The default value is NO. The values are:

Yes

Specifies that objects are to be replaced by the imported objects.

No Specifies that objects are not to be replaced by imported objects.

Example: Import policy information from specific tape volumes

From the server, import the information for all defined policies from tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device assigned to the MENU1 device class.

```
import policy devclass=menu1
volumenames=tape01,tape02,tape03
```

Example: Import policy information from tape volumes listed in a file

From the server, import the information for all defined policies from tape volumes listed in a file named thus:

```
TAPEVOL
TAPEVOL.DATA
```

Specify that these tape volumes be read by a device assigned to the MENU1 device class. The file contains the following lines:

```
TAPE01
TAPE02
TAPE03
import policy devclass=menu1 volumenames=file:tapevol
```

Related commands

Table 181. Commands related to IMPORT POLICY

Command	Description
CANCEL PROCESS	Cancels a background server process.
EXPORT POLICY	Copies policy information to external media or directly to another server.
IMPORT ADMIN	Restores administrative information from external media.
IMPORT NODE	Restores client node information from external media.
IMPORT SERVER	Restores all or part of the server from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY PROCESS	Displays information about background processes.

IMPORT SERVER (Import server information)

Use this command to copy all or part of the server control information and client file data (if specified) from export media to the Tivoli Storage Manager server.

Important: For commands that import administrators or nodes, you must consider the method of authentication. The Tivoli Storage Manager server cannot export or import passwords for nodes or administrators that are authenticating with LDAP directory servers. If the current authentication method uses an LDAP directory server and the password is not already synchronized by that server, you must update the password. After issuing the **IMPORT** command, set the password by issuing the **UPDATE ADMIN** or **UPDATE NODE** command.

IBM Tivoli Storage Manager servers with retention protection enabled do not allow import operations.

Restriction:

- If target and source server levels are not compatible, the operation might not work.

For server compatibility requirements, see the *Administrator's Guide*.

- If you use an LDAP directory server to authenticate passwords, any target servers must be configured for LDAP passwords. Server data that is exported from a node that authenticates with an LDAP directory server is inaccessible if the target server is not properly configured. If your target server is not configured, exported data from an LDAP node can still go there. But the target server must be configured to use LDAP in order for you to access the data.

You can also initiate an import of server information and client file data directly from the originating server. For more information, see the **EXPORT** commands.


This command generates a background process that can be canceled with the **CANCEL PROCESS** command. If an **IMPORT SERVER** background process is canceled, some of the data is already imported. To display information about background processes, use the **QUERY PROCESS** command.

Limitation: The Tivoli Storage Manager server does not convert code pages during export, import, and node replication operations. If servers are running in different locales, some information in databases or system output might become unreadable. Invalid characters might be displayed, for example, in the contact information for the administrator and client nodes, and in descriptions of policy domains. Any field that is stored in the server character set and that includes extended ASCII characters can be affected. To resolve the issue, after the import or node replication operation, update the fields with the appropriate **UPDATE** commands. This server limitation does not affect client data. Any client data that was exported, imported, or replicated can be restored, retrieved, and recalled.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

FILEData

Specifies the type of files that can be imported for all nodes defined to the server. This parameter is optional. The default value is NONE.

The device class used to access the file data is determined by the device class for the storage pool. If it is the same device class specified in this command, two drives are needed to import information. The mount limit for the device class must be set to at least 2.

The following descriptions mention active and inactive backup file copies. An active backup file copy is the most recent backup copy for a file that still exists on the client workstation. All other file copies are called inactive copies. The values are:

ALL

Tivoli Storage Manager imports all backup versions of files, all archived files, and all files that were migrated by a Tivoli Storage Manager for Space Management client.

None

Tivoli Storage Manager does not import files, only node definitions.

ARchive

Tivoli Storage Manager imports only archived files.

Backup

Tivoli Storage Manager imports only backup versions, whether the versions are active or inactive.

IMPORT SERVER

BACKUPActive

Tivoli Storage Manager imports only active backup versions. These active backup versions are the active versions in the Tivoli Storage Manager database at the time that the **IMPORT** command is issued.

ALLActive

Tivoli Storage Manager imports all active backup versions of files, all archived files, and all files that were migrated by a Tivoli Storage Manager for Space Management client. The active backup versions are the active versions in the Tivoli Storage Manager database at the time that the **IMPORT** command is issued.

SPacemanaged

Tivoli Storage Manager imports only files that were migrated by a Tivoli Storage Manager for Space Management client.

Preview

Specifies whether to preview the results of the import operation, without importing information. The **PREVIEW=YES** option still requires that you mount the export volumes. This parameter is optional. The default value is **NO**. Possible values are:

No Specifies that the server information is to be imported.

Yes

Specifies that the operation will be previewed but not performed. Information is transferred to the server console and the activity log.

DEVclass (Required)

Specifies the device class from which import data is to be read. You cannot specify the **DISK** device class or a device class with a device type of **NAS**.

If all drives for the device class are busy when the import runs, Tivoli Storage Manager cancels lower priority operations, such as reclamation, to make a drive available.

Dates

Specifies whether the dates of the file copies are set as specified when the files were exported (**ABSOLUTE**), or are adjusted to the date of import (**RELATIVE**). The default value is **ABSOLUTE**.

If the import media was idle for some time after export (sitting on a shelf for six months, for example), the original backup, or archive dates might be old enough to trigger immediate expiration of file copies when the data is imported into a server. The **RELATIVE** specification for this value adjusts for time elapsed since export so that the file copies are not immediately expired.

For example, assume that an import tape contains an archive file copy that was archived five days before the export operation. If the export media are saved for six months and then imported, the archive file appears to have been inserted six months and five days ago by default (**DATES=ABSOLUTE**) and might be subject to immediate expiration depending upon the retention value that is specified in the file's management class. Specifying **DATES=RELATIVE** results in resetting the archive date for the file to five days ago during import. **DATES=RELATIVE** parameter thus adjusts file backup and archive dates for the time that elapsed since the export operation occurred.

Possible values are:

Absolute

The dates for file copies are set to the values specified when the files were exported.

Relative

The date for file copies are adjusted to the date of import.

VOLumentnames (Required)

Specifies the volumes to be used for the import operation. Volumes must be imported in the same order as they were exported. The values are:

volume_name

Specifies the volume name. To specify multiple volumes, separate the names with commas and no intervening spaces.

FILE:*file_name*

Specifies the name of a file that contains a list of volumes used for the imported data. In the file, each volume name must be on a separate line. Blank and comment lines that begin with an asterisk are ignored.

Use these naming conventions when specifying volumes associated with the following device types:

For this device	Specify
Tape	1 – 6 alphanumeric characters.
FILE	Any fully qualified file name string. An example is /imdata/mt1.
REMOVABLEFILE	1–6 alphanumeric characters.
SERVER	1–250 alphanumeric characters.

Replacedefs

Specifies whether to replace objects on the server. Existing file spaces are not replaced. New file spaces are created when identical names are encountered. The default value is NO. Possible values are:

No Specifies that objects are not to be replaced by imported objects.

Yes

Specifies that objects are to be replaced by the imported objects.

MERGEfilespace

Specifies whether Tivoli Storage Manager merges client files into existing file spaces on the target server (if they exist), or if Tivoli Storage Manager generates new file space names. You cannot merge non-Unicode and Unicode file spaces together. The default is NO. Possible values are:

No Specifies that Tivoli Storage Manager generates a new file space name for imported data on the target server if file spaces with the same name already exist.

Yes

Specifies that imported data on the target server is merged with the existing file space, if a file space with the same name already exists on the target server.

PROXynodeassoc

Specifies if proxy node associations are imported. This parameter is optional. The default value is NO.

Example: Import the information for all defined servers from specific tapes

From the server, import the information for all defined servers from tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device assigned to the MENU1 device class.

```
import server devclass=menu1 volumenames=tape01,tape02,tape03
```

Example: Import information for all defined servers from specific tapes and specify files are merged into existing file spaces

From the server, import the information for all defined servers from tape volumes TAPE01, TAPE02, and TAPE03. Specify that these tape volumes be read by a device assigned to the MENU1 device class and that client files be merged into file spaces on the target server if file spaces of the same names already exist.

```
import server devclass=menu1 volumenames=tape01,tape02,tape03 mergefilespace=yes
```

Example: Import information for all defined servers from tapes listed in a file

From the server, import the information for all defined servers from tape volumes listed in a file named TAPEVOL. Specify that the tape volumes are read by a device assigned to the MENU1 device class. The input file contains these lines:

```
TAPE01
TAPE02
TAPE03
```

```
import server devclass=menu1 volumenames=file:tapevol
```

Related commands

Table 182. Commands related to **IMPORT SERVER**

Command	Description
CANCEL PROCESS	Cancels a background server process.
COPY ACTIVATEDATA	Copies active backup data.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
IMPORT ADMIN	Restores administrative information from external media.
IMPORT NODE	Restores client node information from external media.
IMPORT POLICY	Restores policy information from external media.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY PROCESS	Displays information about background processes.

INSERT MACHINE (Insert machine characteristics information or recovery instructions)

Use this command to add client machine characteristics or recovery instructions to existing machine information in the database.

You can write a program to read files containing the information and generate the appropriate **INSERT MACHINE** commands. See the *Administrator's Guide* for details.

You can use **QUERY** commands to retrieve the information if a disaster occurs.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

▶▶—INsert MACHine—machine_name—sequence_number—————▶
|
|—Characteristics—==—text—————▶
|—RECOVERYInstructions—==—text—————▶

```

Parameters

machine_name (Required)

Specifies the name of the client machine.

sequence_number (Required)

Specifies the sequence number for the line of text in the database.

Characteristics

Specifies machine characteristics information. You must specify the characteristics or recovery instructions, but not both. Enclose the text in quotation marks if it contains blank characters. The text can be up to 1024 characters.

RECOVERYInstructions

Specifies recovery instructions. You must specify the characteristics or recovery instructions, but not both. Enclose the text in quotation marks if it contains blank characters. The text can be up to 1024 characters.

Example: Update a machine's information

For the machine **DISTRICT5**, insert this characteristics text on line 1: "Machine owner is Mary Smith".

```

insert machine district5 1
characteristics="Machine owner is Mary Smith"

```

Related commands

Table 183. Commands related to **INSERT MACHINE**

Command	Description
DEFINE MACHINE	Defines a machine for DRM.
DELETE MACHINE	Deletes a machine.
QUERY MACHINE	Displays information about machines.

ISSUE MESSAGE (Issue a message from a server script)

Use this command with return code processing in a script to issue a message from a server script to determine where the problem is with a command in the script.

Privilege class

Any administrator can issue this command.

Syntax

►—ISSUE MESSAGE—*message_severity*—*message_text*—◄

Parameters

message_severity (Required)

Specifies the severity of the message. The message severity indicators are:

- I** Information. ANR1496I is displayed in the message text.
- W** Warning. ANR1497W is displayed in the message text.
- E** Error. ANR1498E is displayed in the message text.
- S** Severe. ANR1499S is displayed in the message text.

message_text (Required)

Specifies the description of the message.

Example: Issue a message from a server script

Assume you have a script called `backupsript` that quiesces a client's database, takes a backup of that database, and then restarts the client's database. For illustration, your script results in a non-zero return code. Use the **ISSUE MESSAGE** command with the message severity and message text. The following is an example of a server script that calls `backupsript` on the client machine and issues messages based on the return code from `backupsript`.

```
issue message i "Starting backup"
define clientaction nodename action=command objects="c:\backupsript" wait=yes
if (101) goto qfail
if (102) goto qwarn
if (103) goto backupf
if (104) goto restartf
issue message i "Backup of database complete"
exit
qfail: issue message e "Quiesce of database failed"
exit
qwarn: issue message w "Quiesce of database failed, taking fuzzy backup"

exit
backupf: issue message e "Backup of database failed"
exit
restartf: issue message s "Database restart failed"
exit
```

Command

```
issue message e "quiesce of database failed"
```

Related commands

Table 184. Commands related to ISSUE MESSAGE

Command	Description
COPY SCRIPT	Creates a copy of a script.
DEFINE SCRIPT	Defines a script to the Tivoli Storage Manager server.
DELETE SCRIPT	Deletes the script or individual lines from the script.
RENAME SCRIPT	Renames a script to a new name.
RUN	Runs a script.
UPDATE SCRIPT	Changes or adds lines to a script.

LABEL LIBVOLUME (Label a library volume)

Use this command to label tape volumes or, in an automated library, to label the volumes automatically as they are checked in. With this command, the server uses the full-length label with which the volumes are often pre-labeled.

Important: Use this command only for MANUAL, SCSI, ACSLS, and 349X libraries. The command processing will not wait for a drive to become available, even if the drive is only in the IDLE state. If necessary, you can make a library drive available by issuing the **DISMOUNT VOLUME** command to dismount the volume in that particular drive. When the library drive becomes available, you can reissue the **LABEL LIBVOLUME** command.


For detailed and current drive and library support information, see the Supported Devices website for your operating system:

http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html


Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.


Syntax for a manual library



Syntax for a SCSI library




A (SEARCH=Yes, SEARCH=Bulk):




B (LABELSource=Vollist):




Syntax for a 349X library



A (SEARCH=Yes):




Syntax for an ACSLS library



LABEL LIBVOLUME

A (SEARCH=Yes):



Parameters

library_name (Required)

Specifies the name of the library that contains the storage volume.

volume_name

Specifies the name of the volume to be labeled.

- For SCSI libraries: The server requests that the volume be inserted into a slot in the library or, if available, into an entry/exit port. The server identifies a slot by the slot's element address. If you are labeling a volume in a SCSI library with multiple entry/exit ports, the volume in the lowest numbered slot will be labeled.
- For MANUAL libraries: The server requests that the volume be inserted into a drive.
- For 349X libraries: The volume might already be in the library, or you might be prompted to put it into the I/O station.

Remember: If the specified volume name is already defined in a storage pool or in a volume history file, the volume will not be labeled, and a message will be displayed.

CHECKIN

Specifies whether the server checks in the volume. This parameter is optional. The following are possible values:

SCRatch

Specifies that the server checks in the volumes and adds them to the library's scratch pool. If a volume has an entry in volume history, you cannot check it in as a scratch volume.

PRivate

Specifies that the server checks in the volumes and designates them as private. Private volumes are available only when you request them by name.

If you do not specify a value for this parameter, then the command will only label the volume but will not check it in. If you do not specify a value for this parameter and you want to check in the volume, you must issue the **CHECKIN LIBVOLUME** command.

SEARCH

Specifies that the server searches the library for usable volumes to label. This parameter applies to SCSI, 349X, and ACSLS libraries.

Possible values are:

Yes

Specifies that the server labels only volumes that are stored in the library, unless the volume is already labeled or its bar code cannot be read.

If you specify the LABELSOURCE=PROMPT option, the volume is moved into the drive from its location in the library or entry and exit ports. The server prompts you to issue the **REPLY** command containing the label string, and that label is written to the tape.

Bulk

Specifies that the server searches the library entry/exit ports for usable volumes to label. This option is only valid for SCSI libraries.

If you specify LABELSOURCE=BARCODE, the volume bar code is read, and the tape is moved from its location in the library or in the entry/exit ports to a drive where the bar code label is written. After the tape is labeled, it is moved back to its location in the library, to the entry/exit ports, or to a storage slot if the CHECKIN option is specified. For bar code support to work correctly for libraries supported by Tivoli Storage Manager, the Tivoli Storage Manager server and the device driver must be at the same level. Bar code support is available for libraries supported by Tivoli Storage Manager and which use the Tivoli Storage Manager device driver or the IBM Magstar® or LTO Ultrium device driver.

Tip: You can use the VOLRANGE or VOLLIST parameter to limit the search.

VOLRange

Specifies a range of volume names separated by a comma. Use this parameter to limit the search for volumes to be labeled when you specify SEARCH=YES (349X, ACSLS, and SCSI libraries) or SEARCH=BULK (SCSI libraries only). If there are no volumes in the library that are within the specified range, the command completes without errors.

You can specify only volume names that can be numerically incremented. In addition to the incremental area, a volume name can include an alphanumeric prefix and an alphanumeric suffix, for example:

Parameter	Description
volrange=bar110,bar130	The 21 volumes are labeled: bar110, bar111, bar112, ...,bar129, bar130.
volrange=bar11a,bar13a	The 3 volumes are labeled: bar11a, bar12a, bar13a.
volrange=123400,123410	The 11 volumes are labeled: 123400, 123401, ...123409, 123410.

VOLList

Specifies a list of volumes. Use this parameter to limit the search for volumes to be labeled when you specify SEARCH=YES (349X, ACSLS, and SCSI libraries) or SEARCH=BULK (SCSI libraries only). If there are no volumes in the library that are in the list, the command completes without errors. The VOLLIST parameter can also be the source of names to be used to label volumes if the LABELSOURCE parameter is set to VOLLIST. If LABELSOURCE=VOLLIST, you must specify the VOLLIST parameter.

Possible values are:

volume_name

Specifies the names of one or more values that are used for the command. For example: VOLLIST=TAPE01,TAPE02.

FILE:*file_name*

Specifies the name of a file that contains a list of volumes for the

LABEL LIBVOLUME

command. In the file, each volume name must be on a separate line. Blank lines and comment lines that begin with an asterisk are ignored. For example, to use volume *TAPE01*, *TAPE02* and *TAPE03*, create a file named *TAPEVOL* that contains these lines:

```
TAPE01  
TAPE02  
TAPE03
```

You can specify the volumes for the command as follows:
VOLLIST=FILE:TAPEVOL.

Remember: The file name is case-sensitive.

LABELSource

Specifies how or whether the server reads sequential media labels of volumes. This option is only valid for SCSI libraries. Specify this parameter only when SEARCH=YES or SEARCH=BULK.

Possible values are:

Prompt

The server prompts for volume names as necessary.

Barcode

The server attempts to read the bar code label. If the attempt fails, the server will not label the volume and will display a message.

Important: For bar code support to work properly, the appropriate device drivers must be installed for the libraries.

Vollist

This option only applies to SCSI libraries. The server attempts to read the specified file or list of files. If the attempt fails, the server will not label the volumes and will display a message.

OVERWRITE

Specifies whether the server attempts to overwrite existing labels. This parameter is optional. The default is NO. Possible values are:

No Specifies that the server only labels unlabeled volumes. For StorageTek VolSafe volumes and WORM optical volumes, the value must be NO.

Yes

Specifies that the server overwrites existing labels *only* if both the existing label and the prompted or bar code label are not already defined in either the server storage pool or volume history list.

WAITTime

Specifies the number of minutes that the server will wait for you to reply or respond to a request. Specify a value in the range 0-9999. If you want to be prompted by the server, specify a wait time greater than zero. The default value is 60 minutes. For example, suppose the server prompts you to insert a tape into the entry/exit port of a library. If you specified a wait time of 60 minutes, the server will issue a request and wait 60 minutes for you to reply. Suppose, on the other hand, you specify a wait time of 0. If you have already inserted a tape, a wait time of zero will cause the operation to continue without prompting. If you have not inserted a tape, a wait time of zero will cause the operation to fail.

Example: Automatically label library volumes

Label tapes in a SCSI library named AUTO automatically as you are checking the volumes in.

```
label libvolume auto checkin=scratch search=yes labelsource=barcode
overwrite=yes
```

Example: Label sequential library volumes

Label 3 volumes from bar11a to bar13a in a SCSI library named ABC. When you issue the following command, the 3 volumes are labeled: bar11a, bar12a, bar13a.

```
label libvolume abc checkin=scratch search=yes volrange=bar11a,bar13a
labelsource=barcode
```

Related commands

Table 185. Commands related to LABEL LIBVOLUME

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
CANCEL PROCESS	Cancels a background server process.
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY LIBVOLUME	Displays information about a library volume.
QUERY PROCESS	Displays information about background processes.
REPLY	Allows a request to continue processing.
UPDATE LIBVOLUME	Changes the status of a storage volume.

LOCK commands

Use the **LOCK** command to prevent users from accessing the server.

- “LOCK ADMIN (Lock out an administrator)” on page 599
- “LOCK NODE (Lock out a client node)” on page 601
- “LOCK PROFILE (Lock a profile)” on page 603

LOCK ADMIN (Lock out an administrator)

Use this command to prevent an administrator from accessing the server. The administrator is locked out until a system administrator uses the **UNLOCK ADMIN** command to reestablish access for the administrator.

You can use the authentication filter to lock all administrators, excluding console administrators. After configuring an LDAP directory server for password authentication, you can lock administrators to force them to create passwords that authenticate with an LDAP server.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

lock admin [*] authentication=[local|ldap]

```

Parameters

admin_name (Required)

Specifies the name of the administrator to be locked out. You can use wildcard characters to specify the administrator name. You do not have to enter an administrator name if you want to lock all of the administrators according to their authentication method. Use the wildcard with an authentication method to lock multiple administrators.

AUTHentication

Specifies the method of authentication that the administrator uses to log in.

LOcal

Specifies to lock administrators who authenticate to the IBM Tivoli Storage Manager server.

LDap

Specifies to lock administrators who authenticate to the LDAP directory server.

Example: Lock out an administrator

Lock out the administrator CLAUDIA. Issue the command:

```
lock admin claudia
```

Example: Lock out all administrators who authenticate to the Tivoli Storage Manager server database

Use the wildcard character (*) to lock all the administrators who authenticate their passwords locally. Console administrators are not affected by this command. Issue the following command:

```
lock admin * authentication=local
```

LOCK ADMIN

Related commands

Table 186. Commands related to LOCK ADMIN

Command	Description
QUERY ADMIN	Displays information about one or more Tivoli Storage Manager administrators.
UNLOCK ADMIN	Enables a locked administrator to access Tivoli Storage Manager.

LOCK NODE (Lock out a client node)

Use this command to prevent a client node from accessing the server. A locked client node cannot perform any Tivoli Storage Manager operations, even if the operations are scheduled.

After configuring an LDAP directory server for password authentication, you can lock nodes to force them to use passwords that authenticate with an LDAP server.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node belongs.

Syntax

```

lock node [*node_name] AUTHentication=[Local|LDap]

```

Parameters

node_name

Specifies the name of the client node to lock out. You can use a wildcard character instead of a node name if you want to lock all of the nodes according to their method of authentication.

AUTHentication

Specifies the method of password authentication that is needed to log into a node.

Local

Specifies to lock nodes that authenticate with the Tivoli Storage Manager server.

LDap

Specifies to lock nodes that authenticate with an LDAP directory server.

Example: Lock a specific client node

Lock the client node SMITH.

```
lock node smith
```

Example: Lock all nodes that authenticate to the local Tivoli Storage Manager database

Issue the following command to lock all nodes that authenticate with the Tivoli Storage Manager server:

```
lock node * authentication=local
```

Related commands

Table 187. Commands related to LOCK NODE

Command	Description
QUERY NODE	Displays partial or complete information about one or more clients.

LOCK NODE

Table 187. Commands related to LOCK NODE (continued)

Command	Description
UNLOCK NODE	Enables a locked user in a specific policy domain to access the server.

LOCK PROFILE (Lock a profile)

Use this command on a configuration manager to temporarily lock a profile so that configuration information is not distributed to subscribing managed servers.

You can use this command when you are making multiple updates to your configuration and do not want to distribute this information until the changes are completed.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

▶▶ LOCK PROFILE profile_name [ minutes ]

```

Parameters

profile_name (Required)

Specifies the profile to lock. You can use wildcard characters to indicate multiple names.

minutes

Specifies the time, in minutes, before Tivoli Storage Manager unlocks the configuration profile. Specify an integer from 0 to 10000. The default is 60 minutes. If you specify 0, the configuration profile will not unlock automatically. Use the **UNLOCK PROFILE** command to unlock the profile before the time period elapses, or to unlock it if you have specified a value of 0. This parameter is optional.

Example: Lock a profile for a specific amount of time

Lock a profile named DELTA for 30 minutes.

```
lock profile delta 30
```

Related commands

Table 188. Commands related to LOCK PROFILE

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE PROFILE	Deletes a profile from a configuration manager.
QUERY PROFILE	Displays information about configuration profiles.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.

LOCK PROFILE

Table 188. Commands related to LOCK PROFILE (continued)

Command	Description
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.

MACRO (Invoke a macro)

Use this command to invoke a file from the administrative command line that contains one or more Tivoli Storage Manager administrative commands to be performed.

Restriction: Use this command with administrative command-line clients only.

A macro is a file that contains one or more Tivoli Storage Manager administrative commands. You can only issue a macro from the administrative client in batch or interactive mode. A macro is stored as a file on the administrative client machine (or system). Macros are not distributed across servers and cannot be scheduled on the server.


Creating a macro to enter commands can be helpful when you want to issue commands that are used repeatedly, to issue commands that contain several parameters, or to process related commands in a specific order. After you create a macro, you can update the information it contains and use it again, or you can copy the macro file, make changes to the copy, and then run the copy.

For detailed information on macros, how to use them, and the differences between command scripts and administrative command-line client macros, refer to the *Administrator's Guide*.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

macro_name (Required)

Specifies the name of the macro.

substitution_value

Specifies the value for a substitution variable in a macro. When you use a substitution variable, you can reuse a macro whenever you need to perform the same task for different objects or with different parameter values. To specify a value that contains blanks, you must enclose the value in quotation marks. This parameter is optional.

Example: Create a macro to register a new administrator

Create a macro file named REGNG. Use the macro to register and grant authority to a new administrator. Write the macro as follows:

```
/* Register and grant authority to a new administrator */
REGister Admin jones passwd      -
CONtactinfo="x1235"              -
GRant AUTHority jones            -
CLasses=Policy
```

MACRO

Issue the following command to run the macro:

```
macro regng.mac
```

Example: Write a macro using substitution variables

Create a macro file named AUTHRG, containing substitution variables, to register and grant authority to a new administrator. Write the macro as follows:

```
/* Register and grant authority to a new administrator */  
REGister Admin %1 %2 - /* Enter userid and password */  
CONtact=%3 /* Enter contact info (in quotes if nec.) */  
GRant AUTHority %1 - /* Server uses variable already */  
- /* defined by you */  
Classes=%4 /* Enter the privilege class */
```

Issue a command similar to the following, entering the values you want to pass to the server to process the command when you run the macro.

```
macro authrg.mac jones passwd x1235 Policy
```

Related commands

Table 189. Commands related to MACRO

Command	Description
COMMIT	Makes changes to the database permanent.
ROLLBACK	Discards any uncommitted changes to the database since the last COMMIT was executed.

MIGRATE STGPOOL (Migrate storage pool to next storage pool)

Use this command to migrate files from one storage pool to the next storage pool in the storage hierarchy.

This command can only be used with primary storage pools. The storage pool data format cannot be NETAPPDUMP, CELERRADUMP, or NDMPDUMP.

Only one migration or reclamation process for a given storage pool is allowed at any given time. If a migration or reclamation process is already running for the storage pool, you cannot start another migration process for the storage pool.

You should only use this command if you are not going to use automatic migration for the storage pool. To prevent automatic migration from running, set the HIGHMIG attribute of the storage pool definition to 100.

If you use this command to start a migration process, but the storage pool does not have a next storage pool identified in the hierarchy, a reclamation process is triggered for the source storage pool. To prevent the reclamation process, define the next storage pool in the hierarchy. For details, see the topic about setting up a storage pool hierarchy in the *IBM Tivoli Storage Manager Administrator's Guide*. Then, start the migration process.

The **MIGRATE STGPOOL** command will honor the values of the following parameters on the **DEFINE STGPOOL** and **UPDATE STGPOOL** commands:

- MIGPROCESS
- MIGDELAY
- MIGCONTINUE
- NEXTPOOL
- LOWMIG

Tip: You can override the value of the **LOWMIG** parameter on **DEFINE STGPOOL** and **UPDATE STGPOOL** by specifying a value for the **LOWMIG** parameter on the **MIGRATE STGPOOL** command.

The **MIGRATE STGPOOL** command will ignore the value of the **HIGHMIG** parameter of the storage pool definition. Migration will occur regardless of the value of the **HIGHMIG** parameter.

This command creates one or more migration processes that can be canceled with the **CANCEL PROCESS** command. The number of processes is limited by the **MIGPROCESS** attribute of the storage pool definition. To display information about background processes, use the **QUERY PROCESS** command.


Remember: Migrating data from a primary storage pool that is set up for data deduplication to another primary storage pool that is also set up for data deduplication removes duplicate data.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for both the storage pool from which the files are to be migrated and the next storage pool to which files are to be migrated.

MIGRATE STGPOOL

Syntax



Parameters

pool_name (Required)

Specifies the primary storage pool from which files are to be migrated.

DUration

Specifies the maximum number of minutes the migration will run before being automatically canceled. When the specified number of minutes elapses, the server will automatically cancel all migration processes for this storage pool. As soon as the processes recognize the automatic cancellation, they will end. As a result, the migration may run longer than the value you specified for this parameter. You can specify a number from 1 to 9999. This parameter is optional. If not specified, the server will stop only after the low migration threshold is reached.

LOWmig

For random-access and sequential-access disk storage pools, specifies that the server stops migration when the amount of data in the pool is at or below this percentage of the pool's estimated capacity. The calculation for sequential-access disk storage pools includes the capacity of all the scratch volumes specified for the pool. For other types of sequential-access storage pools, this parameter specifies that the server stops migration when the ratio of volumes containing data to the total number of volumes in the storage pool is at or below this percentage. The total number of volumes includes the maximum number of scratch volumes. This parameter is optional. You can specify a number from 0 to 99. The default value is the LOWMIG attribute of the storage pool definition.

When the storage pool reaches the low migration threshold, the server does not start migration of another node's files. However, because all file spaces that belong to a node or all file spaces for all nodes belonging to a collocation group are migrated together, the occupancy of the storage pool can fall below the value you specified for this parameter. To permit migration to empty the storage pool, you can set LOWMIG=0.

If the storage pool is already below the low migration threshold, no data will be migrated.

RECLAIM

Specifies whether reclamation is attempted for the storage pool prior to performing the migration. This parameter can only be specified for a sequential-access storage pool. This parameter is optional. The default is No. Possible values are:

No Specifies that the server will not attempt a reclamation prior to starting the migration.

Yes

Specifies that the server will attempt reclamation prior to starting the migration. Any volumes in the storage pool that meet the reclamation threshold as specified by the RECLAIM attribute of the storage pool definition will be reclaimed prior to performing the migration. If no volumes meet the reclamation threshold or if, after reclamation, the LOWMIG threshold has not been reached, the server will begin the migration. Before reclaiming space for storage pools defined with RECLAMATIONTYPE=SNAPLOCK, the server will delete all empty WORM FILE volumes during reclamation processing that have exceeded their reclaim period.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. This default is No. Possible values are:

No Specifies that the server processes this command in the background.

You can continue with other tasks while the command is being processed. Messages created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the **CANCEL PROCESS** command. If you cancel this process, some files may have already been migrated prior to the cancellation.

Yes

Specifies that the server processes this command in the foreground. The operation must complete before you can continue with other tasks. The server then displays the output messages to the administrative client when the operation completes. Messages are also displayed either in the activity log or the server console, or both, depending on where the messages are logged.

Note: You cannot specify WAIT=YES from the server console.

Example: Migrate a storage pool to the next storage pool

Migrate data from the storage pool named BACKUPPOOL to the next storage pool. Specify that the server should end the migration as soon as possible after 90 minutes.

```
migrate stgpool backuppool duration=90
```

Related commands

Table 190. Commands related to MIGRATE STGPOOL

Command	Description
CANCEL PROCESS	Cancels a background server process.
QUERY PROCESS	Displays information about background process.
QUERY STGPOOL	Displays information about storage pools.
RECLAIM STGPOOL	Performs reclamation for the storage pool.

MOVE commands

Use the **MOVE** commands to either transfer backup or archive data between storage pools, or to move disaster recovery media on and off site.

- “MOVE DATA (Move files on a storage pool volume)” on page 611
- “MOVE DRMEDIA (Move disaster recovery media offsite and back onsite)” on page 615
- “MOVE GRPMEMBER (Move a server group member)” on page 630
- “MOVE MEDIA (Move sequential access storage pool media)” on page 631
- “MOVE NODEDATA (Move data by node in a sequential access storage pool)” on page 639

MOVE DATA (Move files on a storage pool volume)

Use this command to move files from one storage pool volume to other storage pool volumes.

You can move files from a primary storage pool volume only to volumes in the same or a different primary storage pool. You can move files from a copy storage pool volume only to volumes in the same copy storage pool. You can move files from an active-data pool volume only to volumes in the same active-data pool.

In addition to moving data from volumes in storage pools that have NATIVE or NONBLOCK data formats, this command also lets you move data from volumes in storage pools that have NDMP data formats (NETAPPDUMP, CELERRADUMP, or NDMPDUMP). The target storage pool must have the same data format as the source storage pool. If you are moving data out of a storage pool for the purpose of upgrading to new tape technology, the target primary storage pool must be associated with a library that has the new device for the tape drives. Tivoli Storage Manager supports backend data movement for NDMP images. For details, see the *Administrator's Guide*.

If you are moving files to volumes in the same storage pool, sufficient space must be available on the volumes. Otherwise, the operation fails.

When you move files from a sequential access volume, multiple sequential access volume mounts are required to move files that span volumes.

When you move files from a random access volume, the server erases any cached copies of files on the volume.

After a move data operation completes, a volume may not be empty if one or more files cannot be relocated to another volume because of input/output errors on the device or because errors were found in the file. If needed, you can delete the volume using the option to discard any data. The files with I/O or other errors are then deleted.

You can use this command to move files from an offsite volume in a copy storage pool or active-data pool. Because the offsite volume cannot be mounted, the server obtains the files that are on the offsite volume from either a primary storage pool, another copy storage pool, or another active-data pool. These files are then written to the destination volumes in the original copy storage pool or active-data pool.

Do not use the MOVE DATA command if a restore process (**RESTORE STGPOOL** or **RESTORE VOLUME**) is running. The **MOVE DATA** command could cause the restore to be incomplete. If you issue the **MOVE DATA** command during a restore operation and you receive an error message indicating that one or more files are locked and cannot be moved, you need to reissue the **MOVE DATA** command after the restore operation has completed in order to move any remaining files.

Remember:


Issuing this command removes duplicate data when:

- Moving data from a primary storage pool that is set up for data deduplication to another primary storage pool that is also set up for data deduplication.
- Moving data within a copy storage pool that is set up for data deduplication.
- Moving data within an active-data pool that is set up for data deduplication.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to which the volume belongs and also for the new storage pool, if one is specified.

Syntax



Notes:

- 1 The default is NO if either the source or target storage pool is random access. The default is YES if both the source and target storage pools are sequential access.
- 2 This parameter is not available or is ignored if the data format is NETAPPDUMP, CELERRADUMP, or NDMPDUMP data.

Parameters

volume_name (Required)

Specifies the storage pool volume from which to move files.

STGpool

Specifies the primary storage pool to which you want to move files (the target storage pool). This parameter is optional and applies only to moving data from primary storage pool volumes. If you do not specify a value for this parameter, files are moved to other volumes within the same storage pool.

SHREDTONOshred

Specifies whether data will be moved from a storage pool that enforces shredding to a storage pool that does not enforce shredding. This parameter is optional. The default value is NO. Possible values are:

No Specifies that the server will not allow data to be moved from a storage pool that enforces shredding to a storage pool that does not enforce shredding. If the source storage pool enforces shredding and the target storage pool does not, the operation will fail.

Yes

Specifies that the server will allow data to be moved from a storage pool that enforces shredding to a storage pool that does not enforce shredding. The source data will be shredded when the operation is complete. The target data will not be shredded when it is deleted.

RECONSTRUCT

Specifies whether to reconstruct file aggregates during data movement. Reconstruction removes empty space that has accumulated during deletion of logical files from an aggregate. This parameter is optional. If both the source and target storage pools are sequential access, the default value is YES. If either the source or target storage pool is random access, the default is NO.

The parameter is not available or is ignored if any of the following conditions are true:

- The data format is NETAPPDUMP, CELERRADUMP, or NDMPDUMP.
- The data is in a storage pool that is configured for data deduplication.
- The target storage pool for the data movement is configured for data deduplication.

Attention: Reconstruction removes inactive backup files in active-data pools. If you specify RECONSTRUCT=NO when moving the data in an active-data pool that is not configured for data deduplication, inactive backup files remain in the storage pool.

Possible values are:

No Specifies that reconstruction of file aggregates is not performed during data movement.

Yes

Specifies that reconstruction of file aggregates is performed during data movement. You can only specify this option when both the source and the target storage pools are sequential-access.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is No. Possible values are:

No Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed.

The server displays messages that are created from the background process either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the **CANCEL PROCESS** command. If a MOVE DATA background process is canceled, some files may have already moved prior to the cancellation.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify WAIT=YES from the server console.

Example: Move files on a storage pool volume

Move files from storage pool volume STGVOL.1 to any available volumes assigned to the 8MMPPOOL storage pool.

```
move data stgvol.1 stgpool=8mmpool
```

Related commands

Table 191. Commands related to MOVE DATA

Command	Description
CANCEL PROCESS	Cancels a background server process.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DELETE VOLUME	Deletes a volume from a storage pool.
MOVE DRMEDIA	Moves DRM media on-site and off-site.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY PROCESS	Displays information about background processes.
QUERY SHREDSTATUS	Displays information about data waiting to be shredded.
SHRED DATA	Manually starts the process of shredding deleted data.

MOVE DRMEDIA (Move disaster recovery media offsite and back onsite)

Use this command to track database backup and copy storage pool volumes that are to be moved offsite and to identify the expired or empty volumes that are to be moved onsite.

The database backup volumes can be for full plus incremental or snapshot backups. You cannot specify virtual volumes (backup objects stored on another server). You can change volumes through each state, or you can use the TOSTATE parameter and skip states to simplify the movements.

Remember: The **MOVE DRMEDIA** command always processes copy storage-pool volumes. (For more information, see the description of the **COPYSTGPOOL** parameter on this command). By default, volumes in active-data pools are not eligible for processing by the disaster recovery manager. To process active-data pool volumes, you must issue the **SET DRMACTIVEDATASTGPOOL** command, or you must use the **ACTIVEDATASTGPOOL** parameter on the **MOVE DRMEDIA** command. To control whether the command processes database backup volumes, you can use the **SOURCE** parameter on this command.

You can use the **QUERY ACTLOG** command to see if the **MOVE DRMEDIA** command was successful. You can also view this information from the server console.


Restriction: Do not run the **MOVE DRMEDIA** and **BACKUP STGPOOL** commands concurrently. Ensure that the storage pool backup processes are complete before you issue the **MOVE DRMEDIA** command.

Privilege class


To issue this command, you must have one of the following privilege classes:

- If the **CMD** parameter is specified and the **REQSYSAUTHOUTFILE** server option is set to **NO**: operator, unrestricted storage, or system privilege.
- If the **CMD** parameter is specified and the **REQSYSAUTHOUTFILE** server option is set to **YES** (the default): system privilege.

Syntax



MOVE DRMEDIA



Parameters

volume_name (Required)

Specifies the name of the database backup or copy storage pool volume to be processed. If you use wildcard characters to specify this name, you must also specify `WHERESTATE`.

`WHERESTATE`

Specifies the state of volumes to be processed. This parameter is required if the `TOSTATE` is not specified or if you use a wildcard character in the volume name. For more information, see Table 196 on page 625 and Table 197 on page 626. Possible values are:

MUnmountable

These volumes contain valid data and are available for onsite processing. They will change to `NOTMOUNTABLE` if the `TOSTATE` is not specified.

Depending on the behavior of the `REMOVE` parameter, Tivoli Storage Manager might eject volumes in an automated library before changing the destination state.

For external libraries, the server sends requests to the external library manager to eject the volumes. It depends on the external library manager as to whether or not the volumes are ejected from the library. Refer to the external library documentation for information about the procedures that you should follow when using the **MOVE DRMEDIA** command to track the volumes.

NOTMUnmountable

These volumes are onsite, contain valid data, and are not available for onsite processing. They will change to `COURIER` if the `TOSTATE` is not specified.

COUrier

These volumes are with the courier and being moved offsite. They will change only to VAULT.

VAULTRetrieve

These volumes are located at the offsite vault and do not contain valid data. They will change to COURIERRETRIEVE if the TOSTATE is not specified.

COURIERRetrieve

These volumes are with the courier and being moved onsite. They will change only to ONSITERETRIEVE. Tivoli Storage Manager deletes from the database the volume records of the database backup and scratch copy storage pool volumes.

BEGINDate

Specifies the beginning date used to select volumes. This parameter is optional. Volumes are considered eligible if the **MOVE DRMEDIA** command has changed the volume to its current state on or after the specified date. The default is the earliest date for which volume information exists.

You can specify the date using one of the following values:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998
TODAY	The current date	TODAY
<i>TODAY-days or -days</i>	The current date minus days specified	TODAY-7 or -7. To identify volumes that were changed to their current state a week ago, you can specify TODAY-7 or simply -7.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
<i>EOLM-days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
<i>BOTM+days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDDate

Specifies the ending date used to select volumes. This parameter is optional. Volumes are considered eligible if the **MOVE DRMEDIA** command has changed the volume to its current state on or before the specified date. The default is the current date.

You can specify the date using one of the following values:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998

MOVE DRMEDIA

Value	Description	Example
TODAY	The current date	TODAY To identify volumes that were changed to their current state today, specify TODAY.
TODAY-days or -days	The current date minus days specified. The maximum number of days is 9999.	TODAY-1 or -1. To identify volumes that were changed to their current state a week ago, you can specify TODAY-1 or simply -1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time used to select volumes for processing. This parameter is optional. Volumes are considered eligible if the **MOVE DRMEDIA** command has changed the volume to its current state on or after the specified time and date. The default is midnight (00:00:00) on the date specified with the **BEGINDATE** parameter.

You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified begin date	12:33:28
NOW	The current time on the specified begin date	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified begin date	NOW+03:00 or +03:00.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified begin date	NOW-03:30 or -03:30. If you issue the MOVE DRMEDIA command at 9:00 with BEGINTIME=NOW-03:30 or BEGINTIME= -03:30 , Tivoli Storage Manager identifies the volumes that were changed to their current state at 5:30 on the begin date you specify.

ENDTime

Specifies the ending time used to select volumes for processing. This parameter

is optional. Volumes are considered eligible if the **MOVE DRMEDIA** command has changed the volume to its current state on or after the specified time and date. The default is 23:59:59.

You can specify the time using one of the following values:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time on the specified end date	12:33:28
NOW	The current time on the specified end date	NOW
<i>NOW+HH:MM</i> <i>or +HH:MM</i>	The current time plus hours and minutes on the specified end date	<i>NOW+03:00 or +03:00.</i> If you issue the MOVE DRMEDIA command at 9:00 with <i>ENDTIME=NOW+03:30</i> or <i>ENDTIME=+03:30</i> , Tivoli Storage Manager identifies the volumes that were changed to their current state at 12:30 on the end date you specify.
<i>NOW-HH:MM</i> <i>or -HH:MM</i>	The current time minus hours and minutes on the specified end date	<i>NOW-03:30 or -03:30.</i>

COPYstgpool

Specifies the name of the copy storage pool whose volumes are to be processed. This parameter is optional. You can use wildcard characters to specify this name. If you use wildcard characters to specify this name, you must also specify **WHERESTATE**.

The copy storage pools specified with this parameter override those specified with the **SET DRMCOPYSTGPOOL** command. If this parameter is not specified, Tivoli Storage Manager selects the storage pools as follows:

- If the **SET DRMCOPYSTGPOOL** command was previously issued with valid copy storage pool names, Tivoli Storage Manager processes only those storage pools.
- If the **SET DRMCOPYSTGPOOL** command has not been issued, or if all of the copy storage pools have been removed using the **SET DRMCOPYSTGPOOL** command, Tivoli Storage Manager processes all copy storage pool volumes in the specified state (**MOUNTABLE**, **NOTMOUNTABLE**, **COURIER**, **VAULTRETRIEVE**, or **COURIERRETRIEVE**).

Source

Specifies whether to include database backup volumes for processing. This parameter is optional. The default is **DBBACKUP**. The choices are:

DBBackup

Specifies that Tivoli Storage Manager includes full and incremental database backup volumes for processing.

DBSnapshot

Specifies that Tivoli Storage Manager includes database snapshot backup volumes for processing.

DBNone

Specifies that Tivoli Storage Manager does not include any database backup volumes for processing.

ACTIVEDatastgpool

Specifies the name of the active-data pool whose volumes are to be processed.

MOVE DRMEDIA

This parameter is optional. You can use wildcard characters to specify this name. If you use wildcard characters to specify this name, you must also specify the `WHERESTATE` parameter.

The active-data pools that are specified with this parameter override those that are specified with the `SET DRMACTIVEDATASTGPOOL` command. If this parameter is not specified, Tivoli Storage Manager selects the storage pools in the following way:

- If the `SET ACTIVEDATASTGPOOL` command was previously issued with valid active-data pool names, Tivoli Storage Manager processes only those storage pools.
- If the `SET ACTIVEDATASTGPOOL` command has not been issued, or all of the active-data pools have been removed using the `SET ACTIVEDATASTGPOOL` command, Tivoli Storage Manager processes all active-data pool volumes in the specified state (`NOTMOUNTABLE`, `COURIER`, `VAULTRETRIEVE`, or `COURIERRETRIEVE`). Volumes in the `MOUNTABLE` state are not processed.

Remember: You can use the `CMDFILE` parameter with the `MOVE DRMEDIA` command to create a file of executable commands to process active-data pool volumes. For example, you can use `CMDFILE` with the `MOVE DRMEDIA * WHERESTATE=COURIERRETRIEVE TOSTATE=ONSITERETRIEVE` command to generate the `CHECKIN LIBVOL` commands to insert tape volumes into the library.

REMove

Specifies that the Tivoli Storage Manager server tries to move the volume out of the library and into the convenience I/O station or entry/exit ports. This parameter is optional. Possible values are `YES`, `BULK`, `NO`, and `UNTILEEFULL`. The default is `BULK`. The response of the server to each of those options and the default values are described in the following tables.

349X libraries: The following table shows how the server responds for 349X libraries.

Table 192. How the Tivoli Storage Manager Server Responds for 349X Libraries

REMOVE=YES	REMOVE=BULK	REMOVE=NO
The 3494 Library Manager ejects the cartridge to the convenience I/O station.	The 3494 Library Manager ejects the cartridge to the high-capacity output facility.	The 3494 Library Manager does not eject the volume. The server leaves the cartridge in the library in the <code>INSERT</code> category for use by other applications.

Table 193. How the Tivoli Storage Manager server responds for SCSI libraries

If a library . . .	And REMOVE=YES, then...	And REMOVE= BULK, then...	And REMOVE=NO, then...	And REMOVE= UNTILEEFULL then...
<i>Does not have entry/exit ports</i>	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server then prompts you to remove the cartridge from the slot and to issue a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.
<i>Has entry/exit ports and an entry/exit port is available</i>	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server then prompts you to remove the cartridge from the slot and to issue a REPLY command.	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server does not prompt you to remove the cartridge and does not request a REPLY command.	The server specifies the port address in a message. The server does not prompt you to remove the cartridge and does not request a REPLY command.	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server does not prompt you to remove the cartridge and does not request a REPLY command.
<i>Has entry/exit ports, but no ports are available</i>	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server then prompts you to remove the cartridge from the slot and to issue a REPLY command.	The server waits for a port to be made available.	The server specifies the port address in a message. The server does not prompt you to remove the cartridge and does not request a REPLY command.	The command fails and any remaining eligible volumes are not processed. Make the port available and issue the command again.

Table 194. How the Tivoli Storage Manager server responds for ACSLS libraries

REMOVE=YES or REMOVE=BULK	REMOVE=NO
The server ejects the cartridge to the convenience I/O station. The server then deletes the volume entry from the server library inventory. While moving volumes from the MOUNTABLE state with REMOVE=YES specified, the MOVE MEDIA command uses more than one slot in the CAP for a StorageTek library with ACSLS.	The server does not eject the cartridge. The server deletes the volume entry from the server library inventory and leaves the volume in the library.

Table 195. How the Tivoli Storage Manager server responds for external libraries

REMOVE=YES, REMOVE=BULK, or REMOVE=NO

The server requests the external library manager to eject the volume from the library.

It depends on the external library manager as to whether or not the volume is ejected from the library. Refer to the external library documentation for information about the procedures that you should follow when using the **MOVE DRMEDIA** command to track volumes.

TOSTate

Specifies the destination state of the volumes being processed. This parameter is required if **WHERESTATE** is not specified. If you specify **TOSTATE** but not **WHERESTATE**, you must specify the volume name. Wildcard characters are not allowed. For more information, see Table 196 on page 625 and Table 197 on page 626.

Possible values are:

NOTMOUNTABLE

Specifies that volumes are to change to the **NOTMOUNTABLE** state. This value is valid only if the volumes are in the **MOUNTABLE** state.

If volumes are in an automated library, Tivoli Storage Manager might eject, depending on the behavior of the **REMOVE** parameter, the volumes from the library before changing them to the **NOTMOUNTABLE** state.

For external libraries, the server sends requests to the external library manager to eject the volumes. It depends on the external library manager as to whether or not the volumes are ejected from the library. Refer to the external library documentation for information about the procedures that you should follow when using the **MOVE DRMEDIA** command to track the volumes.

COURIER

Specifies that volumes are to change to the **COURIER** state. This value is valid only if the volumes are in the **MOUNTABLE** or **NOTMOUNTABLE** state.

Depending on the behavior of the **REMOVE** parameter and whether volumes are in an automated library, Tivoli Storage Manager might eject the volumes from the library before changing them to the **COURIER** state.

For external libraries, the server sends requests to the external library manager to eject the volumes. It depends on the external library manager as to whether or not the volumes are ejected from the library. Refer to the external library documentation for information about the procedures that you should follow when using the **MOVE DRMEDIA** command to track the volumes.

VAULT

Specifies that volumes are to change to the **VAULT** state. This value is valid only if the volumes are in the **MOUNTABLE**, **NOTMOUNTABLE**, or **COURIER** state.

Depending on the behavior of the **REMOVE** parameter and whether volumes are in an automated library, Tivoli Storage Manager might eject the volumes from the library before changing them to the **VAULT** state.

For external libraries, the server sends requests to the external library manager to eject the volumes. It depends on the external library manager as to whether or not the volumes are ejected from the library. Refer to the

external library documentation for information about the procedures that you should follow when using the **MOVE DRMEDIA** command to track the volumes.

COURIERRetrieve

Specifies that volumes are to change to the COURIERRETRIEVE state. This value is valid only if the volumes are in the VAULTRETRIEVE state.

ONSITERetrieve

Specifies that volumes are to change to the ONSITERETRIEVE state. This value is valid only if the volumes are in the VAULTRETRIEVE or COURIERRETRIEVE state. For database backup and scratch copy storage pool volumes that are changing to the ONSITERETRIEVE state, Tivoli Storage Manager deletes the volume records from the database.

WHERELocation

Specifies the current location of the volumes. This parameter is optional. The maximum length of the location is 255 characters. Enclose the text in quotation marks if it contains any blank characters.

TOLocation

Specifies the destination location of the volumes. This parameter is optional. The maximum length of the location specified is 255 characters. Enclose the text in quotation marks if it contains any blank characters. If you do not specify the destination location, the location defined by the **SET DRMNOTMOUNTABLE** command is used.

CMd

Specifies a command to be issued for each volume that is processed by the MOVE DRMEDIA command. DRM writes the commands to a file that is specified by the CMDFILENAME parameter. After the MOVE DRMEDIA operation is completed, the commands in the file can be issued. The command can contain up to 255 characters. If the command contains more than 240 characters, it is split into multiple lines, and continuation characters (+) are added. You might need to alter the continuation character based on the operating system. This parameter is optional.

command

The command string enclosed in quotation marks. The string must not include embedded quotation marks. For example, this is a valid CMD parameter:

```
cmd="checkin libvol lib8mm &vol status=scratch"
```

This is an example of a CMD parameter that is *not* valid:

```
cmd=""checkin libvol lib8mm" &vol status=scratch"
```

The command can include substitution variables. The variables are not case-sensitive, and must not contain blank spaces after the ampersand (&). The possible variables are:

&VOL

A volume name.

&LOC

A volume location.

&VOLDSN

The file name to be written into the sequential access media labels. For example, if the applicable device class sets TSM as the tape volume

prefix, a copy storage pool tape volume file name might be TSM.BFS and a database backup tape volume file name might be TSM.DBB.

&NL

The new line character. When &NL is specified, the command is split at the &NL variable. If required, you must specify the appropriate continuation character before the &NL. If the &NL is not specified and the command line is greater than 240 characters, the line is split into multiple lines and continuation characters (+) are added.

CMDFilename

Specifies the fully qualified name of the file that will contain the commands specified by CMD parameter. This parameter is optional.

If you do not specify a file name or if you specify a null string (""), DRM uses the file name specified by the SET DRMCMDFILENAME. If you do not specify a file name with the **SET DRMCMDFILENAME** command, DRM generates a file name by appending exec.cmds to the directory path name of the current working directory of the Tivoli Storage Manager server.

If the operation fails after the command file has been created, the file is not deleted.

APPend

Specifies whether to overwrite any existing contents of the command file or append the commands to the file. This parameter is optional. The default is NO. Possible values are:

No DRM overwrites the contents of the file.

Yes

DRM appends the commands to the file.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is NO. Possible values are:

No Specifies that Tivoli Storage Manager processes this command in the background.

Messages created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

To see if the operation was successful, issue the **QUERY ACTLOG** command.

Yes

Specifies that the server processes this command in the foreground. Wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client.

Restriction: You cannot specify WAIT=YES from the server console.

CAP

Specifies which cartridge access port (CAP) to use for ejecting volumes if you specify REMOVE=YES. This parameter applies to volumes in ACSLS libraries only. If the CAP priority value is set to 0 in the library, this parameter is required. If a CAP priority value greater than 0 is set in the library, this parameter is optional. By default, all CAPs initially have a priority value of 0, which means that ACSLS does not automatically select the CAP.

To display valid CAP identifiers (x,y,z), issue the **QUERY CAP** command with ALL specified from the Automated Cartridge System System Administrator (ACSSA) console on the ACSLS server host. The identifiers are as follows:

- x The Automated Cartridge System (ACS) ID. This identifier can be a number in the range 0 - 126.
- y The Library Storage Module (LSM) ID. This identifier can be a number in the range 0 - 23.
- z The CAP ID. This identifier can be a number in the range 0 - 11.

For more information, see the StorageTek documentation.

Rules for destination states and destination locations

The following table shows how DRM determines the destination state and location of a volume.

Destination state

- TOSTATE specified.
- The next state of the WHERESTATE specified, if the TOSTATE is not specified.

Destination location

- TOLOCATION specified.
- The location of the TOSTATE specified, if the TOLOCATION is not specified.
- The location of the next state of the WHERESTATE specified, if the TOLOCATION and TOSTATE are not specified.

Table 196. Volume destination and location

Parameters specified	Destination state	Destination location
WHERESTATE	The next state of the WHERESTATE	Location of the next state
WHERESTATE TOSTATE	TOSTATE	Location of the TOSTATE
WHERESTATE TOLOCATION	The next state of the WHERESTATE	TOLOCATON
WHERESTATE TOSTATE TOLOCATION	TOSTATE	TOLOCATION
TOSTATE	TOSTATE	Location of the TOSTATE
TOSTATE WHERELOCATION	TOSTATE	Location of the TOSTATE
TOSTATE WHERELOCATION TOLOCATION	TOSTATE	TOLOCATION

Rules for state transitions

The following tables show the state transitions that volumes are eligible for, based on their current state.

Table 197. State transitions for volumes

Volume's current state	Destination state			
	MOUNTABLE	NOTMOUNTABLE	COURIER	VAULT
MOUNTABLE	N	Y	Y	Y
NOTMOUNTABLE	N	N	Y	Y
COURIER	N	N	N	Y
VAULT	N	N	N	N
VAULTRETRIEVE	N	N	N	N
COURIERRETRIEVE	N	N	N	N
ONSITERETRIEVE	N	N	N	N

Volume's current state	Destination state		
	VAULT-RETRIEVE	COURIER-RETRIEVE	ONSITE-RETRIEVE
MOUNTABLE	N	N	N
NOTMOUNTABLE	N	N	N
COURIER	N	N	N
VAULT	N	N	N
VAULTRETRIEVE	N	Y	Y
COURIERRETRIEVE	N	N	Y
ONSITERETRIEVE	N	N	N

Example: Move disaster recovery media from the NOTMOUNTABLE state

Move disaster recovery media that is in the NOTMOUNTABLE state to the COURIER state and then query the results.

```
move drmedia * wherestate=notmountable
tostate=courier
```

```
query actlog search="MOVE DRMEDIA"
```

```
08/11/1999 11:12:24 ANR0984I Process 10 for MOVE DRMEDIA started
in the BACKGROUND at 11:12:24.
08/11/1999 11:12:24 ANR0610I MOVE DRMEDIA started by HSIAO as
process 10.
08/11/1999 11:12:25 ANR6683I MOVE DRMEDIA: Volume TAPE0P was moved
from NOTMOUNTABLE state to COURIER.
08/11/1999 11:12:25 ANR6683I MOVE DRMEDIA: Volume TAPE1P was moved
from NOTMOUNTABLE state to COURIER.
08/11/1999 11:12:25 ANR6683I MOVE DRMEDIA: Volume DBTP02 was moved
from NOTMOUNTABLE state to COURIER.
08/11/1999 11:12:25 ANR6683I MOVE DRMEDIA: Volume DBTP01 was moved
from NOTMOUNTABLE state to COURIER.
08/11/1999 11:12:25 ANR6682I MOVE DRMEDIA command ended: 4 volumes
processed.
08/11/1999 11:12:25 ANR0611I MOVE DRMEDIA started by HSIAO as
process 10 has ended.
08/11/1999 11:12:25 ANR0985I Process 10 for MOVE DRMEDIA running in
the BACKGROUND processed 4 items with a
completion state of SUCCESS at 11:12:25.
```

Example: Move disaster recovery media from the MOUNTABLE state

Move disaster recovery media from the MOUNTABLE state to the COURIER state. If the media is in an automated library, MOVE DRMEDIA ejects the media before changing the state.

```
move drmedia * wherestate=mountable tostate=courier wait=yes
```

```
ANR0984I Process 12 for MOVE DRMEDIA started
in the FOREGROUND at 09:57:17.
ANR0609I MOVE DRMEDIA started as process 12.
ANR0610I MOVE DRMEDIA started by HSIAO as
process 12.
ANR6696I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
volume TAPE01 in library LIB8MM starting.
ANR6697I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
volume TAPE01 in library LIB8MM completed
successful.
ANR6683I MOVE DRMEDIA: Volume TAPE01 was moved
from MOUNTABLE state to COURIER.
ANR6696I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
volume TAPE02 in library LIB8MM starting.
ANR6697I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
volume TAPE02 in library LIB8MM completed
successful.
ANR6683I MOVE DRMEDIA: Volume TAPE02 was moved
from MOUNTABLE state to COURIER.
ANR6696I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
volume DBTP05 in library LIB8MM starting.
ANR6697I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
volume DBTP05 in library LIB8MM completed
successful.
ANR6683I MOVE DRMEDIA: Volume DBTP05 was moved
from MOUNTABLE state to COURIER.
ANR6696I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
volume DBTP04 in library LIB8MM starting.
ANR6697I MOVE DRMEDIA: CHECKOUT LIBVOLUME for
volume DBTP04 in library LIB8MM completed
successful.
ANR6683I MOVE DRMEDIA: Volume DBTP04 was moved
from MOUNTABLE state to COURIER.
ANR6682I MOVE DRMEDIA command ended: 4 volumes
processed.
ANR0611I MOVE DRMEDIA started by HSIAO as
process 12 has ended.
ANR0985I Process 12 for MOVE DRMEDIA running
in the FOREGROUND processed 4 items with a
completion state of SUCCESS at 10:12:25.
```

Example: Move disaster recovery media from the VALUTRETRIEVE state

Move disaster recovery media that is in the VAULTRETRIEVE state to the ONSITERETRIEVE state. Generate a **CHECKIN LIBVOLUME** command for each volume that is successfully processed and store the commands in the file:

```
/drm/move/exec.cmds
```

```
move drmedia * wherestate=vaultretrieve tostate=onsiteretrieve
cmdfilename=/drm/move/exec.cmds
cmd="checkin libvol lib8mm &vol status=scratch"
```

Query the results:

```
query actlog search="MOVE DRMEDIA"
```

MOVE DRMEDIA

```
08/13/1999 09:12:24 ANR0984I Process 15 for MOVE DRMEDIA started in
the BACKGROUND at 09:12:24.
08/13/1999 09:12:24 ANR0610I MOVE DRMEDIA started by HSIAO as
process 15.
08/13/1999 09:12:24 ANR6684I MOVE DRMEDIA: Volume CSTEP01 was deleted.
08/13/1999 09:12:24 ANR6684I MOVE DRMEDIA: Volume CSTEP02 was deleted.
08/13/1999 09:12:24 ANR6684I MOVE DRMEDIA: Volume DBTP10 was deleted.
08/13/1999 09:12:24 ANR6684I MOVE DRMEDIA: Volume DBTP11 was deleted.
08/13/1999 09:12:27 ANR6682I MOVE DRMEDIA command ended: 4 volumes
processed.
08/13/1999 09:12:42 ANR0611I MOVE DRMEDIA started by HSIAO as process
15 has ended.
08/13/1997 09:12:42 ANR0985I Process 15 for MOVE DRMEDIA running in
the BACKGROUND processed 4 items with a
completion state of SUCCESS at 09:12:42.
```

The volume check-in commands were also created in the file:

```
/drm/move/exec.cmds
```

The file contains these lines:

```
checkin libvol lib8mm CSTEP01 status=scratch
checkin libvol lib8mm CSTEP02 status=scratch
checkin libvol lib8mm DBTP10 status=scratch
checkin libvol lib8mm DBTP11 status=scratch
```

Remember: To process the **CHECKIN LIBVOLUME** commands, issue the **MACRO** command with the file name as the macro name.

Related commands

Table 198. Commands related to MOVE DRMEDIA

Command	Description
BACKUP DB	Backs up the Tivoli Storage Manager database to sequential access volumes.
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
CANCEL PROCESS	Cancels a background server process.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DISMOUNT VOLUME	Dismounts a sequential, removable volume by the volume name.
PREPARE	Creates a recovery plan file.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY DRMSTATUS	Displays DRM system parameters.
QUERY PROCESS	Displays information about background processes.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
SET DRMACTIVEDATASTGPOOL	Specifies that active-data storage pools are managed by DRM.
SET DRMCOPYSTGPOOL	Specifies that copy storage pools are managed by DRM.

Table 198. Commands related to MOVE DRMEDIA (continued)

Command	Description
SET DRMCOURIERNAME	Specifies the name of the courier for the disaster recovery media.
SET DRMDBBACKUPEXPIREDAYS	Specifies criteria for database backup series expiration.
SET DRMVAULTNAME	Specifies the name of the vault where DRM media is stored.
SET DRMCMDFILENAME	Specifies a file name for containing DRM executable commands.
SET DRMFILEPROCESS	Specifies whether the MOVE DRMEDIA or QUERY DRMEDIA command processes files associated with a device type of file.
SET DRMNOTMOUNTABLENAME	Specifies the location name of the DRM media to be sent offsite.

MOVE GRPMEMBER (Move a server group member)

Use this command to move a member from one server group to another server group. The command fails if the member you are moving has the same name as a current member of the group.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—MOVE GRPMEMBER—*member_name*—*from_group*—*to_group*—◄◄

Parameters

member_name **(Required)**

Specifies the member (a server or a server group) to move.

from_group **(Required)**

Specifies the server group with which the member is currently associated.

to_group **(Required)**

Specifies the new server group for the member.

Example: Move a server to another server group

Move member PAYSON from REGION1 group to REGION2 group.

```
move grpmember payson region1 region2
```

Related commands

Table 199. Commands related to MOVE GRPMEMBER

Command	Description
DEFINE GRPMEMBER	Defines a server as a member of a server group.
DEFINE SERVERGROUP	Defines a new server group.
DELETE GRPMEMBER	Deletes a server from a server group.
DELETE SERVERGROUP	Deletes a server group.
QUERY SERVER	Displays information about servers.
QUERY SERVERGROUP	Displays information about server groups.
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVERGROUP	Updates a server group.

MOVE MEDIA (Move sequential access storage pool media)

Use this command to manage overflow storage pools. The database tracks media that has been moved using this command.

This command applies to sequential access primary and copy storage pool volumes that are managed by an automated library (including an external library). The library does not have to be full. One or more sequential access storage pool volumes can be processed at the same time.

Use the **DAYS** parameter to identify eligible volumes to be moved. Use the **OVERFLOW LOCATION** parameter to record the storage location for the moved media.

This command generates a background process that you can view using the **QUERY PROCESS** command. To cancel, issue the **CANCEL PROCESS** command.

To determine if the command was successful, issue the **QUERY ACTLOG** command or use the server console.

The volumes moved by the **MOVE DRMEDIA** command for offsite recovery are not processed by the **MOVE MEDIA** command.


For additional information see the *Administrator's Guide*.

Privilege class


To issue this command, you must have one of the following privilege classes:

- If the **CMD** parameter is *NOT* specified: Operator or system privilege.
- If the **CMD** parameter is specified and the **REQSYSAUTHOUTFILE** server option is set to **NO**: Operator, unrestricted storage, or system privilege.
- If the **CMD** parameter is specified and the **REQSYSAUTHOUTFILE** server option is set to **YES** (the default): System privilege.

Syntax



MOVE MEDIA



Parameters

volume_name (Required)

Specifies the name of the sequential access primary or copy storage pool volume to be processed. You can use a wildcard character to specify the name. All matching volumes are considered for processing.

STGpool (Required)

Specifies the name of the sequential access primary or copy storage pool that is used to select the volumes for processing. You can use a wildcard character to specify the name. All matching storage pools are processed. If the storage pool specified is not managed by an automated library, no volumes are processed.

Days

Specifies the number of days that must elapse after the volume has been written or read, before the volume is eligible for processing by the command. This parameter is optional. You can specify a number from 0 to 9999. The default value is 0. The most recent of the volumes' last written date or last read date is used to calculate the number of days elapsed.

WHEREState

Specifies the current state of the volumes to be processed. This parameter is used to restrict processing to the volumes that are in the specified state. This parameter is optional. The default value is MOUNTABLEINLIB.

Possible values are:

MOUNTABLEInlib

Specifies that storage pool volumes are to transition from the MOUNTABLEINLIB state to the MOUNTABLENOTINLIB state. Volumes in the MOUNTABLEINLIB state contain valid data and are in the library.

MOUNTABLENotinlib

Specifies that storage pool volumes are to change from the MOUNTABLENOTINLIB state back to the MOUNTABLEINLIB state. Volumes in the MOUNTABLENOTINLIB state may contain valid data and are in the overflow location.

- For *empty scratch volumes*, the **MOVE MEDIA** command deletes the volume records so that they can be used again.
- For *private volumes*, the **MOVE MEDIA** command resets the volume location to blank, changes the volumes' state to CHECKIN, and changes the last update date to the current date.

- For *scratch volumes with data*, the **MOVE MEDIA** command resets the volume location to blank, changes the volumes' state to CHECKIN, and changes the last update date to the current date.

Attention: Volumes in the CHECKIN state may contain valid data and need to be checked into the library.

WHERESTATUS

Specifies that the move process should be restricted by volume status. This parameter is optional. You can specify more than one status in a list by separating each status with a comma and no intervening spaces. If you do not specify this parameter, volumes moved from the MOUNTABLEINLIB state to the MOUNTABLENOTINLIB state are restricted to only full volumes, and volumes moved from the MOUNTABLENOTINLIB state to the MOUNTABLEINLIB state are restricted to only empty volumes.

Possible values are:

FULL

Moves volumes with a status of FULL.

FILLing

Moves volumes with a status of FILLING.

EMPTy

Moves volumes with a status of EMPTY.

ACCess

Specifies how users and system processes access files in the storage pool volume that is moved out from an automated library and stored in an overflow location by the **MOVE MEDIA** command. This parameter is optional. If you do not specify this parameter, moving volumes from the MOUNTABLEINLIB state to the MOUNTABLENOTINLIB process updates the volumes' access mode to READONLY, and moving volumes from the MOUNTABLENOTINLIB state to the MOUNTABLEINLIB process updates the volumes' access mode to READWRITE.

Possible values are:

READWrite

Specifies that users and system processes can read from and write to files stored on the volume that is in the overflow location. If this value is specified, Tivoli Storage Manager requests the volume to be checked into the library when the volume is needed for a read or write operation.

READOnly

Specifies that users and system processes can only read files that are stored on the volume that is in the overflow location. The server requests the volume to be checked into the library only when the volume is needed for a read operation.

OVFL0cation

Specifies the overflow location that is the destination of the volumes being processed. The maximum length of the location name is 255 characters. The location name information must be enclosed in quotation marks if it contains any blank characters. If you do not specify an overflow location and the storage pool also has no overflow location identified, the server changes the location of the ejected volume to a null string ("").

REMove

MOVE MEDIA

Specifies that the Tivoli Storage Manager server tries to move the volume out of the library and into the convenience I/O station or entry/exit ports. This parameter is optional. Possible values are YES, BULK, and NO. The default is BULK. The response of the server to each of those options and the default values are described in the following tables.

349X libraries: The following table shows how the server responds for 349X libraries.

Table 200. How the Tivoli Storage Manager Server Responds for 349X Libraries

REMOVE=YES	REMOVE=BULK	REMOVE=NO
The 3494 Library Manager ejects the cartridge to the convenience I/O station.	The 3494 Library Manager ejects the cartridge to the high-capacity output facility.	The 3494 Library Manager does not eject the volume. The server leaves the cartridge in the library in the INSERT category for use by other applications.

SCSI libraries: The following table shows how the server responds to YES, BULK, and NO for SCSI libraries.

Table 201. How the Tivoli Storage Manager Server Responds for SCSI Libraries

If a library . . .	And REMOVE=YES. . .	And REMOVE=BULK...	And REMOVE=NO
<i>Does not</i> have entry/exit ports	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server then prompts you to remove the cartridge from the slot and issue a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.
<i>Has</i> entry/exit ports and an entry/exit port is available	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server then prompts you to remove the cartridge from the slot and issue a REPLY command.	The server moves the cartridge to the available entry/exit port and specifies the port address in a message. The server does not prompt you to remove the cartridge and does not request a REPLY command.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.
<i>Has</i> entry/exit ports, but no ports are available	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server then prompts you to remove the cartridge from the slot and issue a REPLY command.	The server waits for an entry/exit port to be made available.	The server leaves the cartridge in its current slot within the library and specifies the slot address in a message. The server does not prompt you to remove the cartridge and does not require a REPLY command.

ACSLs libraries: The following table shows how the server responds for ACSLS libraries.

Table 202. How the Tivoli Storage Manager Server Responds for ACSLS Libraries

REMOVE=YES or REMOVE=BULK	REMOVE=NO
The server ejects the cartridge to the convenience I/O station.	The server does not eject the cartridge.
The server then deletes the volume entry from the server library inventory.	The server deletes the volume entry from the server library inventory and leaves the volume in the library.
While moving volumes from the MOUNTABLE state with REMOVE=YES specified, the MOVE MEDIA command uses more than one slot in the CAP for a StorageTek library with ACSLS.	

External libraries: The following table shows how the server responds for external libraries.

Table 203. How the Tivoli Storage Manager Server Responds for External Libraries

REMOVE=YES or REMOVE=BULK	REMOVE=NO
The server ejects the cartridge to the convenience I/O station. The server then deletes the volume entry from the server library inventory.	The server does not eject the cartridge.
	The server deletes the volume entry from the server library inventory and leaves the volume in the library.

Cmd

Specifies the creation of executable commands. This parameter is optional. You must enclose your command specification in quotation marks. The maximum length of the command specification is 255 characters. For each volume successfully processed by the **MOVE MEDIA** command, Tivoli Storage Manager writes the associated commands to a file. Specify the file name with the **CMDFILENAME** parameter.

If you do not specify the file name, the **MOVE MEDIA** command will generate a default filename by appending the string `exec.cmds.media` to the Tivoli Storage Manager server directory.

If the length of the command written to the file exceeds 255 characters, it is split into multiple lines and a continuation character, `+`, is added to all but the last line of the command. You may need to alter the continuation character according to the requirements of the product that runs the commands.

If you do not specify **CMD**, the **MOVE MEDIA** command will not generate any executable commands.

string

Specifies the string to build an executable command. You can specify any free form text for the string. Enclose the full string in quotation marks. For example, the following is a valid executable command specification:

```
CMD="UPDATE VOLUME &VOL"
```

The following is an invalid executable command specification:

```
CMD=""UPDATE VOLUME" &VOL"
```

substitution

Specifies a variable for which you want the command to substitute a value. The possible substitution variables are:

&VOL

Substitute the volume name for **&VOL**. You can specify lowercase

MOVE MEDIA

characters, &vol. No spaces or blanks are allowed between ampersand, &, and VOL. If there are spaces or blanks between ampersand and VOL, the MOVE MEDIA command will treat them as strings and no substitution will be set. If &VOL is not specified, no volume name is set in the executable command.

&LOC

Substitute the volume location for &LOC. You can specify lowercase characters, &loc. No spaces or blanks are allowed between ampersand, &, and LOC. If there are spaces or blanks between ampersand and LOC, the MOVE MEDIA command will treat them as strings and no substitution will be set. If &LOC is not specified, no location name is set in the executable command.

&VOLDSN

Substitute the volume file name for &VOLDSN. An example of a storage pool tape volume file name using the default prefix ADSM is ADSM.BFS. If &VOLDSN is not specified, no volume file name is set in the executable command.

&NL

Substitute a new line character for &NL. When &NL is specified, the **MOVE MEDIA** command splits the command at the position where the &NL is and will not append any continuation character. The user is responsible for specifying the proper continuation character before the &NL if one is required. The user is also responsible for the length of the line written. If the &NL is not specified and the length of the command line exceeds 255, the command line is split into multiple lines and a continuation character, +, is added to all but the last line of the command.

CMDFilename

Specifies the full path name of a file that will contain the commands specified with CMD. This parameter is optional. The maximum length of the file name is 1279 characters.

If you do not specify a file name, the **MOVE MEDIA** command will generate a default file name by appending the string *exec.cmds.media* to the Tivoli Storage Manager server directory. The server directory is the current working directory of the Tivoli Storage Manager server process.

The **MOVE MEDIA** command automatically allocates the file name specified or generated. If the file name exists, you can use the APPEND=YES parameter to add to the file. Otherwise, the file is overwritten. If a file is accidentally overwritten and you need to run the commands that had been in the file, issue the **QUERY MEDIA** command to rebuild the executable commands for the desired volumes. If the **MOVE MEDIA** command fails after the command file is allocated, the file is not deleted.

APPend

Specifies to write at the beginning or ending of the command file data. The default is NO. Possible values are:

No Specifies to write the data from the beginning of the command file. If the given command file exists, its contents are overwritten.

Yes

Specifies to append the command file by writing at the end of the command file data.

CHECKLabel

Specifies whether Tivoli Storage Manager should read volume labels for sequential media. For SCSI devices you can suppress label checking by setting the CHECKLabel to NO. This parameter is not applicable to 349X libraries. This parameter is optional. The default is YES. Possible values are:

Yes

Specifies that Tivoli Storage Manager attempts to read the media label. Reading the media label verifies that the correct volume is being checked out.

No Specifies that Tivoli Storage Manager does not attempt to read media label. This increases performance because the read process does not occur.

CAP

Specifies which cartridge access port (CAP) to use for ejecting volumes if you specify REMOVE=YES. This parameter applies to volumes in ACSLS libraries only. If the CAP priority value is set to 0 in the library, this parameter is required. If a CAP priority value greater than 0 is set in the library, this parameter is optional. By default, all CAPs initially have a priority value of 0, which means that ACSLS does not automatically select the CAP.

To display valid CAP identifiers (*x,y,z*), issue the **QUERY CAP** command with ALL specified from the Automated Cartridge System System Administrator (ACSSA) console on the ACSLS server host. The identifiers are as follows:

- x* The Automated Cartridge System (ACS) ID. This identifier can be a number in the range 0 - 126.
- y* The Library Storage Module (LSM) ID. This identifier can be a number in the range 0 - 23.
- z* The CAP ID. This identifier can be a number in the range 0 - 11.

For more information, see the StorageTek documentation.

Example: Move all full volumes out of the library

Move all full volumes that are in the ARCHIVE sequential primary storage pool out of the library.

```
move media * stgpool=archive
```

Example: Generate the checkin commands

Generate the CHECKIN LIBVOLUME commands for full and partially full volumes that are in the ONSITE.ARCHIVE primary storage pool and stored in the overflow location, Room 2948/Bldg31.

MOVE MEDIA creates the executable commands in /tsm/move/media/checkin.vols

```
move media * stgpool=onsite.archive
wherestate=mountablenotinlib wherestatus=full,filling
ovflocation=room2948/bldg31
cmd="checkin libvol lib3494 &vol status=private"
cmdfilename=/tsm/move/media/checkin.vols
```

```
checkin libvolume lib3494 TAPE04 status=private
checkin libvolume lib3494 TAPE13 status=private
checkin libvolume lib3494 TAPE14 status=private
```

MOVE MEDIA

Tip: Run the **CHECKIN LIBVOLUME** commands by issuing the MACRO command with the following as the macro name:

- /tsm/move/media/checkin.vols

Related commands

Table 204. Commands related to MOVE MEDIA

Command	Description
CANCEL PROCESS	Cancels a background server process.
QUERY MEDIA	Displays information about storage pool volumes moved by the MOVE MEDIA command.
QUERY PROCESS	Displays information about background processes.

MOVE NODEDATA (Move data by node in a sequential access storage pool)

Use this command to move data located in a sequential-access storage pool. You can move data for one or more nodes or for a group of collocated nodes. You can also move selected file spaces for a single node. The data can be located in a primary storage pool, a copy storage pool, or an active-data pool.

This command is helpful for reducing the number of volume mounts during client restore or retrieve operations by consolidating data for a specific node within a storage pool, or to move data to another storage pool. For example, you can use this command for moving data to a random-access storage pool in preparation for client restore processing.

Ensure that the access mode of the volumes *from which you are moving the node data* is read/write or read-only and that the access mode of the volumes *to which you are moving the node data* is set to read/write. This operation will not move data on volumes with access modes of offsite, unavailable, or destroyed.

The **MOVE NODEDATA** command takes two forms, depending on whether you are moving data only for selected filespaces. The syntax and parameters for each form are defined separately.

- “MOVE NODEDATA (Move data in file spaces for one or more nodes or a collocation group)” on page 641
- “MOVE NODEDATA (Move data from selected file spaces of a single node)” on page 644

Restriction: You cannot move node data into or out of a storage pool defined with a CENTERA device class.

Table 205. Commands related to MOVE NODEDATA

Command	Description
CANCEL PROCESS	Cancels a background server process.
COPY ACTIVATEDATA	Copies active backup data.
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node to a collocation group.
DELETE COLLOGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node from a collocation group.
MOVE DATA	Moves data from a specified storage pool volume to another storage pool volume.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY COLLOGROUP	Displays information about collocation groups.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY OCCUPANCY	Displays file space information by storage pool.

MOVE NODEDATA

Table 205. Commands related to MOVE NODEDATA (continued)

Command	Description
QUERY PROCESS	Displays information about background processes.
QUERY STGPOOL	Displays information about storage pools.
QUERY VOLUME	Displays information about storage pool volumes.
UPDATE COLLOCGROUP	Updates the description of a collocation group.


MOVE NODEDATA (Move data in file spaces for one or more nodes or a collocation group)

Use this command to move data in file spaces belonging to one or more nodes or to move data in file spaces belonging to a collocation group. A *collocation group* is a group of nodes whose data is collocated on a minimal number of sequential access volumes.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the source storage pool. If your authorization is restricted storage privilege and you intend to move data to another storage pool, you must also have the appropriate authority for the destination storage pool.

Syntax



Notes:

- 1 The default is NO if either the source or target storage pool is random access. The default is YES if both the source and target storage pools are sequential access.

Parameters

node_name (Required unless the COLLOCGroup parameter is specified)

Specifies the node name related to the data that is moved with this command. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names.

COLLOCGroup (Required unless the node_name parameter is specified)

Specifies the name of the collocation group whose data is to be moved. Data for all the nodes belonging to the collocation group will be moved.

MOVE NODEDATA

FROMstgpool (Required)

Specifies the name of a sequential-access storage pool that contains data to be moved. This storage pool must be in the NATIVE or NONBLOCK data format.

T0stgpool

Specifies the name of a storage pool to which data will be moved. This storage pool must be in the NATIVE or NONBLOCK data format. This parameter is optional and does not apply when the source storage pool is a copy storage pool or an active-data pool. That is, if the source storage pool is a copy storage pool the destination must be the same copy storage pool. Similarly, if the source storage pool is an active-data pool, the destination must be the same active-data pool. If a value is not specified, data is moved to other volumes within the source pool.

Important: If you are moving data within the same storage pool, there must be volumes available that do not contain the node data you are moving. That is, the server cannot use volumes that contain the data to be moved as destination volumes.

Type

Specifies the type of files to be moved. This parameter is optional. The default value is ANY. If the source storage pool is an active-data pool, the only valid values are ANY and BACKUP. However, only the active versions of backup data are moved if TYPE=ANY. Possible values are:

ANY

Specifies that all types of files are moved.

Backup

Specifies that backup files are moved.

ARchive

Specifies that archive files are moved. This value is not valid for active-data pools.

SPacemanaged

Specifies that space-managed files (files that were migrated by a Tivoli Storage Manager for Space Management client) are moved. This value is not valid for active-data pools.

MAXPRocess

Specifies the maximum number of parallel processes to use for moving data. This parameter is optional. You can specify a value from 1–999, inclusive. The default value is 1. Increasing the number of parallel processes should improve throughput.

When determining this value, consider the number of logical and physical drives that can be dedicated to this operation. To access a sequential access volume, Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the move. Each process needs a mount point for storage pool volumes, and, if the device type is not FILE, each process also needs a drive.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is No. Possible values are:

No Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed.

The server displays messages that are created from the background process either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the **CANCEL PROCESS** command. If a background process is canceled, some files may have already been moved before the cancellation.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify WAIT=YES from the server console.

RECONSTRUCT

Specifies whether to reconstruct file aggregates during data movement. Reconstruction removes empty space that has accumulated during deletion of logical files from an aggregate. This parameter is optional. If both the source and target storage pools are sequential access, the default value is YES. If either the source or target storage pool is random access, the default is NO.

The parameter is not available or is ignored if any of the following conditions are true:

- The data format is NETAPPDUMP, CELERRADUMP, or NDMPDUMP.
- The data is in a storage pool that is configured for data deduplication.
- The target storage pool for the data movement is configured for data deduplication.

Attention: Reconstruction removes inactive backup files in active-data pools. If you specify RECONSTRUCT=NO when moving the data in an active-data pool that is not configured for data deduplication, inactive backup files remain in the storage pool.

Possible values are:

No Specifies that reconstruction of file aggregates will not be performed during the move.

Yes

Specifies that reconstruction of file aggregates will be performed during the move. You may only specify this option when both the source and the target storage pools are sequential-access.

Example: Move a specific node's data from a tape storage pool to a disk storage pool

Move all data belonging to node MARY that is stored in storage pool TAPEPOOL. Data should be moved to disk storage pool BACKUPPOOL.

```
move nodedata mary
  fromstgpool=tapepool tostgpool=backuppool
```

MOVE NODEDATA


MOVE NODEDATA (Move data from selected file spaces of a single node)

Use this command to move data for selected file spaces belonging to a single node.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the source storage pool. If your authorization is restricted storage privilege and you intend to move data to another storage pool, you must also have the appropriate authority for the destination storage pool.

Syntax



Notes:

- 1 The default is NO if either the source or target storage pool is random access. The default is YES if both the source and target storage pools are sequential access.

Parameters***node_name* (Required)**

Specifies the node name related to the data that is moved with this command. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names.

FROMstgpool (Required)

Specifies the name of a sequential-access storage pool that contains data to be moved. This storage pool must be in the NATIVE or NONBLOCK data format.

T0stgpool

Specifies the name of a storage pool to which data will be moved. This storage pool must be in the NATIVE or NONBLOCK data format. This parameter is optional and does not apply when the source storage pool is a copy storage pool or an active-data pool. That is, if the source storage pool is a copy storage pool the destination must be the same copy storage pool. Similarly, if the source storage pool is an active-data pool, the destination must be the same active-data pool. If a value is not specified, data is moved to other volumes within the source pool.

Important: If you are moving data within the same storage pool, there must be volumes available that do not contain the node data you are moving. That is, the server cannot use volumes that contain the data to be moved as destination volumes.

FILESpace

Specifies the name of the non-Unicode file space that contains data to be moved. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. This parameter is optional. If you do not specify a value for this parameter and values for UNIFILESPACE or the FSID or both, non-Unicode file spaces are not moved.

UNIFILESpace

Specifies the name of the Unicode file space that contains data to be moved. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. This parameter is optional. If you do not specify a value for this parameter and values for FILESPACE or the FSID or both, non-Unicode file spaces are not moved.

FSID

Specifies file space identifiers (FSIDs) for the file spaces to be moved. Separate multiple names with commas and no intervening spaces. This parameter is optional.

Type

Specifies the type of files to be moved. This parameter is optional. The default value is ANY. If the source storage pool is an active-data pool, the only valid values are ANY and BACKUP. However, only the active versions of backup data are moved if TYPE=ANY. Possible values are:

ANY

Specifies that all types of files are moved.

MOVE NODEDATA

Backup

Specifies that backup files are moved.

ARchive

Specifies that archive files are moved. This value is not valid for active-data pools.

SPacemanaged

Specifies that space-managed files (files that were migrated by a Tivoli Storage Manager for Space Management client) are moved. This value is not valid for active-data pools.

MAXPRocess

Specifies the maximum number of parallel processes to use for moving data. This parameter is optional. You can specify a value from 1–999, inclusive. The default value is 1. Increasing the number of parallel processes should improve throughput.

When determining this value, consider the number of logical and physical drives that can be dedicated to this operation. To access a sequential access volume, Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the move. Each process needs a mount point for storage pool volumes, and, if the device type is not FILE, each process also needs a drive.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is No. Possible values are:

No Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed.

The server displays messages that are created from the background process either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the CANCEL PROCESS command. If a background process is canceled, some files may have already moved before the cancellation.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify WAIT=YES from the server console.

RECONStruct

Specifies whether to reconstruct file aggregates during data movement. Reconstruction removes empty space that has accumulated during deletion of logical files from an aggregate. This parameter is optional. If both the source and target storage pools are sequential access, the default value is YES. If either the source or target storage pool is random access, the default is NO.

The parameter is not available or is ignored if any of the following conditions are true:

- The data format is NETAPPDUMP, CELERRADUMP, or NDMPDUMP.

- The data is in a storage pool that is configured for data deduplication.
- The target storage pool for the data movement is configured for data deduplication.

Attention: Reconstruction removes inactive backup files in active-data pools. If you specify RECONSTRUCT=NO when moving the data in an active-data pool that is not configured for data deduplication, inactive backup files remain in the storage pool.

Possible values are:

No Specifies that reconstruction of file aggregates will not be performed during the move.

Yes

Specifies that reconstruction of file aggregates will be performed during the move. You may only specify this option when both the source and the target storage pools are sequential-access.

Example: Move a node's non-Unicode and Unicode data

Move data for node TOM in storage pool TAPEPOOL. Restrict movement of data to files in non-Unicode file spaces as well as Unicode file spaces, \\jane\d\$. Data should be moved to disk storage pool BACKUPPOOL.

```
move nodedata tom
  fromstgpool=tapepool tostgpool=backuppool
  filespace=* unifilespace=\\jane\d$
```

Example: Move all node data from tape storage pools to a disk storage pool

Move all data for node SARAH, from all primary sequential-access storage pools (for this example, TAPEPOOL*) to DISKPOOL. To obtain a list of storage pools that contain data for node SARAH, issue either of the following **QUERY OCCUPANCY** or **SELECT** commands:

```
query occupancy sarah
SELECT * from OCCUPANCY where node_name='sarah'
```

Attention: For this example assume that the results were TAPEPOOL1, TAPEPOOL4, and TAPEPOOL5.

```
move nodedata sarah
  fromstgpool=tapepool1 tostgpool=DISKPOOL
```

```
move nodedata sarah
  fromstgpool=tapepool4 tostgpool=DISKPOOL
```

```
move nodedata sarah
  fromstgpool=tapepool5 tostgpool=DISKPOOL
```

Example: Move a node's non-Unicode and Unicode file spaces

The following is an example of moving non-Unicode and Unicode file spaces for a node. For node NOAH move non-Unicode filespace \\servtuc\d\$ and Unicode file space \\tsmserv1\e\$ that has a filespace ID of 2 from sequential access storage pool TAPEPOOL to random access storage pool DISKPOOL.

```
move nodedata noah
  fromstgpool=tapepool tostgpool=diskpool
  filespace=\\tsmserv1\d$ fsid=2
```


NOTIFY SUBSCRIBERS (Notify managed servers to update profiles)

Use this command on a configuration manager to notify one or more managed servers to request that their configuration information be immediately refreshed.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

PROFILE (Required)

Specifies the name of the profile. Any managed servers that subscribe to the profile are notified. You can use wildcard characters to specify multiple profiles. To specify multiple profiles, separate the names with commas and no intervening spaces. The default is to notify all subscribers.

Example: Notify managed servers to update profiles

Notify all managed servers that subscribe to a profile named DELTA to request updated configuration information.

```
notify subscribers profile=delta
```

Related commands

Table 206. Commands related to NOTIFY SUBSCRIBERS

Command	Description
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE SUBSCRIBER	Deletes obsolete managed server subscriptions.
DELETE SUBSCRIPTION	Deletes a specified profile subscription.
QUERY SUBSCRIBER	Displays information about subscribers and their subscriptions to profiles.
QUERY SUBSCRIPTION	Displays information about profile subscriptions.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
SET CONFIGREFRESH	Specifies a time interval for managed servers to contact configuration managers.

PERFORM LIBACTION (Define or delete all drives and paths for a library)

Use this command to define or delete all drives and their paths for a single library in one step.

This command can be used when you set up a library environment or modify an existing hardware setup that requires changes to many drive definitions. After you define a library, issue **PERFORM LIBACTION** to define drives and their paths for the library. You can also delete all drives and paths for a library by issuing the command with ACTION=DELETE.

This command is only valid for library types of SCSI and VTL. To use this command with ACTION=DEFINE, the SANDISCOVERY option must be supported and enabled.


For detailed and current library support information, see the Supported Devices website for your operating system:

http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html


Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



A (DEFINE):



Parameters

library_name (Required)

Specifies the name of the library to be defined or deleted. The maximum length of this name is 30 characters unless you are issuing **PERFORM LIBACTION** with ACTION=DEFINE and using the default PREFIX value. In that case, the maximum length of the name is 25 characters.

ACTION

Specifies the action for the PERFORM LIBACTION command. Possible values are:

PERFORM LIBACTION

DEFine

Specifies that drives and their paths are defined for the specified library. SAN discovery must be enabled before you specify this parameter value.

DElete

Specifies that drives and their paths are deleted for the specified library.

DEVIce

Specifies the library device name that is used when you define paths if a path to the library is not already defined. If a path is already defined, the DEVICE parameter is ignored. The maximum length for this value is 64 characters. This parameter is optional.

PREFix

Specifies the prefix that is used for all drive definitions. For example, a PREFIX value of *DR* creates drives *DR0*, *DR1*, *DR2*, for as many drives as are created. If a value is not specified for the PREFIX parameter, the library name is used as the prefix for drive definitions. The maximum length for this value is 25 characters.

SOURCe

Specifies the source server name to be used when you define or delete drive path definitions on a library client or LAN-free client. Use this parameter only if the drives in the library are set up for the local server. If no value is specified for the SOURCE parameter, the local server name, which is the default, is used. The maximum length for the source name is 64 characters.

If a source name other than the local server name is specified with ACTION=DEFINE, drive path definitions are defined with the token value of UNDISCOVERED. The path definitions are then updated dynamically by library clients that support SAN Discovery the first time the drive is mounted.

PREView

Specifies the output of all commands that are processed for **PERFORM LIBACTION** before the command is issued. The PREVIEW parameter is not compatible with the DEVICE parameter. If you are issuing the **PERFORM LIBACTION** command to define a library, you cannot specify both the PREVIEW and the DEVICE parameter.

Possible values are:

No Specifies that a preview of the commands that are issued for **PERFORM LIBACTION** is not displayed.

Yes

Specifies that a preview of the commands that are issued for **PERFORM LIBACTION** is displayed.

Example: Define a shared library

In a SAN, define a library named SHAREDTSM to a library client server named LIBMGR1:

```
define library sharedtsm libtype=shared primarylibmanager=libmgr1
```

Then, issue the **PERFORM LIBACTION** command to define the drives and paths for the library:

```
perform libaction sharedtsm action=define source=libmgr1
```

Example: Define a library with four drives

Define a SCSI library named KONA:

```
define library kona libtype=scsi
```

Then, issue the **PERFORM LIBACTION** command to define drives and paths for the library:

```
perform libaction kona action=define device=/dev/tsm SCSI/lb3
prefix=dr
```

The server then runs the following commands:

```
define path server1 kona srct=server destt=library
device=/dev/tsm SCSI/lb3
define drive kona dr0
define path server1 dr0 srct=server destt=drive library=kona
device=/dev/tsm SCSI/mt1
define drive kona dr1
define path server1 dr1 srct=server destt=drive library=kona
device=/dev/tsm SCSI/mt2
define drive kona dr2
define path server1 dr2 srct=server destt=drive library=kona
device=/dev/tsm SCSI/mt3
define drive kona dr3
define path server1 dr3 srct=server destt=drive library=kona
device=/dev/tsm SCSI/mt4
```

Related commands

Table 207. Commands related to **PERFORM LIBACTION**

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE PATH	Defines a path from a source to a destination.
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE DRIVE	Deletes a drive from a library.
DELETE LIBRARY	Deletes a library.
DELETE PATH	Deletes a path from a source to a destination.
QUERY DRIVE	Displays information about drives.
QUERY LIBRARY	Displays information about one or more libraries.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE DRIVE	Changes the attributes of a drive.
UPDATE LIBRARY	Changes the attributes of a library.
UPDATE PATH	Changes the attributes associated with a path.

PING SERVER (Test the connection between servers)

Use this command to test the connection between the local server and a remote server.

Important: The name and password of the administrator client issuing this command must also be defined on the remote server.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—PING SERVER—*server_name*—————►►

Parameters

server_name **(Required)**

Specifies the name of the remote server.

Example: Ping a server

Test the connection to server FRED.

```
ping server fred
```

Related commands

Table 208. Commands related to PING SERVER

Command	Description
DEFINE SERVER	Defines a server for server-to-server communications.
QUERY SERVER	Displays information about servers.

PREPARE (Create a recovery plan file)

Use this command to create a recovery plan file, which contains the information needed to recover a Tivoli Storage Manager server. You can store a recovery plan file on a file system that is accessible to the source Tivoli Storage Manager server or on a target Tivoli Storage Manager server.


You can use the **QUERY ACTLOG** command to view whether the **PREPARE** command was successful.

You can also view this information from the server console or, if the **WAIT** parameter equals **YES**, an administrative client session.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

Source

Specifies the type of database backup series that Tivoli Storage Manager assumes when generating the recovery plan file. This parameter is optional. The default is **DBBACKUP**. The choices are:

DBBackup

Specifies that Tivoli Storage Manager assumes the latest full with plus database backup series.

DBSnapshot

Specifies that Tivoli Storage Manager assumes the latest database snapshot backup series.

DEVclass

Specifies the device class name that is used to create a recovery plan file object on a target server. The device class must have a device type of **SERVER**.

PREPARE

Important: The maximum capacity for the device class must be larger than the size of the recovery plan file. If the size of the recovery plan file exceeds the maximum capacity, the command fails.

The naming convention for the archive object that contains the recovery plan file on the target server is:

- **Filespace name:**
ADSM.SERVER
- **High-level qualifier:**
devclassprefix/servername.yyyymmdd.hhmmss
- **Low-level qualifier:**
RPF.OBJ.1

The recovery plan file virtual volume name as recorded in the volume history table on the source server is in the format *servername.yyyymmdd.hhmmss*.

If the DEVCLASS parameter is not specified, the recovery plan file is written to a file based on the plan prefix.

If SOURCE=DBBACKUP is specified or is defaulted to, the volume history entry for the recovery plan file object specifies a volume type of RPFFILE. If SOURCE=DBSNAPSHOT is specified, the volume history entry specifies a volume type of RPFsnapshot.

PLANPrefix

Specifies the path name prefix that is used in the recovery plan file name. This parameter is optional.

The maximum length is 250 characters.

Tivoli Storage Manager appends to the prefix the sortable date and time format *yyymmdd.hhmmss*. For example: 20081115.051421.

The prefix can be one of the following:

Directory path

End the prefix with the forward slash (/). For example:

```
PLANPREFIX=/admsrv/recplans/
```

The resulting file name would look like this:

```
/admsrv/recplans/20081115.051421
```

Directory path followed by a string

Tivoli Storage Manager treats the string as part of the file name. For example:

```
PLANPREFIX=/admsrv/recplans/accounting
```

The resulting file name looks like this:

```
/admsrv/recplans/accounting.20081115.051421
```

Note the period before the date and time.

String only

Tivoli Storage Manager specifies the directory path. Tivoli Storage Manager uses the name of the current working directory. For example, the current working directory is /opt/tivoli/tsm/server/bin and you specify the following parameter:

```
PLANPREFIX=shipping
```

The resulting file name looks like this:

```
/opt/tivoli/tsm/server/bin/shipping.20081115.051421
```

Note the period before the date and time.

If the PLANPREFIX parameter is not specified, Tivoli Storage Manager selects the prefix in one of these ways:

- If the **SET DRMPLANPREFIX** command has been issued, Tivoli Storage Manager uses the prefix specified in that command.
- If the **SET DRMPLANPREFIX** command has not been issued, Tivoli Storage Manager uses the directory path name of the current working directory. For example, the current working directory is the following:

```
/opt/tivoli/tsm/server/bin
```

The resulting file name looks like this:

```
/opt/tivoli/txm/server/bin/20081115.051421
```

INSTRPrefix

Specifies the prefix of the path name used by Tivoli Storage Manager to locate the files that contain the recovery instructions. The maximum length is 250 characters.

The prefix can be one of the following:

Directory path

End the prefix with the forward slash (/). For example:

```
INSTRPREFIX=/admsrv/recinstr/  
/admsrv/recinstr/RECOVERY.INSTRUCTIONS.GENERAL
```

Directory path followed by a string

Tivoli Storage Manager treats the string as part of the file name. For example:

```
INSTRPREFIX=/admsrv/recinstr/accounts
```

Tivoli Storage Manager appends the appropriate recovery plan file stanza name. For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name is:

```
/admsrv/recinstr/accounts.RECOVERY.INSTRUCTIONS.GENERAL
```

String only

- Tivoli Storage Manager specifies the directory path and appends the appropriate recovery plan file stanza name. Tivoli Storage Manager uses the name of the current working directory. For example, the current working directory is /opt/tivoli/tsm/server/bin and you specify the following parameter:

```
INSTRPREFIX=shipping
```

For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name looks like this:

```
/opt/tivoli/tsm/server/bin/shipping.RECOVERY.INSTRUCTIONS.GENERAL
```

If you do not specify the INSTRPREFIX parameter, Tivoli Storage Manager selects the prefix in one of these ways:

- If the **SET DRMINSTRPREFIX** command has been issued, Tivoli Storage Manager uses the prefix specified in that command.
- If the **SET DRMINSTRPREFIX** command has not been issued, Tivoli Storage Manager uses the current working directory.

PREPARE

For example, if the current working directory is `/opt/tivoli/tsm/server/bin`, for the `RECOVERY.INSTRUCTIONS.GENERAL` file, the resulting file name would be:

```
/opt/tivoli/tsm/server/bin/RECOVERY.INSTRUCTIONS.GENERAL
```

PRIMstgpool

Specifies the names of the primary storage pools that you want to restore. Separate the storage pool names with commas and no intervening spaces. You can use wildcard characters. If this parameter is not specified, Tivoli Storage Manager selects the storage pools as follows:

- If the **SET DRMPRIMSTGPOOL** command has been issued, Tivoli Storage Manager includes the primary storage pools named in that command.
- If the **SET DRMPRIMSTGPOOL** command has not been issued, Tivoli Storage Manager includes all the primary storage pools.

COPYstgpool

Specifies the names of the copy storage pools used to back up the primary storage pools that you want to restore (see the `PRIMSTGPOOL` parameter). Separate storage pool names with commas and no intervening spaces. You can use wildcard characters. If this parameter is not specified, Tivoli Storage Manager selects the storage pools as follows:

- If the **SET DRMCOPYSTGPOOL** command has been issued, Tivoli Storage Manager includes those copy storage pools.
- If the **SET DRMCOPYSTGPOOL** command has not been issued, Tivoli Storage Manager includes all copy storage pools.

ACTIVEDatastgpool

Specifies the names of the active-data storage pools that you want to have available for offsite access. Separate active-data storage-pool names with commas and no intervening spaces. You can use wildcard characters. If this parameter is not specified, Tivoli Storage Manager selects the storage pools as follows:

- If the **SET ACTIVEDATASTGPOOL** command has been previously issued with valid active-data storage pool names, Tivoli Storage Manager processes those storage pools.
- If the **SET ACTIVEDATASTGPOOL** command has not been issued, or all of the active-data storage pools have been removed using the **SET ACTIVEDATASTGPOOL** command, Tivoli Storage Manager processes only the active-data pool volumes that were marked on-site at the time the **PREPARE** command is run. Tivoli Storage Manager will mark these volumes as `UNAVAILABLE`.

Wait

Specifies whether this command is processed in the background or foreground.

No Specifies background processing. This is the default.

Yes

Specifies foreground processing.

You cannot specify `YES` from the server console.

Example: Create a recovery plan file

Issue the **PREPARE** command and query the activity log to check the results.

```
prepare
query actlog search=prepare
```

```

05/03/2008 12:01:13 ANR0984I Process 3 for PREPARE started in the
BACKGROUND at 12:01:13.
05/03/2008 12:01:13 ANR6918W PREPARE: Recovery instructions file
/home/guest/drmtest/prepare/tserver/DSM1509/
RECOVERY.INSTRUCTIONS.DATABASE not found.
05/03/2008 12:01:13 ANR6918W PREPARE: Recovery instructions file
/home/guest/drmtest/prepare/tserver/DSM1509/
RECOVERY.INSTRUCTIONS.STGPOOL not found.
05/03/2008 12:01:13 ANR6913W PREPARE: No volumes with backup data
exist in copy storage pool CSTORAGEEP.
05/03/2008 12:01:13 ANR6913W PREPARE: No volumes with backup data
exist in copy storage pool CSTORAGEPSM.
05/03/2008 12:01:14 ANR6920W PREPARE: Generated replacement volume
name BACK4X@ is not valid for device type
8MM. Original volume name: BACK4X. Stanza is
PRIMARY.VOLUMES.REPLACEMENT macro.
05/03/2008 12:01:14 ANR6900I PREPARE: The recovery plan file
/home/guest/drmtest/prepare/plandir/DSM1509/
r.p.20080503.120113 was created.
05/03/2008 12:01:14 ANR0985I Process 3 for PREPARE running in the
BACKGROUND completed with completion state
SUCCESS at 12:01:14.

```

Related commands

Table 209. Commands related to PREPARE

Command	Description
DELETE VOLHISTORY	Removes sequential volume history information from the volume history file.
QUERY DRMSTATUS	Displays DRM system parameters.
QUERY RPFCONTENT	Displays the contents of a recovery plan file.
QUERY RPFFILE	Displays information about recovery plan files.
QUERY SERVER	Displays information about servers.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
SET DRMACTIVEDATASTGPOOL	Specifies that active-data storage pools are managed by DRM.
SET DRMCOPYSTGPOOL	Specifies that copy storage pools are managed by DRM.
SET DRMINSTRPREFIX	Specifies the prefix portion of the path name for the recovery plan instructions.
SET DRMPPLANVPOSTFIX	Specifies the replacement volume names in the recovery plan file.
SET DRMPPLANPREFIX	Specifies the prefix portion of the path name for the recovery plan.
SET DRMPRIMSTGPOOL	Specifies that primary storage pools are managed by DRM.
SET DRMRPFEXPIREDAYS	Set criteria for recovery plan file expiration.
UPDATE VOLHISTORY	Adds or changes location information for a volume in the volume history file.

QUERY commands

Use the **QUERY** commands to request or display information about Tivoli Storage Manager objects.

- “QUERY ACTLOG (Query the activity log)” on page 661
- “QUERY ADMIN (Display administrator information)” on page 667
- “QUERY ALERTTRIGGER (Query the list of defined alert triggers)” on page 672
- “QUERY ALERTSTATUS (Query the status of an alert)” on page 674
- “QUERY ASSOCIATION (Query client node associations with a schedule)” on page 679
- “QUERY AUDITOCUPANCY (Query client node storage utilization)” on page 681
- “QUERY BACKUPSET (Query a backup set)” on page 684
- “QUERY BACKUPSETCONTENTS (Query contents of a backup set)” on page 689
- “QUERY CLOPTSET (Query a client option set)” on page 692
- “QUERY COLLOCGROUP (Query a collocation group)” on page 694
- “QUERY CONTENT (Query the contents of a storage pool volume)” on page 697
- “QUERY COPYGROUP (Query copy groups)” on page 705
- “QUERY DATAMOVER (Display data mover definitions)” on page 710
- “QUERY DB (Display database information)” on page 713
- “QUERY DBSPACE (Display database storage space)” on page 716
- “QUERY DEVCLASS (Display information on one or more device classes)” on page 718
- “QUERY DIRSPACE (Query storage utilization of FILE directories)” on page 723
- “QUERY DOMAIN (Query a policy domain)” on page 724
- “QUERY DRIVE (Query information about a drive)” on page 727
- “QUERY DRMEDIA (Query disaster recovery media)” on page 731
- “QUERY DRMSTATUS (Query disaster recovery manager system parameters)” on page 740
- “QUERY ENABLED (Query enabled events)” on page 743
- “QUERY EVENT (Query scheduled and completed events)” on page 745
- “QUERY EVENTRULES (Query rules for server or client events)” on page 757
- “QUERY EVENTSERVER (Query the event server)” on page 759
- “QUERY EXPORT (Query for active or suspended export operations)” on page 760
- “QUERY FILESPACE (Query one or more file spaces)” on page 767
- “QUERY LIBRARY (Query a library)” on page 774
- “QUERY LIBVOLUME (Query a library volume)” on page 777
- “QUERY LICENSE (Display license information)” on page 780
- “QUERY LOG (Display information on the recovery log)” on page 782
- “QUERY MACHINE (Query machine information)” on page 784
- “QUERY MEDIA (Query sequential access storage pool media)” on page 787
- “QUERY MGMTCLASS (Query a management class)” on page 793
- “QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)” on page 796

- “QUERY MONITORSTATUS (Query the monitoring status)” on page 799
- “QUERY MOUNT (Display information on mounted sequential access volumes)” on page 803
- “QUERY NASBACKUP (Query NAS backup images)” on page 805
- “QUERY NODE (Query nodes)” on page 809
- “QUERY NODEDATA (Query client data in volumes)” on page 820
- “QUERY NODEGROUP (Query a node group)” on page 823
- “QUERY OCCUPANCY (Query client file spaces in storage pools)” on page 825
- “QUERY OPTION (Query server options)” on page 829
- “QUERY PATH (Display a path definition)” on page 831
- “QUERY POLICYSET (Query a policy set)” on page 835
- “QUERY PROCESS (Query one or more server processes)” on page 838
- “QUERY PROFILE (Query a profile)” on page 841
- “QUERY PROXYNODE (Query proxy authority for a client node)” on page 844
- “QUERY PVUESTIMATE (Display processor value unit estimate)” on page 845
- “QUERY RECOVERYMEDIA (Query recovery media)” on page 849
- “QUERY REPLICATION (Query node replication processes)” on page 852
- “QUERY REPLNODE (Display information about replication status for a client node)” on page 862
- “QUERY REPLRULE (Query replication rules)” on page 866
- “QUERY REQUEST (Query one or more pending mount requests)” on page 868
- “QUERY RESTORE (Query restartable restore sessions)” on page 869
- “QUERY RPFCONTENT (Query recovery plan file contents stored on a target server)” on page 872
- “QUERY RPFFILE (Query recovery plan file information stored on a target server)” on page 874
- “QUERY SAN (Query the devices on the SAN)” on page 877
- “QUERY SCHEDULE (Query schedules)” on page 880
- “QUERY SCRIPT (Query Tivoli Storage Manager scripts)” on page 888
- “QUERY SERVER (Query a server)” on page 891
- “QUERY SERVERGROUP (Query a server group)” on page 895
- “QUERY SESSION (Query client sessions)” on page 897
- “QUERY SHREDSTATUS (Query shredding status)” on page 901
- “QUERY SPACETRIGGER (Query the space triggers)” on page 903
- “QUERY STATUS (Query system parameters)” on page 906
- “QUERY STATUSTHRESHOLD (Query status monitoring thresholds)” on page 914
- “QUERY SSLKEYRINGPW (Query SSL key database file password)” on page 905
- “QUERY STGPOOL (Query storage pools)” on page 917
- “QUERY SUBSCRIBER (Display subscriber information)” on page 928
- “QUERY SUBSCRIPTION (Display subscription information)” on page 930
- “QUERY SYSTEM (Query the system configuration and capacity)” on page 932
- “QUERY TAPEALERTMSG (Display status of SET TAPEALERTMSG command)” on page 934
- “QUERY TOC (Display table of contents for a backup image)” on page 935

QUERY commands

- “QUERY VIRTUALFSMAPPING (Query a virtual file space mapping)” on page 938
- “QUERY VOLHISTORY (Display sequential volume history information)” on page 940
- “QUERY VOLUME (Query storage pool volumes)” on page 948

QUERY ACTLOG (Query the activity log)

Use this command to display messages generated by the server and client. This command provides filtering options that can be used to limit the number of messages displayed and the time that it takes to process this query. If you do not specify any parameters with this command, all messages generated in the previous hour are displayed.


The activity log contains all messages that are sent to the server console under normal operation. The results of commands entered at the server console are not recorded in the activity log unless the command affects or starts a background process or client session. Error messages are displayed in the activity log.

Restriction: You cannot schedule the **QUERY ACTLOG** command by using the **DEFINE SCHEDULE** command.

Privilege class

Any administrator can issue this command.

Syntax



A:

Parameters

BEGINDate

Specifies the beginning date of the range for messages to be displayed. All

QUERY ACTLOG

messages meeting the time range criteria that occurred after this date are displayed. The default is the current date. This parameter is optional.

You can specify the date using one of the following values:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY- <i>days</i> or <i>-days</i>	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -7 or -7. To display information beginning with messages created a week ago, you can specify <code>BEGINDATE=TODAY-7</code> or <code>BEGINDATE= -7</code> .
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time of the range for messages to be displayed. All messages meeting the time range criteria that occurred after this time are displayed. If you do not specify time, all messages that occurred in the last hour are displayed.

You can specify the time using one of the following values:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time on the specified begin date	10:30:08
NOW	The current time on the specified begin date	NOW
NOW+ <i>HH:MM</i> or <i>+HH:MM</i>	The current time plus hours and minutes on the specified begin date	NOW+03:00 or +03:00. If you issue this command at 9:00 with <code>BEGINTIME=NOW+3</code> or <code>BEGINTIME=+3</code> , Tivoli Storage Manager displays messages with a time of 12:00 or later on the begin date.

Value	Description	Example
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified begin date	NOW-04:00 or -04:00. If you issue the QUERY ACTLOG command at 9:00 with BEGINTime=NOW-3:30 or BEGINTime= -3:30, Tivoli Storage Manager displays messages with a time of 5:30 or later on the begin date.

ENDDate

Specifies the ending date of the range for messages to be displayed. All messages meeting the time range criteria that occurred before this date are displayed. If you do not specify a value, the current date is used. This parameter is optional.

You can specify the date using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY-1 or -1. To display information created up to yesterday, you can specify ENDDATE=TODAY-1 or simply ENDDATE= -1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+days	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time of the range for messages to be displayed. All messages meeting this time range criteria that occurred before this time are displayed. If you do not specify a value, all messages are displayed up to the time when you issued this command. This parameter is optional.

You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified end date	10:30:08
NOW	The current time on the specified end date	NOW

QUERY ACTLOG

Value	Description	Example
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified end date	NOW+03:00 or +03:00. If you issue this command at 9:00 with ENDTIME=NOW+3:00 or ENDTIME= +3:00, Tivoli Storage Manager displays messages with a time of 12:00 or earlier on the end date you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified end date	NOW-03:30 or -03:30. If you issue this command at 9:00 with ENDTIME=NOW-3:30 or ENDTIME= -3:30, IBM Tivoli Storage Manager displays messages with a time of 5:30 or earlier on the end date you specify.

MSGno

Specifies an integer that defines the number of the message to be displayed from the activity log. This integer is just the numeric part of the message. This parameter is optional.

Search

Specifies a text string that you want to search for in the activity log. Enclose the string expression in quotation marks if it contains blanks. You can use text and a wildcard character to specify this string. This parameter is optional.

Note: Do not enter as a text string either the IBM Tivoli Storage Manager server name or text and a wildcard character that would find the server name. If you do so, the output includes messages that do not include the search string.

NODENAME

Specifies that the query displays messages logged for this node. If you do not specify a value for this parameter, messages for all nodes are displayed.

ORIGINATOR

Specifies that the query displays messages logged by the server, client, or both. The default is ALL. Possible values are:

ALL

Specifies that the query displays messages that originated from the client and the server.

SERVER

Specifies that the query displays messages that originated from the server.

CLIENT

Specifies that the query displays messages that originated from the client.

You can specify one of the following values to minimize processing time when querying the activity log for messages logged by the client:

OWNERNAME

Specifies that the query displays messages logged for a particular owner. If you do not specify a value for this parameter, messages for all owners are displayed.

SCHEDNAME

Specifies that the query displays messages logged by a particular

scheduled client activity. If you do not specify a value for this parameter, messages for all schedules are displayed.

Domainname

Specifies that the query displays messages logged for a particular policy domain to which a named schedule belongs. This parameter is optional, unless you are specifying a schedule name.

SESsnum

Specifies that the query displays messages logged from a particular client session number. If you do not specify a value for this parameter, messages for all client sessions are displayed.

Example: Search activity log for messages with specific text

Search the activity log for any message that contains the string “delete”. The output includes only messages produced during the past hour. Issue the command:
`query actlog search=delete`

Date/Time	Message
08/27/1998 15:19:43	ANR0812I Inventory client file expiration complete: 0 files deleted.

Example: Search activity log for messages within a specific time frame

Display messages that occurred yesterday between 9:30 and 12:30. Issue the command:

```
query actlog begindate=today-1
beginntime=09:30:00 endtime=12:30:00
```

Date/Time	Message
10/21/1998 10:52:36	ANR0407I Session 3921 started for administrator ADMIN (WebBrowser) (HTTP 9.115.20.100(2315)).
10/21/1998 11:06:08	ANR0405I Session 3922 ended for administrator ADMIN (WebBrowser).
10/21/1998 12:16:50	ANR0405I Session 3934 ended for administrator ADMIN (WebBrowser).

Example: Search activity log for messages from a specific client node

Search the activity log for Tivoli Storage Manager messages from the client for node JEE. Issue the command:

```
query actlog originator=client node=jee
```

Date/Time	Message
06/10/1998 15:46:22	ANE4007E (Session No: 3 Node: JEE) Error processing '/jee/report.out': access to the object is denied
06/11/1998 15:56:56	ANE4009E (Session No: 4 Node: JEE) Error processing '/jee/work.lst': disk full condition

Example: Search activity log for client and server messages from a specific client node and session

Search the activity log for Tivoli Storage Manager messages from the client and server for node A associated with Session 1. The output includes all messages with the defined text string, "SESSION: 1". Issue the command:

```
query actlog search="(SESSION:1)"
```

Date/Time	Message
02/13/2012 12:13:42	ANR0406I Session 1 started for node A (WinNT) (Tcp/Ip colind(2463)). (SESSION: 1)
02/13/2012 12:13:56	ANE4952I (ANE4985I Session: 1, ANE4986I Node: A) Total number of objects inspected: 34 (SESSION: 1)
02/13/2012 12:13:56	ANE4954I (ANE4985I Session: 1, ANE4986I Node: A) Total number of objects backed up: 34 (SESSION: 1)
02/13/2012 12:13:56	ANE4958I (ANE4985I Session: 1, ANE4986I Node: A) Total number of objects updated: 0 (SESSION: 1)
02/13/2012 12:13:56	ANE4964I (ANE4985I Session: 1, ANE4986I Node: A) Elapsed processing time: 00:00:02 (SESSION: 1)
02/13/2012 12:13:59	ANR0403I Session 1 ended for node A (WinNT). (SESSION: 1)

Example: Search activity log for client-generated messages from a client session

Search the activity log for Tivoli Storage Manager messages from a specific client session. The output includes only messages generated by the client. Issue the command:

```
query actlog sessnum=1
```

Date/Time	Message
02/13/2012 12:13:56	ANE4952I (ANE4985I Session: 1, ANE4986I Node: A) Total number of objects inspected: 34 (SESSION: 1)
02/13/2012 12:13:56	ANE4954I (ANE4985I Session: 1, ANE4986I Node: A) Total number of objects backed up: 34 (SESSION: 1)
02/13/2012 12:13:56	ANE4958I (ANE4985I Session: 1, ANE4986I Node: A) Total number of objects updated: 0 (SESSION: 1)
02/13/2012 12:13:56	ANE4964I (ANE4985I Session: 1, ANE4986I Node: A) Elapsed processing time: 00:00:02 (SESSION: 1)

Related commands

Table 210. Command related to QUERY ACTLOG

Command	Description
SET ACTLOGRETENTION	Specifies the number of days to retain log records in the activity log.


QUERY ADMIN (Display administrator information)

Use this command to display information about one or more administrators.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

admin_name

Specifies the name of the administrator for which you want to display information. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a value for this parameter, all administrators are displayed.

Classes

Specifies that you want to restrict output to those administrators that have privilege classes that you specify. This parameter is optional. You can specify multiple privilege classes in a list by separating the names with commas and no intervening spaces. If you do not specify a value for this parameter, information about all administrators is displayed, regardless of privilege class. Possible values are:

System

Display information on administrators with system privilege.

Policy

Display information on administrators with policy privilege.

Storage

Display information on administrators with storage privilege.

Operator

Display information on administrators with operator privilege.

Node

Display information on users with client node privilege.

QUERY ADMIN

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified administrators.

Detailed

Specifies that complete information is displayed for the specified administrators.

Authentication

Specifies the password authentication method for the administrator.

Local

Display those administrators authenticating to the IBM Tivoli Storage Manager server.

Ldap

Display those administrators authenticating to an LDAP directory server. The administrator password is case-sensitive.

Alert

Specifies whether alerts are sent to an administrators email address.

Yes

Specifies that alerts are sent to the specified administrators email address.

No

Specifies that alerts are not sent to the specified administrators email address. This is the default value.

Tip: Alert monitoring must be enabled, and email settings must be correctly defined to successfully receive alerts by email. To view the current settings, issue the **QUERY MONITORSETTINGS** command.

Example: Display information about all administrators

Display partial information on all administrators. Issue the command:

```
query admin
```

Administrator Name	Days Since Last Access	Days Since Password Set	Locked?	Privilege Classes
ADMIN	<1	<1	No	System
SERVER_CONSOLE			No	System

See “Field descriptions” on page 669 for field descriptions.

Example: Display complete information about one administrator

From a managed server, display complete information for the administrator named ADMIN. Issue the command:

```
query admin admin format=detailed
```

```

Administrator Name: ADMIN
Last Access Date/Time: 1998.06.04 17.10.52
Days Since Last Access: <1
Password Set Date/Time: 1998.06.04 17.10.52
Days Since Password Set: 26
Invalid Sign-on Count: 0
    Locked?: No
    Contact:
        System Privilege: Yes
        Policy Privilege: **Included with system privilege**
        Storage Privilege: **Included with system privilege**
        Operator Privilege: **Included with system privilege**
        Client Access Privilege: **Included with system privilege**
        Client Owner Privilege: **Included with system privilege**
Registration Date/Time: 05/09/1998 23:54:20
Registering Administrator: SERVER_CONSOLE
Managing profile:
Password Expiration Period: 90 Day (s)
    Email Address:
        Email Aerts: Yes
    Authentication: Local
    SSL Required: No
    
```

See “Field descriptions” for field descriptions.

Field descriptions

Administrator Name

Specifies the name of the administrator.

Last Access Date/Time

Specifies the date and time that the administrator last accessed the server.

Days Since Last Access

Specifies the number of days since the administrator last accessed the server.

Password Set Date/Time

Specifies the date and time that the administrator’s password was defined or most recently updated.

Days Since Password Set

Specifies the number of days since the administrator’s password was defined or most recently updated.

Invalid Sign-on Count

Specifies the number of invalid sign-on attempts that have been made since the last successful sign-on. This count can only be non-zero when an invalid password limit (SET INVALIDPWLIMIT) is greater than zero. When the number of invalid attempts equals the limit set by the SET INVALIDPWLIMIT command, the administrator is locked out of the system.

Locked?

Specifies whether the administrator is locked out of the system.

Contact

Specifies any contact information for the administrator.

System Privilege

Specifies whether the administrator has been granted system privilege.

Policy Privilege

Specifies whether the administrator has been granted unrestricted policy privilege or the names of any policy domains that the restricted policy administrator can manage.

Storage Privilege

Specifies whether the administrator has been granted unrestricted storage privilege or the names of any storage pools that the restricted storage administrator can manage.

Operator Privilege

Specifies whether the administrator has been granted operator privilege.

Client Access Privilege

Specifies that client access authority has been granted to a user with node privilege.

Client Owner Privilege

Specifies that client owner authority has been granted to a user with node privilege.

Registration Date/Time

Specifies the date and time that the administrator was registered.

Registering Administrator

Specifies the name of the administrator who registered the administrator. If this field contains \$\$CONFIG_MANAGER\$\$, the administrator is associated with a profile that is managed by the configuration manager.

Managing Profile

Specifies the profiles to which the managed server subscribed to get the definition of this administrator.

Password Expiration Period

Specifies the administrator's password expiration period.

Email Address

Applies only to IBM Tivoli Storage Manager Express and is shown for compatibility purposes.

Email Alerts

Specifies whether alerts are sent to the specified administrator by email.

Authentication

Specifies the password authentication method: LOCAL, LDAP, or LDAP (pending).

Authentication Target	Authentication Method
Tivoli Storage Manager server	LOCAL
LDAP directory server	LDAP
This administrator is configured to authenticate with an LDAP directory server, but the administrator did not yet authenticate through a client node.	LDAP (pending)

SSL Required

Specifies if the security setting for the administrator user ID requires Secure Sockets Layer (SSL). Values can be YES, NO, or Default. You must have system level authority to update the administrator **SSLREQUIRED** setting.

Related commands

Table 211. Commands related to QUERY ADMIN

Command	Description
GRANT AUTHORITY	Assigns privilege classes to an administrator.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER ADMIN	Defines a new administrator without granting administrative authority.
REMOVE ADMIN	Removes an administrator from the list of registered administrators.
RENAME ADMIN	Changes a Tivoli Storage Manager administrator's name.
RESET PASSEXP	Resets the password expiration for nodes or administrators.
REVOKE AUTHORITY	Revokes one or more privilege classes or restricts access to policy domains and storage pools.
SET INVALIDPWLIMIT	Sets the number of invalid logon attempts before a node is locked.
SET MINPWLENGTH	Sets the minimum length for client passwords.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.


QUERY ALERTTRIGGER (Query the list of defined alert triggers)

Use this command to display which server messages are defined as alerts.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

message_number

Specifies the message number that you want to query. Specify multiple message numbers, which are separated by commas, and no intervening spaces. Message numbers are a maximum of eight characters in length. Wildcard characters can be used to specify message numbers. If you do not specify a message number, all alert triggers are displayed.

Query alert triggers to display which messages are designated as alerts

Display all messages that are designated as alerts by issuing the following command:

```
query alertrigger
```

Example output:

Alert Trigger	Category	Administrator
ANR1067E	SERVER	HARRYH
ANR1073E	SERVER	CSDADMIN, DJADMIN, HARRYH
ANR1074E	STORAGE	CSDADMIN, DJADMIN, HARRYH
ANR1096E	STORAGE	CSDADMIN, DJADMIN, HARRYH, MHAYE

Query alert triggers for a specific message number

Display all alert triggers that have message number ANR1067E designated to them by issuing the following command:

```
query alertrigger ANR1067E
```

Example output:

Alert Trigger	Category	Administrator
ANR1067E	SERVER	HARRYH

Related commands

Table 212. Commands related to QUERY ALERTTRIGGER

Command	Description
“DEFINE ALERTTRIGGER (Define an alert trigger)” on page 112	Associates specified messages to an alert trigger.

Table 212. Commands related to QUERY ALERTTRIGGER (continued)

Command	Description
"DELETE ALERTTRIGGER (Remove a message from an alert trigger)" on page 400	Removes a message number that can trigger an alert.
"QUERY ALERTSTATUS (Query the status of an alert)" on page 674	Displays information about alerts that have been issued on the server.
"UPDATE ALERTTRIGGER (Update a defined alert trigger)" on page 1165	Updates the attributes of one or more alert triggers.
"UPDATE ALERTSTATUS (Update the status of an alert)" on page 1168	Updates the status of a reported alert.

QUERY ALERTSTATUS


QUERY ALERTSTATUS (Query the status of an alert)

Use this command to display information about alerts that are reported on the Tivoli Storage Manager server.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

Status

Specifies the status type that you want to display. If you do not specify a status, all alerts are queried and displayed. Specify one of the following values:

Active

Displays alerts that are specified in the Tivoli Storage Manager server database as active.

INactive

Displays alerts that are in the inactive state.

CLosed

Displays alerts that are in the closed state.

ANy

Displays all alerts, without regard to state.

MSGnum

Specifies the message number that you want to display. Specify the numerical portion of a Tivoli Storage Manager server message. Values range from 0 to 9999. For example, the message number in message ANR2044E would be 2044. Specify multiple message numbers by separating them with commas and no intervening spaces.

CATegory

Specifies the category type for the alert, which is determined by the message types. The default value is SERVER. Specify one of the following values:

APplication

Alert is classified as application category. For example, you can specify this category for messages that are associated with application (TDP) clients.

CATalog

Alert is classified as catalog category. For example, you can specify this category for messages that are associated with the database, active log file, or archive log file.

Note: Changing the category of an alert trigger does not change the category of existing alerts on the server. New alerts are categorized with the new category.

CLient

Alert is classified as client category. For example, you can specify this category for messages that are associated with general client activities.

DEvice

Alert is classified as device category. For example, you can specify this category for messages that are associated with device classes, libraries, drives, or paths.

SErver

Alert is classified as general server category. For example, you can specify this category for messages that are associated with general server activities or events.

STorage

Alert is classified as storage category. For example, you can specify this category for messages that are associated with storage pools.

SYstems

Alert is classified under system clients category. For example, you can specify this category for messages that are associated with system backup and archive or hierarchical storage management (HSM) backup-archive clients.

VMcIient

Alert is classified under VMclient category. For example, you can specify this category for messages that are associated with virtual machine clients.

SOURCEType

Specifies the source type that is being queried. Specify one of the following values:

LOcal

Displays alerts that originated from the local Tivoli Storage Manager server.

QUERY ALERTSTATUS

CLient

Displays alerts that originated from the Tivoli Storage Manager client.

REmote

Displays alerts that originated from another Tivoli Storage Manager server.

SOURCEName

Specifies the name of the source where the alert originated. **SOURCENAME** can be the name of a local or remote Tivoli Storage Manager server, or a Tivoli Storage Manager client.

ID This optional parameter specifies the unique ID of the alert that you want to display. Specify a value from 1 to 9223372036854775807.

ASSigned

Specifies the administrator name that is assigned the alert that you want to query.

RESolvedby

Specifies the administrator name that resolved the alert that you want to query.

Query active alerts

Display only alerts that are active in the server database by issuing the following command:

```
query alertstatus status=active
```

Query active alerts for two messages issued by the local server

Issue the following command to display only active alerts for message numbers ANE4958I and ANR4952E that were issued by the local server:

```
query alertstatus msgnum=4958,4952 status=active sourcetype=local
```

Query active alerts for messages ANR4958I and ANR4952E issued by a client

Issue the following command to display only active alerts for message numbers ANE4958I and ANE4952I that were issued by a client:

```
query alertstatus msgnum=4958,4952 status=active sourcetype=client
```

Query all alerts on a server

Issue the following command to display all alerts that are on the server:

```
query alertstatus
```

Example output: Display all the alerts that are on the server:

```

Alert Identifier: 83
Alert Message Number: 293
  Source Name: SEDONA
  Source Type: LOCAL
  First Occurrence: 03/07/2013 17:08:35
Most Recent Occurrence: 03/07/2013 17:08:35
  Count: 1
  Status: ACTIVE
  Last Status Change: 12/31/1969 17:00:00
  Category: INVENTORY
  Message: ANR0293I Reorganization for table AF_BITFILES started.
  Assigned:
  Resolved By:
  Remark:

Alert Identifier: 85
Alert Message Number: 293
  Source Name: SEDONA
  Source Type: LOCAL
  First Occurrence: 03/08/2013 05:45:00
Most Recent Occurrence: 03/08/2013 05:45:00
  Count: 1
  Status: ACTIVE
  Last Status Change: 12/31/1969 17:00:00
  Category: INVENTORY
  Message: ANR0293I Reorganization for table BF_AGGREGATED_BITFILES started.
  Assigned:
  Resolved By:
  Remark:

Alert Identifier: 1282
Alert Message Number: 293
  Source Name: ALPINE
  Source Type: LOCAL
  First Occurrence: 02/13/2013 15:47:50
Most Recent Occurrence: 02/13/2013 15:47:50
  Count: 1
  Status: CLOSED
  Last Status Change: 02/26/2013 09:46:39
  Category: INVENTORY
  Message: ANR0293I Reorganization for table TSMMON_ALERT started.
  Assigned:
  Resolved By:
  Remark:

Alert Identifier: 1792
Alert Message Number: 293
  Source Name: ALPINE
  Source Type: LOCAL
  First Occurrence: 02/19/2013 08:58:14
Most Recent Occurrence: 02/19/2013 08:58:14
  Count: 1
  Status: CLOSED
  Last Status Change: 03/01/2013 12:39:21
  Category: INVENTORY
  Message: ANR0293I Reorganization for table ACTIVITY_LOG started.
  Assigned:
  Resolved By:
  Remark:

```

Related commands

Table 213. Commands related to QUERY ALERTSTATUS

Command	Description
"DEFINE ALERTRIGGER (Define an alert trigger)" on page 112	Associates specified messages to an alert trigger.

QUERY ALERTSTATUS

Table 213. Commands related to **QUERY ALERTSTATUS** (continued)

Command	Description
"DELETE ALERTTRIGGER (Remove a message from an alert trigger)" on page 400	Removes a message number that can trigger an alert.
"QUERY ALERTTRIGGER (Query the list of defined alert triggers)" on page 672	Displays message numbers that trigger an alert.
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"UPDATE ALERTTRIGGER (Update a defined alert trigger)" on page 1165	Updates the attributes of one or more alert triggers.
"UPDATE ALERTSTATUS (Update the status of an alert)" on page 1168	Updates the status of a reported alert.


QUERY ASSOCIATION (Query client node associations with a schedule)

Use this command to display information about which client nodes are associated with one or more schedules. Client nodes associated with a schedule perform operations such as backup or archive according to that schedule.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

domain_name

Specifies the name of the policy domain to display. You can use a wildcard character to specify this name. All matching policy domain names are displayed. If you do not specify a value for this parameter, all existing policy domains are queried. If you specify a domain name, you do not have to specify a schedule name.

schedule_name

Specifies the name of the schedule to display. You can use a wildcard character to specify this name. All matching schedule names are displayed. If you do not specify a value for this parameter, all existing schedules are queried. If you specify a schedule name, you must also specify a policy domain name.

Example: Display client nodes that are associated with a schedule

Display all the client nodes that are associated with each schedule that belongs to the EMPLOYEE_RECORDS policy domain. Issue the command:

```
query association employee_records *
```

```

Policy Domain Name: EMPLOYEE_RECORDS
Schedule Name: WEEKLY_BACKUP
Associated Nodes: JOE JOHNSON LARRY SMITH SMITHERS TOM
  
```

See “Field descriptions” for field descriptions.

Field descriptions

Policy Domain Name

Specifies the name of the policy domain to which the schedule belongs.

Schedule Name

Specifies the name of the schedule.

Associated Nodes

Specifies the names of the client nodes that are associated with the specified schedule.

QUERY ASSOCIATION

Related commands

Table 214. Commands related to QUERY ASSOCIATION

Command	Description
DEFINE ASSOCIATION	Associates clients with a schedule.
DELETE ASSOCIATION	Deletes the association between clients and a schedule.

QUERY AUDITOCUPANCY (Query client node storage utilization)

Use this command to display information about client node server storage utilization. To display current license audit information from the server, use the **AUDIT LICENSE** command before you issue the **QUERY AUDITOCUPANCY** command.


As part of a license audit operation, the server calculates, by node, the amount of backup, archive, and space management storage in use. For servers that manage large amounts of data, this calculation can take a great deal of processor time and can stall other server activity. You can use the **AUDITSTORAGE** server option to specify that storage is not to be calculated as part of a license audit.

You can use the information from this query to determine if and where client node storage utilization must be balanced. This information can also assist you with billing clients for storage usage.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

node_name

Specifies a list of nodes for which to display server storage use information. Specify more than one node by separating the node names with commas, with no intervening spaces. You can use wildcard characters to specify names. The default (*) is to query all client nodes. Use the **DOMAIN** parameter to limit this list by policy domain. This parameter is optional.

D0main

Specifies a list of policy domains to restrict which nodes are displayed. Nodes belonging to the specified policy domains are displayed. Specify more than one policy domain by separating the policy domain names with commas, with no intervening spaces. You can use wildcard characters to specify names. This parameter is optional.

P0o1type

Specifies the type of storage pool to display. This parameter is optional. The default is **ANY**. Possible values are:

ANY

Specifies both primary and copy storage pools. The value that is presented is the total for the two pools.

QUERY AUDITOCUPANCY

PRimary

Specifies primary storage pools only.

COpy

Specifies copy storage pools only.

Example: Display storage usage

Display combined storage use in primary and copy storage pools. Issue the command:

```
query auditoccupancy
```

```
License information as of last audit on 05/22/1996 14:49:51.
```

Client Node Name	Backup Storage Used (MB)	Archive Storage Used (MB)	Space-Managed Storage Used (MB)	Total Storage Used (MB)
CLIENT	245	20	0	265
SMITH	245	20	0	265
SMITHERS	245	20	0	265
JOHNSON	300	15	0	320
JOE	245	20	0	265
TOM	300	15	0	320
LARRY	245	20	0	265

See “Field descriptions” for field descriptions.

Field descriptions

Client Node Name

Specifies the name of the client node.

Backup Storage Used (MB)

| Specifies the total backup storage use for the node. For this value, one MB
| = 1048576 bytes.

Archive Storage Used (MB)

| Specifies the total archive storage use for the node. For this value, one MB
| = 1048576 bytes.

Space-Managed Storage Used (MB)

| Specifies the amount of server storage that is used to store files that are
| migrated from the client node by a Tivoli Storage Manager for Space
| Management client. For this value, one MB = 1048576 bytes.

Total Storage Used (MB)

| Specifies the total storage use for the node. For this value, one MB =
| 1048576 bytes.

Related commands

Table 215. Commands related to QUERY AUDITOCUPANCY

Command	Description
AUDIT LICENSES	Verifies compliance with defined licenses.
QUERY LICENSE	Displays information about licenses and audits.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

Table 215. Commands related to QUERY AUDIT OCCUPANCY (continued)

Command	Description
REGISTER LICENSE	Registers a license with the Tivoli Storage Manager server.
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.


QUERY BACKUPSET (Query a backup set)

Use this command to display information about one or more backup sets.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

node_name or *node_group_name* (Required)

Specifies the name of the client node and node groups whose data is contained in the backup set to be displayed. To specify multiple node names and node group names, separate the names with commas and no intervening spaces. You can use wildcard characters with node names but not with node group names.

backup_set_name

Specifies the name of the backup set whose information is to be displayed. The backup set name you specify can contain wildcard characters. You can specify more than one backup set name by separating the names with commas and no intervening spaces.

BEGINDate

Specifies the beginning date of the range in which the point-in-time date of the backup set to be displayed must fall. This parameter is optional. You can use

this parameter with the **BEGINTIME** parameter to specify a range for the date and time. If you specify a begin date without a begin time, the time will be at 12:00 a.m. (midnight) on the date you specify.

You can specify the date by using one of the following values:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1999
TODAY	The current date	TODAY
TODAY+ <i>days</i> or + <i>days</i>	The current date plus days specified.	TODAY +3 or +3.
TODAY- <i>days</i> or - <i>days</i>	The current date minus days specified.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time of the range in which the point-in-time date of the backup set to be displayed must fall. This parameter is optional. You can use this parameter with the **BEGINDATE** parameter to specify a range for the date and time. If you specify a begin time without a begin date, the date will be the current date at the time you specify.

You can specify the time by using one of the following values:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08
NOW	The current time	NOW
NOW+ <i>HH:MM</i> or + <i>HH:MM</i>	The current time plus hours and minutes specified	NOW+02:00 or +02:00.
NOW- <i>HH:MM</i> or - <i>HH:MM</i>	The current time minus hours and minutes specified	NOW-02:00 or -02:00.

ENDDate

Specifies the ending date of the range in which the point-in-time date of the backup set to be displayed must fall. This parameter is optional. You can use this parameter with the **ENDTIME** parameter to specify an ending date and time. If you specify an end date without an end time, the time will be at 11:59:59 p.m. on the specified end date.

You can specify the date by using one of the following values:

QUERY BACKUPSET

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1999
TODAY	The current date	TODAY
TODAY+ <i>days</i> or + <i>days</i>	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
TODAY- <i>days</i> or - <i>days</i>	The current date minus days specified.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time of the range in which the point-in-time date of the backup set to be displayed must fall. This parameter is optional. You can use this parameter with the ENDDATE parameter to specify a date and time. If you specify an end time without an end date, the date will be the current date at the time you specify.

You can specify the time by using one of the following values:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08
NOW	The current time	NOW
NOW+ <i>HH:MM</i> or + <i>HH:MM</i>	The current time plus hours and minutes specified	NOW+02:00 or +02:00.
NOW- <i>HH:MM</i> or - <i>HH:MM</i>	The current time minus hours and minutes specified	NOW-02:00 or -02:00.

WHERERetention

Specifies the retention value, specified in days, that must be associated with the backup sets to be displayed. You can specify an integer from 0 to 30000. The values are:

days

Specifies that backup sets that are retained this number of days are displayed.

NOLimit

Specifies that backup sets that are retained indefinitely are displayed.

WHEREDescription

Specifies the description that must be associated with the backup set to be

displayed. The description you specify can contain wildcard characters. This parameter is optional. Enclose the description in quotation marks if it contains any blank characters.

WHEREDevice

Specifies the name of the device class that must be associated with the backup set to be displayed. You can use wildcard characters to specify a device class name. This parameter is optional.

WHEREToCexists

Specifies whether a backup set must have a table of contents in order to be displayed. This parameter is optional. The default is to display all backup sets whether or not they have a table of contents.

WHEREDataType

Specifies the data type of a backup set to be displayed. This parameter is optional. The default is to display all types of backup sets. To specify multiple data types, separate data types with commas and no intervening spaces.

FILE

Specifies that a file level backup set is to be displayed. File level backup sets contain files and directories backed up by the backup-archive client.

IMAGE

Specifies that an image backup set is to be displayed. Image backup sets contain images created by the backup-archive client **BACKUP IMAGE** command.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified backup sets.

Detailed

Specifies that complete information is displayed for the specified backup sets.

Example: Query a backup set

Display information for backup sets whose names begin with PERS_DATA. The backup sets belong to the node JANE and are assigned to the DVLMENT device class.

query backupset jane pers_data*

```

Node Name: JANE
Backup Set Name: PERS_DATA.3089
Data Type: File
Date/Time: 03/17/2007 16:17:47
Retention Period: 60
Device Class Name: DVLMENT
Description: backupset created from /srvr
Has Table of Contents (TOC)?: Yes
    
```

Field descriptions

Node Name

Specifies the name of the client node whose data is contained in the backup set.

QUERY BACKUPSET

Backup Set Name

Specifies the name of the backup set.

Data Type:

Displays the data type of the backup sets. Possible types are file, image, and application.

Date/Time

Specifies the date and time (PITDate and PITTime) of the **GENERATE BACKUPSET** command. The PITDate and PITTime specify that files that were active on the specified date and time and that are still stored on the Tivoli Storage Manager server are to be included in the backup set, even if they are inactive at the time you issue the **GENERATE BACKUPSET** command. The default is the date on which the **GENERATE BACKUPSET** command is run.

Retention Period

Specifies the number of days that the backup set is retained on the server.

Device Class Name

Specifies the name of the device class for which the volumes containing the backup set is assigned.

Description

Specifies the description associated with the backup set.

Has Table of Contents (TOC)?

Specifies whether the backup set has a table of contents.

Related commands

Table 216. Commands related to **QUERY BACKUPSET**

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
QUERY BACKUPSETCONTENTS	Displays contents contained in backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.

QUERY BACKUPSETCONTENTS (Query contents of a backup set)


Use this command to display information about the files and directories contained in a backup set for a client node.

Remember: Processing this command can use considerable network resources and mount points.

Privilege class

To issue this command, you must have system privilege or policy privilege for the domain to which the client node is assigned.

Syntax



Parameters

node_name (Required)

Specifies the name of the client node whose data is contained in the backup set to display. The name you specify cannot contain wildcard characters nor can it be a list of node names separated by commas.

backup_set_name (Required)

Specifies the name of the backup set to display. The name that you specify cannot contain wildcard characters nor can it be a list of node names that are separated by commas.

DATATYPE

Specifies that the backup set containing the specified types of data is to be queried. This parameter is optional. The default is that a file level backup set is to be queried. Possible values are:

FILE

Specifies that a file level backup set is to be queried. File level backup sets contain files and directories backed up by the backup-archive client.

IMAGE

Specifies that an image backup set is to be queried. Image backup sets contain images created by the backup-archive client **BACKUP IMAGE** command.

Example: Query contents of a backup set for a specific node

Display the contents from backup set named PERS_DATA.3099 belonging to client node JANE. Issue the command:

```
query backupsetcontents jane pers_data.3099
```

QUERY BACKUPSETCONTENTS

Node Name	Filespace Name	Client's Name for File
JANE	/srvr	/deblock
JANE	/srvr	/deblock.c
JANE	/srvr	/dsmerror.log
JANE	/srvr	/dsmxxxx.log
JANE

Field descriptions

Node Name

Specifies the name of the client node whose data is contained in the backup set.

Filespace Name

Specifies the name of the file space to which the specified file belongs.

File space names can be in a different code page or locale than the server. If they are, the names in the Administration Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, Tivoli Storage Manager converts the name to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

Client's Name for File

Specifies the name of the file.

File space names and file names that can be in a different code page or locale than the server do not display correctly on the Administration Center or the administrative command-line interface. The data itself is backed up and can be restored properly, but the file space or file name may display with a combination of invalid characters or blank spaces.

If the file space name is Unicode enabled, the name is converted to the server's code page for display. The results of the conversion for characters not supported by the current code page depends on the operating system. For names that Tivoli Storage Manager is able to partially convert, you may see question marks (?), blanks, unprintable characters, or "...". These characters indicate to the administrator that files do exist. If the conversion is not successful, the name is displayed as "...". Conversion can fail if the string includes characters that are not available in the server code page, or if the server has a problem accessing system conversion routines.

A file name that is displayed as "....." indicates that both the file path and file name were not successfully converted. An example of the path and name could be:

```
my\dir\...
```

Related commands

Table 217. Commands related to QUERY BACKUPSETCONTENTS

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
GENERATE BACKUPSET	Generates a backup set of a client's data.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
DELETE BACKUPSET	Deletes a backup set.
QUERY BACKUPSET	Displays backup sets.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.


QUERY CLOPTSET (Query a client option set)

Use this command to query a client option set.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

option_set_name

Specifies the name of the client option set to be queried. You can use wildcard characters to specify this name. This parameter is optional. The default is option set names.

DESCRiption

Specifies the description used on the **DEFINE** or **UPDATE CLOPTSET** commands to be used as a filter. If the description contains spaces, enclose it in quotation marks. This parameter is optional.

Example: Query a client option set

From a managed server, query a client option set named ENG. Issue the following command:

```
query cloptset eng
```

```

      Optionset:  ENG
      Description:
Last Update by (administrator): $$CONFIG_MANAGER$$
      Managing profile:

      Option:  SCROLLINES
      Sequence number: 0
Use Option Set Value (FORCE): No
      Option Value: 40

      Option:  SCROLLPROMPT
      Sequence number: 0
Use Option Set Value (FORCE): No
      Option Value: yes
  
```

Field descriptions

Optionset

Specifies the name of the option set.

Description

Specifies the description of the client option set.

Last Update by (administrator)

Specifies the name of the administrator that most recently updated the option set. If this field contains \$\$CONFIG_MANAGER\$\$, the client option set is associated with a profile that is managed by the configuration manager.

Managing Profile

Specifies the profile to which the managed server subscribed to get the definition of the client option set.

Option

Specifies the name of the option.

Sequence Number

Specifies the sequence number of the option.

Use Option Set Value (FORCE)

Specifies whether the server option setting overrides the client's option setting. NO indicates the server option setting does not override the client option. YES indicates the server's option setting overrides the client option setting. This option is set with the **FORCE** parameter on the **DEFINE CLIENTOPT** command.

Option Value

Specifies the value of the option.

Related commands

Table 218. Commands related to QUERY CLOPTSET

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLIENTOPT	Adds a client option to a client option set.
DEFINE CLOPTSET	Defines a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
DELETE CLOPTSET	Deletes a client option set.
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.
UPDATE CLOPTSET	Updates the description of a client option set.
DEFINE PROFASSOCIATION	Associates objects with a profile.


QUERY COLLOGROUP (Query a collocation group)

Use this command to display the collocation groups defined on the server.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

group_name

Specifies the name of the collocation group to display. To specify multiple names, use a wildcard character. This parameter is optional. The default is to display all collocation groups.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed. To display the members of the collocation group, you must specify FORMAT=DETAILED.

Example: Display defined collocation groups

Display the collocation groups defined on the server. Issue the following command:

```
query collogroup
```

Collocation Group Name	Collocation Group Description
DEPT ED	Education department
GROUP1	Low cap client nodes.

See “Field descriptions” on page 695 for field descriptions.

Example: Display detailed information for collocation groups

Display complete information about all collocation groups and determine which client nodes belong to which collocation groups. Issue the following command:

```
query collogroup format=detailed
```

```

Collocation Group Name: DEPT_ED
Collocation Group Description: Education department
Last Update by (administrator): SERVER_CONSOLE
    Last Update Date/Time: 04/21/2004 10:59:03
Collocation Group Member(s): EDU_1 EDU_7

Collocation Group Name: GROUP1
Collocation Group Description: Low cap client nodes.
Last Update by (administrator): SERVER_CONSOLE
    Last Update Date/Time: 04/21/2004 10:59:16
Collocation Group Member(s): CHESTER REX NOAH JARED
    
```

See “Field descriptions” for field descriptions.

Field descriptions

Collocation Group Name

The name of the collocation group.

Collocation Group Description

The description for the collocation group.

Last Update by (administrator)

The name of the administrator that defined or most recently updated the collocation group.

Last Update Date/Time

The date and time that an administrator defined or most recently updated the collocation group.

Collocation Group Member(s)

The members of the collocation group.

Related commands

Table 219. Commands related to QUERY COLLOGROUP

Command	Description
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node to a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node from a collocation group.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY STGPOOL	Displays information about storage pools.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
UPDATE COLLOGROUP	Updates the description of a collocation group.
UPDATE STGPOOL	Changes the attributes of a storage pool.

QUERY COLLOGROUP

QUERY CONTENT (Query the contents of a storage pool volume)

Use this command to display information about files in a storage pool volume, as well as the names of client files that link to a deduplicated group of files.

You can use this command to identify files that the server has found to be damaged and files that have been backed up to a copy storage pool or copied to an active-data pool. This command is useful when a volume is damaged or before you:

- Request the server to fix inconsistencies between a volume and the database
- Move files from one volume to another volume
- Delete a volume from a storage pool


Because this command can take a long time to run and the results can be very large, consider using the **COUNT** parameter to limit the number of files displayed.

Note: Files that are cached in a disk volume and that are marked as damaged are not included in the results.


Privilege class

Any administrator can issue this command.

Syntax



QUERY CONTENT



Notes:

- 1 Use this parameter only for volumes in primary storage pools.

Parameters

volume_name (Required)

Specifies the volume to be queried.

NODE

Specifies the backup-archive client or the Tivoli Storage Manager for Space Management associated with the file space to query. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a name, all backup-archive and Tivoli Storage Manager for Space Management clients are included.

FILESPACE

Specifies the file space to query. This parameter is optional. You can use wildcard characters to specify this name. File space names are case-sensitive. If you do not specify a file space name, all file spaces are included.

For a server that has clients with Unicode support, you may need to have the server convert the file space name that you enter. For example, you may need to have the server convert the name you enter from the server's code page to Unicode. See the **NAMETYPE** parameter for details. If you do not specify a file space name or specify only a single wildcard character for the name, you can use the **CODETYPE** parameter to limit the operation to Unicode file spaces or non-Unicode file spaces.

COUNT

Specifies the number of files to be displayed. This parameter is optional. You can specify either a positive integer or a negative integer. If you specify a positive integer, *n*, the first *n* files are displayed. If you specify a negative integer, *-n*, the last *n* files are displayed in *reverse* order. You cannot specify COUNT=0. If you do not specify a value for this parameter, all files are displayed.

Type

Specifies the types of files to query. This parameter is optional. The default value is ANY. If the volume being queried is assigned to an active-data pool, the only valid values are ANY and BACKUP. Possible values are:

ANY

Specifies that all types of files in the storage pool volume are queried: backup versions of files, archived copies of files, and files migrated by Tivoli Storage Manager for Space Management clients from client nodes.

Backup

Specifies that only backup files are queried.

Archive

Specifies that only archive files are queried. This value is not valid for active-data pools.

SPacemanaged

Specifies that only space-managed files (files that were migrated by a Tivoli Storage Manager for Space Management client) are queried. This value is not valid for active-data pools.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed. Unicode names are converted to the server code page.

Detailed

Specifies that complete information is displayed. Unicode names are displayed in hexadecimal.

DAmaged

Specifies criteria to restrict the query output based on whether files have been marked as damaged. For purposes of this criteria, the server examines only physical files (a file that may be a single logical file or an aggregate that consists of logical files). This parameter is optional. The default value is ANY. Possible values are:

ANY

Specifies that files are displayed regardless of whether the server has found the files to be damaged.

Yes

Specifies that only files that have been marked as damaged are displayed. These are files in which the server found errors when a user attempted to restore, retrieve, or recall the file, or when an **AUDIT VOLUME** command was run.

No Specifies that only files *not* known to have been damaged are displayed.

COPIed

Specifies criteria to restrict the query output based on whether files have been backed up to a copy storage pool. Whether files are stored in an active-data pool does not effect the output. This parameter is optional. The default value is ANY. Possible values are:

ANY

Specifies that files are displayed regardless of whether the files are backed up to a copy storage pool. Primary and cached file copies are displayed.

Yes

Specifies that the files displayed are only those for which at least one usable backup copy exists in a copy storage pool. A file is not displayed if its copy in the copy storage pool is known to have errors. Cached file copies are not displayed because these files are never restored.

Use **COPIED=YES** to identify primary files that can be restored using the **RESTORE VOLUME** or **RESTORE STGPOOL** command.

QUERY CONTENT

No Specifies that the files displayed are only those for which no usable backup copies exist in a copy storage pool. Cached file copies are not displayed because these files are never restored.

Use **COPIED=NO** to identify primary files that cannot be restored using the **RESTORE VOLUME** or **RESTORE STGPOOL** command.

NAMEType

Specify how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with Unicode support. A backup-archive client with Unicode support is currently available only for Windows, Macintosh OS 9, Macintosh OS X, and NetWare. Use this parameter only when you specify a partly or fully qualified file space name.

The default value is **SERVER**. Possible values are:

SERVER

The server uses the server's code page to interpret the file space names.

UNICODE

The server converts the file space names from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the names and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server has a problem accessing system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

CODEType

Specify how you want the server to interpret the file space names that you enter. Use this parameter only when you enter a single wildcard character for the file space name.

The default value is **BOTH**, which means that the file spaces are included regardless of code page type. Possible values are:

UNICODE

Include only file spaces that are in Unicode.

NONUNICODE

Include only file spaces that are not in Unicode.

BOTH

Include file spaces regardless of code page type.

FOLLOWLinks

Specifies whether to display only the files that are stored on the volume or only files that are linked to the volume. You can also display both stored files and linked files. The default is **NO**. Possible values are:

No Display only the files that are stored in the volume. Do not display files that have links to the volume.

Yes

Display all files, including files that are stored on the volume and any files that have links to the volume.

JUSTLinks

Display only the files that have links to the volume. Do not display files that are stored on the volume.

Example: Display the contents of a volume for a specific client node

Query the contents of a volume and limit the results to files backed up from the PEGASUS client node.

For the volume /tsmstg/diskvol1.dsm, issue the command:

```
query content /tsmstg/diskvol1.dsm node=pegasus
type=backup
```

Results of the command include all logical files that make up any aggregate that is on the volume, even if the aggregate is stored on more than this volume. For aggregates, the query does not determine which logical files are actually stored on the volume for which the query is performed.

Node Name	Type	Filespace Name	FSID	Client's Name for File
PEGASUS	Bkup	\\pegasus\e\$	1	\UNI_TEST\ SM01.DAT
PEGASUS	Bkup	\\pegasus\e\$	1	\UNI_TEST\ SM02.DAT

See "Field descriptions" on page 702 for field descriptions.

Example: Display detailed information for a tape volume

Query the contents of the tape volume named WPD001. Display only files that are backed up by the node MARK, and files that are either stored on the volume or linked to the volume. Display only the first four files on the volume.

```
query content wpd001 node=mark count=4 type=backup followlinks=yes
format=detailed
```

QUERY CONTENT

```
Node Name: MARK
Type: Bkup
Filespace Name: \\mark\e$
Hexadecimal Filespace Name:
FSID: 1
Client's Name for File: \UNI_TEST\ SM01.DAT
Hexadecimal Client's Name for File:
Aggregated?: 1/3
Stored Size: 2,746
Segment Number:
Cached Copy?: No
Linked?: No

Node Name: MARK
Type: Bkup
Filespace Name: \\mark\e$
Hexadecimal Filespace Name:
FSID: 1
Client's Name for File: \UNI_TEST\ SM02.DAT
Hexadecimal Client's Name for File:
Aggregated?: 2/3
Stored Size: 2,746
Segment Number:
Cached Copy?: No
Linked?: No

Node Name: MARK
Type: Bkup
Filespace Name: \\mark\e$
Hexadecimal Filespace Name:
FSID: 1
Client's Name for File: \UNI_TEST\ SM03.DAT
Hexadecimal Client's Name for File:
Aggregated?: 3/3
Stored Size: 2,746
Segment Number:
Cached Copy?: No
Linked?: No

Node Name: HOWARD
Type: Bkup
Filespace Name: \\howard\f$
Hexadecimal Filespace Name:
FSID: 2
Client's Name for File: \UNI_TEST\ HOWARD.C
Hexadecimal Client's Name for File:
Aggregated?: 1/3
Stored Size: 2,372
Segment Number:
Cached Copy?: No
Linked?: Yes
```

See “Field descriptions” for field descriptions.

Field descriptions

Node Name

The node to which the file belongs.

Type The type of file: archive (Arch), backup (Bkup), or space-managed (SpMg) by a Tivoli Storage Manager for Space Management client.

Filespace Name

The file space to which the file belongs.

File space names can be in a different code page or locale than the server. If they are, the names in the Administration Center and the administrative command-line interface might not be displayed correctly. Data is backed

up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, Tivoli Storage Manager converts the name to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

Hexadecimal Filespace Name

The file space to which the file belongs. If the file space name is in Unicode, the name is displayed in hexadecimal format.

FSID The file space ID (FSID) for the file space. The server assigns a unique FSID when a file space is first stored on the server.

Client's Name for File

The client's name for the file.

File space names and file names that can be in a different code page or locale than the server do not display correctly on the Administration Center or the administrative command-line interface. The data itself is backed up and can be restored properly, but the file space or file name may display with a combination of invalid characters or blank spaces. The results of the conversion for characters not supported by the current code page depends on the operating system. For names that Tivoli Storage Manager is able to partially convert, you may see question marks (??), blanks, unprintable characters, or "...". These characters indicate to the administrator that files do exist.

Hexadecimal Client's Name for File

The client's name for the file displayed in hexadecimal format.

Aggregated?

Whether the file is a logical file that is stored as part of an aggregate. If the file is part of an aggregate, the sequence of this file within the aggregate and the total number of logical files in the aggregate are displayed. Results of the command include all logical files that make up any aggregate that is on the volume, even if the aggregate is stored on more than this volume. The query does not determine which logical files are actually stored on the volume for which the query is performed.

If the file is not part of an aggregate, the field simply displays "no."

Stored Size

The size of the physical file, in bytes. If the file is a logical file that is stored as part of an aggregate, this value indicates the size of the entire aggregate.

Segment Number

For volumes in sequential-access storage pools, specifies whether the physical file (either a single logical file or an aggregate of logical files) is stored across multiple volumes. For example, if the logical file is stored in an aggregate that spans two volumes, the segment number indicates 1/2 (the first part of the physical file is stored on the volume) or 2/2 (the second part of the physical file is stored on the volume). If the segment

QUERY CONTENT

number is 1/1, the physical file is completely stored on the volume. For volumes in random-access storage pools, no value is displayed for this field.

Cached Copy?

Whether the physical file is a cached copy of a file migrated to the next storage pool. If the file is part of an aggregate, this value pertains to the aggregate.

Linked?

Indicates whether the file is stored on the volume or whether the file is linked to the volume.

Related commands

Table 220. Commands related to QUERY CONTENT

Command	Description
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
COPY ACTIVEDATA	Copies active backup data.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DELETE VOLUME	Deletes a volume from a storage pool.
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.
UPDATE VOLUME	Updates the attributes of storage pool volumes.


QUERY COPYGROUP (Query copy groups)

Use this command to display information about one or more copy groups.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

domain_name

Specifies the policy domain associated with the copy group to query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a value for this parameter, all policy domains are queried. You must specify this parameter when querying an explicitly named copy group.

policy_set_name

Specifies the policy set associated with the copy group to query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a value for this parameter, all policy sets are queried. You must specify this parameter when querying an explicitly named copy group.

class_name

Specifies the management class associated with the copy group to query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a value for this parameter, all management classes are queried. You must specify this parameter when querying an explicitly named copy group.

STANDARD

Specifies the name of the copy group. This parameter is optional. The name of the copy group must be STANDARD. The default is STANDARD.

Type

Specifies the type of copy group to be queried. This parameter is optional. The default value is BACKUP. Possible values are:

Backup

Specifies that you want to query backup copy groups.

QUERY COPYGROUP

Archive

Specifies that you want to query archive copy groups.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display detailed information on one backup copy group

Display complete information on the backup copy group assigned to the ACTIVEFILES management class in the VACATION policy set of the EMPLOYEE_RECORDS policy domain. Issue the command:

```
query copygroup employee_records vacation
activefiles format=detailed
```

```
Policy Domain Name: EMPLOYEE_RECORDS
Policy Set Name: VACATION
Mgmt Class Name: ACTIVEFILES
Copy Group Name: STANDARD
Copy Group Type: Backup
Versions Data Exists: 5
Versions Data Deleted: 1
Retain Extra Versions: 30
Retain Only Version: 60
Copy Mode: Absolute
Copy Serialization: Shared Static
Copy Frequency: 3
Copy Destination: BACKUPPOOL
Table of Contents (TOC) Destination:
Last Update by (administrator): ADMIN
Last Update Date/Time: 10/02/2002 17.51.49
```

Example: Display information on the backup copy group in the STANDARD management class and policy set

From a managed server, display complete information on the backup copy group assigned to the STANDARD management class in the STANDARD policy set of the ADMIN_RECORDS policy domain. Issue the command:

```
query copygroup admin_records
standard standard format=detailed
```

```

Policy Domain Name: ADMIN_RECORDS
Policy Set Name: STANDARD
Mgmt Class Name: STANDARD
Copy Group Name: STANDARD
Copy Group Type: Backup
Versions Data Exists: 2
Versions Data Deleted: 1
Retain Extra Versions: 30
Retain Only Version: 60
Copy Mode: Modified
Copy Serialization: Shared Static
Copy Frequency: 0
Copy Destination: BACKUPPOOL
Table of Contents (TOC) Destination:
Last Update by (administrator): $$CONFIG_MANAGER$$
Last Update Date/Time: 2002.10.02 17.51.49
Managing profile: ADMIN_INFO
Changes Pending: Yes

```

Example: Display information on an archive copy group

From a managed server, display complete information on the archive copy group STANDARD that is assigned to the MCLASS1 management class in the SUMMER policy set of the PROG1 policy domain. Issue the command:

```

query copygroup prog1 summer mclass1
type=archive format=detailed

```

```

Policy Domain Name: PROG1
Policy Set Name: SUMMER
Mgmt Class Name: MCLASS1
Copy Group Name: STANDARD
Copy Group Type: Archive
Retain Version: 730
Retention Initiation: Creation
Minimum Retention:
Copy Serialization: Shared Static
Copy Frequency: Cmd
Copy Mode: Absolute
Copy Destination: ARCHPOOL
Last Update by (administrator): $$CONFIG_MANAGER$$
Last Update Date/Time: 2002.10.02 17.42.49
Managing profile: ADMIN_INFO

```

Example: Display information on the copy group for a NAS backup

Query the copy group for the NAS backup. Issue the command:

```

query copygroup nasdomain
type=backup

```

QUERY COPYGROUP

```
Policy Domain Name: NASDOMAIN
Policy Set Name: ACTIVE
Mgmt Class Name: STANDARD
Copy Group Name: STANDARD
Copy Group Type: Backup
Versions Data Exists: 2
Versions Data Deleted: 1
Retain Extra Versions: 30
Retain Only Version: 60
Copy Mode: Modified
Copy Serialization: Shared Static
Copy Frequency: 0
Copy Destination: NASPOOL
Table of Contents (TOC) Destination: BACKUPOOL
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 10/02/2002 12:16:52
Managing profile:
Changes Pending: Yes
```

Field descriptions

Policy Domain Name

The name of the policy domain.

Policy Set Name

The name of the policy set.

Mgmt Class Name

The name of the management class.

Copy Group Name

The name of the copy group. This name is always STANDARD.

Copy Group Type

The type of the copy group.

Versions Data Exists

The maximum number of backup versions to retain for files that are currently on the client file system.

Versions Data Deleted

The maximum number of backup versions to retain for files that have been deleted from the client file system after being backed up using Tivoli Storage Manager.

Retain Extra Versions

The number of days to retain a backup version after that version becomes inactive.

Retain Only Versions

The number of days to retain the last backup version of a file that has been deleted from the client file system.

Retain Version

The number of days to keep an archive copy.

Retention Initiation

The time which the server initiates the retention time specified by the RETAIN VERSION parameter. CREATION specifies that the retention time is initiated at the time an archive copy is stored on the server. EVENT specifies that the retention time is initiated at the time the server is notified of a retention-initiating event for the archive copy.

Minimum Retention

The minimum number of days to keep an archive copy when Retention Initiation is EVENT. The value of this parameter is not displayed when Retention Initiation is CREATION.

Copy Serialization

Whether a file can be in use during an archive operation.

Copy Frequency

The copy frequency of the copy group. For archive copy groups, this value is always CMD.

Copy Mode

Specifies that files in the copy group are archived regardless of whether they have been modified. For archive copy groups, this value is always ABSOLUTE.

Copy Destination

The name of the storage pool where the server initially stores files associated with this archive copy group.

Table of Contents (TOC) Destination

The name of the primary storage pool in which TOCs are initially stored for image backup operations in which TOC generation is requested.

Last Update by (administrator)

The name of the administrator or server that most recently updated the copy group. If this field contains \$\$CONFIG_MANAGER\$\$, the copy group is associated with a domain that is managed by the configuration manager.

Last Update Date/Time

The date and time when the copy group was most recently defined or updated.

Managing Profile

The profile or profiles to which the managed server subscribed to get the definition of this policy copy group.

Changes Pending

Whether or not changes are being made but not activated. Once the changes are activated, the field resets to No.

Related commands

Table 221. Commands related to QUERY COPYGROUP

Command	Description
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DELETE COPYGROUP	Deletes a backup or archive copy group from a policy domain and policy set.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.


QUERY DATAMOVER (Display data mover definitions)

Use this command to display data mover definitions.

Privilege class

Any administrator can issue this command.

Syntax



Notes:

- 1 The type must be specified if FORMAT=DETAILED

Parameters

data_mover_name

Specifies the name of the data mover to display. You can specify multiple names with a wildcard character. The default displays all data movers.

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that the assigned name and addressing information is specified.

Detailed

Specifies that complete information is displayed.

Type

Specifies the type of data mover to be displayed. This parameter must be specified if the FORMAT=DETAILED.

NAS

Specifies a NAS device.

Example: Display information on all data movers

Display the data movers on the server. Issue the command:

```
query datamover
```

Data Mover Name	Type	Online
NASMOVER1	NAS	Yes
NASMOVER2	NAS	No

See “Field descriptions” on page 711 for field descriptions.

Example: Display information for one data mover

Display partial information on data mover DATAMOVER6. Issue the command:
`query datamover datamover6 type=nas`

Source Name	Type	Online
-----	-----	-----
DATAMOVER6	NAS	Yes

See “Field descriptions” for field descriptions.

Example: Display detailed information on one data mover

Display detailed information on data mover DATAMOVER6. The TYPE parameter is required when FORMAT=DETAILED. Issue the command:

`query datamover datamover6 format=detailed type=nas`

Data Mover Name:	DataMover6
Type:	NAS
HLAddress:	9.115.38.5
LLAddress:	10000
USERid:	NDMPadmin
Storage Pool Data Format:	NDMPDUMP
Online:	Yes
Last Update by (administrator):	ADMIN
Last Update Date/Time:	04/28/2003 09:26:33

See “Field descriptions” for field descriptions.

Field descriptions

Data Mover Name

Specifies the name of the data mover.

Type Specifies the type of the data mover.

HLAddress

Specifies the IP address for the data mover.

LLAddress

Specifies the TCP port number for the data mover.

User ID

Specifies the user ID that the server uses to get access to the data mover.

Storage Pool Data Format

Specifies the data format that is used by this data mover.

Online

Specifies whether the data mover is online and available for use.

Last Update by (administrator)

Specifies the ID of the administrator who performed the last update.

Last Update Date/Time

Specifies the date and time when the last update occurred.

QUERY DATAMOVER

Related commands

Table 222. Commands related to QUERY DATAMOVER

Command	Description
DEFINE DATAMOVER	Defines a data mover to the Tivoli Storage Manager server.
DELETE DATAMOVER	Deletes a data mover.
UPDATE DATAMOVER	Changes the definition for a data mover.


QUERY DB (Display database information)

Use this command to display information about the database.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display summary statistics about the database

Display statistical information about the database. Issue the command:

```
query db
```

Database Name	Total Pages	Usable Pages	Used Pages	Free Pages
TSMDB1	32,776	32,504	24,220	8,284

See “Field descriptions” on page 714 for field descriptions.

Example: Display detailed database information

Display detailed statistical information about the database. Issue the command:

```
query db format=detailed
```

```

Database Name: TSM_DB2
Total Size of File System (MB): 1,748,800
Space Used by Database(MB): 448
Free Space Available (MB): 235,609
Total Pages: 32,776
Usable Pages: 32,504
Used Pages: 24,220
Free Pages: 8,284
Buffer Pool Hit Ratio: 99.3
Total Buffer Requests: 204,121
Sort Overflows: 0
Package Cache Hit Ratio: 89.8
Last Database Reorganization: 05/25/2009 16:44:06
Full Device Class Name: FILE
Number of Backup Streams: 4
Incrementals Since Last Full: 0
Last Complete Backup Date/Time: 05/18/2009 22:55:19
  
```

See “Field descriptions” for field descriptions.

Field descriptions

Database Name

The name of the database that is defined and configured for use by the Tivoli Storage Manager server.

Total Size of File System (MB)

The total size, in megabytes, of the file systems in which the database is located.

Space Used by Database (MB)

The amount of database space, in megabytes, that is in use.

Free Space Available (MB)

The amount of database space, in megabytes, that is not in use.

Total Pages

The total number of pages in the table space.

Usable Pages

The number of usable pages in the table space.

Used Pages

The number of used pages in the table space.

Free Pages

The total number of free pages in all table spaces. The Tivoli Storage Manager database has up to 10 table spaces.

Buffer Pool Hit Ratio

The total hit ratio percent.

Total Buffer Requests

The total number of buffer pool data logical reads and index logical reads since the last time the database was started or since the database monitor was reset.

Sort Overflows

The total number of sorts that ran out of the sort heap and might have required disk space for temporary storage.

Package Cache Hit Ratio

A percentage indicating how well the package cache is helping to avoid reloading packages and sections for static SQL from the system catalogs. It also indicates how well the package cache is helping to avoid recompiling dynamic SQL statements. A high ratio indicates it is successful in avoiding these activities.

Last Database Reorganization

The last time that the database manager performed an automatic reorganization activity.

Full Device Class Name

The name of the device class this is used for full database backups.

Number of Backup Streams

The number of concurrent data movement streams that were used during the database backup.

Incrementals Since Last Full

The number of incremental backups that were performed since the last full backup.

Last Complete Backup Date/Time

The date and time of the last full backup.

Related commands

Table 223. Commands related to QUERY DB

Command	Description
BACKUP DB	Backs up the Tivoli Storage Manager database to sequential access volumes.
EXTEND DBSPACE	Adds directories to increase space for use by the database.
QUERY DBSPACE	Displays information about the storage space defined for the database.

QUERY DBSPACE (Display database storage space)

Use this command to display information about the directories used by the database to store data.

Privilege class

Any administrator can issue this command.

Syntax

►►—QUERY DBSpace—◄◄

Parameters

None.

Example: Display database storage space information

Display information about database storage space. Issue the command:

```
query dbspace
```

Location	Total Size of File System (MB)	Space Used on File System (MB)	Free Space Available (MB)
/tsmdb001	1,748,800	1,513,191.125	117,804.422
/tsmdb002	1,748,800	1,513,191.125	117,804.422

See “Field descriptions” for field descriptions.

Field descriptions

Location

Specifies the locations of database directories.

Total Size of File System (MB)

The total size, in megabytes, of the file system in which the database is located.

Space Used on File System (MB)

The amount of storage space, in megabytes, that is in use.

When you run the **QUERY DBSPACE** command, the value in the output might be greater than the value that is obtained by running the `df` system command. The output from the `df` system command does not include the amount of space that is reserved for the root user.

If you run the `df` system command, the default percentage of space that is reserved for the root user is 5%. You can change this default value.

Free Space Available(MB)

The amount of space, in megabytes, that is not in use.

Related commands*Table 224. Commands related to QUERY DBSPACE*

Command	Description
BACKUP DB	Backs up the Tivoli Storage Manager database to sequential access volumes.
EXTEND DBSPACE	Adds directories to increase space for use by the database.
QUERY DB	Displays allocation information about the database.


QUERY DEVCLASS (Display information on one or more device classes)

Use this command to display information on one or more device classes.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

device_class_name

Specifies the name of the device class to be queried. This parameter is optional. You can use wildcard characters to specify this name. All matching device classes are displayed. If you do not specify a value for this parameter, all device classes are displayed.

Format

Specifies how the information is displayed. This parameter is optional. The default value is **STANDARD**. Possible values are:

Standard

Specifies that partial information is displayed for the specified device class.

Detailed

Specifies that complete information is displayed for the specified device class.

Example: List all device classes

Display information on all device classes.

```
query devclass
```

Device Class Name	Device Access Strategy	Storage Pool Count	Device Type	Format	Est/Max Capacity (MB)	Mount Limit
8MMTAPE	Sequential	1	8MM	DRIVE	6,144.0	2
DISK	Random	4				
PLAINFILES	Sequential	1	FILE		50.0	1
8MMSP2	Sequential	2	8MM	DRIVE	44.4	DRIVES

See “Field descriptions” on page 720 for field descriptions.

Example: Display detailed information for a specific FILE device class

Display information in full detail on the PLAINFILES device class.

```
query devclass plainfiles format=detailed
```

```

Device Class Name: PLAINFILES
Device Access Strategy: Sequential
Storage Pool Count: 1
Device Type: FILE
Format:
Est/Max Capacity (MB): 50.0
Mount Limit: 1
Mount Wait (min):
Mount Retention (min):
Label Prefix:

Library:
Directory:
Server Name:
Retry Period:

Retry Interval:
Shared:
Primary Allocation (MB):
Secondary Allocation (MB):
Compression:
Retention:
Protection:
Expiration Date:
Unit:
Logical Block Protection:
Last Update by (administrator): ADMIN
Last Update Date/Time: 05/31/2000 13:15:36

```

See “Field descriptions” on page 720 for field descriptions.

Example: Display detailed information for a specific 3592 device class

Display full details on the 3592 device class.

```
query devclass 3592 format=detailed
```

QUERY DEVCLASS

```
Device Class Name: 3592
Device Access Strategy: Sequential
Storage Pool Count: 1
  Device Type: 3592
  Format: 3592
Est/Max Capacity (MB):
  Mount Limit: DRIVES
  Mount Wait (min): 60
  Mount Retention (min): 60
  Label Prefix: ADSM

Library: MANLIB
Directory:
Server Name:
Retry Period:

Retry Interval:
  Shared:
  HLAddr:
  WORM: No
  Scaled Capacity: 90
  Drive Encryption: On
Primary Allocation (MB):
Secondary Allocation (MB):
  Compression:
  Retention:
  Protection:
  Expiration Date:
  Unit:
Logical Block Protection: Read/Write
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 08/04/03 14:28:31
```

See “Field descriptions” for field descriptions.

Field descriptions

Device Class Name

The name of the device class.

Device Access Strategy

How data is written to the device class.

Storage Pool Count

The number of storage pools that are assigned to the device class.

Device Type

The device type of the device class.

Format

The recording format.

Est/Max Capacity (MB)

The estimated or maximum capacity of a volume associated with the device class.

Mount Limit

The maximum number of sequential access volumes that can be mounted concurrently or specifies that DRIVES is the mount limit.

Mount Wait (min)

The maximum number of minutes to wait for a sequential access volume to be mounted.

Mount Retention (min)

The number of minutes to retain an idle sequential access volume before dismounting it.

Label Prefix

The high-level qualifier of the data set name that the server writes into the sequential access media labels.

Library

The name of the defined library object that contains the drives used by the device class.

Directory

The directory or directories for a shared FILE device class.

Server Name

The name of a defined server.

Retry Period

The interval over which the server should attempt to contact a target server if communications failure is suspected.

Retry Interval

How often the retries are done within a given retry period.

Two-Sided

Whether a removable file is two-sided.

Shared

Whether this FILE device class will be shared between the server and one or more storage agents.

HLAddress

The IP address of the device in dotted decimal format.

WORM

Whether this drive is a write once, read many (WORM) device.

Scaled Capacity

The percentage of the media capacity that can be used to store data.

Drive Encryption

Whether drive encryption is permitted. This field applies only to volumes in a storage pool associated with a device type of 3592, LTO, or ECARTRIDGE.

Primary Allocation

For **FILE** device classes that represent storage managed by a z/OS media server, this specifies the initial amount of space that is dynamically allocated when a new volume is opened.

Secondary Allocation

For **FILE** device classes that represent storage managed by a z/OS media server, this specifies the amount of space by which a file volume is extended when space that is already allocated to the file volume is used up.

Compression

For tape device classes that represent storage managed by a z/OS media server, this specifies whether the data is compressed.

Retention

For tape device classes that represent storage managed by a z/OS media server, this specifies the number of days to retain the tape, if retention is used.

QUERY DEVCLASS

Protection

For tape device classes that represent storage managed by a z/OS media server, this specifies whether the volumes are protected by the RACF program.

Expiration Date

For tape device classes that represent storage managed by a z/OS media server, this specifies the expiration date that is placed on the tape labels for this device class, if expiration is used.

Unit For tape device classes that represent storage managed by a z/OS media server, this specifies the esoteric unit name for the group of tape devices.

Logical Block Protection

Specifies whether logical block protection is enabled and, if it is, the mode. Possible values are Read/Write, Write-only, and No. You can use logical block protection only with the following types of drives and media:

- IBM LTO5 and later
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later
- Oracle StorageTek T10000C drives

Last Update by (administrator)

The administrator that made the last update to the device class.

Last Update Date/Time

The date and time of the last update.

Related commands

Table 225. Commands related to QUERY DEVCLASS

Command	Description
DEFINE DEVCLASS	Defines a device class.
DEFINE DEVCLASS (z/OS media server)	Defines a device class to use storage managed by a z/OS media server.
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE DEVCLASS	Deletes a device class.
QUERY DIRSPACE	Displays information about FILE directories.
QUERY SERVER	Displays information about servers.
UPDATE DEVCLASS	Changes the attributes of a device class.
UPDATE DEVCLASS (z/OS media server)	Changes the attributes of a device class for storage managed by a z/OS media server.

QUERY DIRSPACE (Query storage utilization of FILE directories)

Use this command to display information about free space in the directories associated with a device class with a device type of FILE.

Privilege class

Any administrator can issue this command.

Syntax

```

>> Query DIRSPace [device_class_name]
  
```

Parameters

device_class_name

Specifies the name of the device class to be queried. This parameter is optional. You can use wildcard characters to specify this name. All matching device classes of device type FILE are displayed. If you do not specify a value for this parameter, all device classes of device type FILE are displayed.

Example: List FILE type device classes

Display information for all device classes with a device type of FILE. In the following example the unit M is equivalent to megabytes, and the unit G is equivalent to gigabytes.

```
query dirspace
```

Device Class	Directory	Estimated Capacity	Estimated Available
DBBKUP	/This/is/a/large/directory	13,000 M	5,543 M
DBBKUP	/This/is/directory2	13,000 M	7,123 M
DBBKUP2	/This/is/a/huge/directory	2,256 G	2,200 G

Related commands

Table 226. Commands related to QUERY DIRSPACE

Command	Description
DEFINE DEVCLASS	Defines a device class.
DELETE DEVCLASS	Deletes a device class.
QUERY DEVCLASS	Displays information about device classes.
UPDATE DEVCLASS	Changes the attributes of a device class.


QUERY DOMAIN (Query a policy domain)

Use this command to display information on one or more policy domains.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

domain_name

Specifies the policy domain to query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a value for this parameter, all policy domains are displayed.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display a summary of policy domains

Display partial information for all policy domains on the server. Issue the command:

```
query domain
```

Policy Domain Name	Activated Policy Set	Activated Default Mgmt Class	Number of Registered Nodes	Description
EMPLOYEE-RECORDS	VACATION	ACTIVEFI-LES	6	Employee Records Domain
PROG1			0	Programming Group Test Domain
PROG2			0	Programming Group Test Domain
STANDARD	STANDARD	STANDARD	1	Installed default policy domain

See “Field descriptions” on page 725 for field descriptions.

Example: Display the list of active-data pools

Display the active-data pool list. Issue the command:

```
query domain format=detailed
```

```

Policy Domain Name: STANDARD
Activated Policy Set: STANDARD
Activation Date/Time: 05/16/2006 16:18:05
Days Since Activation: 15
Activated Default Mgmt Class: STANDARD
Number of Registered Nodes: 1
Description: Installed default policy domain.
Backup Retention (Grace Period): 30
Archive Retention (Grace Period): 365
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 05/31/2006 15:17:48
Managing profile:
Changes Pending: Yes
Active Data Pool List: ADPPPOOL

```

See “Field descriptions” for field descriptions.

Field descriptions

Policy Domain Name

The name of the policy domain.

Activated Policy Set

The name of the policy set that was last activated in the domain.

The definitions in the last activated policy set and the ACTIVE policy set are not necessarily identical. When you activate a policy set, the server copies the contents of the policy set to the policy set with the special name ACTIVE. The copied definitions in the ACTIVE policy set can be modified only by activating another policy set. You can modify the original policy set without affecting the ACTIVE policy set. Therefore, definitions in the policy set that was last activated might not be the same as those in the ACTIVE policy set.

Activation Date/Time

The date and time that the policy set was activated.

Days Since Activation

The number of days since the policy set was activated.

Activated Default Mgmt Class

The assigned default management class for the policy set.

Number of Registered Nodes

The number of client nodes registered to the policy domain.

Description

The description of the policy domain.

Backup Retention (Grace Period)

The number of days to retain inactive backup versions of files when any of the following conditions occur:

- A file is rebound to a new management class, but neither the new management class nor default management class contains a backup copy group.
- The management class to which a file is bound no longer exists, and the default management class does not contain a backup copy group.
- The backup copy group is deleted from the management class to which a file is bound and the default management class does not contain a backup copy group.

QUERY DOMAIN

Archive Retention (Grace Period)

The number of days to retain an archive file that meets either of the following conditions:

- The management class to which a file is bound no longer exists, and the default management class does not contain an archive copy group.
- The archive copy group is deleted from the management class to which a file is bound and the default management class does not contain an archive copy group.

Last Update by (administrator)

The administrator that defined or most recently updated the policy domain. If this field contains \$\$CONFIG_MANAGER\$\$, the policy domain is associated with a profile that is managed by the configuration manager.

Last Update Date/Time

When the administrator defined or most recently updated the policy domain.

Managing Profile

The profile or profiles to which the managed server subscribed to get the definition of this policy domain.

Changes Pending

Whether or not changes are being made but not activated. Once the changes are activated, the field resets to No.

Active Data Pool List

The list of active-data pools in the domain.

Related commands

Table 227. Commands related to QUERY DOMAIN

Command	Description
COPY DOMAIN	Creates a copy of a policy domain.
DEFINE DOMAIN	Defines a policy domain that clients can be assigned to.
DELETE DOMAIN	Deletes a policy domain along with any policy objects in the policy domain.
UPDATE DOMAIN	Changes the attributes of a policy domain.


QUERY DRIVE (Query information about a drive)

Use this command to display information about the drives associated with a library.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

library_name

Specifies the name of the library where the queried drive is located. This parameter is optional. You can use a wildcard character to specify this name.

You must specify a value for this parameter if you specify a drive name.

drive_name

Specifies the name assigned to the drive. This parameter is optional. You can use a wildcard character to specify this name. If you specify a drive name, you must also specify a *library_name*.

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the drive.

Detailed

Specifies that complete information is displayed for the drive.

Example: List drives associated with the server

Display information about all drives associated with the server. Issue the command:

```
query drive
```

Library Name	Drive Name	Device Type	Online
LIB1	DRIVE01	3590	Yes
LIB2	DRIVE02	3590	Yes

See “Field descriptions” on page 728 for field descriptions.

Example: Display detailed information on a specific drive and library

Display detailed information about the drive named DRIVE02 that is associated with the library LIB2. Issue the command:

QUERY DRIVE

```
query drive lib2 drive02 format=detailed
```

```
Library Name: LIB2
Drive Name: DRIVE02
Device Type: 3590
On Line: Yes
Drive State: Empty
ACS DriveID:
Allocated to:
Last Update by (administrator): ADMIN
Last Update Date/Time: 02/29/2002 09:26:23
Cleaning Frequency (Gigabytes/ASNEEDED/NONE): NONE
```

See “Field descriptions” for field descriptions.

Field descriptions

Library Name

The name of the library to which the drive is assigned.

Drive Name

The name assigned to the drive.

Device Type

The device type as specified in the associated device class. The server must have a path defined from the server to the drive in order to determine the true device type. As long as there is a path defined from the server to the drive, the server will display the true device type of the drive even if there are other paths defined to this drive. Exceptions to this occur if the device type is remote or unknown.

REMOTE

The server does not have a path to the device. The only defined paths to the device are from data movers.

UNKNOWN

No path exists.

Tip: Review the output of the **QUERY PATH** command to determine if the desired paths are defined. If they are not defined, define those desired paths using the **DEFINE PATH** command. Also, if using a data mover device, review the output of the **QUERY DATAMOVER** command to determine the type of the data mover device. If you are using a path from the server to a drive, the device type of the device class and the drive need to match. If you are using a path from a data mover device to a drive, review the documentation for your type of data mover to ensure the device type of the device class is compatible with the type of data mover device.

On Line

Specifies the status of the drive:

Yes The drive is online and available for server operations.

No The drive is offline and was put in this state by an administrator updating the status.

Unavailable Since

Specifies that the drive has been unavailable since *mm/dd/yy hh:mm:ss*. Output shows the time the server marked the drive as unavailable.

Polling Since

Specifies that the server is polling the drive because the drive

stopped responding. Output shows the time the server detected a problem and began polling. The server polls a drive before stating it is unavailable. The time output follows the format: mm/dd/yy hh:mm:ss.

For more information on conditions under which the server will poll, see the *Administrator's Guide*.

Drive State

This specifies the current state of this particular drive based on the result of the last SCSI command to the drive or library. The server tracks the state of the drive to improve its selection of a drive for an operation and its drive recovery operations. The values are:

Unavailable

The drive is not available to the library for operations.

Empty The drive is empty and ready for operations.

Loaded

The drive is currently loaded, and the server is performing operations to the drive.

Unloaded

The media has been ejected from the drive.

Reserved

The drive is reserved for a mount request.

Unknown

The drive begins in drive state unknown as a result of being defined, as a result of server initialization, or as a result of having its status updated to online.

ACS DriveId

The ID that specifies the physical location of the drive in an ACSLS library. This drive ID is specified as *a,l,p,d*, where *a* is the ACSID, *l* is the LSM (library storage module), *p* is the panel number, and *d* is the drive ID.

Allocated To

The name of the library client that is currently using the drive. This applies to shared SCSI libraries only; the field is left blank for all other libraries.

Last Update by (administrator)

Who performed the last update to the drive.

Last Update Date/Time

The date and time when the last update occurred.

Cleaning Frequency

How often the server activates drive cleaning. This value can be the number of gigabytes, NONE, or ASNEEDED.

Related commands

Table 228. Commands related to QUERY DRIVE

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.

QUERY DRIVE

Table 228. Commands related to **QUERY DRIVE** (continued)

Command	Description
DEFINE PATH	Defines a path from a source to a destination.
DELETE DRIVE	Deletes a drive from a library.
DELETE LIBRARY	Deletes a library.
QUERY LIBRARY	Displays information about one or more libraries.
UPDATE DRIVE	Changes the attributes of a drive.

QUERY DRMEDIA (Query disaster recovery media)

Use this command to display information about database backup and copy storage pool volumes. You can also use the command to create a file of executable commands to process the volumes.

Remember: The **QUERY DRMEDIA** command always processes eligible copy storage-pool volumes. (For details about eligible copy storage pool volumes, see the description of the **COPYSTGPOOL** parameter in this command). By default, volumes in active-data pools are not eligible for processing by the disaster recovery manager. To process active-data pool volumes, you must issue the **SET DRMACTIVEDATASTGPOOL** command, or you must use the **ACTIVEDATASTGPOOL** parameter on the **QUERY DRMEDIA** command. Depending on the setting for the **SOURCE** parameter in this command, database backups (full and incremental types, or snapshot types) can also be processed. The other parameters of the command (such as **WHERESTATE**) can further limit the results of the query.


If you are using an external library and have moved a volume to the **NOTMOUNTBLE** state using the **MOVE DRMEDIA** command, the **QUERY DRMEDIA** command might still report the volume state as **MOUNTABLE** if it detects that the volume is in the library. Refer to the external library documentation for information about the procedures that you should follow when using the **MOVE DRMEDIA** and the **QUERY DRMEDIA** commands.

Privilege class


To issue this command, you must have one of the following privilege classes:

- If the **CMD** parameter is **NOT** specified: operator or system privilege.
- If the **CMD** parameter is specified and the **REQSYSAUTHOUTFILE** server option is set to **NO**: operator, unrestricted storage, or system privilege.
- If the **CMD** parameter is specified and the **REQSYSAUTHOUTFILE** server option is set to **YES** (the default): system privilege.

Syntax



QUERY DRMEDIA



Parameters

volume_name

Specifies the names of the database backup and copy storage pool volumes to be queried. You can use wildcard characters to specify multiple names. This parameter is optional. The server looks for matching names among the following eligible volumes:

- Database backup volumes, as selected by the `SOURCE` parameter of this command.
- Copy storage pool volumes from copy storage pools named in the `COPYSTGPOOL` parameter. If you do not use the `COPYSTGPOOL` parameter, the server queries volumes from copy storage pools previously selected by the **SET DRMCOPYSTGPOOL** command.

Other parameters can also limit the results of the query.

WHEREState

Specifies the state of volumes to be processed. This parameter is optional. The default is `ALL`. Possible values are:

ALL

Specifies all volumes in all states.

MOnstable

Volumes in this state contain valid data and are accessible for onsite processing.

NOTMOnstable

Volumes in this state are onsite, contain valid data, and not accessible for onsite processing.

COUrier

Volumes in this state are being moved to an offsite location.

VAult

Volumes in this state are offsite, contain valid data, and are not accessible for onsite processing.

VAULTRetrieve

Volumes in this state are located at the offsite vault, do not contain valid data and can be moved back onsite for reuse or disposal:

- A copy storage pool volume is considered to be in the VAULTRETRIEVE state if it has been empty for at least the number of days specified with the REUSEDELAY parameter on the **DEFINE STGPOOL** command.
- A database backup volume is considered to be in the VAULTRETRIEVE state if it is associated with a database backup series that was expired based on the value specified using the **SET DRMDBBACKUPEXPIREDAYS** command.

Important: When you issue **QUERY DRMEDIA WHERESTATE=VAULTRETRIEVE**, the server dynamically determines which volumes can be moved back onsite for reuse or disposal. Therefore, to ensure that you identify all volumes that are in a VAULTRETRIEVE state, issue **QUERY DRMEDIA WHERESTATE=VAULTRETRIEVE** without the **BEGINDATE**, **ENDDATE**, **BEGINTIME** or **ENDTIME** parameters. The Last Update Date/Time field in the output for **QUERY DRMEDIA WHERESTATE=VAULTRETRIEVE** displays the date and time that a volume was moved to the VAULT state, not VAULTRETRIEVE.

COURIERRetrieve

Volumes in this state are being moved back to the onsite location.

REmote

Volumes in this state contain valid data and are located at the offsite remote server.

BEGINDate

Specifies the beginning date used to select volumes. This parameter is optional. Volumes are considered eligible if the **MOVE DRMEDIA** command has changed the volume to its current state on or after the specified date. The default is the earliest date for which volume information exists.

You can specify the date using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY-days or -days	The current date minus days specified. The maximum number of days is 9999.	TODAY"7 or "7. To query volumes beginning with records changed to their current state a week ago, you can specify BEGINDATE=TODAY-7 or BEGINDATE=-7 .
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM-days	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM

QUERY DRMEDIA

Value	Description	Example
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDDate

Specifies the ending date used to select volumes. This parameter is optional. Volumes are considered eligible if the MOVE DRMEDIA command has changed the volume to its current state on or before the specified date. The default is the current date.

You can specify the date using one of the following values:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY- <i>days</i> or - <i>days</i>	The current date minus days specified. The maximum number of days is 9999.	TODAY"7 or "7. To query volumes beginning with records changed to their current state a week ago, you can specify BEGINDATE=TODAY-7 or BEGINDATE=-7.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time used to select volumes. This parameter is optional. Volumes are considered eligible if the **MOVE DRMEDIA** command has changed the volume to its current state on or after the specified time and date. The default is midnight (00:00:00) on the date specified with the BEGINDATE parameter.

You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified begin date	12:33:28
NOW	The current time on the specified begin date	NOW

Value	Description	Example
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified begin date	NOW+03:00 or +03:00. If you issue QUERY DRMEDIA command at 9:00 with BEGINTIME=NOW+03:00 or BEGINTIME=+03:00, Tivoli Storage Manager displays volumes that were changed to their current state at 12:00 on the begin date you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified begin date	NOW-03:30 or -03:30. If you issue QUERY DRMEDIA command at 9:00 with BEGINTIME=NOW-03:30 or BEGINTIME=-03:30, Tivoli Storage Manager displays volumes that were changed to their current state at 5:30 on the begin date you specify.

ENDTime

Specifies the ending time used to select volumes. This parameter is optional. Volumes are considered eligible if the **MOVE DRMEDIA** command has changed the volume to its current state on or before the specified time and date. The default is 23:59:59.

You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified end date	10:30:08
NOW	The current time on the specified end date	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified end date	NOW+03:00 or +03:00. If you issue QUERY DRMEDIA command at 9:00 with ENDTIME=NOW+03:00 or ENDTIME=+03:00, Tivoli Storage Manager processes volumes that were changed to their current state at 12:00 on the end date you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified end date	NOW-03:30 or -03:30 If you issue QUERY DRMEDIA command at 9:00 with ENDTIME=NOW-03:00 or ENDTIME=-03:00, Tivoli Storage Manager processes volumes that were changed to their current state at 6:00 on the end date you specify.

COPYstgpool

Specifies the name of the copy storage pool whose volumes are to be processed. This parameter is optional. You can use wildcard characters to specify this name. The copy storage pools specified with this parameter override those specified with the **SET DRMCOPYSTGPOOL** command.

If this parameter is not specified, Tivoli Storage Manager selects the storage pools as follows:

- If the **SET DRMCOPYSTGPOOL** command was previously issued with valid copy storage pool names, Tivoli Storage Manager processes only those storage pools.
- If the **SET DRMCOPYSTGPOOL** command has not been issued, or if all of the copy storage pools have been removed using the **SET DRMCOPYSTGPOOL** command, Tivoli Storage Manager processes all copy storage pool volumes in the specified state (ALL, MOUNTABLE, NOTMOUNTABLE, COURIER, VAULT, VAULTRETRIEVE, COURIERRETRIEVE, or REMOTE).

Source

Specifies whether any database backup volumes are selected. This parameter is optional. The default is DBBACKUP. Possible values are:

DBBackup

Full and incremental database backup volumes are selected.

DBSnapshot

Snapshot database backup volumes are selected.

DBNone

No database backup volumes are selected.

ACTIVEDatastgpool

Specifies the name of the active-data storage pool whose volumes are to be processed. This parameter is optional. You can use wildcard characters to specify this name. The active-data pools that are specified with this parameter override those specified with the **SET DRMACTIVEDATASTGPOOL** command.

If this parameter is not specified, Tivoli Storage Manager selects the storage pools as follows:

- If the **SET ACTIVEDATASTGPOOL** command was previously issued with valid active-data pool names, Tivoli Storage Manager processes only those storage pools.
- If the **SET ACTIVEDATASTGPOOL** command has not been issued, or all of the active-data pools have been removed using the **SET ACTIVEDATASTGPOOL** command, Tivoli Storage Manager processes all active-data pool volumes in the specified state (ALL, NOTMOUNTABLE, COURIER, VAULT, VAULTRETRIEVE, COURIERRETRIEVE, or REMOTE). Volumes in the MOUNTABLE state are not processed.

Format

Specifies the information to be displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that detailed information is displayed.

Cmd

Specifies that executable commands are built for the selected volumes. If you specify **FORMAT=CMD**, you must also specify the **CMD** parameter.

WHERELocation

Specifies the location of the volumes to be queried. This parameter is optional. The maximum length of the location is 255 characters. Enclose the text in quotation marks if it contains any blank characters. If you specify a target

server name, the disaster recovery manager displays all database backup volumes and copy storage pool volumes located on the target server.

CMD

Specifies the creation of executable commands to process the volume name and location obtained by this command. This parameter is optional. You must enclose the command specification in quotation marks. The maximum length of this parameter is 255 characters. The disaster recovery manager writes the commands to a file specified by the **CMDFILENAME** parameter or the **SET DRMCMDFILENAME** command, or generated by the **QUERY DRMEDIA** command. If the command length is greater than 240 characters, it is split into multiple lines and continuation characters (+) are added. You may need to alter the continuation character according to the product that runs the commands.

If you do not specify the **FORMAT=CMD** parameter, this command will not create any command lines.

string

The command string. The string must not include embedded quotation marks. For example, this is a valid **CMD** parameter:

```
cmd="checkin libvol lib8mm &vol status=scratch"
```

This is an example of a **CMD** parameter that is *not* valid:

```
cmd="checkin libvolume lib8mm" &vol status=scratch"
```

substitution

Specifies a substitution variable to tell **QUERY DRMEDIA** to substitute a value for the variable. The variables are not case-sensitive, and must not contain blank spaces after the ampersand (&). The possible variables are:

&VOL

A volume name variable.

&LOC

A volume location.

&VOLDSN

The name of the file the server writes into the sequential access media labels. An example of a copy storage pool tape volume file name using the default prefix **TSM** is **TSM.BFS**. An example of a database backup tape volume file name using a prefix **TSM310** defined with the device class is **TSM310.DBB**.

&NL

The new line character. When **&NL** is specified, **QUERY DRMEDIA** command splits the command at the **&NL** variable and does not append a continuation character. You must specify the proper continuation character before the **&NL** if one is required. If the **&NL** is not specified and the command line is greater than 240 characters, the line is split into multiple lines and continuation characters (+) are added.

CMDFilename

Specifies the fully qualified name of the file to contain the commands specified with **CMD** parameter. This parameter is optional.

If you do not specify a name with the **SET DRMCMDFILENAME** command, Tivoli Storage Manager creates a file name by appending **exec.cmds** to the absolute directory path name of the Tivoli Storage Manager instance directory.

QUERY DRMEDIA

If you specify a null string (" ") the commands are displayed on the console only. You can redirect the commands to a file using the redirection character for the operating system.

If the operation fails after the command file has been created, the file is not deleted.

APPend

Specifies whether to overwrite any existing contents of the command file or append the commands to the file. This parameter is optional. The default is NO. Possible values are:

No The disaster recovery manager overwrites the contents of the file.

Yes

The disaster recovery manager appends the commands to the file.

Example: List volumes to be sent to offsite storage

Display all volumes to be given to a courier for offsite storage.

```
query drmedia wherestate=notmountable
format=standard
```

Volume Name	State	Last Update Date/Time	Automated LibName
TAPE01	Not mountable	01/20/1998 14:25:22	
DBTP01	Not mountable	01/20/1998 14:25:22	
DBTP03	Not mountable	01/20/1998 14:31:53	

See "Field descriptions" for field descriptions.

Example: Display information on volumes at the vault

Display detailed information about all volumes at the vault.

```
query drmedia wherestate=vault format=detailed
```

Volume Name: DBTP02
State: Vault
Last Update Date/Time: 01/20/1998 13:29:02
Location: Ironmnt
Volume Type: DBBackup
Copy Storage Pool Name:
Active-Data Storage Pool Name: TSMACTIVEPOOL
Automated LibName:

See "Field descriptions" for field descriptions.

Field descriptions

Volume Name

The name of the database backup or copy storage pool volume.

State The state of the volume.

Last Update Date/Time

The date and time that the volume state was last updated. For volumes in the VAULTRETRIEVE state, this field displays the date and time that a volume was moved to the VAULT state, not VAULTRETRIEVE. The server does not "update" volumes to VAULTRETRIEVE. At the time the **QUERY DRMEDIA** command is issued, the server dynamically determines whether

the data in copy storage pool volumes and database backup volumes is no longer valid and whether the volume can be brought back onsite for reuse or disposal.

Location

The volume location.

Volume Type

The type of volume. Possible values are:

DBBackup

A full or incremental database backup volume.

DBSnapshot

A database snapshot backup volume.

CopyStgPool

A copy storage pool volume.

Copy Storage Pool Name

For a copy storage pool volume, the name of the copy storage pool.

Active-Data Storage Pool Name

For an active-data pool volume, the name of the active-data pool.

Automated LibName

The name of the automated library if the volume is in a library.

Related commands

Table 229. Commands related to QUERY DRMEDIA

Command	Description
BACKUP DB	Backs up the Tivoli Storage Manager database to sequential access volumes.
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
MOVE DRMEDIA	Moves DRM media on-site and off-site.
SET DRMACTIVEDATASTGPOOL	Specifies that active-data storage pools are managed by DRM.
SET DRMCOPYSTGPOOL	Specifies that copy storage pools are managed by DRM.
SET DRMDBBACKUPEXPIREDAYS	Specifies criteria for database backup series expiration.
SET DRMCMDFILENAME	Specifies a file name for containing DRM executable commands.
SET DRMFILEPROCESS	Specifies whether the MOVE DRMEDIA or QUERY DRMEDIA command processes files associated with a device type of file.

QUERY DRMSTATUS (Query disaster recovery manager system parameters)

Use this command to display information about the system parameters defined for disaster recovery manager (DRM).

Privilege class

Any administrator can issue this command.

Syntax

►►—Query DRMStatus—◄◄

Parameters

None.

Example: Display DRM system parameter information

Display information about the DRM system parameters:

```
query drmstatus
```

```
Recovery Plan Prefix:
Plan Instructions Prefix:
Replacement Volume Postfix: @
Primary Storage Pools: PRIM1 PRIM2
Copy Storage Pools: COPY*
Active-Data Storage Pools: TSMACTIVEPOOL
Not Mountable Location name: Local
Courier Name: Fedex
Vault Site Name: Ironmnt
DB Backup Series expiration days: 30 Day(s)
Recovery Plan File Expiration Days: 30 Days(s)
Check Label?: No
Process FILE Device Type?: No
Command file name:
```

Field descriptions

Recovery Plan Prefix

User-specified prefix portion of the file name for the recovery plan file.

Plan Instructions Prefix

User-specified prefix portion of the file names for the server recovery instructions files.

Replacement Volume Postfix

The character added to the end of the replacement volume names in the recovery plan file.

Primary Storage Pools

The primary storage pools that are eligible for processing by the **PREPARE** command. If this field is blank, all primary storage pools are eligible.

Copy Storage Pools

The copy storage pools that are eligible for processing by the **MOVE DRMEDIA**, **PREPARE**, and **QUERY DRMEDIA** commands. If this field is blank, all copy storage pools are eligible.

Active-data Storage Pools

The active-data pools that are eligible for processing by the **MOVE DRMEDIA**, **PREPARE**, and **QUERY DRMEDIA** commands. If this field is blank, active-data pools are not eligible.

Not Mountable Location Name

The name of the offsite location where the media to be shipped are stored.

Courier Name

The name of the courier used to carry the media to the vault.

Vault Site Name

The name of the vault where the media is stored.

DB Backup Series Expiration Days

The minimum number of days that must elapse since a database series has been created before it is eligible to be expired. See the SET DRMDBBACKUPEXPIREDDAYS command for information about the criteria that must be met for database backup series expiration.

Recovery Plan File Expiration Days

The minimum number of days that must elapse since a recovery plan file, stored on a target server, has been created before it is eligible to be expired. See the SET DRMRPFEXPIREDDAYS command for information about the criteria that must be met for recovery plan file expiration.

Check Label?

Whether media labels are read for sequential media volumes checked out by the **MOVE DRMEDIA** command. Possible values are Yes or No.

Process FILE Device Type?

Whether **MOVE DRMEDIA** or **QUERY DRMEDIA** commands process database backup and copy storage pool volumes associated with a device class with a FILE device type. Possible values are Yes or No.

Command File Name

The full path file name that contains the executable commands generated by the **MOVE DRMEDIA** or **QUERY DRMEDIA** command.

Related commands

Table 230. Commands related to QUERY DRMSTATUS

Command	Description
SET DRMCHECKLABEL	Specifies whether Tivoli Storage Manager should read volume labels during MOVE DRMEDIA command processing.
SET DRMACTIVEDATASTGPOOL	Specifies that active-data storage pools are managed by DRM.
SET DRMCOPYSTGPOOL	Specifies that copy storage pools are managed by DRM.
SET DRMCMDFILENAME	Specifies a file name for containing DRM executable commands.
SET DRMCOURIERNAME	Specifies the name of the courier for the disaster recovery media.
SET DRMDBBACKUPEXPIREDDAYS	Specifies criteria for database backup series expiration.

QUERY DRMSTATUS

Table 230. Commands related to **QUERY DRMSTATUS** (continued)

Command	Description
SET DRMFILEPROCESS	Specifies whether the MOVE DRMEDIA or QUERY DRMEDIA command processes files associated with a device type of file.
SET DRMINSTRPREFIX	Specifies the prefix portion of the path name for the recovery plan instructions.
SET DRMPLANVPOSTFIX	Specifies the replacement volume names in the recovery plan file.
SET DRMPLANPREFIX	Specifies the prefix portion of the path name for the recovery plan.
SET DRMPRIMSTGPOOL	Specifies that primary storage pools are managed by DRM.
SET DRMRPFEXPIREDAYS	Set criteria for recovery plan file expiration.
SET DRMVAULTNAME	Specifies the name of the vault where DRM media is stored.
SET DRMNOTMOUNTABLENAME	Specifies the location name of the DRM media to be sent offsite.


QUERY ENABLED (Query enabled events)

Use this command to display either a list of enabled events or a list of disabled events, whichever is shorter.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

receiver

Specifies a type of receiver for enabled events. This is a required parameter. Valid values are:

ACTLOG

Specifies the Tivoli Storage Manager activity log as a receiver.

CONSOLE

Specifies the standard server console as a receiver.

EVENTSERVER

Specifies the event server as a receiver.

FILE

Specifies a user file as a receiver. Each logged event is a record in the file and a person cannot read each logged event easily.

FILETEXT

Specifies a user file as a receiver. Each logged event is a fixed-size, readable line.

SNMP

Specifies the simple network management protocol (SNMP) as a receiver.

TIVOLI

Specifies the Tivoli Management Environment (TME) as a receiver.

USEREXIT

Specifies a user-written routine to which Tivoli Storage Manager writes information as a receiver.

NODename

Specifies a node name to be queried. You can specify NODENAME or SERVERNAME. If neither parameter is specified, the query is for events enabled for the server running this command.

SERVername

Specifies a server name to be queried. You can specify NODENAME or

QUERY ENABLED

SERVERNAME. If neither parameter is specified, the query is for events enabled for the server running this command.

Example: Query the server for console events

Query the server for server events that are enabled for the console. There are 10000 possible server events. Either a list of enabled events or disabled events is displayed (whichever list is shorter).

```
query enabled console
```

9998 events are enabled for the CONSOLE receiver. The following events are DISABLED for the CONSOLE receiver:

```
ANR8409, ANR8410
```

Related commands

Table 231. Commands related to QUERY ENABLED

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
DISABLE EVENTS	Disables specific events for receivers.
ENABLE EVENTS	Enables specific events for receivers.
END EVENTLOGGING	Ends event logging to a specified receiver.
QUERY EVENTRULES	Displays information about rules for server and client events.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

QUERY EVENT (Query scheduled and completed events)

Use this command to display the status of scheduled events. The time and date parameters allow you to limit the query to events that were scheduled to occur within the specified times and dates. Limiting the output to events whose scheduled start times fall within a date and time range also minimizes the time it takes to process this query.

The command syntax differs for queries that apply to scheduled client operations and to scheduled administrative commands.

- “QUERY EVENT (Display administrative event schedules)” on page 753
- “QUERY EVENT (Display client schedules)” on page 746

Table 232. Commands related to QUERY EVENT

Command	Description
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.
DELETE EVENT	Deletes event records before a specified date and time.
QUERY ACTLOG	Displays messages from the server activity log.
SET EVENTRETENTION	Specifies the number of days to retain records for scheduled operations.
SET RANDOMIZE	Specifies the randomization of start times within a window for schedules in client-polling mode.

QUERY EVENT – client schedules

QUERY EVENT (Display client schedules)


Use the **QUERY EVENT** command to display scheduled and completed events for selected clients.

Tivoli Storage Manager keeps only one version of an event record in the database. If a client schedule is changed, all previous event records for that schedule are removed from the database.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

domain_name (Required)

Specifies the name of the policy domain to which the schedules belong. You can use a wildcard character to specify this name.

schedule_name (Required)

Specifies the name of the schedule for which events are displayed. You can use a wildcard character to specify this name.

Type=Client

Specifies that the query displays events for client schedules. This parameter is optional. The default is CLIENT.

Nodes

Specifies the name of the client node that belongs to the specified policy domain for which events are displayed. You can specify multiple client nodes by separating the names with commas and no intervening spaces. You can use wildcard characters to specify nodes. If you do not specify a client name, events display for all clients that match the domain name and the schedule name.

BEGINDate

Specifies the beginning date of the time range for events to be displayed. All events scheduled to start during this time are displayed. This parameter is optional. The default is the current date.

You can specify the date using one of the values below:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+days or <i>+days</i>	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
TODAY-days or <i>-days</i>	The current date minus days specified	TODAY-7 or -7. To query events scheduled to start during the past seven days, specify BEGINDATE=TODAY-7 ENDDATE=TODAY or BEGINDATE=-7 ENDDATE=TODAY.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
<i>EOLM-days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
<i>BOTM+days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time of the range for events to be displayed. All events scheduled to start during this time are displayed. This parameter is optional. The default value is 00:00.

You can specify the time using one of the values below:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time on the specified begin date	10:30:08
NOW	The current time on the specified begin date	NOW
NOW+HH:MM or <i>+HH:MM</i>	The current time plus hours and minutes on the specified begin date	NOW+03:00 or +03:00. If you issue this command at 9:00 to query events scheduled to start 3 hours from now, you can specify either BEGINTIME=NOW+03:00 or BEGINTIME=+03:00. Tivoli Storage Manager displays events at 12:00 on the specified begin date.

QUERY EVENT – client schedules

Value	Description	Example
NOW- <i>HH:MM</i> or - <i>HH:MM</i>	The current time minus hours and minutes on the specified begin date	NOW-04:00 or -04:00. If you issue this command at 9:00 to query events scheduled to start during the last 4 hours, you can specify either BEGINTIME=NOW-04:00 ENDTIME=NOW or BEGINTIME=-04:00 ENDTIME=NOW. Tivoli Storage Manager displays events at 5:00 on the specified begin date.

ENDDate

Specifies the ending date of the time range for events to be displayed. All events that were schedule to start during this time are displayed. This parameter is optional. The default is the value used for the BEGINDATE.

You can specify the date using one of the values below:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+ <i>days</i> or + <i>days</i>	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
TODAY- <i>days</i> or - <i>days</i>	The current date minus days specified	TODAY-8 or -8. To query events scheduled to start during a one-week period that ended yesterday, you can specify either BEGINDATE=TODAY-8 ENDDATE=TODAY-1 or BEGINDATE=-8 ENDDATE=-1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time of the range for events to be displayed. All events that were scheduled to start during this time are displayed. This parameter is optional. The default value is 23:59.

You can specify the time using one of the values below:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time on the specified end date	10:30:08
NOW	The current time on the specified end date	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified end date	NOW+03:00 or +03:00. If you issue this command at 9:00 to query events scheduled to start 3 hours from now, you can specify either BEGINTIME=NOW ENDTIME=NOW+03:00 or BEGINTIME=NOW ENDTIME=+03:00.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified end date	NOW-04:00 or -04:00

EXceptiononly

Specifies the type of information you want on scheduled or completed events. This parameter is optional. The default is NO. Possible values are:

No Specifies that the information on past and projected events is displayed.

Yes

Specifies that the events that failed or did not process as scheduled are displayed.

Format

Specifies how information displays. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information for events displays.

Detailed

Specifies that complete information for events displays.

Display partial information for unsuccessful events

Display partial information for all events scheduled for DOMAIN1 that did not run successfully. Limit the search to the client named JOE. Limit the events displayed to those that were scheduled to occur from February 11, 2001 (02/11/2001) to February 12, 2001 (02/12/2001).

```
query event domain1 * nodes=joe begindate=02/11/2001
enddate=02/12/2001 exceptiononly=yes
```

Scheduled Start	Actual Start	Schedule Name	Node Name	Status
02/11/1999 01:00:00	02/11/1999 01:13:55	BACK1	JOE	Failed
02/12/1999 01:00:00		DAILYBKP	JOE	Missed

See “Field descriptions” on page 751 for field descriptions.

Display partial information for scheduled events for a client

Display complete information for all events that are scheduled for processing. Use the start time as 10 days previous to today, and the finish time includes today.

QUERY EVENT – client schedules

```
query event * * begindate=today-10 enddate=today
```

Scheduled Start	Actual Start	Schedule Name	Node Name	Status
02/04/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/04/2013 14:00:00	02/04/2013 14:12:49	VDATAMVR1-IN1	VDATAMVR1-T1	Completed
02/04/2013 14:30:00	02/04/2013 14:33:10	VDATAMVR1-IN2	VDATAMVR1-T2	Completed
02/04/2013 15:00:00	02/04/2013 15:01:49	VDATAMVR1-IN3	VDATAMVR1-T3	Completed
02/04/2013 15:30:00	02/04/2013 15:42:00	VDATAMVR1-IN4	VDATAMVR1-T4	Completed
02/05/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/05/2013 14:00:00	02/05/2013 14:05:22	VDATAMVR1-F1	VDATAMVR1-F1	Completed
02/05/2013 14:30:00	02/05/2013 14:32:53	VDATAMVR1-F2	VDATAMVR1-F2	Failed 12
02/05/2013 15:00:00	02/05/2013 15:00:38	VDATAMVR1-F3	VDATAMVR1-F3	Completed
02/05/2013 15:30:00	02/05/2013 15:36:41	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/06/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/06/2013 14:00:00	02/06/2013 14:06:42	VDATAMVR1-F1	VDATAMVR1-F1	Completed
02/06/2013 14:30:00	02/06/2013 14:35:41	VDATAMVR1-F2	VDATAMVR1-F2	Completed
02/06/2013 15:00:00	02/06/2013 15:08:56	VDATAMVR1-F3	VDATAMVR1-F3	Completed
02/06/2013 15:30:00	02/06/2013 15:40:49	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/07/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/07/2013 14:00:00	02/07/2013 14:03:43	VDATAMVR1-F1	VDATAMVR1-F1	Completed
02/07/2013 14:30:00	02/07/2013 14:35:10	VDATAMVR1-F2	VDATAMVR1-F2	Completed
02/07/2013 15:00:00	02/07/2013 15:09:12	VDATAMVR1-F3	VDATAMVR1-F3	Completed
02/07/2013 15:30:00	02/07/2013 15:40:21	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/08/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/08/2013 14:00:00	02/08/2013 14:10:17	VDATAMVR1-F1	VDATAMVR1-F1	Completed
02/08/2013 14:30:00	02/08/2013 14:39:16	VDATAMVR1-F2	VDATAMVR1-F2	Completed
02/08/2013 15:00:00	02/08/2013 15:08:17	VDATAMVR1-F3	VDATAMVR1-F3	Completed
02/08/2013 15:30:00	02/08/2013 15:41:16	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/09/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/09/2013 14:02:16		VDATAMVR1-F1	VDATAMVR1-F1	Failed 12
02/09/2013 14:30:00	02/09/2013 14:44:26	VDATAMVR1-F2	VDATAMVR1-F2	Failed 12
02/09/2013 15:00:00	02/09/2013 15:06:24	VDATAMVR1-F3	VDATAMVR1-F3	Failed 12
02/09/2013 15:30:00	02/09/2013 15:32:18	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/11/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/11/2013 14:00:00	02/11/2013 14:01:05	VDATAMVR1-F1	VDATAMVR1-F1	Failed 12
02/11/2013 14:30:00	02/11/2013 14:31:42	VDATAMVR1-F2	VDATAMVR1-F2	Failed 12
02/11/2013 15:00:00	02/11/2013 15:06:17	VDATAMVR1-F3	VDATAMVR1-F3	Failed 12
02/11/2013 15:30:00	02/11/2013 15:30:19	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/12/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/12/2013 14:00:00	02/12/2013 14:03:37	VDATAMVR1-F1	VDATAMVR1-F1	Completed
02/12/2013 14:30:00	02/12/2013 14:33:07	VDATAMVR1-F2	VDATAMVR1-F2	Completed
02/12/2013 15:00:00	02/12/2013 15:03:56	VDATAMVR1-F3	VDATAMVR1-F3	Completed
02/12/2013 15:30:00	02/12/2013 15:36:44	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/13/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Missed
02/13/2013 14:00:00	02/13/2013 14:06:24	VDATAMVR1-F1	VDATAMVR1-F1	Completed
02/13/2013 14:30:00	02/13/2013 14:34:50	VDATAMVR1-F2	VDATAMVR1-F2	Completed
02/13/2013 15:00:00	02/13/2013 15:15:01	VDATAMVR1-F3	VDATAMVR1-F3	Completed
02/13/2013 15:30:00	02/13/2013 15:30:18	VDATAMVR1-F4	VDATAMVR1-F4	Completed
02/14/2013 14:00:00		SCHD_INCR-DM1	TSM_CET_DM1	Future
02/14/2013 14:00:00		VDATAMVR1-F1	VDATAMVR1-F1	Future
02/14/2013 14:30:00		VDATAMVR1-F2	VDATAMVR1-F2	Future
02/14/2013 15:00:00		VDATAMVR1-F3	VDATAMVR1-F3	Future

See “Field descriptions” on page 751 for field descriptions.

Display detailed information for scheduled events for a client

Display the detailed information for events scheduled for processing by client DOC between the hours of 10:00 a.m. and 11:00 a.m. on November 1, 2005 (11/01/2005). Notice that when the status is FAILED, the result code is displayed.

```
query event domain1 * nodes=doc begindate=11/01/2005
begintime=10:00 endtime=11:00 enddate=11/01/2005
exceptionsonly=yes format=detailed
```

Scheduled Start	Actual Start	Schedule Name	Node Name	Status
11/01/2005 10:01:01	11/01/2005 10:03:46	T1	DOC	Failed 8
11/01/2005 10:16:01	11/01/2005 10:16:10	T1	DOC	Failed 4
11/01/2005 10:31:01	11/01/2005 10:33:08	T1	DOC	Completed
11/01/2005 10:46:01		T1	DOC	Missed
11/01/2005 10:57:49	11/01/2005 10:58:07	T0	DOC	Failed 12

Field descriptions

Policy Domain Name

Specifies the name of the policy domain to which the schedule is assigned.

Schedule Name

Specifies the name of the schedule that initiated this event.

Node Name

Specifies the client scheduled to perform the operation.

Scheduled Start

Specifies the scheduled starting date and time for the event.

Actual Start

Specifies the date and time at which the client began processing the scheduled operation. No information is displayed if the scheduled operation has not started.

Completed

Specifies the date and time the scheduled event is completed.

Status Specifies the status of the event at the time the **QUERY EVENT** command is issued. Possible values are:

Completed

Specifies that the scheduled event has been completed.

Failed Specifies that the client reports a failure running the scheduled operation and successive retries have not succeeded.

Failed - no restart

Specifies an intermediate status, when a client session is interrupted by a communications error or timeout on the server. This status can be changed to a final status of "Completed" or "Failed" when the event completes.

Future Specifies that the beginning of the startup window for the event is in the future. This status also indicates that an event record has not been created for this event.

In Progress

Specifies that the scheduled event is running and has not yet reported the completion state to the server.

Periodically check the status for completion of the scheduled event. If this status is not updated in a reasonable amount of time, review your client `dsmsched.log` and `dsmerror.log` to determine why the client did not report the outcome of this event to the server. If the scheduled backup failed, rerun the scheduled event or perform a manual incremental backup to ensure the data backup.

QUERY EVENT – client schedules

Missed

Specifies that the scheduled startup window for this event has passed and the schedule has not begun.

Pending

Specifies that the **QUERY EVENT** command was issued during the startup window for the event, but processing the scheduled operation has not yet begun.

Restarted

Specifies that the client has retried processing the scheduled operation.

Severed

Specifies that the communications with the client is severed before the event can complete.

Started

Specifies that the event has begun processing.

Uncertain

Specifies that the state of the event cannot be determined. The server specifies Uncertain if the **QUERY EVENT** command does not find an event record. An event record will not be found if the record was deleted or if the server was unavailable during the scheduled startup window (the schedule was never started). Records with Uncertain status do not reside in the database. If you do not want these records to display, either specify **EXCEPTIONSONLY=YES** or delete the schedule if it is no longer needed.

Attention: When a scheduled operation is processing, without having been restarted within its specified duration, the Status field shows Started. If the operation continues beyond the specified duration, no event record is created. If a query is issued after the specified duration has passed, the Status shows as Failed even if the operation is still running. After the operation completes, an event record is created, and a subsequent query shows the result in the Status field.

Result Specifies the return code from the client that identifies whether the schedule has processed successfully. If the return code is other than a zero, examine the client's error log and schedule log.

Return code	Explanation
0	All operations completed successfully.
4	The operation completed successfully, but some files were not processed.
8	The operation completed with at least one warning message.
12	The operation completed with at least one error message (except for error messages for skipped files).

If a schedule has **ACTION=COMMAND** as a parameter, and the command is not a Tivoli Storage Manager command, the command can produce other values in the Result field. See the *Backup-Archive Clients Installation and User's Guide* for details.

Reason

Specifies the reason for the return code.


QUERY EVENT (Display administrative event schedules)

Use the **QUERY EVENT** command to display scheduled and completed events for selected administrative command schedules.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

schedule_name (Required)

Specifies the name of the schedule for which events display. You can use wildcard characters to specify names.

Type=Administrative (Required)

Specifies that the query displays events for administrative command schedules.

BEGINDate

Specifies the beginning date of the time range for events to be displayed. All events scheduled to start during this time are displayed. This parameter is optional. The default is the current date.

You can specify the date using one of the values below:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+days or <i>+days</i>	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
TODAY-days or <i>-days</i>	The current date minus days specified	TODAY-7 or -7. To query events scheduled to start during the past seven days, specify BEGINDATE=TODAY-7 ENDDATE=TODAY or BEGINDATE=-7 ENDDATE=TODAY.

QUERY EVENT — administrative schedules

Value	Description	Example
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time of the range for events to be displayed. All events scheduled to start during this time are displayed. This parameter is optional. The default value is 00:00.

You can specify the time using one of the values below:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time on the specified begin date	10:30:08
NOW	The current time on the specified begin date	NOW
NOW+ <i>HH:MM</i> or + <i>HH:MM</i>	The current time plus hours and minutes on the specified begin date	NOW+03:00 or +03:00. If you issue this command at 9:00 to query events scheduled to start 3 hours from now, you can specify either BEGINTIME=NOW+03:00 or BEGINTIME=+03:00. Tivoli Storage Manager displays events at 12:00 on the specified begin date.
NOW- <i>HH:MM</i> or - <i>HH:MM</i>	The current time minus hours and minutes on the specified begin date	NOW-04:00 or -04:00. If you issue this command at 9:00 to query events scheduled to start during the last 4 hours, you can specify either BEGINTIME=NOW-04:00 ENDTIME=NOW or BEGINTIME=-04:00 ENDTIME=NOW. Tivoli Storage Manager displays events at 5:00 on the specified begin date.

ENDDate

Specifies the ending date of the time range for events to be displayed. All events that were schedule to start during this time are displayed. This parameter is optional. The default is the value used for the BEGINDATE.

You can specify the date using one of the values below:

QUERY EVENT — administrative schedules

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY+days or <i>+days</i>	The current date plus days specified. The maximum number of days you can specify is 9999.	TODAY +3 or +3.
TODAY-days or <i>-days</i>	The current date minus days specified	TODAY-8 or -8. To query events scheduled to start during a one-week period that ended yesterday, you can specify either BEGINDATE=TODAY-8 ENDDATE=TODAY-1 or BEGINDATE=-8 ENDDATE=-1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
<i>EOLM-days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
<i>BOTM+days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time of the range for events to be displayed. All events that were scheduled to start during this time are displayed. This parameter is optional. The default value is 23:59.

You can specify the time using one of the values below:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time on the specified end date	10:30:08
NOW	The current time on the specified end date	NOW
NOW+HH:MM or <i>+HH:MM</i>	The current time plus hours and minutes on the specified end date	NOW+03:00 or +03:00. If you issue this command at 9:00 to query events scheduled to start 3 hours from now, you can specify either BEGINTIME=NOW ENDTIME=NOW+03:00 or BEGINTIME=NOW ENDTIME=+03:00.
NOW-HH:MM or <i>-HH:MM</i>	The current time minus hours and minutes on the specified end date	NOW-04:00 or -04:00

QUERY EVENT — administrative schedules

EXceptiononly

Specifies the type of information you want on scheduled or completed events. This parameter is optional. The default is NO. Possible values are:

No Specifies that the information on past and projected events is displayed.

Yes

Specifies that the events that failed or did not process as scheduled are displayed.

Format

Specifies how the information displays. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information for events displays.

Detailed

Specifies that complete information for events displays.

Example: List events for a specific administrative schedule

Display partial information for all events scheduled for an administrative schedule named DOSADMIN. Limit the query to events that are scheduled for March 30, 1999 (03/30/1999). Issue the command:

```
query event dosadmin type=administrative  
begindate=03/30/1999  
enddate=03/30/1999
```

Scheduled Start	Actual Start	Schedule Name	Status
03/30/1999 00:00:00	03/30/1999 00:00:01	DOSADMIN	Completed
03/30/1999 04:00:00	03/30/1999 04:00:01	DOSADMIN	Completed
03/30/1999 12:00:00		DOSADMIN	Future
03/30/1999 16:00:00		DOSADMIN	Future

Field descriptions

Scheduled Start

Specifies the scheduled starting date and time for the event.

Actual Start

Specifies the date and time at which the client began processing the scheduled operation. No information displays if the schedule has not started executing.

Schedule Name

Specifies the name of the schedule that initiated this event.

Status For administrative commands or scripts that specify WAIT=YES, the status of a scheduled event is STARTED until the operation specified by the command or script is completed. The final status of the scheduled event depends on the return code of the operation. However, if WAIT=YES and if the schedule is running a script that specifies PREVIEW=YES, the final status is COMPLETED, unless the script contained a syntax error.

For administrative commands or scripts that specify WAIT=NO, the status of a scheduled event is COMPLETED if the scheduled command or script started. The success of the schedule is independent of the success of the operation performed by the command or script.


QUERY EVENTRULES (Query rules for server or client events)

Use this command to display the history of events that are enabled or disabled by a specified receiver for the server or for a client node.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

receivers

Specifies the name of one or more receivers for enabled events. This parameter is optional.

You can use a wildcard character to specify all receivers.

Valid values are:

CONSOLE

Specifies the standard console as a receiver.

ACTLOG

Specifies the Tivoli Storage Manager activity log as a receiver.

EVENTSERVER

Specifies the event server as a receiver.

FILE

Specifies a user file as a receiver. Each logged event is a record in the file and a person cannot read each logged event easily.

FILETEXT

Specifies a user file as a receiver. Each logged event is a fixed-size, readable line.

SNMP

Specifies the simple network management protocol (SNMP) as a receiver.

TIVOLI

Specifies the Tivoli Management Environment (TME) as a receiver.

USEREXIT

Specifies a user-written routine to which Tivoli Storage Manager writes information as a receiver.

QUERY EVENTRULES

NODENAME

Specifies a node name to be queried. You can use a wildcard character to specify a name. You can specify NODENAME or SERVERNAME. If neither parameter is specified, the query is for event rules for the server running this command.

SERVER

Specifies a server name to be queried. You can use a wildcard character to specify a name. You can specify NODENAME or SERVERNAME. If neither parameter is specified, the query is for event rules for the server running this command.

Example: Display the history of client events for the server console

Display the history of client events enabled or disabled for the server console and activity log receivers.

```
query eventrules console,actlog nodename=*
```

Date/Time	Client Event Rules
05/29/97 13:39:58	ENABLE EVENTS CONSOLE ANE4001 NODENAMES=JEE
05/30/97 13:46:25	DISABLE EVENTS ACTLOG ANE4962 NODENAMES=JEE
05/30/97 13:46:25	DISABLE EVENTS ACTLOG ANE4963 NODENAMES=JEE
05/30/97 13:46:25	DISABLE EVENTS ACTLOG ANE4965 NODENAMES=JEE
05/30/97 13:46:25	DISABLE EVENTS ACTLOG ANE4966 NODENAMES=JEE
05/30/97 13:46:25	DISABLE EVENTS ACTLOG ANE4967 NODENAMES=JEE
05/30/97 13:46:25	DISABLE EVENTS ACTLOG ANE4968 NODENAMES=JEE
05/30/97 14:24:20	ENABLE EVENTS CONSOLE ANE4015 NODENAMES=RON
05/30/97 14:24:50	ENABLE EVENTS CONSOLE ANE4026 NODENAMES=DONNA
05/30/97 14:25:59	ENABLE EVENTS CONSOLE ANE4015 NODENAMES=DONNA

Example: Display the history of client events for all receivers

Display the history of server events enabled or disabled for all receivers.

```
query eventrules
```

Date/Time	Server Event Rules
05/22/97 14:35:13	ENABLE EVENTS CONSOLE ANR2578
05/30/97 14:29:31	ENABLE EVENTS CONSOLE ANR0272
05/30/97 14:31:46	ENABLE EVENTS USEREXIT ANR0130
05/30/97 14:31:54	ENABLE EVENTS USEREXIT ANR0131
05/30/97 14:50:28	ENABLE EVENTS USEREXIT ANR0266

Related commands

Table 233. Commands related to QUERY ENABLED

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
DISABLE EVENTS	Disables specific events for receivers.
ENABLE EVENTS	Enables specific events for receivers.
END EVENTLOGGING	Ends event logging to a specified receiver.
QUERY ENABLED	Displays enabled or disabled events for a specific receiver.

QUERY EVENTSERVER (Query the event server)

Use this command to display the name of the event server.

Privilege class

Any administrator can issue this command.

Syntax

►►—Query EVENTSERVER—◄◄

Example: Display the event server name

Display the name of the event server.

```
query eventserver
```

ANR1669I Server EVENT is defined as the event server.

Related commands

Table 234. Commands related to QUERY EVENTSERVER

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
DEFINE EVENTSERVER	Defines a server as an event server.
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE EVENTSERVER	Deletes reference to the event server.
DELETE SERVER	Deletes the definition of a server.
END EVENTLOGGING	Ends event logging to a specified receiver.

QUERY EXPORT (Query for active or suspended export operations)


Use this command to list all restartable export operations. A restartable export is a server-to-server export operation whose FILEDATA value is not NONE. Only active server-to-server export operations that can be suspended are displayed.

Any **EXPORT NODE** or **EXPORT SERVER** operation with FILEDATA=NONE are not displayed. Additionally, the **QUERY EXPORT** command does not show export operations where the target device is either sequential media or virtual volumes.

Privilege class

An administrator can issue this command.

Syntax



Parameters

export_identifier

This optional parameter is the unique string identifier for the server-to-server export operation. Wildcard characters can be used to specify this name, and all matching export operations are queried. If you do not specify a value for this parameter and you also do not specify a PROCESS identifier, then all export operations are queried.

STate

This optional parameter queries the state of the valid server-to-server export operations. The default value is ALL. The possible values are:

ALL

Lists all running and suspended server-to-server export operations.

RUNning

Lists all active server-to-server export operations that are identifying eligible files or exporting files to the target server.

SUSPended

Lists all suspended server-to-server export operations. These suspended operations stopped running because of a failure or by the **SUSPEND EXPORT** command being issued.

PROcEss

This optional parameter specifies the number of a running server-to-server export operation that you want to query. If PROCESS is specified, Tivoli Storage Manager only displays the running server-to-server export operation associated with the process number. If PROCESS is not specified, Tivoli Storage Manager displays information on all server-to-server export operations. You

cannot specify this parameter if you specify an export identifier or if you specify the STATE parameter with a value of SUSPENDED.

Format

This optional parameter specifies how the information is displayed. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified export operations.

Detailed

When specified, displays all available information for the export operations.

Example: Display running and suspended export operations

List information for all currently running and suspended export operations. Issue the following command:

query export state=all

Export Identifier	Start Time	State	Process ID	Command
MYEXPORTNODE	01/24/2007 10:30:03	Suspended	--	Export NODE me,you,them filesystem=c\$ nametype=unicode filedata=all durunits= indefinite toserver=athens exportid=MYEXPORTNODE
EXPORT_HOME_DIRS	01/25/2007 09:30:03	Running	11	Export NODE n2,n3,n4 filesystem=/home nametype= server filedata=all durunits= indefinite toserver=athens exportid=EXPORT_HOME_DIRS
EXPORT_NODE_0001	01/25/2007 14:30:33	Running Not Suspendible	--	Export NODE n5,n6,n7 filesystem=d\$ nametype=unicode filedata=archive durunits= indefinite toserver=athens

See "Field descriptions" on page 763 for field descriptions.

Example: Display information about a running export operation

List information for the currently running export operation with process number "7." Issue the following command:

query export process=7

Export Identifier	Start Time	State	Process	Command
MYEXPORTNODE	01/24/2007 10:30:03	Running	7	Export NODE me,you,them filesystem=c\$ nametype=unicode filedata=all toserver=athens exportid=MYEXPORTNODE

See "Field descriptions" on page 763 for field descriptions.

Example: Display detailed information about all suspended export operations

List information for all currently suspended export operations. Issue the following command:

```
query export state=suspended format=detailed
```

```
Export Identifier : MyExportNode
Start Time : 01/24/2007 10:30:03
State : Suspended
Process Id : --
Command : Export NODE m* filespace=c$ nametype=unicode
filedata=all durunits=indefinite
toserver=athens
Phase : File list complete. Exporting eligible files
Total Running Time : 3 Days 0 Hours 24 Minutes
Current Process Running Time :
Export Operation Restart Count: 0
Date and Time of Last Restart : --
Date and Time of Last Suspend : 01/25/2007 08:30:11
Policy Domains Exported : 0
Policy Sets Exported : 0
Schedules Exported : 0
Mgmt Classes Exported : 0
Copy Groups Exported : 0
Administrators Exported : 1
Option Sets Exported : 0
Node Definitions Exported : 3
Filespace Definitions Exported : 7
Archive Files Exported : 50,000
Backup Files Exported : 150,000
Space Managed Files Exported : 0
Archive Files Skipped : 0
Backup Files Skipped : 25
Space Managed Files Skipped : 0
Total bytes Transferred (MB) : 7,000
Total Files to be Transferred : 900,000
Files Remaining : 700,000
```

See “Field descriptions” on page 763 for field descriptions.

Example: Display information for server-to-server export operations

List detailed information for all currently running server-to-server export operations. Issue the following command:

```
query export state=running format=detailed
```

```

Export Identifier : export_HOME_Dirs
Start Time : 01/25/2007 09:30:03
State : Running
Process Id : 11
Command : Export NODE n2,n3,n4
         filespace=/home nametype=
         server filedata=all
         toserver=athens
Phase : Identifying and exporting
        eligible files
Total Running Time : 0 Days 22 Hours 0 Minutes
Current Process Running Time : 01:30:00
Export Operation Restart Count: 4
Date and Time of last Restart : 02/01/2007 11:00:03
Date and Time of last Suspend : 01/31/2007 05:01:00
Policy Domains Exported : 0
Policy Sets Exported : 0
Schedules Exported : 0
Mgmt Classes Exported : 0
Copy Groups Exported : 0
Administrators Exported : 1
Option Sets Exported : 0
Node Definitions Exported : 3
Filespace Definitions Exported : 7
Archive Files Exported : 0
Backup Files Exported : 1000
Space Managed Files Exported : 0
Archive Files Skipped : 0
Backup Files Skipped : 0
Space Managed Files Skipped : 0
Total bytes Transferred (MB) : 50
Total Files to be Transferred : 400,000
Files Remaining : 399,000

```

See “Field descriptions” for field descriptions.

Field descriptions

Export identifier

The unique identifier assigned to this server-to-server export operation.

Start time

The time and date that this export operation was first initiated.

State The current state of this export operation. The value is one of the following:

Running - Not Suspending

The operation is active and is transmitting definitions to the target server. The process cannot be suspended, and if the process fails while in this state, you cannot restart it.

Running

The operation is active and is either searching for eligible files or transmitting file data to the target server.

Running - Suspend in Progress

The operation is in the process of being suspended as a result of a **SUSPEND EXPORT** command. The export operation is fully suspended when all of the data from the export operation is saved. An export operation in this state does not respond to the following commands:

- **CANCEL PROCESS**
- **CANCEL EXPORT**
- **RESTART EXPORT**

- **SUSPEND EXPORT**

Suspended

The operation stopped running due to a failure or was suspended with the **SUSPEND EXPORT** command.

Process ID

The process ID for the export operation when the status is either "Initializing" or "Running".

Command

The full command issued to start this server-to-server export.

Phase The current step that the operation is performing. The possible phases are shown in the order in which they are performed:

Creating definitions on target server

The operation is exporting definitions. The process cannot be suspended. Should the process fail in this phase, it cannot be restarted.

Identifying and exporting eligible files

The operation is building a list of eligible files for export. Some files may also be transmitted to the target during this phase. A process in this phase can be suspended. Should the process fail in this phase, it can be restarted.

File list complete. Exporting eligible files

The operation has completed building the list of eligible files for export and it is now transmitting the files to the target. A process in this phase can be suspended. Should the process fail in this phase, it can be restarted.

Total running time

The overall running time for this server-to-server export operation. For example, if this operation started and was then suspended and restarted two times, this value is the total running time of all three active processes of the export operation.

Current process running time

The running time of the active process of a server-to-server export operation. No value is displayed for a suspended operation because no active process exists.

Export operation restart count

The number of times the server-to-server export operation was restarted.

Date and time of last restart

The last date and time at which this server-to-server export operation was restarted.

Date and time of last suspend

The last date and time at which this server-to-server export operation was suspended.

Policy domains exported

The number of policy domain definitions successfully exported to the target server.

Policy sets exported

The number of policy set definitions successfully exported to the target server.

Schedules exported

The number of schedule definitions successfully exported to the target server.

Mgmt classes exported

The number of management class definitions successfully exported to the target server.

Copy groups exported

The number of copy group definitions successfully exported to the target server.

Administrators exported

The number of administrator definitions successfully exported to the target server.

Option sets exported

The number of option set definitions successfully exported to the target server.

Node definitions exported

The number of node definitions successfully exported to the target server.

File space definitions exported

The number of file space definitions successfully exported to the target server.

Archive files exported

The number of archive files successfully exported to the target server.

Backup files exported

The number of backup files successfully exported to the target server.

Space managed files exported

The number of space managed files successfully exported to the target server.

Archive files skipped

The number of archive files that were eligible for export but were skipped.

Backup files skipped

The number of backup files that were eligible for export but were skipped.

Space managed files skipped

The number of space managed files that were eligible for export but were skipped.

Total bytes transferred (MB)

The total number of bytes transmitted so far to the target server for this export operation.

Total files to be transferred

The total number of files transmitted so far to the target server for this export operation.

Files remaining

The total number of files remaining to be transmitted to the target server for this export operation.

QUERY EXPORT

Related commands

Table 235. Commands related to **QUERY EXPORT**

Command	Description
CANCEL PROCESS	Cancels a background server process.
CANCEL EXPORT	Deletes a suspended export operation.
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
IMPORT NODE	Restores client node information from external media.
IMPORT SERVER	Restores all or part of the server from external media.
QUERY PROCESS	Displays information about background processes.
RESTART EXPORT	Restarts a suspended export operation.
SUSPEND EXPORT	Suspends a running export operation.

QUERY FILESPACE (Query one or more file spaces)


Use this command to display information about file spaces that belong to a client node. The information includes the results of the last incremental backup or replication.

Important: If a node has more than one file space, you can issue a **DELETE FILESPACE** command for one of those file spaces. However, if you issue a **QUERY FILESPACE** command for the node during the deletion process, the output shows no file spaces. To obtain accurate information about remaining file spaces, issue the **QUERY FILESPACE** command after the deletion process ends.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

node_name

Specifies the client node to which the file space belongs. You can use wildcard characters to specify this name. This parameter is optional. The default is all client node names. You must specify a value for this parameter if you specify a file name.

file_space_name

Specifies the name of the file space to be queried. You can use wildcard characters to specify this name. This parameter is optional. If a value is not specified, all file spaces are queried.

For a server that has clients with support for Unicode-enabled file spaces, you might have to have the server convert the file space name that you enter. For example, you might have to have the server convert the name you enter from the server code page to Unicode. See the **NAMETYPE** parameter for details. If you do not specify a file space name, or if you specify only a single wildcard character for the name, you can use the **CODETYPE** parameter to limit the operation to Unicode file spaces or to non-Unicode file spaces.

File space names are case-sensitive. You can use the **QUERY FILESPACE** command to determine the correct capitalization for the file space to be queried.

QUERY FILESPACE

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. You can specify one of the following values:

Standard

Specifies that partial information is displayed for the specified file space.

Detailed

Specifies that complete information is displayed for the specified file space.

NAMETYPE

Specify how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with support for Unicode. You can use this parameter for Unicode-enabled Tivoli Storage Manager clients that have Windows, Macintosh OS X, and NetWare operating systems.

Use this parameter only when you enter a partly-qualified or fully-qualified file-space name. The default value is SERVER. You can specify one of the following values:

SERVER

The server uses the server code page to interpret the file space names.

UNICODE

The server converts the file space name entered from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server has problems accessing system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

CODETYPE

Specify what type of file spaces are to be included in the operation. The default is BOTH, meaning that file spaces are included regardless of code page type. Use this parameter only when you enter a single wildcard character for the file space name. You can specify one of the following values:

UNICODE

Include only file spaces that are in Unicode.

NONUNICODE

Include only file spaces that are not in Unicode.

BOTH

Include file spaces regardless of code page type.

Example: List all file spaces

Query all file spaces associated with all client nodes.

```
query filesystem
```

Node Name	Filespace Name	FSID	Platform	Filespace Type	Is Filespace Unicode?	Capacity (MB)	Pct Util
JOE	\\joe\c\$	1	WinNT	NTFS	Yes	2,502.3	75.2
JOE	\\joe\d\$	2	WinNT	NTFS	Yes	6,173.4	59.6

See "Field descriptions" on page 770 for field descriptions.

Example: Display detailed file space information for a virtual file space

Display detailed information for the file space /HomeDir which is a virtual file space mapping and belongs to the NAS node NAS1.

```
query filesystem nas1 /HomeDir
```

Node Name	Filespace Name	FSID	Platform	Filespace Type	Is Filespace Unicode?	Capacity (MB)	Pct Util
NAS1	/HomeDir	1	NetApp	WAFL (VFS)	No	2,502.3	75.2

See “Field descriptions” on page 770 for field descriptions.

Important: You might not see the expected results after requesting a detailed format because several fields must be completed by the API application. These fields include:

- Filespace type
- Platform
- Capacity (MB)
- Pct Util
- Last backup start Date/Time
- Last backup completion Date/Time

For more information about specific fields that are updated by the API, see the *IBM Tivoli Storage Manager: Using the Application Programming Interface*.

Example: Display detailed file space information for a specific file space and node

Display detailed information about the \\joe\c\$ file space that belongs to the client node JOE.

```
query filesystem joe \\joe\c$ nametype=unicode format=detailed
```

```

Node Name: JOE
Filespace Name: \\joe\c$
Hexadecimal Filespace Name: 5c5c6a6f655c6324
FSID: 1
Platform: WinNT
Filespace Type: NTFS
Is Filespace Unicode?: Yes
Capacity (MB): 2,502.3
Pct Util: 75.2
Last Backup Start Date/Time:
Days Since Last Backup Started:
Last Backup Completion Date/Time:
Days Since Last Backup Completed:
Last Replication Start Date/Time: 12/02/2010, 12:42:00
Days Since Last Node Replication Started: 30
Last Replication End Date/Time: 12/02/2010, 12:42:00
Days Since Last Replication Ended: 30
Backup Replication Rule: ACTIVE_DATA
Backup Data-Type Replication State: ENABLED
Archive Replication Rule: DEFAULT
Archive Data-Type Replication State: ENABLED
Space-Managed Replication Rule: NONE
Space-Managed Data-Type Replication State: DISABLED

```

See “Field descriptions” on page 770 for field descriptions.

Field descriptions

Important: You might not see the expected results after requesting a detailed format because several fields must be completed by the API application. These fields include:

- Filespace Type
- Platform
- Capacity (MB)
- Pct Util
- Last Backup Start Date/Time
- Last Backup Completion Date/Time

For more information about specific fields that are updated by the API, see the *IBM Tivoli Storage Manager: Using the Application Programming Interface*.

Node Name

Specifies the name of the client node.

Filespace Name

The name of the file space that belongs to the node.

File space names can be in a different code page or locale than the server. If they are, the names in the Administration Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, Tivoli Storage Manager converts the name to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

Hexadecimal Filespace Name

Specifies the hexadecimal name of the file space for the client node in UTF-8 format.

FSID Specifies the file space ID of the file space.

Platform

Specifies the platform for the client node.

Filespace Type

Specifies the type of file space.

A file space type that is appended with "(VFS)" denotes that this file space name is a virtual file space mapping for a directory path on a NAS device.

Is Filespace Unicode?

Indicates whether the file space is Unicode.

Capacity (MB)

Specifies the amount of space, in megabytes, assigned to this file space on the client node.

For a file space that is a virtual file space mapping for a directory path, this field represents the capacity of the file space on which the directory path is located.

Pct Util

Specifies the percentage of the file space that is occupied.

For a file space that is a virtual file space mapping for a directory path, the percentage utilized is calculated as the percentage of the capacity of the file space that was occupied by the directory at the time of the last full backup.

Last Backup Start Date/Time

Specifies the start date and time of the last incremental backup of the file space.

Days Since Last Backup Started

Specifies the number of days since the start of the last incremental backup of the file space.

Last Backup Completion Date/Time

Specifies the completion date and time of the last incremental backup of the file space.

Days Since Last Backup Completed

Specifies the number of days since the completion of the last incremental backup of the file space.

Last Replication Start Date/Time

Specifies the date and time that the last replication of file space data started.

Days Since Last Replication Started

Specifies the number of days since the last replication of file space data started.

Last Replication End Date/Time

Specifies the date and time that the last replication of file space data ended.

Days Since Last Replication Ended

Specifies the number of days since the last replication of file space data ended.

Backup Replication Rule

Specifies the replication rule that applies to backup data in the file space. The following values are possible:

ALL_DATA

Replicates active and inactive backup data. The data is replicated with a normal priority.

ACTIVE_DATA

Replicates only active backup data. The data is replicated with a normal priority.

ALL_DATA_HIGH_PRIORITY

Replicates active and inactive backup data. The data is replicated with a high priority.

ACTIVE_DATA_HIGH_PRIORITY

Replicates only active backup data. The data is replicated with a high priority.

QUERY FILESPACE

DEFAULT

Replicates backup data according to the client node rule for backup data. If the client node rule for backup data is DEFAULT, backup data is replicated according to the server rule for backup data.

NONE

Backup data in the file space is not replicated.

Backup Data-Type Replication State

Specifies whether replication of backup data in the file space is enabled or disabled. If the state is ENABLED, backup files are eligible for replication. If the state is DISABLED, backup files are not eligible for replication.

Archive Replication Rule

Specifies the replication rule that applies to archive data in the file space. The following values are possible:

ALL_DATA

Replicates archive data. The data is replicated with a normal priority.

ALL_DATA_HIGH_PRIORITY

Replicates archive data. The data is replicated with a high priority.

DEFAULT

Replicates archive data according to the client rule for archive data. If the client rule for archive data is DEFAULT, archive data is replicated according to the server rule for archive data.

NONE

Archive data in the file space is not replicated.

Archive Data-Type Replication State

Specifies whether replication of archive data in the file space is enabled or disabled. If the state is ENABLED, archive files are eligible for replication. If the state is DISABLED, archive files are not eligible for replication.

Space-Managed Replication Rule

Specifies the replication rule that applies to space-managed data in the file space. The following values are possible:

ALL_DATA

Replicates space-managed data. The data is replicated with a normal priority.

ALL_DATA_HIGH_PRIORITY

Replicates space-managed data. The data is replicated with a high priority.

DEFAULT

Replicates space-managed data according to the client rule for space-managed data. If the client rule for space-managed data is DEFAULT, space-managed data is replicated according to the server rule for space-managed data.

NONE

Space-managed data in the file space is not replicated.

Space-Managed Data-Type Replication State

Specifies whether replication of space-managed data in the file space is enabled or disabled. If the state is ENABLED, space-managed files are eligible for replication. If the state is DISABLED, space-managed files are not eligible for replication.

Related commands

Table 236. Commands related to QUERY FILESPACE

Command	Description
DEFINE VIRTUALFSMAPPING	Define a virtual file space mapping.
DELETE FILESPACE	Deletes data associated with clients file spaces.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
RENAME FILESPACE	Renames a client filesystem on the server.
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE NODE	Changes the attributes associated with a client node.


QUERY LIBRARY (Query a library)

Use this command to display information about libraries.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

library_name

Specifies the name of the library to be queried. You can use wildcards to specify names. This parameter is optional.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the library.

Detailed

Specifies that complete information is displayed for the library.

Example: Display summary information about a specific library

Display information about the library named AUTO. Issue the command:

```
query library auto
```

```

Library Name: AUTO
Library Type: SCSI
ACS Id:
Private Category:
Scratch Category:
WORM Scratch Category:
External Manager:
Shared: No
LanFree:
ObeyMountRetention:
  
```

See “Field descriptions” on page 775 for field descriptions.

Example: Display detailed library information about a specific library

Display information in full detail about the library named EZLIFE. Issue the command:

```
query library ezlife format=detailed
```

```

Library Name: EZLIFE
Library Type: SCSI
ACS Id:
Private Category:
Scratch Category:
WORM Scratch Category:
External Manager:
Shared: Yes
LanFree:
ObeyMountRetention:
Primary Library Manager: EZSERVER
WWN:
Serial Number:
AutoLabel: OVERWRITE
Relabel Scratch: Yes
Last Update by (administrator): DOCTOR_MIKE
Last Update Date/Time: 2002-12-05 15:24:53

```

See “Field descriptions” for field descriptions.

Field descriptions

Library Name

The name of the library.

Library Type

The type of library.

ACS Id

Specifies that the library is a StorageTek library that is controlled by StorageTek Automated Cartridge System Library Software (ACSL).

Private Category

The category number for private volumes that must be mounted by name.

The information displayed in this field only applies to an IBM 3494 or 3495 Tape Library Dataserver.

Scratch Category

The category number to use for scratch volumes in the library.

The information displayed in this field only applies to an IBM 3494 or 3495 Tape Library Dataserver.

WORM Scratch Category

The category number used for WORM scratch volumes in the library.

The information displayed in this field only applies to an IBM 3494 or 3495 Tape Library Dataserver.

External Manager

The location of the external library manager where the server can send media access requests.

Shared

Whether this library is shared with other Tivoli Storage Manager servers in a storage area network (SAN).

LanFree

Whether an external library is used for LAN-free operations.

ObeyMountRetention

Whether the server uses the value set for mount retention in the device class associated with this external library.

QUERY LIBRARY

Primary Library Manager

The name of the server that is responsible for controlling access to library resources.

WWN The Fibre Channel worldwide name for the library.

Serial Number

Specifies the serial number for the library being queried.

AutoLabel

Specifies whether the server attempts to automatically label tape volumes.

Relabel Scratch

Specifies whether the server relabels volumes that have been deleted and returned to scratch.

Last Update by (administrator)

Who performed the last update to the library.

Last Update Date/Time

The date and time when the last update occurred.

Related commands

Table 237. Commands related to QUERY LIBRARY

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
DEFINE LIBRARY	Defines an automated or manual library.
DEFINE PATH	Defines a path from a source to a destination.
DELETE LIBRARY	Deletes a library.
QUERY PATH	Displays information about the path from a source to a destination.
UPDATE LIBRARY	Changes the attributes of a library.
UPDATE PATH	Changes the attributes associated with a path.


QUERY LIBVOLUME (Query a library volume)

Use this command to display information about one or more volumes that are checked into an automated library for use by the Tivoli Storage Manager server.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

library_name

Specifies the name of the library. You can use wildcard characters to specify this name. This parameter is optional. The default is all libraries.

volume_name

Specifies the volume name. You can use wildcard characters to specify this name. This parameter is optional. The default is all volumes.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: List checked in volumes for a specific library

Display information about all of the volumes that are checked into the library named TAPE. See “Field descriptions” on page 778 for field descriptions.

```
query libvolume tape
```

Library Name	Volume Name	Status	Owner	Last Use	Home Element	Device Type
TAPE	000114	Scratch			1,000	LTO
TAPE	NY1602	Scratch			1,001	DLT

Example: Display detailed information for a specific library

Display detailed information about a volume named JJY008. See “Field descriptions” on page 778 for field descriptions.

```
query libvolume jjy008 format=detailed
```

QUERY LIBVOLUME

```
Library Name: HPW3494
Volume Name: JJY008
Status: Private
Owner: SUNSET
Last Use: Data
Home Element:
Device Type:
Cleanings Left:
Media Type:
```

Field descriptions

Library Name

The name of the library where the storage volume is located.

Volume Name

The name of the storage volume.

Status The status of the storage volume according to the library inventory. If the status is Private, the volume is being used by Tivoli Storage Manager. If the status is Scratch, the volume is available for use.

Owner

The owner server of the volume, if the volume is private.

Last Use

The type of data on the volume. This field applies only to volumes in Private status. For storage pool volumes, this field shows **Data**. For database backup volumes (full, incremental, or snapshot), this field shows **DbBackup**.

Home Element

The element address of the library slot containing the volume.

Device Type

The type of device that the volume is used on. This field will display a value only for volumes checked into a library that has mixed media capabilities.

Cleanings Left

For cleaner cartridges, the number of cleanings left.

Media Type

The type of media the volume represents (for example, 8mm tape).

Related commands

Table 238. Commands related to **QUERY LIBVOLUME**

Command	Description
AUDIT LIBRARY	Ensures that an automated library is in a consistent state.
CHECKIN LIBVOLUME	Checks a storage volume into an automated library.
CHECKOUT LIBVOLUME	Checks a storage volume out of an automated library.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.

Table 238. Commands related to QUERY LIBVOLUME (continued)

Command	Description
LABEL LIBVOLUME	Labels volumes in manual or automated libraries.
QUERY LIBRARY	Displays information about one or more libraries.
UPDATE LIBVOLUME	Changes the status of a storage volume.

QUERY LICENSE (Display license information)

Use this command to display license audit, license terms, and compliance information.

Privilege class

Any administrator can issue this command.

Syntax

►►—Query LICense—◄◄

Parameters

None.

To display the license information, issue the following command:

```
query license
```

The following example output is displayed:

```
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY LICENSE

                Last License Audit: 02/14/2013 14:28:08
                Number of TDP for Oracle in use: 0
                Number of TDP for Oracle in try buy mode: 0
                Number of TDP for MS SQL Server in use: 0
Number of TDP for MS SQL Server in try buy mode: 0
                Number of TDP for MS Exchange in use: 0
                Number of TDP for MS Exchange in try buy mode: 0
                Number of TDP for Lotus Notes in use: 12
                Number of TDP for Lotus Notes in try buy mode: 0
                Number of TDP for Lotus Domino in use: 0
                Number of TDP for Lotus Domino in try buy mode: 0
                Number of TDP for Informix in use: 1
                Number of TDP for Informix in try buy mode: 0
                Number of TDP for SAP R/3 in use: 0
                Number of TDP for SAP R/3 in try buy mode: 0
                Number of TDP for ESS in use: 0
                Number of TDP for ESS in try buy mode: 0
                Number of TDP for ESS R/3 in use: 0
                Number of TDP for ESS R/3 in try buy mode: 0
                Number of TDP for EMC Symmetrix in use: 0
                Number of TDP for EMC Symmetrix in try buy mode: 0
                Number of TDP for EMC Symmetrix R/3 in use: 6
                Number of TDP for EMC Symmetrix R/3 in try buy mode: 0
                Number of TDP for WAS in use: 0
                Number of TDP for WAS in try buy mode: 0
                Is IBM System Storage Archive Manager in use?: No
                Is IBM System Storage Archive Manager licensed?: Yes
                Is Tivoli Storage Manager Basic Edition in use: Yes
                Is Tivoli Storage Manager Basic Edition licensed: Yes
                Is Tivoli Storage Manager Extended Edition in use: No
                Is Tivoli Storage Manager Extended Edition licensed: Yes
                Server License Compliance: Valid
```

A product is in use if you purchased the product and registered the license.

Related commands

Table 239. Commands related to QUERY LICENSE

Command	Description
AUDIT LICENSES	Verifies compliance with defined licenses.
QUERY AUDITOCCUPANCY	Displays the server storage utilization for a client node.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY PVUESTIMATE	Displays processor value unit estimates. Remember: The QUERY PVUESTIMATE command reports licenses by providing PVU information on a per-node basis for server devices.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER LICENSE	Registers a license with the Tivoli Storage Manager server.
REGISTER NODE	Defines a client node to the server and sets options for that user.
SET CPUINFOREFRESH	Specifies the number of days between client scans for workstation information used for PVU estimates.
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.
UPDATE NODE	Changes the attributes associated with a client node.


QUERY LOG (Display information on the recovery log)

Use this command to display information about the recovery log.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display summary information about the recovery log

Display summary information about the recovery log. See “Field descriptions” for field descriptions.

```
query log
```

Total Space (MB)	Used Space (MB)	Free Space (MB)
38,912	401.3	38,368.7

Example: Display detailed information about the recovery log

Display detailed information about the recovery log. See “Field descriptions” for field descriptions.

```
query log format=detailed
```

Total Space (MB): 38,912
Used Space (MB): 401.3
Free Space (MB): 38,368.7
Active Log Directory : /activelog
Archive Log Directory : /archivelog
Mirror Log Directory : /mirrorlog
Archive Failover Log Directory : /archfailoverlog

Field descriptions

Total Space

Specifies the maximum size of the active log, in megabytes.

Used Space

Specifies the amount of used active log space, in megabytes.

Free Space

Specifies the amount of active log space that is not being used by uncommitted transactions, in megabytes.

Active Log Directory

Specifies the location where active log files are stored. When you change the active log directory, the server moves all archived logs to the archive log directory and all active logs to a new active log directory.

Mirror Log Directory

Specifies the location where the mirror for the active log is maintained.

Archive Failover Log Directory

Specifies the location into which the server saves archive logs if the logs cannot be archived to the archive log directory.

Archive Log Directory

Specifies the location into which the server can archive a log file after all the transactions that are represented in that log file are completed.

QUERY MACHINE (Query machine information)

Use this command to display information for one or more machines. You can use this information to recover Tivoli Storage Manager client machines in case of a disaster.

Attention: Tivoli Storage Manager does not use the information in any way. It is available only to help you plan for the disaster recovery of client machines.


Tivoli Storage Manager displays information for multiple machines in the following order:

- According to the priority specified.
- Within a priority, according to the specified location and machine name.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

machine_name

Specifies the name of one or more machines to be queried. You can use wildcard characters to specify this name. This parameter is optional. The default is all machines that meet the specified criteria.

BUilding

Specifies the name or number of the building that the machines are in. This parameter is optional. Enclose the text in quotation marks if it contains any blank characters.

FLoor

Specifies the name or number of the floor that the machines are on. This parameter is optional. Enclose the text in quotation marks if it contains any blank characters.

R0om

Specifies the name or number of the room that the machines are in. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters.

PRIority

Specifies the priority number of the machines. This parameter is optional.

ADSMServer

Specifies if the machine contains a Tivoli Storage Manager server. This parameter is optional. The default is to display any machines that meet the other criteria. Possible values are:

Yes

The machine contains a Tivoli Storage Manager server.

No

The machines do not contain a Tivoli Storage Manager server.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Tivoli Storage Manager displays partial information for the machines.

Detailed

Tivoli Storage Manager displays all information for the machines.

RECOVERYInstructions

Tivoli Storage Manager displays only machine recovery instructions. This option is valid only when querying a specific machine.

CHaracteristics

Tivoli Storage Manager displays only machine characteristics. This option is valid only when querying a specific machine.

Example: Display information for a specific machine

Display information for a machine named MACH1. See “Field descriptions” on page 786 for field descriptions.

query machine MACH1

Machine Name	Machine Priority	Building	Floor	Room	Node Name	Recovery Media Name
MACH1	1	21	2	2929	VIRGINIA	RECMED1

Example: Display detailed information for priority 1 machines

Display detailed information for all priority 1 machines on the second floor of building 21. See “Field descriptions” on page 786 for field descriptions.

query machine * building=21 floor=2 priority=1
format=detailed

```

Machine Name: MACH1
Machine Priority: 1
Building: 21
Floor: 2
Room: 2929
Server?: Yes
Description: TSM server machine
Node Name: VIRGINIA
Recovery Media Name: RECMED1
Characteristics?: Yes
Recovery Instructions?: Yes
    
```

QUERY MACHINE

Field descriptions

Machine Name

The name of the machine.

Machine Priority

The recovery priority of the machine.

Building

The building in which the machine is located.

Floor The floor on which the machine is located.

Room The room in which the machine is located.

Server?

Whether the machine contains a Tivoli Storage Manager server.

Description

A description of the machine.

Node Name

The Tivoli Storage Manager client nodes associated with this machine.

Recovery Media Name

The recovery media associated with this machine.

Characteristics?

Whether the characteristics text of the machine is stored in the database.

Recovery Instructions?

Specifies whether recovery instructions text for a machine is stored in the Tivoli Storage Manager database.

Related commands

Table 240. Commands related to QUERY MACHINE

Command	Description
DEFINE MACHINE	Defines a machine for DRM.
DEFINE MACHNODEASSOCIATION	Associates a Tivoli Storage Manager node with a machine.
DEFINE RECMEDMACHASSOCIATION	Associates recovery media with a machine.
DELETE MACHINE	Deletes a machine.
INSERT MACHINE	Inserts machine characteristics or recovery instructions into the Tivoli Storage Manager database.
UPDATE MACHINE	Changes the information for a machine.


QUERY MEDIA (Query sequential access storage pool media)

Use this command to display information about the sequential access primary and copy storage pool volumes moved by the **MOVE MEDIA** command.

Privilege class

Any administrator with system or operator privilege can issue this command unless it includes the **CMD** parameter. If the **CMD** parameter is specified and the **REQSYSAUTHOUTFILE** server option is set to **NO**, the administrator must have operator, unrestricted storage, or system privilege. If the **CMD** parameter is specified and the **REQSYSAUTHOUTFILE** server option is set to **YES** (the default), the administrator must have system privilege.

Syntax



Parameters

volume_name

Specifies the name of the sequential access primary or copy storage pool volume to display. This parameter is optional. You can use a wildcard character to specify the name. All matching volumes are considered for processing. If you do not specify this parameter, all volumes defined in the storage pool specified with the **STGPOOL** parameter display.

STGpool (Required)

Specifies the name of the sequential access primary or copy storage pool that is used to select the volumes for processing. You can use wildcard characters to specify the name. All matching storage pools are processed. If the storage pool specified is not managed by an automated library, no volumes display.

QUERY MEDIA

Days

Specifies the number of days that must elapse, after the volume has been written to or read from, before the volume is eligible for processing. This parameter is optional. You can specify a number from 0 to 9999. The default value is 0. The most recent of the volume's last written date or last read date is used to calculate the number of days elapsed.

WHERESTATUS

Specifies that the output of the query should be restricted by volume status. This parameter is optional. You can specify more than one status in a list by separating each status with a comma and no intervening spaces. If you do not specify a value for this parameter, all volumes in the specified storage pool, regardless of their status, are displayed.

Possible values are:

FULL

Specifies that volumes with a status of FULL display.

FILLing

Specifies that volumes with a status of FILLING display.

EMPtY

Specifies that volumes with a status of EMPTY display.

WHEREACCEss

Specifies that output should be restricted by volume access mode. This parameter is optional. If you do not specify a value for this parameter, output is not restricted by access mode.

Possible values are:

READWrite

Specifies that volumes with an access mode of READWRITE display.

READOnly

Specifies that volumes with an access mode of READONLY display.

Format

Specifies how information displays. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information displays for the specified sequential access storage pool volumes.

Detailed

Specifies that complete information displays for the specified sequential access storage pool volumes.

Cmd

Specifies that executable commands are built for the storage pool volumes processed by the **QUERY MEDIA** command. These commands will be in the file specified with the **CMDFILENAME** parameter on the **QUERY MEDIA** command. If you want the commands to display on the console only, specify a null string ("") for the **CMDFILENAME**. If **FORMAT=CMD** is specified but no command string is specified with the **CMD** parameter, the **QUERY MEDIA** command will fail.

WHEREState

Specifies the state of volumes to process. This parameter restricts processing to volumes that have the specified state. This parameter is optional. The default is ALL. Possible values are:

A11

Specifies that volumes in all states are queried. The valid states are: MOUNTABLEINLIB and MOUNTABLENOTINLIB.

MOUNTABLEInlib

Specifies that volumes that are currently in the MOUNTABLEINLIB state are queried. Volumes in the MOUNTABLEINLIB state are in the library, and are onsite, contain valid data, and are available for onsite processing.

MOUNTABLENotinlib

Specifies that volumes that are currently in the MOUNTABLENOTINLIB state are queried. Volumes in the MOUNTABLENOTINLIB state are not in the library, do not contain valid data, and are not available for onsite processing.

WHEREOVFLocation

Specifies the overflow location of the volumes to display. This parameter is optional. This parameter restricts processing to volumes that are in the specified location. The maximum length of the location is 255 characters. The location must be enclosed in quotation marks if it contains any blank characters.

Cmd

Specifies the creation of executable commands. Enclose the command specification in quotation marks. The maximum length of the command specification is 255 characters. This parameter is optional.

For each volume successfully processed by the **QUERY MEDIA** command, the server writes the associated commands to a file. Specify the file name with the **CMDFILENAME** parameter.

If you do not specify a filename, the command will generate a default filename by appending the string `exec.cmds.media` to the server directory.

Remember:

1. If the command written to the file exceeds 255 characters, it is split into multiple lines, and a continuation character (+) is added to all but the last line. You may need to alter the continuation character according to the requirements of the product that runs the commands.
2. If an executable command is specified with any value for **FORMAT** other than **CMD**, the command string is ignored, and the **QUERY MEDIA** command will not write any command line.

Specify a command string and any substitution variables:

string

Specifies the string to build an executable command to process the volume name or volume location or both. You can specify any free form text for the string. Do not use embedded quotation marks. For example, the following is a valid executable command specification:

```
cmd="checkin libvolume &vol"
```

The following is an invalid executable command specification:

```
cmd="checkin libvolume "&vol"
```

substitution

Specifies a variable for which you want the **QUERY MEDIA** command to substitute a value. The possible substitution variables are:

QUERY MEDIA

&VOL

Substitute the volume name for &VOL. You can specify lowercase characters, &vol. No spaces or blanks are allowed between ampersand, &, and VOL. If there are spaces or blanks between ampersand and VOL, the **QUERY MEDIA** command will treat them as strings and no substitution will be set. If &VOL is not specified, no volume name is set in the executable command.

&LOC

Substitute the volume location for &LOC. You can specify lowercase characters, &loc. No spaces or blanks are allowed between ampersand, &, and LOC. If there are spaces or blanks between ampersand and LOC, the **QUERY MEDIA** command will treat them as strings and no substitution will be set. If &LOC is not specified, no location name is set in the executable command.

&VOLDSN

Substitute the volume file name for &VOLDSN. An example of a copy storage pool tape volume file name using the defined prefix IBM Tivoli Storage Manager310 is IBM Tivoli Storage Manager310.BFS. If &VOLDSN is not specified, no volume file name is set in the executable command.

&NL

Substitute the new line character for &NL. When &NL is specified, the **QUERY MEDIA** command will split the command at the position where the &NL is and will not append any continuation character. The user is responsible for specifying the proper continuation character before the &NL if one is required. The user is also responsible for the length of the line written. If the &NL is not specified and the command exceeds 255 characters, the command is split into multiple lines, and a continuation character (+) is added to all but the last line.

CMDFilename

Specifies the full path name that will contain the commands specified with CMD parameter when **FORMAT=CMD** is specified. This parameter is optional. The maximum length of the file name is 1279 characters.

If you specify "" with the **CMDFILENAME** parameter, the **QUERY MEDIA** command will generate a file name by appending the "exec.cmds.media" to the server directory. The server directory is the current working directory of the server process.

If you specify a null string (""), the commands built are displayed on the console only. You can redirect the commands displayed to a file by using the redirection characters for the operating system (> or >>).

If the filename is not specified, the command will generate a default filename by appending the string "exec.cmds.media" to the server directory.

The **QUERY MEDIA** command automatically allocates the file name specified or generated. If the file name exists, the **QUERY MEDIA** command will attempt to use it and the existing data, if any, in the file to be overwritten. You can specify **APPEND=YES** to prevent the existing data from being overwritten. If the **QUERY MEDIA** command fails after the command file is allocated, the file is not deleted.

APPend

Specifies to write at the beginning or the ending of the command file data. This parameter is optional. The default is **NO**. Possible values are:

No Specifies to write the data from the beginning of the command file. If the given command file exists, its contents are overwritten.

Yes

Specifies to append the command file by writing at the end of the command file data.

Example: Display information on a specific sequential access storage pool

Display all full and partial full volumes that are in the sequential access primary storage pool, ARCHIVE. See “Field descriptions” on page 792 for field descriptions.

```
query media * stgpool=archive wherestatus=full, filling
```

Volume Name	State	Location	Automated LibName
TAPE01	Mountable in Library		LIB3494
TAPE03	Mountable not in Lib.	Room1234/Bldg31	
TAPE07	Mountable in Library		LIB3494
TAPE09	Mountable not in Lib.	Room1234/Bldg31	

Example: Display information on sequential access storage pool with a specific prefix

Display in detail all full volumes in MOUNTABLENOTINLIB state for sequential access storage pools that have a prefix name of ONSITE. See “Field descriptions” on page 792 for field descriptions.

```
query media wherestate=mountablenotinlib stgpool=onsite* wherestatus=full format=detailed
```

```

Volume Name: TAPE21
State: Mountable not in library
Status: Full
Access: ReadOnly
Last Reference Date: 01/30/98
Last Update Date/Time: 08/20/1996 13:29:02
Location: Rm569/bldg31
Storage Pool Name: ONSITE.ARCHIVE
Automated Libname:

Volume Name: TAPE22
State: Mountable not in library
Status: Full
Access: ReadOnly
Last Reference Date: 01/30/98
Last Update Date/Time: 08/20/1996 15:29:02
Location: Rm569/bldg31
Storage Pool Name: ONSITE.ARCHIVEPOOL
Automated Libname:
    
```

Example: Generate checkin commands

Generate the **CHECKIN LIBVOLUME** commands for full and partially full volumes that are in the ONSITE.ARCHIVE primary storage pool and stored in the overflow location Room 2948/Bldg31.

```
query media * stgpool=onsite.archive format=cmd wherestatus=full,filling wherestate=mountablenotinlib whereovflocation=room2948/bldg31 cmd="checkin libvol lib3494 &vol status=private" cmdfilename=/tsm/move/media/checkin.vols
```

QUERY MEDIA

The **QUERY MEDIA** command created the **CHECKIN LIBVOLUME** executable commands in `/tsm/move/media/checkin.vols`, which can be run by issuing the **MACRO** command with `/tsm/move/media/checkin.vols` as the macro name.

```
checkin libvol lib3494 TAPE04 status=private
checkin libvol lib3494 TAPE13 status=private
checkin libvol lib3494 TAPE14 status=private
```

Field descriptions

Volume Name

Specifies the name of the primary sequential access storage pool volume.

State Specifies the state of the volume.

Status Specifies the status of the volume.

Access

Specifies the access mode of the volume.

Last Reference Date

Specifies the volume's last written date or last read date, whichever is more recent.

Last Update Date/Time

Specifies the date and time when the volume was most recently updated.

Location

Specifies where the volume is stored. If the volume is ejected from the library and its location is not specified or defined, a question mark (?) is displayed for the location.

Storage Pool Name

Specifies the name of the sequential access storage pool where the volume is defined.

Automated LibName

Specifies the automated library name if the volume is in the library.

Related commands

Table 241. Commands related to **QUERY MEDIA**

Command	Description
MOVE MEDIA	Moves storage pool volumes that are managed by an automated library.


QUERY MGMTCLASS (Query a management class)

Use this command to display information about management classes.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

domain_name

Specifies the policy domain associated with the management class to query. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a value for this parameter, management classes in all policy domains are queried. You must specify this parameter when querying an explicitly named management class.

policy_set_name

Specifies the policy set associated with the management class to query. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a value for this parameter, management classes in all policy sets are queried. You must specify this parameter when querying an explicitly named management class.

class_name

Specifies the management class to query. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a value for this parameter, all management classes are queried.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display information for all management classes

Query all management classes for all policy domains. Create the output in standard format. See “Field descriptions” on page 794 for field descriptions.

```
query mgmtclass
```

QUERY MGMTCLASS

Policy Domain Name	Policy Set Name	Mgmt Class Name	Default Mgmt Class ?	Description
EMPLOYEE-RECORDS	ACTIVE	ACTIVEFILES	Yes	Modified default management class
EMPLOYEE-RECORDS	HOLIDAY	ACTIVEFILES	Yes	Modified default management class
EMPLOYEE-RECORDS	HOLIDAY	FILEHISTORY	No	Test modified management class
EMPLOYEE-RECORDS	VACATION	ACTIVEFILES	Yes	Original default management class
EMPLOYEE-RECORDS	VACATION	FILEHISTORY	No	Test modified management class
PROG1	SUMMER	MCLASS1	No	Technical Support Mgmt Class
PROG2	SUMMER	MCLASS1	No	Technical Support Mgmt Class
STANDARD	ACTIVE	STANDARD	Yes	Installed default management class
STANDARD	STANDARD	STANDARD	Yes	Installed default management class

Example: Display detailed information for a specific management class

Query the ACTIVEFILES management class that is assigned to the VACATION policy set of the EMPLOYEE_RECORDS policy domain. Create the output in detailed format. See “Field descriptions” for field descriptions.

```
query mgmtclass employee_records vacation
activefiles format=detailed
```

```
Policy Domain Name: EMPLOYEE_RECORDS
Policy Set Name: VACATION
Mgmt Class Name: ACTIVEFILES
Default Mgmt Class ? : Yes
Description: Installed default management class
Space Management Technique: None
Auto-Migrate on Non-Use: 0
Migration Requires Backup?: Yes
Migration Destination: SPACEMGPOOL
Last Update by (administrator): $$CONFIG_MANAGER$$
Last Update Date/Time: 05/31/1998 13:15:45
Managing Profile: EMPLOYEE
Changes Pending: Yes
```

Field descriptions

Policy Domain Name

The policy domain.

Policy Set Name

The policy set.

Mgmt Class Name

The management class.

Default Mgmt Class?

Whether the management class is the default management class for the policy set.

Description

The description of the management class.

Space Management Technique

The space management technique for the management class, for Tivoli Storage Manager for Space Management clients.

Auto-Migrate on Non-Use

The number of days that must elapse since a file was last accessed before it is eligible for automatic migration by Tivoli Storage Manager for Space Management clients.

Migration Requires Backup?

Whether a backup version of a file must exist before a file can be migrated by Tivoli Storage Manager for Space Management clients.

Migration Destination

The storage pool that is the destination for files migrated by Tivoli Storage Manager for Space Management clients.

Last Update by (administrator)

The administrator or server that most recently updated the management class. If this field contains \$\$CONFIG_MANAGER\$\$, the management class is associated with a domain that is managed by the configuration manager.

Last Update Date/Time

The date and time when the management class was most recently defined or updated.

Managing Profile

The profile or profiles to which the managed server subscribed to get the definition of this management class.

Changes Pending

Whether or not changes are being made but not activated. Once the changes are activated, the field resets to No.

Related commands

Table 242. Commands related to QUERY MGMTCLASS

Command	Description
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE MGMTCLASS	Defines a management class.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
QUERY DOMAIN	Displays information about policy domains.
UPDATE MGMTCLASS	Changes the attributes of a management class.

QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)

Use this command to display information about alert monitoring and server status settings.

Privilege class

Any administrator can issue this command.

Syntax

►► Query MONITORSettings ◀◀

Display monitoring settings

Display details about the monitoring settings. See Field descriptions for more details.

query monitorsettings

Example output:

```
Monitor Status: On
Status Refresh Interval (Minutes): 5
Status Retention (Hours): 48
Monitor Message Alerts: On
Alert Update Interval (Minutes): 10
Alert to Email: On
Send Alert Summary to Administrators: On
Alert from Email Address: DJADMIN@MYDOMAIN.COM
Alert SMTP Host: DJHOST.MYDOMAIN.COM
Alert SMTP Port: 25
Alert Active Duration (Minutes): 480
Alert Inactive Duration (Minutes): 480
Alert Closed Duration (Minutes): 60
Monitoring Admin: ADMIN
Monitored Group: MONGROUP
Monitored Servers: SERVER2
At-Risk Interval for Applications: 24
Skipped files as At-Risk for Applications?: Yes
At-Risk Interval for Virtual Machines: 24
Skipped files as At-Risk for Virtual Machines?: Yes
At-Risk Interval for Systems: 24
Skipped files as At-Risk for Systems?: Yes
```

Field descriptions

Monitor status

Specifies whether alert monitoring on the server is enabled or disabled.

Status refresh interval (minutes)

Specifies the number of minutes between intervals that the monitoring server gathers event data.

Status retention (hours)

Specifies the number of hours that status monitoring indicators are retained.

Monitor message alerts

Specifies whether alerts are sent to administrators by email.

| **Alert update interval (minutes)**

| Specifies the length of time, in minutes, that the alert monitor waits before
| the alert is updated and pruned on the server.

| **Alert to email**

| Specifies whether alerts are sent to administrators by email.

| **Send alert summary to administrators**

| Specifies the administrators that receive a summary of existing alerts on
| the server in an email.

| **Alert from email address**

| Specifies the email address of the sender.

| **Alert SMTP host**

| Specifies the Simple Mail Transfer Protocol (SMTP) host mail server that is
| used to send alerts by email.

| **Alert SMTP port**

| Specifies the SMTP mail server port that is used to send alerts by email.

| **Alert active duration (minutes)**

| Specifies how long, in minutes, an alert remains active.

| **Alert inactive duration (minutes)**

| Specifies how long, in minutes, an alert remains inactive.

| **Alert closed duration (minutes)**

| Specifies how long, in minutes, an alert remains closed before it is deleted
| from the server.

| **Monitoring admin**

| Specifies the name of the monitoring administrator that is used to connect
| to the servers in the monitored group.

| **Monitored group**

| Specifies the name of the monitored server group.

| **Monitored servers**

| Specifies the names of the servers in the monitored server group. The
| monitor settings might be different on each monitored server. If so, issue
| the query command for each server to display the monitoring settings.

| **At-risk interval for applications**

| Specifies how long, in hours, an applications client can log no activity
| before it is considered at-risk.

| **Skipped files as at risk for applications**

| Specifies that the server considers skipped files, by the client as a failure,
| and marks the client at-risk.

| **At-risk interval for virtual machines**

| Specifies how long, in hours, a virtual client can log no activity before it is
| considered at-risk.

| **Skipped files as at risk for virtual machines**

| Specifies that the server considers skipped files, by the client as a failure
| and marks the client at-risk.

| **At-risk interval for systems**

| Specifies how long, in hours, a systems client can log no activity before it
| is considered at-risk.

QUERY MONITORSETTINGS

Skipped files as at risk for systems

Specifies that the server considers skipped files, by the client as a failure, and marks the client at-risk.

Related commands

Table 243. Commands related to QUERY MONITORSETTINGS

Command	Description
"DEFINE ALERTTRIGGER (Define an alert trigger)" on page 112	Associates specified messages to an alert trigger.
"DELETE ALERTTRIGGER (Remove a message from an alert trigger)" on page 400	Removes a message number that can trigger an alert.
"QUERY ALERTSTATUS (Query the status of an alert)" on page 674	Displays information about alerts that have been issued on the server.
"QUERY ALERTTRIGGER (Query the list of defined alert triggers)" on page 672	Displays message numbers that trigger an alert.
"SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)" on page 1142	Specifies whether to enable client at-risk activity interval evaluation.
"SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)" on page 1147	Specifies whether to use client at-risk skipped files as failure evaluation.
"UPDATE ALERTTRIGGER (Update a defined alert trigger)" on page 1165	Updates the attributes of one or more alert triggers.
"UPDATE ALERTSTATUS (Update the status of an alert)" on page 1168	Updates the status of a reported alert.

QUERY MONITORSTATUS (Query the monitoring status)


Use this command to display monitoring messages that are within the defined status retention period.

You can limit the output to a specified status, such as only messages with a status of active. If you do not specify any parameters, all messages are displayed.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

Format

Specifies the amount of information that is displayed. The default value is STANDARD. Specify one of the following values:

Standard

Specifies that only partial information is displayed for the specified messages.

Detailed

Specifies that all information is displayed for the specified messages.

Type

This parameter restricts the output to only messages with the specified type value. Specify one of the following values:

All

Displays all information.

Active

Displays all active messages. This is the default value.

Inactive

Displays all inactive messages.

QUERY MONITORSTATUS

Activity

Specifies the activity that you want to query. See the **DEFINE STATUSTHRESHOLD** command for details on available activities to query.

Name

Specifies the name that you want to query. The NAME value refers to the name of the element with the specified activity. For example, a status indicator that contains information about a storage pool that is called backuppool has the NAME set to BACKUPPOOL.

Status

Specifies the status of the messages that you want to query. You can specify multiple status values in a list by separating the values with commas and no intervening spaces. If you do not specify a value for this parameter, information for all status values is displayed. Specify one of the following values:

Normal

Displays all messages with a normal status.

Warning

Displays all messages with a warning status.

Error

Displays all messages with an error status.

Display monitoring settings

Display details about the monitoring status.

Query MONITORStatus type=active

Example output:

```
Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: CAPACITY OF PRIMARY DISK AND FILE STORAGE
Element Name: CAPACITY OF PRIMARY DISK AND FILE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL

Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: USED CAPACITY OF PRIMARY DISK AND FILE STORAGE
Element Name: USED CAPACITY OF PRIMARY DISK AND FILE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL

Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: CAPACITY OF PRIMARY TAPE STORAGE
Element Name: CAPACITY OF PRIMARY TAPE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL

Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: USED CAPACITY OF PRIMARY TAPE STORAGE
Element Name: USED CAPACITY OF PRIMARY TAPE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL
```

Display monitoring settings

Display details about the monitoring status.

```
query monitorstatus f=d type=active
```

Example output:

```

Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: CAPACITY OF PRIMARY DISK AND FILE STORAGE
Element Name: CAPACITY OF PRIMARY DISK AND FILE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL
Element Details:
Primary Repair Suggestion:
First Alternate Repair Suggestion:
Second Alternate Repair Suggestion:

Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: USED CAPACITY OF PRIMARY DISK AND FILE STORAGE
Element Name: USED CAPACITY OF PRIMARY DISK AND FILE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL
Element Details:
Primary Repair Suggestion:
First Alternate Repair Suggestion:
Second Alternate Repair Suggestion:

Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: CAPACITY OF PRIMARY TAPE STORAGE
Element Name: CAPACITY OF PRIMARY TAPE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL
Element Details:
Primary Repair Suggestion:
First Alternate Repair Suggestion:
Second Alternate Repair Suggestion:

Server Name: SERVER1
Activity Date: 03/05/2013 15:57:37
Activity Name: USED CAPACITY OF PRIMARY TAPE STORAGE
Element Name: USED CAPACITY OF PRIMARY TAPE STORAGE
Element Numeric Value: 0
Element String Value:
Element State: NORMAL
Element Details:
Primary Repair Suggestion:
First Alternate Repair Suggestion:
Second Alternate Repair Suggestion:

```

Related commands

Table 244. Commands related to QUERY MONITORSTATUS

Command	Description
“DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)” on page 337	Defines a status monitoring threshold.
“DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)” on page 451	Deletes a status monitoring threshold.
“QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)” on page 796	Displays information about monitoring alerts and server status settings.

QUERY MONITORSTATUS

Table 244. Commands related to **QUERY MONITORSTATUS** (continued)

Command	Description
"QUERY STATUSTHRESHOLD (Query status monitoring thresholds)" on page 914	Displays information about a status monitoring thresholds.
"SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)" on page 1142	Specifies whether to enable client at-risk activity interval evaluation.
"SET STATUSMONITOR (Specifies whether to enable status monitoring)" on page 1144	Specifies whether to enable status monitoring.
"SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)" on page 1146	Specifies the refresh interval for status monitoring.
"SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)" on page 1147	Specifies whether to use client at-risk skipped files as failure evaluation.
"UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)" on page 1375	Changes the attributes of an existing status monitoring threshold.


QUERY MOUNT (Display information on mounted sequential access volumes)

Use this command to display information about the status of one or more sequential access volumes that are mounted.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

volume_name

Specifies the name of the mounted sequential access volume. You can use wildcard characters to specify this name. This parameter is optional. The default is *all mounted volumes*.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: List all mounted sequential volumes

Display information on all mounted sequential media volumes.

```
query mount
```

```

ANR8330I 3590 volume D6W992 is mounted R/O
in drive RMT1/dev/IBMtape1, status: IN USE.
ANR8334I 1 volumes found.
ANR8331I 8MMTAPE volume WPD000 is mounted R/W
in drive 8MM.1 (/dev/tmsmcsi/mt0), status: DISMOUNTING.
ANR8334I 1 volumes found.
  
```

Remember:

1. If the status of a volume is full or if its access mode is read-only (R/O), the mount mode of the volume will be R/O. To determine the status and access mode of a volume, issue the **QUERY VOLUME FORMAT=DETAILED** command. If a volume can be written to (that is, the status is filling or empty), the mount mode of the volume will be read/write (R/W), even if it is currently only being read.
2. In a storage pool associated with the FILE device type, the server can perform concurrent multiple read-access and one write-access to the same volume. As a result, a volume in a storage pool with a device type of FILE can appear as mounted more than once.

QUERY MOUNT

3. In the message ANR8448I, the drive name is listed as UNKNOWN for volumes of the FILE device type with a non-shared device class. The reason is that no drive is associated with the volumes; drive names are shown in the file-based library.
4. If you issue the **QUERY MOUNT** command while the drive is being cleaned, the command output continues to show a DISMOUNTING status for the dismounted volume until the cleaning completes.

Example: Display detailed information about mounted sequential volumes

Display details about mounted volumes.

```
query mount format=detailed
```

```
ANR2017I Administrator SERVER_CONSOLE issued command: QUERY
MOUNT format=detailed
ANR8487I Mount point in device class FILE is waiting for the
volume mount to
complete -- owning server: SERVER1, status: WAITING FOR VOLUME
(session: 0, process: 1).
ANR8488I LTO volume 015005L4 is mounted R/W in drive IBMVTL1
(/dev/rmt37) -- owning
server: SERVER1, status: IN USE (session: 0, process: 2).
ANR8486I Mount point in device class FILE is reserved -- owning
server: SERVER1,
status: RESERVED (session: 5, process: 0).
ANR8334I          3 matches found.
```

Related commands

Table 245. Commands related to **QUERY MOUNT**

Command	Description
DISMOUNT VOLUME	Dismounts a sequential, removable volume by the volume name.
REPLY	Allows a request to continue processing.

QUERY NASBACKUP (Query NAS backup images)


Use this command to display information about the file system image objects that have been backed up for a specific NAS node and file space. You can only use this command to display objects that were backed up for a NAS node using NDMP.

The server displays all matching objects, the dates that these objects were backed up, and information about a table of contents (TOC) for the object.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

node_name (Required)

Specifies the name of the NAS node for which backup objects are displayed. You cannot use wildcards to specify this name.

filespace_name (Required)

Specifies the name of the file space for which backup objects are displayed. You can use wildcards to specify this name.

BEGINDate

Specifies the beginning date to select the backup objects to display. All backup objects that were created on or after the specified date are displayed. The default is seven days prior to the current date. You can use this parameter with the `BEGINTIME` parameter to specify a range for the date and time. This parameter is optional.

You can specify the date using one of the following values:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/2002
TODAY	The current date	TODAY
TODAY- <i>days</i> or <i>-days</i>	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY -7 or -7. To display information about the image objects that have been created a week ago, you can specify BEGINDATE=TODAY-7 or BEGINDATE= -7.

QUERY NASBACKUP

Value	Description	Example
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time to select the backup objects to display. All backup objects created on or after the specified time display. This parameter is optional. The default is midnight (00:00:00) on the date specified for the BEGINDATE.

You can specify the time using one of the following values:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time on the specified begin date	10:30:08
NOW	The current time on the specified begin date	NOW
NOW+ <i>HH:MM</i> or + <i>HH:MM</i>	The current time plus hours and minutes on the specified begin date	NOW+03:00 or +03:00. If you issue this command at 9:00 with BEGINTime=NOW+3 or BEGINTime=+3, Tivoli Storage Manager displays image objects with a time of 12:00 or later on the begin date.
NOW- <i>HH:MM</i> or - <i>HH:MM</i>	The current time minus hours and minutes on the specified begin date	NOW-04:00 or -04:00. If you issue this command at 9:00 with BEGINTime=NOW-3:30 or BEGINTime=-3:30, Tivoli Storage Manager displays image objects with a time of 5:30 or later on the begin date.

ENDDate

Specifies the ending date used to select the backup objects to be displayed. All backup objects created on or before the specified date are displayed. This parameter is optional. The default is the current date. You can use this parameter with the ENDTIME parameter to specify an ending date and time.

You can specify the date using one of the following values:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/2002
TODAY	The current date	TODAY

Value	Description	Example
TODAY- <i>days</i> or <i>-days</i>	The current date minus days specified. The maximum number of days you can specify is 9999.	TODAY-1 or -1. To display information created up to yesterday, you can specify ENDDATE=TODAY-1 or simply ENDDATE= -1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time used to select the backup objects to be displayed. All backup objects created on or before the specified time are displayed. This parameter is optional. The default is 23:59:59. You can use this parameter with the ENDDATE parameter to specify a range for the date and time.

You can specify the time using one of the following values:

Value	Description	Example
HH:MM:SS	A specific time on the specified end date	10:30:08
NOW	The current time on the specified end date	NOW
NOW+HH:MM or +HH:MM	The current time plus hours and minutes on the specified end date	NOW+03:00 or +03:00. If you issue this command at 9:00 with ENDTIME=NOW+3:00 or ENDTIME= +3:00, Tivoli Storage Manager displays image objects with a time of 12:00 or later on the end date you specify.
NOW-HH:MM or -HH:MM	The current time minus hours and minutes on the specified end date	NOW-03:30 or -03:30. If you issue this command at 9:00 with ENDTIME=NOW-3:30 or ENDTIME= -3:30, IBM Tivoli Storage Manager displays image objects with a time of 5:30 or later on the end date you specify.

TYPE

Specifies the type of NDMP backup images for which you want to display information. The default value for this parameter is BACKUPIIMAGE. Other image types represent backup methods that might be specific to a particular file server. Possible values are:

QUERY NASBACKUP

BACKUPImage

Specifies that the output should show only the standard NAS base and differential images. This is the default value for this parameter.

SNAPMirror

Specifies whether to display information about NetApp SnapMirror images. SnapMirror images are block-level full-backup images of a file system. A SnapMirror image can only be restored to a file system that has been prepared as a SnapMirror target volume. Refer to the documentation that came with your NetApp file server for more information. This parameter is valid for NetApp and IBM N-Series file servers only.

Example:

Issue the **QUERY NASBACKUP** command to display information about a node, `nas1`, and a filespace, `/vol/vol1`.

```
query nasbackup nas1 /vol/vol1
```

Node Name	Filespace Name	Object Type (MB)	Object Size (MB)	Creation Date Contents	Has Table of Contents (TOC)	Mgmt Class Name	Image Storage Pool Name
NAS1	vol/vol1	Full image	1050.5	10/22/2002 10:50:57	YES	DEFAULT	NASBACKUPS
NAS1	vol/vol1	Differential image	9.1	10/22/2002 11:03:21	YES	DEFAULT	NASBACKUPS
NAS1	vol/vol1	Full image	1050.5	10/22/2006 10:43:00	YES	STANDARD	FILEPOOL
NAS1	vol/vol1	Differential image	9.1	10/25/2006 11:53:21	YES	STANDARD	FILEPOOL

Example:

Issue the **QUERY NASBACKUP** command to display information about all NetApp SnapMirror to Tape images for a node, `nas2`, and a filespace, `/vol/vol2`.

```
query nasbackup nas2 /vol/vol2 type=snapmirror
```

Node Name	Filespace Name	Object Type	Object Size (MB)	Creation Date	Mgmt Class Name	Image Storage Pool Name
NAS2	vol/vol2	SnapMirror	1050.5	04/02/2008 10:50:57	STANDARD	MYPOOL
NAS2	vol/vol2	SnapMirror	1450.5	04/02/2008 11:03:21	STANDARD	MYPOOL

Related commands

Table 246. Commands related to **QUERY NASBACKUP**

Command	Description
BACKUP NODE	Backs up a network-attached storage (NAS) node.
BACKUP NAS (Tivoli Storage Manager client command)	Creates a backup of NAS node data.
QUERY TOC	Displays details about the table of contents for a specified backup image.
RESTORE NODE	Restores a network-attached storage (NAS) node.


QUERY NODE (Query nodes)

Use this command to view information about one or more registered nodes.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

node_name

Specifies the name of the client node to be queried. You can use wildcard characters to specify this name. All matching client nodes are queried. If you do not specify a value for this parameter, all client nodes are queried. The parameter is optional.

D0main

Specifies a list of policy domains that limit the client node query. Only nodes that are assigned to one of the specified policy domains are displayed. This parameter is optional. Separate the items in the list by commas, with no intervening spaces. You can use wildcard characters to specify a domain. All clients assigned to a matching domain are displayed. If you do not specify a value for this parameter, all policy domains are included in the query.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. You can specify one of the following values:

Standard

Specifies that partial information is displayed for the specified client nodes.

Detailed

Specifies that complete information is displayed for the specified client nodes.

Type

Specifies the type of node to include in the query results. The parameter is optional. The default value is CLIENT. You can specify one of the following values:

QUERY NODE

Any

Specifies any type of node.

Client

Specifies client nodes that are backup-archive clients, Tivoli Storage Manager for Space Management clients, or application clients.

NAS

Specifies NAS nodes.

Server

Specifies client nodes that are other Tivoli Storage Manager servers.

Authentication

Specifies the password authentication method for the node.

Local

Display those nodes authenticating to the Tivoli Storage Manager server.

Ldap

Display those nodes authenticating to an LDAP directory server. The node password is case-sensitive.

Example: Display information about registered client nodes

Display information about all registered client nodes.

```
query node
```

Node Name	Platform	Policy Domain Name	Days Since Last Access	Days Since Password Set	Locked?
CLIENT1	AIX	STANDARD	6	6	No
GEORGE	AIX	STANDARD	1	1	No
JANET	AIX	STANDARD	1	1	No
JARED	Linux86	STANDARD	1	1	No
JOE2	Mac	STANDARD	<1	<1	No
TOMC	WinNT	STANDARD	1	1	No

Example: Displayed detailed information about a client node

Display complete information about the client node named Joe.

```
query node joe format=detailed
```

```

Node Name: JOE
Platform: WinNT
Client OS Level: 4.00
Client Version: Version 5, Release 4, Level 0.0
Policy Domain Name: STANDARD
Last Access Date/Time: 09/24/2002 18:55:46
Days Since Last Access: 6

Password Set Date/Time: 09/24/2002 18:26:43
Days Since Password Set: 6

Invalid Sign-on Count: 0
Locked?: No
Contact:
Compression: Client
Archive Delete Allowed?: Yes
Backup Delete Allowed?: No
Registration Date/Time: 09/24/2002 18:26:43
Registering Administrator: SERVER_CONSOLE
Last Communication Method Used: Tcp/Ip
Bytes Received Last Session: 108,731
Bytes Sent Last Session: 698
Duration of Last Session (sec): 0.00
Pct. Idle Wait Last Session: 0.00
Pct. Comm. Wait Last Session: 0.00
Pct. Media Wait Last Session: 0.00
Optionset:
URL: http://joe.host.name:1581
Node Type: Client
Password Expiration Period: 60
Keep Mount Point?: No
Maximum Mount Points Allowed: 2
Auto Filespace Rename: No
Validate Protocol: No
TCP/IP Name:
TCP/IP Address: 9.11.153.39
Globally Unique ID: 11.9c.54.e0.8a.b5.11.d6.b3.c3.00.06.29.45.c1
Transaction Group Max: 0
Data Write Path: ANY
Data Read Path: ANY
Session Initiation: ClientOrServer
High-level Address:
Low-level Address: 1501
Collocation Group Name:
Proxynode Target:
Proxynode Agent:
Node Groups:
Email Address:
Deduplication: ServerOnly
Replication State: Enabled
Replication Mode: Send
Backup Default Replication Rule: DEFAULT
Archive Default Replication Rule: ALL_DATA
Space Managed Default Replication Rule: None
Users allowed to back up: ALL
Client OS Name: WIN: Windows XP
Client Processor Architecture: x86
Client Products Installed: WIN, FCM, VE
Client Target Version: Version 6, Release 2, Level 0.0
Authentication: Local
SSLRequired: No

```

Field descriptions

Node Name

The name of the client node.

Platform

The client node platform (operating system) as of the last time that the

QUERY NODE

client node contacted the server. A question mark (?) is displayed until the client node first accesses the server and reports its platform type.

Client OS Level

The level of the client operating system as of the last time that the client node contacted the server.

Client Version

The version of the client that is installed on the client node.

This field does not apply to NAS nodes.

Policy Domain Name

The assigned policy domain of the client node.

Last Access Date/Time

The last date and time that the client node accessed the server.

Days Since Last Access

The number of days that have elapsed since the last time that the client node accessed the server.

Password Set Date/Time

The date and time that the password was set for the client node.

Days Since Password Set

The number of days that have elapsed since the password was set for the client node.

Invalid Sign-on Count

The number of invalid sign-on attempts that have been made since the last successful sign-on. This count can be non-zero only when the invalid password limit (**SET INVALIDPWLIMIT**) is greater than zero. When the number of invalid attempts equals the limit set by the **SET INVALIDPWLIMIT** command, the node is locked out of the system.

Locked?

Whether the client node is locked out of Tivoli Storage Manager.

Contact

Any contact information for the client node.

Compression

Whether compression is enabled on the client node.

This field does not apply to NAS nodes.

Archive Delete Allowed?

Whether the client node can delete its own archive files.

Backup Delete Allowed?

Whether the client node can delete its own backup files.

Registration Date/Time

The date and time that the client node was registered.

Registering Administrator

The name of the administrator that registered the client node.

Last Communication Method Used

The communication method that was last used by the client node to contact the server.

Bytes Received Last Session

The number of bytes received by the server during the last client node session.

This field does not apply to NAS nodes.

Bytes Sent Last Session

The number of bytes sent to the client node.

This field does not apply to NAS nodes.

Duration of Last Session (sec)

How long the most recent client node session lasted.

This field does not apply to NAS nodes.

Pct. Idle Wait Last Session

The percentage of the total session time that the client was not performing any functions.

This field does not apply to NAS nodes.

Pct. Comm. Wait Last Session

The percentage of the total session time that the client waited for a communication response from the server.

This field does not apply to NAS nodes.

Pct. Media Wait Last Session

The percentage of the total session time that the client waited for a removable volume to be mounted.

This field does not apply to NAS nodes.

Optionset

The name of the client option set.

URL The web address of the client. The administrator can use this address to contact the client from the Administration Center.

Node Type

The type of client node. The value is one of the following:

- Client: a backup-archive client, a Tivoli Storage Manager for Space Management client, or an application client
- Server: a Tivoli Storage Manager server
- NAS: a NAS file server

Password Expiration Period

The password expiration period of the client node.

Keep Mount Point?

Whether the client node retains a mount point during a session.

Maximum Mount Points Allowed

The number of mount points that a client node can use on the server for Tivoli Storage Manager for Space Management migration and for backup and archive operations. This parameter does not apply to nodes with a type of NAS or SERVER. If a client node was registered to a server at Version 3.7 or later, the value is 0-999, depending on the value that is set with the MAXNUMMP parameter of the **REGISTER NODE** command. If the client node was registered under previous versions of the server and the MAXNUMMP parameter was not explicitly set by using the **UPDATE NODE** command, the value is set to NOLIMIT. The MAXNUMMP value is not evaluated or enforced during client data read operations such as restore,

retrieve, and Tivoli Storage Manager for Space Management recall. However, mount points in use for data read operations are evaluated against attempted concurrent data store operations for the same client node and might prevent the data store operations from being able to acquire mount points.

Auto Filespace Rename

Whether Tivoli Storage Manager prompts the client to rename file spaces when the client system upgrades to a client that supports Unicode. This field is valid only for client systems that use Windows, Macintosh OS X, or NetWare operating systems.

Validate Protocol

Whether the client has data validation enabled. If the client has data validation enabled, this field specifies whether Tivoli Storage Manager validates only the file data or all data, which includes file metadata. You can enable data validation by using the **REGISTER NODE** or **UPDATE NODE** commands.

TCP/IP Name

The host name of the client node as of the last time that the client node contacted the server. The field is blank if the client software does not support reporting this information to the server.

TCP/IP Address

The TCP/IP address of the client node as of the last time that the client node contacted the server. The field is blank if the client software does not support reporting this information to the server.

Globally Unique ID

The globally unique identifier (GUID) as of the last time that the client node contacted the server. This GUID identifies the host computer on which the node is located.

Transaction Group Max

Specifies the number of files per transaction committed that are transferred between a client and a server. Client performance might be improved by using a larger value for this option.

Data Write Path

Specifies the transfer path used when the client sends data to the server, storage agent, or both, during storage operations such as backup or archive. If a path is unavailable, the node cannot send any data.

Data transfer path options are ANY, LAN, or LAN-free.

Data Read Path

Specifies the transfer path used when the server, storage agent, or both, read data for a client, during operations such as restore or retrieve. If a path is unavailable, data cannot be read.

Data transfer path options are ANY, LAN, or LAN-free.

Session Initiation

Controls whether the server or client initiates sessions. The two options are:

- ClientOrServer
- Serveronly

High-level Address

Specifies the client IP address that the server contacts to initiate scheduled events when SESSIONINITIATION is set to SERVERONLY.

Low-level Address

Specifies the client port number on which the client listens for sessions from the server when SESSIONINITIATION is set to SERVERONLY.

Collocation Group Name

Specifies the name of the collocation group to which a node belongs. If a node does not belong to a collocation group, this field is blank.

Proxynode Target

Specifies which nodes are proxy nodes (agents) for other nodes, in a space-separated list. If there are no nodes in that type of association, this field is blank.

Proxynode Agent

Specifies the originating (target) node name for a proxy node session, in a space separated list. If there are no nodes in that type of association, this field is blank.

Node Groups

Specifies the name of the node group to which a node belongs. If a node does not belong to a node group, this field is blank.

Email Address

Specifies the email address of the client node.

Deduplication

The location where data is deduplicated. The value ServerOnly specifies that data stored by this node can be deduplicated on the server only. The value Clientorserver specifies that data stored by this node can be deduplicated on either the client or the server.

Replication State

Indicates whether the node is enabled for replication. The following values are possible:

Enabled

The node is configured for replication and ready to replicate.

Disabled

The node is configured for replication but is not ready to replicate.

None The node is not configured for replication.

Replication Mode

Indicates whether the node is configured as the source of or target for replicated data. If this field is blank, the node is not configured for replication. The following values are possible:

Send The node is configured as the source of data for replication.

Receive

The node is configured as the target of data for replication.

SyncSend

The data that belongs to the node is to be synchronized with the node data that is on the target replication server. Synchronization applies only to nodes whose data was imported from a source replication server and imported to the target replication server. Synchronization occurs during replication.

SyncReceive

The data that belongs to the node is to be synchronized with the node data that is on the source replication server. Synchronization applies only to nodes whose data was imported from a source replication server and imported to the target replication server. Synchronization occurs during replication.

None The node is not configured for replication.

Backup Default Replication Rule

Archive Default Replication Rule

Space Managed Default Replication Rule

The replication rule that applies to backup, archive, and space-managed data that belongs to the node. The following values are possible:

ALL_DATA

Replicates backup, archive, or space-managed data. The data is replicated with normal priority.

ACTIVE_DATA

Replicates active backup data. The data is replicated with normal priority.

ALL_DATA_HIGH_PRIORITY

Replicates backup, archive, or space-managed data. The data is replicated with high priority.

ACTIVE_DATA_HIGH_PRIORITY

Replicates active backup data. The data is replicated with high priority.

DEFAULT

Replicates backup, archive, or space-managed data according to the server rule for the data type.

NONE

No data is replicated. For example, if the replication rule for archive data is NONE, archive data that belongs to the node is not replicated.

Users allowed to back up

Specifies whether a non-root user ID or only a root user ID can back up files to the server. ALL indicates all users, while ROOT indicates that just the root user ID can back up files to the server. This output is not available if the client node operating system is considered a single-user operating system.

Client OS Name

The operating system of the client. The client deployment wizard uses this information when deploying a package to the client. This field is reported only for Tivoli Storage Manager clients at version 6.2.0.0 and later.

Client Processor Architecture

The client architecture. The client deployment wizard uses this value to determine which package to deploy when updating the client. This field is reported only for Tivoli Storage Manager clients at version 6.2.0.0 and later.

Client Products Installed

The products that are on the node. The products that might be listed are:

- BA (Backup-Archive Client)
- VE (Virtual Environments)

- FCM (FlashCopy Manager)

Client Target Version

The version of the client that is installed at a time scheduled through the **DEFINE SCHEDULE** or **UPDATE SCHEDULE** command. This field is reported only for Tivoli Storage Manager clients at version 6.2.0.0 and later.

Authentication

Specifies the password authentication method: LOCAL, LDAP, or LDAP (pending).

Authentication Target	Authentication Method
Tivoli Storage Manager server	LOCAL
LDAP directory server	LDAP
This node is configured to authenticate with an LDAP directory server, but the node did not yet authenticate.	LDAP (pending)

SSL Required

Specifies if the security setting for the node requires Secure Sockets Layer (SSL). Values can be YES, NO, or Default. You must have system level authority to update the node **SSLREQUIRED** setting.

Example: Display information about node roles

The example output is only a portion of the full display.

query node alvin f=d

```

Proxynode Agent:
  Node Groups:
  Email Address:
  Deduplication: ServerOnly
Users allowed to back up: All
  Role: Server
  Role Override: UseReported
Processor Vendor: ORACLE
Processor Brand: UltraSPARC-T2
Processor Type: 4
Processor Model:
Processor Count: 1
Hypervisor:
API Application: NO
  Scan Error: NO
MAC Address:
    
```

Field Descriptions

Role The processor role as reported by the client.

Role Override

The override value for Role, specified with the **UPDATE NODE** command.

Processor Vendor

The processor vendor as reported by the client.

Processor Brand

The processor brand as reported by the client.

Processor Type

The processor type as reported by the client. This numeric value specifies the number of processor cores used for PVU calculation.

QUERY NODE

Processor Model

The processor model as reported by the client.

Processor Count

The processor count as reported by the client.

Hypervisor

The hypervisor as reported by the client.

API Application

The client indicator that this is an API application.

Scan Error

The indicator of whether the latest scan for processor information might be failing and needs investigation.

MAC Address

MAC Address as reported by the client.

Example: View all nodes that authenticate to the Tivoli Storage Manager server

If you want to view all nodes that authenticate locally, specify the following command:

```
query node * authentication=local
```

Node Name	Platform	Policy Domain Name	Days Since Last Access	Days Since Password Set	Locked?
NODE1	WinNT	STANDARD	3	3	No
LOCAL	(?)	STANDARD	7	7	No

Related commands

Table 247. Commands related to QUERY NODE

Command	Description
LOCK NODE	Prevents a client from accessing the server.
QUERY ADMIN	Displays information about one or more Tivoli Storage Manager administrators.
QUERY REPLNODE	Displays information about the replication status of a client node.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
REMOVE REPLNODE	Removes a node from replication.
RENAME NODE	Changes the name for a client node.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
RESET PASSEXP	Resets the password expiration for nodes or administrators.
SET INVALIDPWLIMIT	Sets the number of invalid logon attempts before a node is locked.

Table 247. Commands related to QUERY NODE (continued)

Command	Description
SET MINPWLENGTH	Sets the minimum length for client passwords.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.
UNLOCK NODE	Enables a locked user in a specific policy domain to access the server.
UPDATE NODE	Changes the attributes associated with a client node.


QUERY NODEDATA (Query client data in volumes)

Use this command to display information about the data for one or more nodes in a sequential access storage pool. **QUERY NODEDATA** displays the name of the volume on which a node's data is written and the amount of space occupied by the data on that volume. This information is useful when determining how to group nodes into collocated storage pools.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

node_name

Specifies the name of the client node for which you want to locate data. You can specify one or more names. If you specify multiple names, separate the names with commas; do not use intervening spaces. You can also use wildcard characters when specifying multiple names. You must specify either a node name or collocation group name, but not both.

COLLOCGroup

Specifies the name of the collocation group for which you want to locate data. You must specify either a node name or collocation group name, but not both.

Important: If the amount of space needed to complete the query about a collocation group exceeds the SQL buffer limit, the **QUERY NODEDATA** command can fail. If the command fails for this reason, issue the **QUERY COLLOGROUP** command to display a list of nodes in the group. Then issue the **QUERY NODEDATA** command for each node in the group.

STGpool

Specifies the name of the sequential storage pool to query. This parameter is optional. You can use wildcard characters to specify the names. If a wildcard matches the name of a disk storage pool, the name of the disk storage pool will be ignored. If you do not specify a value for this parameter, all sequential storage pools are queried.

VOLume

Specifies the volume that contains the data. This parameter is optional. You can use wildcard characters to specify multiple names. If you do not specify a value for this parameter, all volumes in the storage pool are queried.

Example: Use wildcards to display node data for a sequential access storage pool

Display information about where node data is stored in a sequential storage pool. Use a wildcard character to indicate node names. See “Field descriptions” for field descriptions.

query nodedata e*

Node Name	Volume Name	Storage Pool Name	Physical Space Occupied (MB)
EDU_J2	E:\tsm\server\00000117.BFS	EDU512	0.01
EDU_J2	E:\tsm\server\00000122.BFS	EDU319	0.01
EDU_J3	E:\tsm\server\00000116.BFS	EDU512	0.01
EDU_J3	E:\tsm\server\00000120.BFS	EDU319	0.01
EDU_J7	E:\tsm\server\00000118.BFS	EDU512	0.04
EDU_J7	E:\tsm\server\00000123.BFS	EDU319	0.04
EDU_JJ1	E:\tsm\server\00000116.BFS	EDU512	0.01
EDU_JJ1	E:\tsm\server\00000121.BFS	EDU512	0.01

Example: Display node data information for a specific collocation group

Display information about the location of node data in a sequential storage pool for a particular collocation group. In this example, nodes EDU_J3 and EDU_JJ1 are the only members that belong to collocation group, grp1, and have data in a sequential access storage pool. See “Field descriptions” for field descriptions.

query nodedata collocgroup=grp1

Node Name	Volume Name	Storage Pool Name	Physical Space Occupied (MB)
EDU_J3	E:\tsm\server\00000116.BFS	EDU512	0.01
EDU_J3	E:\tsm\server\00000120.BFS	EDU319	0.01
EDU_JJ1	E:\tsm\server\00000116.BFS	EDU512	0.01
EDU_JJ1	E:\tsm\server\00000121.BFS	EDU512	0.01

Field descriptions

Node Name

Specifies the name of the node.

Volume Name

Specifies the name of the volume that contains the node data.

Storage Pool Name

Specifies the name of the storage pool in which the volume is located.

Physical Space Occupied (MB)

Specifies the amount of physical space occupied by the node's data.

Physical space includes empty space within aggregates, from which files may have been deleted or expired.

QUERY NODEDATA

Related commands

Table 248. Commands related to QUERY NODEDATA

Command	Description
DEFINE COLLOGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node to a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node from a collocation group.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY COLLOGROUP	Displays information about collocation groups.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY STGPOOL	Displays information about storage pools.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
UPDATE COLLOGROUP	Updates the description of a collocation group.
UPDATE STGPOOL	Changes the attributes of a storage pool.


QUERY NODEGROUP (Query a node group)

Use this command to display the node groups defined on the server.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

group_name

Specifies the name of the node group to display. To specify multiple names, use a wildcard character. This parameter is optional. The default is to display all node groups.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed. To display the members of the node group, you must specify `FORMAT=DETAILED`.

Example: List node groups on the server

Display the node groups defined on the server. See “Field descriptions” on page 824 for field descriptions.

```
query nodegroup
```

Node Group Name	Node Group Description
DEPT_ED	Education department
GROUP1	Low cap client nodes.

Example: Display detailed node group information

Display complete information about all node groups and determine which client nodes belong to which node groups. See “Field descriptions” on page 824 for field descriptions.

```
query nodegroup format=detailed
```

QUERY NODEGROUP

```
Node Group Name: DEPT_ED
Node Group Description: Education department
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 04/21/2006 10:59:03
Node Group Member(s): EDU_1 EDU_7

Node Group Name: GROUP1
Node Group Description: Low cap client nodes.
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 04/21/2006 10:59:16
Node Group Member(s): CHESTER REX NOAH JARED
```

Field descriptions

Node Group Name

The name of the node group.

Node Group Description

The description for the node group.

Last Update by (administrator)

The name of the administrator that defined or most recently updated the node group.

Last Update Date/Time

The date and time that an administrator defined or most recently updated the node group.

Node Group Member(s)

The members of the node group.

Related commands

Table 249. Commands related to **QUERY NODEGROUP**

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE BACKUPSET	Deletes a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
QUERY BACKUPSET	Displays backup sets.
UPDATE BACKUPSET	Updates a retention value associated with a backup set.
UPDATE NODEGROUP	Updates the description of a node group.


QUERY OCCUPANCY (Query client file spaces in storage pools)

Use this command to show where client file spaces are stored and how much space they occupy.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

node_name

Specifies the node that owns the file spaces that you want to locate. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a value for this parameter, all nodes are queried.

file_space_name

Specifies the file space that you want to locate. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a value for this parameter, all file spaces are queried. You must specify a node name if you specify a file space name.

For a server that has clients with Unicode support, you might need to have the server convert the file space name that you enter. For example, you might need to have the server convert the name that you enter from the server's code page to Unicode. See the NAMETYPE parameter for details. If you do not specify a file space name or specify only a single wildcard character for the name, you can use the CODETYPE parameter to limit the operation to Unicode file spaces or non-Unicode file spaces.

STGpool

Specifies the storage pool to query for files from the specified file space. This

QUERY OCCUPANCY

parameter is optional. You can use wildcard characters to specify names. If you do not specify a value for this parameter, all storage pools are queried.

DEVclass

Specifies the device class that is associated with the devices where the file spaces are stored. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a value for this parameter, storage pools that are associated with any device class are queried.

Type

Specifies the types of files to query in the file spaces. This parameter is optional. The default value is ANY. Possible values are:

ANY

Specifies that all types of files are queried: back up versions of files, archived copies of files, and files that are migrated from Tivoli Storage Manager for Space Management clients.

Backup

Specifies that backup files are queried.

Archive

Specifies that archive files are queried.

SPacem

Specifies that space-managed files (files that were migrated by a Tivoli Storage Manager for Space Management client) are queried.

NAMETYPE

Specifies how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with Unicode support. A backup-archive client with Unicode support is currently available only for Windows, Macintosh OS 9, Macintosh OS X, and NetWare. Use this parameter only when you specify a partly or fully qualified file space name.

The default value is SERVER. Possible values are:

SERVER

The server uses the server's code page to interpret the file space names.

UNICODE

The server converts the file space names from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the names and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server has a problem accessing system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

CODETYPE

Specifies how you want the server to interpret the file space names that you enter. Use this parameter only when you enter a single wildcard character for the file space name or when you do not specify any file space name.

The default value is BOTH, which means that the file spaces are included regardless of code page type. Possible values are:

UNICODE

Include only file spaces that are Unicode enabled.

NONUNICODE

Include only file spaces that are not Unicode enabled.

BOTH

Include file spaces regardless of code page type.

Example: Display file spaces assigned to a specific node

Display information about where all file spaces assigned to the node named DAISY are stored. See “Field descriptions” for field descriptions.

query occupancy daisy

Node Name	Type	Filespace Name	FSID	Storage Pool Name	Number of Files	Physical Space Occupied (MB)	Logical Space Occupied (MB)
DAISY	Bkup	DRIVED	1	COPYFILE	38	0.45	0.42

Example: Display file spaces assigned to a specific node with a backup file type

Display information about the file spaces that belong to the node WAYNE, and that have a backup file type. See “Field descriptions” for field descriptions.

query occupancy wayne type=backup

Node Name	Type	Filespace Name	FSID	Storage Pool Name	Number of Files	Physical Space Occupied (MB)	Logical Space Occupied (MB)
WAYNE	Bkup	DWG1	1	BACKUPPOOL1	2,330	53.19	50.01
WAYNE	Bkup	OS2C	2	BACKUPPOOL1	1,554	32.00	31.30

Field descriptions

Node Name

The node that owns the file space.

Type The type of data. Possible values are:

Arch Data that has been archived.

Bkup Data that has been backed up.

Spmg Data that has been migrated from a Tivoli Storage Manager for Space Management client.

Filespace Name

The name of the file space that belongs to the node.

File space names can be in a different code page or locale than the server. If they are, the names in the Administration Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, Tivoli Storage Manager converts the name to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server

QUERY OCCUPANCY

cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

Storage Pool Name

The storage pool where the file space is currently located.

Number of Files

The number of logical files that belong to the file space and are stored in this storage pool.

Physical Space Occupied (MB)

The amount of physical space that is occupied by the file space. Physical space includes empty space within aggregates, from which files might have been deleted or expired. For this value, one MB = 1048576 bytes.

Tip: This field does not display a value for storage pools that are set up for data deduplication. If you turn off data deduplication for a storage pool, a value for physical occupancy is not displayed until the storage pool is empty of deduplicated files.

Logical Space Occupied (MB)

The amount of space that is occupied by logical files in the file space. Logical space is the space that is actually used to store files, excluding empty space within aggregates. For this value, one MB = 1048576 bytes.

FSID The file space ID (FSID) for the file space. The server assigns a unique FSID when a file space is first stored on the server.

Related commands

Table 250. Commands related to QUERY OCCUPANCY

Command	Description
DELETE FILESPACE	Deletes data associated with clients file spaces.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.

QUERY OPTION (Query server options)

Use this command to display information about server options.

Change server options by editing the server options file or by issuing the **SETOPT** command. When you edit the server options file, you must restart the server before any changes take effect. Any changes you make by issuing the **SETOPT** command take effect immediately.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

optionname

Specifies the name of an option in the server options file. This parameter is optional. You can use wildcard characters to specify this name. All matching server options display. If you do not specify this parameter, information on all options displays.

Example: Display all server options

Display general information about all server options. The output lists all options with their specified values.

```
query option
```

Example: Display options settings using a wildcard character

View the option settings for all options that begin with L.

```
query option l*
```

Server Option	Option Setting
-----	-----
Language	AMENG

Example: Display LDAP directory servers

View the settings for all LDAP directory servers.

```
query option ldapurl
```

Server Option	Option Setting
-----	-----
LDAP URL	ldap://tophoy.tucson.com/cn=tsmdata
LDAP URL	ldap://krypton.ibm.com/ou=tsmdata,dc=ibm,dc=com

QUERY OPTION

Related commands

Table 251. Commands related to QUERY OPTION

Command	Description
SETOPT	Updates a server option without stopping and restarting the server.


QUERY PATH (Display a path definition)

Use this command to display the path between a source and a destination.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

source_name

Specifies the name of a source for which to display paths. This parameter is optional. You can specify wildcard characters. The default is to display paths for all sources.

A source is a data mover, a server, or a storage agent.

destination_name

Specifies the name of a destination for which to display paths. This parameter is optional. You can specify wildcard characters. The default is to display paths for all destinations.

SRCType

Specifies the type of the source. This parameter is optional. The default is to display paths for all source types. Possible values are:

ANY

Specifies to display paths with any source type.

DATAMover

Specifies to only display paths with the DATAMOVER source type.

SERVer

Specifies to only display paths with the SERVER source type. (A source that has a source type of SERVER is a storage agent.)

QUERY PATH

DESTType

Specifies the type of the destination. This parameter is optional. The default is to display paths for all destination types. Possible values are:

ANY

Specifies to display paths with any destination type.

DRive

Specifies to display only paths with the DRIVE destination type. When the destination type is a drive, you must specify the library name. You can refine which paths are displayed by entering a name in the LIBRARY parameter.

LIBRARY

Specifies that only paths with destination type LIBRARY display.

LIBRARY

Specifies the name of the library to which the drive belongs. This parameter is required when the destination type is a drive (DESTTYPE=DRIVE).

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: Display summary path information

Display information about paths for the source NETAPP1. See “Field descriptions” on page 833 for field descriptions.

```
query path netapp1
```

Source Name	Source Type	Destination Name	Destination Type	Online
NETAPP1	DATAMOVER	DRIVE1	DRIVE	Yes
NETAPP1	DATAMOVER	NASLIB	LIBRARY	Yes

Example: Display detailed path information

Display detailed information about paths for the source NETAPP1. See “Field descriptions” on page 833 for field descriptions.

```
query path netapp1 format=detailed
```

```

Source Name: NETAPP1
Source Type: DATAMOVER
Destination Name: NASLIB
Destination Type: LIBRARY
Library:
Device: /dev/tsm SCSI/mc0
Directory:
On-Line: Yes
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 06/21/2002 20:52:56

Source Name: NETAPP1
Source Type: DATAMOVER
Destination Name: DRIVE1
Destination Type: DRIVE
Library: NASLIB
Device: rst01
Directory:
On-Line: Yes
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 06/21/2002 20:55:23

```

Example: Display detailed path information for a z/OS media server

Display detailed information about a z/OS media server path. See “Field descriptions” for field descriptions.

query path format=detailed

```

Source Name: SERVER1
Source Type: SERVER
Destination Name: ZOSMEDIA
Destination Type: LIBRARY
Library:
Node Name:
Device:
External Manager:
ZOS Media Server: MEDSERV1
Comm Method:
LUN:
Initiator: 0
Directory:
On-Line: Yes
Last Update by (administrator): ADMIN
Last Update Date/Time: 06/08/2011 15:33:39

```

Field descriptions

Source Name
The name of the source.

Destination Name
The name of the destination.

Source Type
The type of the source.

Destination Type
The type of the destination.

Library
The name of the library that contains the drive that is the destination. This field will be blank if the destination type is library. The library name is in destination name field when the destination is a library.

QUERY PATH

Node Name

The name of the device that is the destination.

Device

The name of the device that is the destination.

Directory

Specifies the directory location of a file on the source.

LUN Specifies the logical unit name through which the disk can be accessed by the source.

Online

Whether the path is online and available for use.

Last Update by (administrator)

The ID of the administrator who performed the last update.

Last Update Date/Time

The date and time when the last update occurred.

Related commands

Table 252. Commands related to QUERY PATH

Command	Description
DEFINE PATH	Defines a path from a source to a destination.
DELETE PATH	Deletes a path from a source to a destination.
UPDATE PATH	Changes the attributes associated with a path.


QUERY POLICYSET (Query a policy set)

Use this command to display information about one or more policy sets.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

domain_name

Specifies the policy domain associated with the policy set to query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a value for this parameter, all policy domains are queried. You must specify this parameter when querying an explicitly named policy set.

policy_set_name

Specifies the policy set to query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify either ACTIVE or a policy set name, all policy sets are queried.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: List policy sets for all policy domains

Query all policy sets for all policy domains. Create the output in standard format. See “Field descriptions” on page 836 for field descriptions.

```
query policyset
```

QUERY POLICYSET

Policy Domain Name	Policy Set Name	Default Mgmt Class Name	Description
EMPLOYEE-RECORDS	ACTIVE	ACTIVEFILES	Personnel Department
EMPLOYEE-RECORDS	HOLIDAY	ACTIVEFILES	Personnel Department
EMPLOYEE-RECORDS	VACATION	ACTIVEFILES	Personnel Department
PROG1	SUMMER		Programming Group Policies
PROG2	SUMMER		Programming Group Policies
STANDARD	ACTIVE	STANDARD	Installed default policy set.
STANDARD	STANDARD	STANDARD	Installed default policy set.

Example: Displayed detailed information about a specific policy set

Query the VACATION policy set that is in the EMPLOYEE_RECORDS policy domain. Create the output in detailed format. See "Field descriptions" for field descriptions.

```
query policyset employee_records vacation
format=detailed
```

```
Policy Domain Name: EMPLOYEE_RECORDS
Policy Set Name: VACATION
Default Mgmt Class Name: ACTIVEFILES
Description: Personnel Department
Last Update by (administrator): $$CONFIG_MANAGER$$
Last Update Date/Time: 05/31/1998 13:15:50
Managing profile: ADSM_INFO
Changes Pending: Yes
```

Field descriptions

Policy Domain Name

The name of the policy domain.

Policy Set Name

The name of the policy set.

Default Mgmt Class Name

The management class assigned as the default for the policy set.

Description

The description of the policy set.

Last Update by (administrator)

The name of the administrator or server that most recently updated the policy set. If this field contains \$\$CONFIG_MANAGER\$\$, the policy set is associated with a domain that is managed by the configuration manager.

Last Update Date/Time

The date and time when the policy set was most recently defined or updated.

Managing Profile

The profile or profiles that manage the domain to which this policy set belongs.

Changes Pending

Whether or not changes are being made but not activated. Once the changes are activated, the field resets to No.

Related commands

Table 253. Commands related to QUERY POLICYSET

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
COPY POLICYSET	Creates a copy of a policy set.
DEFINE POLICYSET	Defines a policy set within the specified policy domain.
DELETE POLICYSET	Deletes a policy set, including its management classes and copy groups, from a policy domain.
QUERY DOMAIN	Displays information about policy domains.
UPDATE POLICYSET	Changes the description of a policy set.
VALIDATE POLICYSET	Verifies and reports on conditions the administrator must consider before activating the policy set.

QUERY PROCESS (Query one or more server processes)

Use this command to display information about active background processes.

To cancel background processes, issue the **CANCEL PROCESS** command. To display detailed information about node replication processes, issue the **QUERY REPLICATION** command.

Privilege class

Any administrator can issue this command.

Syntax

```

▶▶ Query Process process_number ▶▶
    
```

Parameters

process_number

Specifies the number of the background process to be queried. This parameter is optional. If not specified, information about all background processes is displayed.

Example: Query a single background process

Display information about background process 202. See “Field descriptions” on page 839 for field descriptions.

```
query process 202
```

Process Number	Process Description	Status
202	EXPORT SERVER	ANR0NNNI EXPORT Identifier MYEXPORTSERVER ANR0648I Have copied the following: 8 Domains 2 Policy Sets 10 Management Classes 4 Copy Groups 1 Administrators 746 Bytes (0 errors have been detected) Current input volume(s): C:\BUILD\540\GA\BUILD\NT\I386\DEBUG\ -00000014.BFS,(6 Seconds)

Example: Query all background processes

Display information about all background processes. See “Field descriptions” on page 839 for field descriptions.

```
query process
```

Process Number	Process Description	Status
304	IDENTIFY DUPLICATES	Storage Pool FILEPOOL, Volume /tspool2/00006664. BFS, Files Processed: 2000, Duplicate Extents Found: 344, Duplicate Bytes Found: 3,238,123, Current Physical File (bytes): 2,626,676,296. Status: Processing
284	IDENTIFY DUPLICATES	Storage Pool FILEPOOL, Volume /tspool2/00006666. BFS, Files Processed: 2000, Duplicate Extents Found: 344, Duplicate Bytes Found: 3,238,123, Current Physical File (bytes): None. Status: Idle
4	Replicate Node	Replicating Node(s) IRONMAN. File spaces complete: 0. File spaces identifying and replicating: 1. File spaces replicating: 0. File spaces not started: 3. Files current: 11,920. Files replicated: 0 of 0. Files updated: 0 of 0. Files deleted: 0 of 0. Amount Replicated: 11,482 KB of 11,482 KB. Amount transferred: 11,482 KB. Elapsed time: 0 Day(s), 0 Hour(s), 1 Minute(s).
37	Expiration	Processed 12 nodes out of 30 total nodes, examined 411 objects, deleting 411 backup objects, 0 archive objects, 0 DB backup volumes, 0 recovery plan files; 0 objects have been retried and 0 errors encountered.

Field descriptions

Process Number

Specifies the number assigned to the active background process.

Process Description

Specifies a description of the active background process.

Status Specifies the status of the active background process.

Related commands

Table 254. Command related to QUERY PROCESS

Command	Description
CANCEL EXPORT	Deletes a suspended export operation.
CANCEL PROCESS	Cancels a background server process.
IDENTIFY DUPLICATES	Identifies duplicate data in a storage pool.
QUERY EXPORT	Displays the export operations that are currently running or suspended.

QUERY PROCESS

Table 254. Command related to QUERY PROCESS (continued)

Command	Description
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLNODE	Displays information about the replication status of a client node.
RESTART EXPORT	Restarts a suspended export operation.
SUSPEND EXPORT	Suspends a running export operation.

QUERY PROFILE (Query a profile)


Use this command to display information about profiles and associated objects. Issue this command from a configuration manager or from a managed server. You can use this command to get profile information from any configuration manager defined to the server, even if the server does not subscribe to any profile.

If you query a locked profile from the configuration manager to which the profile belongs, complete profile information is displayed. If you query a locked profile from another server, the query displays only that the profile is locked.

Privilege class

Any administrator can issue this command.

Syntax



Notes:

- 1 The server name you specify depends on the server from which you issue the command. See the description of the SERVER parameter.

Parameters

profile_name

Specifies the profile to display. To specify multiple names, use a wildcard character. This parameter is optional. The default is to display all profiles.

SERVER

Specifies the configuration manager whose profile information is displayed. The requirements for the name depends on where the query is issued:

- From a configuration manager: This parameter is optional. The default is the configuration manager's name.
- From a managed server: This parameter is optional. The default is the name of the configuration manager for this managed server.
- From a server that is neither a configuration manager nor a managed server: You must specify a name.

Format

Specifies whether partial or detailed information is displayed. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that detailed information is displayed.

QUERY PROFILE

USELocal

When you perform the query from a managed server, this parameter specifies whether the profile information is obtained from the configuration manager or the managed server. If the profile information does not exist on the managed server, the information is obtained from the configuration manager, regardless of the value of this parameter.

If you use this parameter on a server that is not managed by the configuration manager that owns the profile, the parameter is ignored. The default value is YES. Possible values are:

Yes

Specifies that the profile information, if available, is obtained from the managed server. The configuration manager is contacted if information is not available from the managed server.

No Specifies that the profile information is obtained from the configuration manager even if the information is available from the managed server. This ensures that you receive current information about the profile.

Example: List profiles from a configuration manager

Display profile information from a configuration manager. See “Field descriptions” for field descriptions.

```
query profile
```

Configuration manager	Profile name	Locked?
SERVER1	DEFAULT_PROFILE	No
SERVER1	ADMIN_INFO	No
SERVER1	EMPLOYEE	No
SERVER1	PERSONNEL	Yes

Example: Display detailed profile information for a managed server

From a managed server, display current detailed information for profile ADMIN_INFO. See “Field descriptions” for field descriptions.

Note: When the profile is locked, most fields are not displayed.

```
query profile admin_info  
format=detailed useLocal=no
```

Configuration manager: SERVER1
Profile name: ADMIN_INFO
Locked: No
Description: Distributed administrative schedules
Server administrators: DENNIS EMILY ANDREA
Policy domains: ADMIN RECORDS
Administrative command schedules: ** all objects **
Server Command Scripts:
Client Option Sets:
Servers:
Server Groups:

Field descriptions

Configuration manager

The name of the configuration manager that owns the profile.

Profile name

The name of the profile.

Locked?

Whether the profile is locked.

Description

The description of the profile.

Server administrators

The administrators that are associated with the profile.

Policy domains

The policy domains that are associated with the profile.

Administrative command schedules

The administrative schedules that are associated with the profile.

Server Command Scripts

The server command scripts that are associated with the profile.

Client Option Sets

The client option sets that are associated with the profile.

Servers

The servers that are associated with the profile.

Server Groups

The names of server groups that are associated with the profile.

Related commands

Table 255. Commands related to QUERY PROFILE

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE PROFILE	Deletes a profile from a configuration manager.
LOCK PROFILE	Prevents distribution of a configuration profile.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UNLOCK PROFILE	Enables a locked profile to be distributed to managed servers.
UPDATE PROFILE	Changes the description of a profile.


QUERY PROXYNODE (Query proxy authority for a client node)

Use this command to display client nodes with authority to act as proxy to other client nodes in the Tivoli Storage Manager server.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

TArget

Specifies the name of the node targeted by the node with proxy authority. It is optional to specify a target node name. Wildcard names can be used to specify the target node name. A comma-separated list of node names is also allowed.

Example: List client nodes with proxy authority

To display all Tivoli Storage Manager client nodes with proxy authority to the target node named MYCLUSTER, issue the following command.

```
query proxynode target=mycluster
```

Target Node	Agent Node
FRED	MOE MINIE MICKEY
ALPHA	BETA GAMMA DELTA

Related commands

Table 256. Commands related to QUERY PROXYNODE

Command	Description
GRANT PROXYNODE	Grant proxy authority to an agent node.
REVOKE PROXYNODE	Revoke proxy authority from an agent node.

QUERY PVUESTIMATE (Display processor value unit estimate)

Use this command to obtain an estimate of the client devices and server devices that are being managed by the Tivoli Storage Manager server. In addition, this command provides an estimate of the processor value unit (PVU) totals for the server devices.

This command generates a PVU estimate that is based on the number of logical nodes that are defined to the Tivoli Storage Manager server. By contrast, the calculation of license obligations is based on the number of physical computers. There might not be a one-to-one correlation between the number of logical nodes and the number of physical computers. The report generated by the **QUERY PVUESTIMATE** command is an estimate, which is not legally binding.


For purposes of the **QUERY PVUESTIMATE** command, nodes on Microsoft Windows 7, Microsoft Windows XP Professional, and Apple systems are assumed to be client devices. Nodes on all other platforms are considered to be server devices. The server on which Tivoli Storage Manager is running is also classified as a server device. However, you can reclassify server devices as client devices if required. If your system includes retired workstations, test workstations, or others that can be ignored for purposes of PVU calculation, you can specify them as type other. To change a node classification, use the **UPDATE NODE** command or the **REGISTER NODE** command.

Note: The PVU information reported by Tivoli Storage Manager is not considered an acceptable substitute for the IBM License Metric Tool.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

Format

Specifies the output format. This parameter is optional. The default is Standard. The following values can be used:

Standard

Specifies standard output.

Detailed

Specifies detailed output.

Example: Display the estimated number of devices and PVU

Display the estimated number of client devices and server devices, and the estimated PVU for the server devices, for a Tivoli Storage Manager server. Issue the following command:

```
query pvuestimate
```

QUERY PVUESTIMATE

Table 257. Sample output for several products managed by one Tivoli Storage Manager server

Product	Number of Client Devices	Number of Server Devices	PVU of Server Devices
Tivoli Storage Manager Extended Edition	1,000	905	90,500
Tivoli Storage Manager for Storage Area Networks	50	10	1,000
Tivoli Storage Manager for Space Management	0	0	0
Tivoli Storage Manager for Mail	0	25	5,000
Tivoli Storage Manager for Databases	0	1,025	20,500
Tivoli Storage Manager for Enterprise Resource Planning	0	25	5,000
Tivoli Storage Manager for System Backup and Recovery	0	0	0

Other Node Classifications	Number
Nodes earlier than Version 6.3 with no PVU information available at this time	10
Nodes classified by the administrator as other	8
Nodes defined as a non-licensed API application	6

The following list provides details about the example fields:

Product

The Tivoli Storage Manager product name.

Number of Client Devices

The estimated number of client devices that are managed by the product. By default, only nodes on Microsoft Windows 7, Microsoft Windows XP Professional, and Apple systems are assumed to be client devices.

Number of Server Devices

The estimated number of server devices that are managed by the product. By default, nodes on all platforms except for Microsoft Windows 7, Microsoft Windows XP Professional, and Apple systems are assumed to be server devices. This number also includes the server on which Tivoli Storage Manager is running.

PVU of Server Devices

The estimated PVUs of all nodes that are connected as server devices.

Nodes earlier than Version 6.3 with no PVU

Devices that do not report processor information to the server.

Nodes classified by the administrator as other device

Nodes that are excluded from PVU counting by the administrator by using the **update node roleoverride=other** command.

Nodes defined as a non-licensed API application

Nodes such as DB2[®] backup or custom API applications.

Example: Display detailed node information

Display information for individual nodes by specifying the detailed (d) value for the **Format** parameter. Issue the following command:

```
tsm: PATMOS_630> query pvueestimate f=d
```

Table 258. Node classifications for specific products

Product	Number of Client Devices	Number of Server Devices	PVU of Server Devices
Tivoli Storage Manager Extended Edition	1,000	905	90,500
- banode1	1		
- banode2		1	200
- banode3	1		
- banode3		1	100
Tivoli Storage Manager for Storage Area Networks	50	10	1,000
- stagent1		1	50
- stagent2		1	100
Tivoli Storage Manager for Space Management	0	0	0
Tivoli Storage Manager for Mail	0	25	5,000
- mailnode1		1	200
- mailnode2		1	100
Tivoli Storage Manager for Databases	0	1,025	20,500
- dbnode1		1	200
- dbnode2		1	100
Tivoli Storage Manager for Enterprise Resource Planning	0	25	5,000
- erpnode1		1	50
- erpnode2		1	100
Tivoli Storage Manager for System Backup and Recovery	0	0	0
Other Node Classifications			Number
Nodes earlier than Version 6.3 with no PVU information available at this time			10
- oldnode1			1
- oldnode2			1
- mailnote44			1
- erpnode66			1
Nodes at Version 6.3 or later with no PVU match			10
- badcitnode1			1
- badcitnode2			1
- mailnode23			1
- erpnode34			1

QUERY PVUESTIMATE

Other Node Classifications	Number
Nodes classified by administrator as "other-device"	8
- overriddennode1	1
- overriddennode2	1
- mailnode77	
Nodes defined as a non-licensed API application	6
- vendorapinode1	1
- vendorapinode2	1

Related commands

Table 259. Commands related to QUERY PVUESTIMATE

Command	Description
AUDIT LICENSES	Verifies compliance with defined licenses.
QUERY LICENSE	Displays information about licenses and audits.
QUERY NODE	Displays partial or complete information about one or more clients.
REGISTER LICENSE	Registers a license with the Tivoli Storage Manager server.
REGISTER NODE	Defines a client node to the server and sets options for that user.
SET CPUINFOREFRESH	Specifies the number of days between client scans for workstation information used for PVU estimates.
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.
UPDATE NODE	Changes the attributes associated with a client node.

QUERY RECOVERYMEDIA (Query recovery media)


Use this command to display information about the media (for example, boot media) needed to recover a machine. Media are displayed in alphabetical order by name.

Remember: Tivoli Storage Manager does not use the information. It is available only to help you plan for the disaster recovery of client machines.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

media_name

Specifies the name of the recovery media. You can use wildcard characters to specify the name. This parameter is optional. The default is all recovery media.

Type

Specifies the type of media to be queried. This parameter is optional. If this parameter is not specified, all recovery media are queried. Possible values are:

B0ot

Only boot media are queried.

0ther

All media other than boot media are queried.

L0cation

Specifies the location of the recovery media to be queried. This parameter is optional. You can specify up to 255 characters. Enclose the description in quotation marks if it contains any blank characters.

Format

Specifies how the information is displayed. This parameter is optional. Possible values are:

Standard

Tivoli Storage Manager displays partial information. This is the default.

Detailed

Tivoli Storage Manager displays all information.

QUERY RECOVERYMEDIA

Example: Display summary information for a specific recovery media

Display information for the recovery media named RECMED1. See “Field descriptions” for field descriptions.

```
query recoverymedia RECMED1
```

Recovery Media Name	Volume Names	Location	Machine Name
RECMED1	vol1 vol2 vol3 vol4	IRONMOUNTAIN	MACH1

Example: Display detailed information for a specific recovery media

Display detailed information for the recovery media named RECMED1. See “Field descriptions” for field descriptions.

```
query recoverymedia RECMED1 format=detailed
```

```
Recovery Media Name: RECMED1
Type: Boot
Volume Names: vol1 vol2 vol3 vol4
Location: IRONMOUNTAIN
Description:
Product:
Product Information:
Machine Name: MACH1
```

Field descriptions

Recovery Media Name

The name of the recovery media.

Type Whether the recovery media are boot media or another type of media. Possible values are:

Boot The recovery media are boot media.

Other The recovery media are not boot media.

Volume Names

The set of volumes that contain the data needed to recover machines associated with this media.

Location

Where the recovery media is stored.

Description

A description of the recovery media.

Product

The product used to create the boot media.

Product Information

Information about the product that created the boot media. This information may be needed for restoring the machine.

Machine Name

The machines that are associated with this recovery media.

Related commands

Table 260. Commands related to QUERY RECOVERYMEDIA

Command	Description
DEFINE RECMEDMACHASSOCIATION	Associates recovery media with a machine.
DEFINE RECOVERYMEDIA	Defines the media required to recover a machine.
DELETE RECOVERYMEDIA	Deletes recovery media.
UPDATE RECOVERYMEDIA	Changes the attributes of recovery media.

QUERY REPLICATION (Query node replication processes)

Use this command to display information about running and ended node-replication processes.

Issue this command on the server that acts as a source for replicated data.

Important: You cannot display information about running replication processes for client nodes that are being converted from import and export operations to replication operations. The conversion process might run for a long time, but it occurs only once for a client node that is being converted.


By default, records about completed node-replication processes are retained for 30 calendar days. A *calendar day* consists of 24-hours, from midnight to midnight.

To display the retention period, issue the **QUERY STATUS** command. Check the value in the Replication Record Retention Period field. To change the retention period, issue the **SET REPRETENTION** command.

Privilege class

Any administrator can issue this command.

Syntax



Notes:

- 1 Do not mix FSIDs (file space identifiers) and file space names in the same command.
- 2 Do not specify FSID if you use wildcard characters for the client node name.

Parameters

node_name (Required)

Specifies the name of the client node to be queried. You can use wildcard characters when you specify this name, with one exception. If the value of the **NAMETYPE** parameter is FSID, do not specify wildcard characters for the client node name. The FSID value indicates the file space identifier. File spaces with identical names can have different identifiers in different client nodes.

file_space_name or *FSID*

Specifies the name of the file space or the file space identifier (FSID) to be queried. A name or FSID is optional. If you do not specify a name or an FSID, all file spaces are queried.

file_space_name

Specifies the name of the file space that has data to be queried. File space names are case sensitive. To determine the correct capitalization for the file space, issue the **QUERY FILESPACE** command. Separate multiple names with commas with no intervening spaces. When you specify a name, you can use wildcard characters.

A server that has clients with Unicode-enabled file spaces might have to convert the file space name. For example, the server might have to convert a name from the server code page to Unicode. For details, see the **NAMETYPE** parameter. If you do not specify a file space name, or if you specify only a single wildcard character for the name, you can use the **CODETYPE** parameter to limit the operation to Unicode file spaces or to non-Unicode file spaces.

FSID

Specifies the file space identifier for the file space to be queried. The server uses FSIDs to find the file spaces to replicate. To determine the FSID for a file space, issue the **QUERY FILESPACE** command. Separate multiple FSIDs with commas with no intervening spaces. If you specify an FSID, the value of the **NAMETYPE** parameter must be FSID.

NAMETYPE

Specifies how you want the server to interpret the file space names that you enter. You can use this parameter for Tivoli Storage Manager clients that are Unicode-enabled and that have Windows, Macintosh OS X, or NetWare operating systems.

Use this parameter only if you enter a partly-qualified or fully-qualified file space name. The default value is SERVER. You can specify one of the following values:

SERVER

The server uses the server code page to interpret file space names.

UNIcode

The server converts file space names from the server code page to the UTF-8 code page. The success of the conversion depends on the characters in the name and the server code page. Conversion can fail if the string includes characters that are not available in the server code page. Conversion can also fail if the server cannot access system conversion routines.

FSID

The server interprets file space names using their file space identifiers.

QUERY REPLICATION

CODEType

Specifies the type of file spaces to be included in the query. The default value is **BOTH**, which means that file spaces are included regardless of code page type. Use this parameter only if you enter a single wildcard character for the file space name. You can specify one of the following values:

UNICODE

Include only file spaces that are in Unicode.

NONUNICODE

Include only file spaces that are not in Unicode.

BOTH

Include all file spaces regardless of code page type.

DISplay

Specifies the number of days of node replication history to display. The default value is 1, which displays information about running node replication processes and about processes that completed during the current calendar day. The maximum value is 9999.

You can specify a number that is the same as or less than the number of days specified as the retention period for the replication history records. If you specify a value that is more than the value of the replication retention period or more than the number of days that replication records have been collected, the server displays only the number of replication history records that are available. For example, suppose that the replication retention period is 30 days and that the replication process has been running for only 10 days. If you specify **DISPLAY=20**, only 10 days of replication history are displayed.

PROcessid

Specifies the node replication history that is associated with a particular process identified by the process identifier. This parameter is optional. If you do not specify this parameter, all processes are displayed for the number of days specified by the **DISPLAY** parameter.

Restarting the server can cause the server to reuse process IDs. Reuse of process IDs can result in duplicate process IDs for separate processes.

STate

Specifies the state of the node replication processes to query. This parameter is optional. The default value is **ALL**. You can specify one of the following values:

ALL

Specifies all running and ended node-replication processes.

RUNning

Specifies all active node-replication processes that are identifying files or replicating files to the target replication server.

ENDED

Specifies all node replication processes that ended.

Format

Specifies how the information is displayed. This parameter is optional. The default value is **STANDARD**. You can specify one of the following values:

Standard

Specifies that partial information is displayed for node replication processes.

Detailed

Specifies that all available information for the node replication processes is displayed.

Example: Display information about replication processes for a file space

Display information about replication processes for a file space in client node PAYROLL. The file space identifier is 10.

query replication ironman

NodeName	Filespace Name	FSID	Start Time	End Time	Status	Phase
IRONMAN	/space	2	02/08/11 21:44:19	02/08/11 21:48:14	Ended	None

query replication ironman format=detailed

```

Node Name: IRONMAN
Filespace Name: /space
FSID: 2
Start Time: 02/08/11 21:44:19
End Time: 02/08/11 21:48:14
Status: Ended
Process Number: 4
Command: replicate node ironman
Phase: None
Process Running Time: 0 Day(s) 0 Hour(s)
4 Minute(s)
Completion State: Complete
Reason For Incompletion: None
Backup Last Update Date/Time:
Backup Target Server:
Backup Files Needing No Action: 0
Backup Files To Replicate: 0
Backup Files Replicated: 0
Backup Files Not Replicated Due to Errors: 0
Backup Files Not Yet Replicated: 0
Backup Files To Delete: 0
Backup Files Deleted: 0
Backup Files Not Deleted Due To Errors: 0
Backup Files To Update: 0
Backup Files Updated: 0
Backup Files Not Updated Due To Errors: 0
Backup Bytes To Replicate (MB): 0
Backup Bytes Replicated (MB): 0
Backup Bytes Transferred (MB): 0
Backup Bytes Not Replicated Due
To Errors (MB): 0
Backup Bytes Not Yet Replicated (MB): 0

Archive Last Update Date/Time: 02/08/11 21:48:14
Archive Target Server: NIGLINA
Archive Files Needing No Action: 0
Archive Files To Replicate: 39,416
Archive Files Replicated: 39,206
Archive Files Not Replicated Due to Errors: 210
Archive Files Not Yet Replicated: 0
Archive Files To Delete: 0
Archive Files Deleted: 0
Archive Files Not Deleted Due To Errors: 0
Archive Files To Update: 0
Archive Files Updated: 0
    
```

QUERY REPLICATION

```
Archive Files Not Updated Due To Errors: 0
  Archive Bytes To Replicate (MB): 4,335
  Archive Bytes Replicated (MB): 4,335
  Archive Bytes Transferred (MB): 0
  Archive Bytes Not Replicated
    Due To Errors (MB): 0
  Archive Bytes Not Yet Replicated (MB): 0

  Space Managed Last Update Date/Time:
  Space Management Target Server:
  Space Managed Files Needing No Action: 0
  Space Managed Files To Replicate: 0
  Space Managed Files Replicated: 0
  Space Managed Files Not Replicated
    Due to Errors: 0
  Space Managed Files Not Yet Replicated: 0
  Space Managed Files To Delete: 0
  Space Managed Files Deleted: 0
  Space Managed Files Not Deleted
    Due To Errors: 0
  Space Managed Files To Update: 0
  Space Managed Files Updated: 0
  Space Managed Files Not Updated
    Due To Errors: 0
  Space Managed Bytes To Replicate (MB): 0
  Space Managed Bytes Replicated (MB): 0
  Space Managed Bytes Transferred (MB): 0
  Space Managed Bytes Not Replicated
    Due To Errors (MB): 0
  Space Managed Bytes Not Yet Replicated (MB): 0
  Total Files Needing No Action: 0
  Total Files To Replicate: 39,416
  Total Files Replicated: 39,206
  Total Files Not Replicated Due To Errors: 210
  Total Files Not Yet Replicated: 0
  Total Files To Delete: 0
  Total Files Deleted: 0
  Total Files Not Deleted Due To Errors: 0
  Total Files To Update: 0
  Total Files Updated: 0
  Total Files Not Updated Due To Errors: 0
  Total Bytes To Replicate (MB): 4,335
  Total Bytes Replicated (MB): 4,335
  Total Bytes Transferred (MB):
  Total Bytes Not Replicated
    Due to Errors (MB):
  Total Bytes Not Yet Replicated (MB):
  Estimated Percentage Complete: 100
  Estimated Time Remaining:
  Estimated Time of Completion:
```

Field descriptions

Node Name

The name of the client node whose data is displayed.

Filespace Name

The name of the client file space whose data is displayed.

FSID The file space identifier.

Start Time

The date and time that the node replication process started.

End Time

The date and time that the node replication process ended.

Status The status of the node replication process. The following values are possible:

Running

The process is active and is either searching for eligible data or sending data to the target replication server.

Ended The process is ended.

Process Number

The identifier for the node replication process.

The same process number can have different start times. If a replication process starts and the server is restarted, the server begins assigning process numbers beginning with the number 1. Replication processes that start after a server restart can obtain process numbers that are already assigned to other replication processes in the replication history. To identify unique replication processes, use the start time.

Command

The command that was issued to start the node replication process.

Phase The phase of a running node-replication process. The following phases are listed in the order in which they occur:

Identifying

The node replication process started to identify data to be replicated, but data is not yet being sent to the target replication server.

Identifying and replicating

The node replication process is identifying data to be replicated and transferring the data to the target replication server.

Replicating

The node replication process identified the data and is transferring files to the target replication server.

None The node replication process is not running.

Process Running Time

The running time of the node replication process.

Completion State

The state of the node replication process. The following values are possible:

Complete

The node replication process completed.

Incomplete

The node replication process ended without running to completion. To determine the reason, check the value in the Reason for Incompletion field.

Reason For Incompletion

The reason why the node replication process ended without completing. Possible values include *cancelled* and *other*. The value *other* can indicate that the server was halted during replication or that the server failed.

Backup Last Update Date/Time, Archive Last Update Date/Time, Space Managed Last Update Date/Time

The date and time that statistics for backup, archive, or space-managed files were last updated. The specified time is the time that the files in the file space were identified for replication or when each batch of files was sent to the target replication server.

QUERY REPLICATION

Backup Target Server, Archive Target Server, Space Management Target Server
The name of the target replication server for backup, archive, or space-managed files.

Backup Files Needing No Action, Archive Files Needing No Action, Space Managed Files Needing No Action
The number of backup, archive, or space-managed files in the file space that did not need to be replicated, updated or deleted.

Backup Files To Replicate, Archive Files To Replicate, Space Managed Files To Replicate
The number of backup, archive, or space-managed files to replicate to the target replication server.

Backup Files Replicated, Archive Files Replicated, Space Managed Files Replicated
The number of backup, archive, or space-managed files that have been replicated to the target replication server.

Backup Files Not Replicated Due To Errors, Archive Files Not Replicated Due To Errors, Space Managed Files Not Replicated Due To Errors
The number of backup, archive, or space-managed files that were not replicated to the target replication server because of errors.

Backup Files Not Yet Replicated, Archive Files Not Yet Replicated, Space Managed Files Not Yet Replicated
The number of backup, archive, or space-managed files not yet replicated to the target replication server.

Backup Files To Delete, Archive Files To Delete, Space Managed Files To Delete
The number of backup, archive, or space-managed files to be deleted on the target replication server.

Backup Files Deleted, Archive Files Deleted, Space Managed Files Deleted
The number of backup, archive, or space-managed files that have been deleted on the target replication server.

Backup Files Not Deleted Due To Errors, Archive Files Not Deleted Due To Errors, Space Managed Files Not Deleted Due To Errors
The number of backup, archive, or space-managed files that were not deleted from the target replication server because of errors.

Backup Files To Update, Archive Files To Update, Space Managed Files To Update
The number of backup, archive, or space-managed files to update on the target replication server. If the metadata of a file is changed, the changed fields are sent to the target replication server.

Backup Files Updated, Archive Files Updated, Space Managed Files Updated
The number of backup, archive, or space-managed files that have been updated on the target replication server.

Backup Files Not Updated Due To Errors, Archive Files Not Updated Due To Errors, Space Managed Files Not Updated Due To Errors
The number of backup, archive, or space-managed files that were not updated on the target replication server because of errors.

Backup Bytes To Replicate (MB), Archive Bytes To Replicate (MB), Space Managed Bytes To Replicate (MB)
The number of backup, archive, or space-managed bytes to replicate to the target replication server.

Backup Bytes Replicated (MB), Archive Bytes Replicated (MB), Space Managed Bytes Replicated (MB)

The number of backup, archive, or space-managed bytes that have been replicated to the target replication server.

If a file was stored in a deduplicated storage pool, the number of bytes in the stored file might be less than the number of bytes in the original file. This field represents the number of physical bytes in the original file.

Backup Bytes Transferred (MB), Archive Bytes Transferred (MB), Space Managed Bytes Transferred (MB)

The number of backup, archive, or space-managed bytes that were sent to the target replication server.

The value in this field represents the actual number of file bytes sent to the target replication server. This value is calculated by subtracting the number of bytes not sent because of deduplication from the number of bytes to replicate.

Backup Bytes Not Replicated Due to Errors (MB), Archive Bytes Not Replicated Due to Errors (MB), Space Managed Bytes Not Replicated Due to Errors (MB)

The number of backup, archive, or space-managed bytes that were not replicated to the target replication server because of errors.

Backup Bytes Not Yet Replicated (MB), Archive Bytes Not Yet Replicated (MB), Space Managed Bytes Not Yet Replicated (MB)

The number of backup, archive, or space-managed bytes not yet replicated to the target replication server.

Total Files Needing No Action

The total number of files in the file space that did not need to be replicated, updated, or deleted.

Total Files To Replicate

The total number of files to replicate to the target replication server.

Total Files Replicated

The total number of files that have been replicated to the target replication server.

Total Files Not Replicated Due To Errors

The total number of files that were not replicated because of errors.

Total files Not Yet Replicated

The total number of files not yet replicated to the target replication server.

Total Files To Delete

The total number of files that were deleted on the target replication server.

Total Files Deleted

The total number of files that have been deleted on the target replication server.

Total Files Not Deleted Due to Errors

The total number of backup, archive, and space-managed files that were not deleted on the target replication server because of errors.

Total Files To Update

The total number of files to be updated on the target replication server. When the metadata of a file is changed, the changed fields are sent to the target replication server.

QUERY REPLICATION

Total Files Updated

The total number of files that have been updated on the target replication server.

Total Files Not Updated Due to Errors

The total number of backup, archive, and space-managed files that were not updated on the target replication server because of errors.

Total Bytes To Replicate (MB)

The total number of bytes to replicate to the target replication server.

Total Bytes Replicated (MB)

The total number of bytes that have been replicated to the target server.

If a file was stored in a deduplicated storage pool, the number of bytes in the stored file might be less than the number of bytes in the original file. This field represents the number of physical bytes in the original file.

Total Bytes Transferred (MB)

The total number of bytes that were transferred to the target replication server.

For files stored in a deduplicated storage pool, the value in this field includes the number of bytes in the original file before duplicate extents were removed. If duplicate extents were already on the target replication server, the number of bytes in the original file is more than the number of bytes transferred.

Total Bytes Not Replicated Due to Errors (MB)

The total number of bytes that were skipped because the source replication server was unable to transfer them to the target replication server.

Total Bytes Not Yet Replicated (MB)

The total number of bytes not yet transferred to the target replication server.

Estimated Percentage Complete

The estimated completion percentage based on the number of bytes.

Estimated Time Remaining

The estimated time remaining before the node replication process is complete.

Estimated Time Of Completion

The estimated time when the node replication process will end.

Table 261. Commands related to QUERY REPLICATION

Command	Description
CANCEL REPLICATION	Cancels node replication processes.
QUERY ACTLOG	Displays messages from the server activity log.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY PROCESS	Displays information about background processes.
QUERY REPLNODE	Displays information about the replication status of a client node.

Table 261. Commands related to QUERY REPLICATION (continued)

Command	Description
QUERY REPLRULE	Displays information about node replication rules.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
SET REPLRETENTION	Specifies the retention period for replication history records.

QUERY REPLNODE (Display information about replication status for a client node)


Use this command to display the number of files that are stored for each replicated file space. Information is displayed about file spaces for every client node that is configured for replication.

A client node is configured for replication if it is enabled or disabled.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

node_name (Required)

Specifies the client node that owns the files about which you want information. You can specify one or more names. If you specify multiple names, separate the names with commas. Do not use intervening spaces. You can use wildcard characters to specify multiple names.

Information about client nodes that match the file criteria, but that are not configured for replication, is not displayed.

server_name

Specifies the name of the replication server to query for replication information. This parameter is optional. If you do not specify a value for this parameter, the server that is the default target for replicated data is queried.

As the value for this parameter, you can also specify a server that was formerly a target for replicated data.

The client nodes that are defined to a replication server can be the source or the target of replicated data. To determine whether a particular client node is sending or receiving data, issue the **QUERY NODE** command. Look for the value *Send* or *Receive* in the Replication Mode field of the output.

To display the name of the active target replication server, issue the **QUERY STATUS** command, and look for the name in the Target Replication Server field.

Example: List client node files on a source and a target replication server

The name of the client node is NODE1.

```
query replnode *
```

Node Name	Type	Filespace Name	FSID	Files on Server	Replication Server (1)	Files on Server (1)
NODE1	SpMg	/hmsmfs	1	1		
NODE1	Bkup	/1space2	2	27		
NODE1	Arch	/1space2	2	22	TGTSRV	22
NODE1	Bkup	/1space	3	18,096		
NODE1	Arch	/1space	3	61,150	TGTSRV	61,150
NODE2						

The number of files that are displayed for the replication servers might be different for the following reasons:

- The output of the **QUERY REPLNODE** command displays the number of files obtained from the occupancy table. The occupancy table contains only files that have a length greater than zero. Files that have a length of 0 and have been replicated are not reflected in this output.
- If only active data is replicated to the target server, the number of files that are displayed for the source server will be larger than the number of files that are displayed on the target server. The reason for the difference is that the source replication server has both active and inactive data, and the target server has only active data.
- A client node might have data that was exported from the source replication server and imported to the target replication server. If that data was synchronized and if the client node also stored data to the target replication server, then the number of files on the target replication server will be greater than the number of files stored as a result of export-and-import operations and replication.

Field descriptions

Node Name

The name of the client node that owns the files.

Type The type of data. If this field is blank, the client node is configured for replication, but it does not have data on the replication server. In the example output, NODE2 is configured for replication, but it does not have backup, archive, or space-managed data.

The following values are possible:

Arch Archive data

Bkup Backup data

SpMg Data that was migrated by IBM Tivoli Storage Manager for Space Management clients

Filespace Name

The name of the file space that belongs to the node.

If this field is blank, the client node is configured for replication, but it does not have data on the replication server.

File space names can be in a different code page or locale than the server. If they are, the names in the Administration Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

QUERY REPLNODE

If the file space name is Unicode-enabled, Tivoli Storage Manager converts the name to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

FSID The file space identifier for the file space. The server assigns a unique FSID when a file space is initially stored on the server. If this field is blank, the client node is configured for replication, but it does not have data on the replication server.

Files on Server

The number of backup, archive, and space-managed files on the server on which this command is issued. If this field is blank, the client node is configured for replication, but it does not have data on the replication server.

Replication Server (1)

The name of the replication server that is being queried for information. If this field is blank, one or more of the following conditions might exist:

- The file space of the node on the replication server where the command was issued does not have any data.
- The client node is not defined on replication server (1).
- The client node is defined on replication server (1), but the node is not configured for replication.
- The corresponding file space on replication server (1) does not have data or the file space is not defined.

Files on Server (1)

The number of files for the data type that are stored on the target replication server. This field can be blank. If it is, one or more of the following conditions might exist:

- Replication server (1) does not have any data.
- The client node is not defined on replication server (1).
- The client node is defined on replication server (1), but the node is not configured for replication.
- The corresponding file space on replication server (1) does not have data or the file space is not defined.

Related commands

Table 262. Commands related to QUERY REPLNODE

Command	Description
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLRULE	Displays information about node replication rules.

Table 262. Commands related to QUERY REPLNODE (continued)

Command	Description
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
UPDATE REPLRULE	Enables or disables replication rules.

QUERY REPLRULE (Query replication rules)


Use this command to display information about replication rules.

Issue this command on the server that acts as a source for replicated data.

Privilege class

Any administrator can issue this command.

Syntax



Privilege class

Any administrator can issue this command.

Parameters

rule_name

Specifies the name of the replication rule that you want to display information about. This parameter is optional. You can use wildcard characters to specify one or more rules. If you do not specify this parameter, information about all rules is displayed in the query output. You can specify the following values:

ALL_DATA

Displays information about the ALL_DATA replication rule. This rule replicates backup, archive, or space-managed data. The data is replicated with a normal priority.

ACTIVE_DATA

Displays information about ACTIVE_DATA replication rule. This rule replicates only active backup data. The data is replicated with a normal priority. This rule is not valid for archive or space-managed data.

ALL_DATA_HIGH_PRIORITY

Displays information about the ALL_DATA_HIGH_PRIORITY rule. This rule replicates backup, archive, or space-managed data. The data is replicated with a normal priority. In a replication process, high-priority data is replicated before normal-priority data.

ACTIVE_DATA_HIGH_PRIORITY

Displays information about the ACTIVE_DATA_HIGH_PRIORITY rule. This rule replicates only active backup data. The data is replicated with a high priority. In a replication process, high-priority active data is replicated before normal priority data. This rule is not valid for archive or space-managed data.

Example: Display information about a server replication rule

The name of the rule is ALL_DATA_HIGH_PRIORITY

```
query replrule all_data_high_priority
```

Rule name: ALL_DATA_HIGH_PRIORITY
 Priority: High
 Active-data only?: No
 Enabled: Yes

Field descriptions

Rule name

Specifies the name of the rule that was queried.

Priority

Specifies the preference given to data during replication processing. In a replication process that includes high-priority and normal-priority data, high-priority data is replicated first.

Active data only?

Specifies whether the rule applies only to active backup data. The following values are possible:

- Yes** Specifies that only active backup data is replicated for file spaces to which this rule is assigned.
- No** Specifies that all backup data is replicated for file spaces to which this rule is assigned.

Enabled

Specifies whether the rule is enabled or disabled. The following values are possible:

- Yes** Specifies that the rule is enabled for replication. Data in file spaces to which the rule is assigned is replicated.
- No** Specifies that the rule is not enabled for replication. Data in file spaces to which the rule is assigned is not replicated.

Related commands

Table 263. Commands related to QUERY REPLRULE

Command	Description
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLNODE	Displays information about the replication status of a client node.
UPDATE REPLRULE	Enables or disables replication rules.

QUERY REQUEST (Query one or more pending mount requests)

Use this command to display information about one or more pending mount requests. The server makes requests for the administrator to perform an action, like inserting a tape volume in a library after a **CHECKIN LIBVOL** has been issued.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

request_number

Specifies the identification number of the pending mount request. This parameter is optional. The default is all pending mount requests.

Example: List all pending mount requests

Display information on all pending mount requests after a **CHECKIN LIBVOL** has been issued.

```
query request
```

Output for a Manual Library

```
ANR8352I Requests outstanding:
ANR8326I 001: Mount 8MM volume EXP001 R/W
in drive 8MM.1 (/dev/mt0) of library
MANUALLIB within 60 minute(s).
```

```
ANR8352I Requests outstanding:
ANR8326I 001: Mount 8MM volume EXP001 R/W
in drive 8MM.1 (/dev/mt0) of library
MANUALLIB within 60 minute(s).
```

Output for an Automated Library

```
ANR8352I Requests outstanding:
ANR8306I 001: Insert 3590 volume 133540 R/W into the slot with element
number 31 of library 3590LIB within 60 minutes; issue 'REPLY'
along with the request ID when ready.
```

Related commands

Table 264. Related commands for QUERY REQUEST

Command	Description
CANCEL REQUEST	Cancels pending volume mount requests.
REPLY	Allows a request to continue processing.


QUERY RESTORE (Query restartable restore sessions)

Use this command to display information about the restartable restore sessions.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

node_name

Specifies the client node to be queried. This parameter is optional. If you do not specify a value, all client nodes with restartable restore sessions are displayed. You must specify a value for this parameter if you specify a file space name.

file_space_name

Specifies the file space to be queried. This parameter is optional. If you do not specify a value, all file spaces are matched for the specified node.

For a server that has clients with support for Unicode, you may need to have the server convert the file space name that you enter. For example, you may need to have the server convert the name you enter from the server's code page to Unicode. See the NAMETYPE parameter for details.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

NAMEType

Specify how you want the server to interpret the file space names that you enter. This parameter is useful when the server has clients with support for Unicode. You can use this parameter for Unicode-enabled Tivoli Storage Manager clients using Windows, Macintosh OS 9, Macintosh OS X, and NetWare operating systems.

Use this parameter only when you enter a partly or fully qualified file space name. The default value is SERVER. Possible values are:

SERVER

The server uses the server's code page to interpret the file space names.

QUERY RESTORE

UNICODE

The server converts the file space name entered from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server has a problem accessing system conversion routines.

FSID

The server interprets the file space names as their file space IDs (FSIDs).

Example: Display a restartable restore session on a specific client node

Display detailed information about client node JAMES associated with file space DRIVE_F_R. See "Field descriptions" for field descriptions.

```
query restore james drive_f_r format=detailed
```

```
Sess Number: -1
Restore State: Restartable
Elapsed Minutes: 2
Node Name: JAMES
FSID: 1
Filespace Name: DRIVE_F_R:
File Spec: /RESTORE/TESTDIR\
```

Field descriptions

Sess Number

Specifies the session number for the restartable restore session. The number for active restore sessions is the same number displayed on the **QUERY SESSION** command. For restore sessions in the restartable state, a negative number is displayed for the session number. Any session number displayed in the **QUERY RESTORE** output may be specified from the **QUERY RESTORE** output.

Restore State

- Active: Specifies the restore session is actively restoring files to the client.
- Restartable: Specifies the restore session failed and can be restarted from where it left off.

Elapsed Minutes

Specifies the number of minutes since the restore session started. Any restartable restore session with an elapsed time greater than the **RESTOREINTERVAL** server option can be automatically deleted from the database when needed or during expiration processing. If the elapsed time is less than the **RESTOREINTERVAL**, you can delete this entry (and unlock the filesystem) only by issuing the **CANCEL RESTORE** command lowering the **RESTOREINTERVAL** value.

Node Name

Specifies the node associated with the restartable restore session.

FSID Specifies the file space ID of the file space.

Filespace Name

Specifies the file space associated with the restartable restore session.

File space names can be in a different code page or locale than the server. If they are, the names in the Administration Center and the administrative command-line interface might not be displayed correctly. Data is backed

up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, Tivoli Storage Manager converts the name to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

File Spec

Specifies the file specification used on the restore operation. The same file specification must be specified if a failed restore operation is to be restarted from where it stopped.

Related commands

Table 265. Commands related to QUERY RESTORE

Command	Description
CANCEL RESTORE	Cancels a restartable restore session.

Example: Display the source server recovery plan

On the source server, display the contents of a recovery plan file that was created for this server on March 19, 1998, at 6:10 A.M. The **PREPARE** command specifies the device class REMOTE. The output of this command is the entire contents of the recovery plan file. See the *Administrator's Guide* for an example of a disaster recovery plan file.

```
query rpfcontent branch1.19980319.061000 devclass=remote
```

Example: Display the target server recovery plan

On the target server, display the contents of a recovery plan file that was stored in this server on March 19, 1998, at 6:10 A.M. The server that created the file is registered on the target server as a node named POLARIS with a node type of SERVER. The output of this command is the entire contents of the recovery plan file. See the *Administrator's Guide* for an example of a disaster recovery plan file.

```
query rpfcontent branch1.19980319.061000 nodename=polaris
```

Related commands

Table 266. Commands related to QUERY RPFCONTENT

Command	Description
PREPARE	Creates a recovery plan file.
QUERY RPFFILE	Displays information about recovery plan files.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.


QUERY RPFIL (Query recovery plan file information stored on a target server)

Use this command to display information about recovery plan files stored on a target server. You can issue this command from either the server that created the file (the source server) or the server that stores the recovery plan file (the target server).

Privilege class

Any administrator can issue this command.

Syntax



Parameters

DEVclass

Specifies the name of the device class that was used to create the recovery plan files. Use this parameter when logged on to the server that created the recovery plan file. You can use wildcard characters in the device class name. All recovery plan files that are created with the device class specified are included in the query.

NODENAME

Specifies the node name, registered on the target server, of the source server that created the recovery plan files. Use this parameter when logged on to the target server. You can use this parameter when the source server is not available. You can use wildcard characters to specify the node name. All file objects that are stored with the node name specified are included in this query.

Source

Specifies the type of database backup that was specified when the recovery plan file was prepared. This parameter is optional. The default is DBBACKUP. Possible values are:

DBBackup

The recovery plan file was prepared with full and incremental database backups specified.

DBSnapshot

The recovery plan file was prepared with snapshot database backups specified.

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Tivoli Storage Manager displays partial information for the recovery plan file.

Detailed

Tivoli Storage Manager displays all information for the recovery plan file.

Example: Display detailed information about the recovery plans

Display recovery plan files that were created for this server using the specified device class. See "Field descriptions" for field descriptions.

```
query rpfiler devclass=* format=detailed
```

```
Recovery Plan File Name: ALASKA.20000406.170423
      Node Name: BRANCH1
      Device Class Name: REMOTE
Recovery Plan File Type: RPFIL
      Mgmt Class Name: STANDARD
Recovery Plan File Size: 16,255 Bytes
      Marked for Deletion: Yes
      Deletion Date: 06/12/2000 13:05:31

Recovery Plan File Name: ALASKA.20000407.170845
      Node Name: BRANCH1
      Device Class Name: REMOTE
Recovery Plan File Type: RPFNSAPSHOT
      Mgmt Class Name: STANDARD
Recovery Plan File Size: 16,425 Bytes
      Marked for Deletion: No
      Deletion Date:
```

Example: Display a list of recovery plans for a specific node name

Display a list of all recovery plan file objects that are stored with the specified node name (TYPE=SERVER). See "Field descriptions" for field descriptions.

```
query rpfiler nodename=branch1
```

Recovery Plan File Name	Node Name	Device Class Name
ALASKA.19980406.170423	BRANCH1	REMOTE
ALASKA.19980407.170845	BRANCH1	REMOTE

Field descriptions

Recovery Plan File Name

The recovery plan file name.

Node Name

The node name that is registered with the target server and used to store the recovery plan file objects.

Device Class Name

The device class name that is defined in the source server and used to create the recovery plan files.

Recovery Plan File Type

The type of recovery plan file:

RPFIL

The plan assumes full plus incremental database backups.

RPFNSAPSHOT

The plan assumes snapshot database backups.

QUERY RPFIL

Mgmt Class Name

The management class name that the recovery plan file is associated with in the target server.

Recovery Plan File Size

Estimated size of the recovery plan file object on the target server.

Marked For Deletion

Whether the object that contains the recovery plan file has been deleted from the source server and marked for deletion on the target server if the grace period has not expired. Possible values are:

Yes The object is marked for deletion.

No The object is not marked for deletion.

Deletion Date

Date that the object has been deleted from the source server and marked for deletion on the target server. This field is blank if the object has not been marked for deletion.

Related commands

Table 267. Commands related to QUERY RPFIL

Command	Description
PREPARE	Creates a recovery plan file.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
QUERY RPFCONTENT	Displays the contents of a recovery plan file.

QUERY SAN (Query the devices on the SAN)


Use this command to obtain information about devices that can be detected on a storage area network (SAN), so that you can configure Tivoli Storage Manager for LAN-free data movement.

The **QUERY SAN** command requires the libhaapi.so that supports SNIA common Host Bus Adapter (HBA) API. With this library object, Tivoli Storage Manager can call the hbaapi functions specified in the SNIA common HBAAPI standard.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

Type

Specifies the type of device that will be displayed. This parameter is optional. The default value is Any. Possible values are:

Any

Specifies that any device detected on the SAN is displayed.

DISK

Specifies that only disk devices are displayed.

DRive

Specifies that only drive devices are displayed.

LIBRARY

Specifies that only library devices are displayed.

OTHER

Specifies that devices that are not disks, drives, or libraries are displayed.

Format

Specifies the type of information that will be displayed. This parameter is optional. The default value is Standard. Possible values are:

Standard

Specifies that the information displayed will be summarized.

Detailed

Specifies that complete information will be displayed.

Tip: The output might not display the serial number of the device. If this happens, look on the back of the device or contact the manufacturer of the device.

Example: List drive devices

Display summary information for drive devices on a SAN. See “Field descriptions” for field descriptions.

```
query san type=drive
```

Device Type	Vendor	Product	Serial	Device
LIBRARY	STK	L180	MPC01000128	/dev/smc1
DRIVE	STK	9840	331001017229	/dev/rmt3
DRIVE	Quantum	DLT4000	JF62806275	/dev/rmt4
DRIVE	Quantum	DLT4000	JP73213185	/dev/rmt5
DRIVE	STK	9840	331000028779	/dev/rmt6

Example: Display drive device information

Display detailed information for all drive devices on a SAN. See “Field descriptions” for field descriptions.

```
query san type=drive format=detailed
```

Device Type:	DRIVE
Vendor:	IBM
Product:	03570B02
Serial:	
Device:	mt10.2.0.3
DataMover:	No
Node WWN:	5005076206039E05
Port WWN:	5005076206439E05
Lun:	0
SCSI Port:	3
SCSI Bus:	0
SCSI Target:	10

Field descriptions

Device Type

The type of device that is being displayed.

Vendor

The name of the device's vendor.

Product

The name of the product assigned by the vendor.

Serial The serial number of the device.

Device

The device special file name.

Data Mover

Whether the device is a data mover.

Node WWN

The worldwide name for the device.

Port WWN

The worldwide name for the device, which is specific to the port that the device is connected to.

Lun The Logical Unit Number of the device.

SCSI Port

The port of the Fibre Channel (or SCSI) Host Bus Adapter.

SCSI Bus

The bus of the Host Bus Adapter card.

SCSI Target

The target number of the device.

Related commands

Table 268. Commands related to QUERY SAN

Command	Description
DEFINE DATAMOVER	Defines a data mover to the Tivoli Storage Manager server.
DEFINE DRIVE	Assigns a drive to a library.
DEFINE LIBRARY	Defines an automated or manual library.

QUERY SCHEDULE (Query schedules)

Use this command to display information about one or more schedules.

The QUERY SCHEDULE command takes two forms, depending on whether the schedule applies to client operations or administrative commands. The syntax and parameters for each operation are defined separately. Some options in the query display will be blank depending on whether the schedule style is classic or enhanced.

- “QUERY SCHEDULE (Query an administrative schedule)” on page 885
- “QUERY SCHEDULE (Query client schedules)” on page 881

Table 269. Commands related to QUERY SCHEDULE

Command	Description
COPY SCHEDULE	Creates a copy of a schedule.
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.
UPDATE SCHEDULE	Changes the attributes of a schedule.


QUERY SCHEDULE (Query client schedules)

Use this command to display information about one or more client schedules.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

domain_name

Specifies the name of the policy domain to which the schedule belongs. You can use a wildcard character to specify this name. If you specify a domain name, you do not have to specify a schedule name.

schedule_name

Specifies the name of the schedule that belongs to the specified policy domain. You can use a wildcard character to specify this name. If you specify a schedule name, you must also specify a policy domain name.

Type=Client

Specifies that the query displays client schedules. This parameter is optional. The default is CLIENT.

Nodes

Specifies the name of one or more client nodes that are associated with the schedules to be displayed. This parameter is optional. You can use a wildcard character to specify client nodes. If you do not specify a client name, all schedules matching the DOMAINNAME and SCHEDULENAME parameters are displayed. You can specify multiple client nodes by separating the names with commas and no intervening spaces.

Format

Specifies how information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the schedules.

Detailed

Specifies that detailed information is displayed for the schedules.

The standard format displays a blank in the period column and an asterisk in the day column for enhanced schedules. To display complete information about an enhanced schedule, issue FORMAT=DETAILED.

QUERY SCHEDULE

Example: List schedules for a specific policy domain

Display all schedules that belong to the EMPLOYEE_RECORDS policy domain. See “Field descriptions: Schedules for a specific policy domain” for field descriptions.

```
query schedule employee_records
```

The standard format displays a blank in the period column and an asterisk in the day column for enhanced schedules. To display complete information about an enhanced schedule, issue `FORMAT=DETAILED`.

Domain	* Schedule Name	Action	Start Date/Time	Duration	Period	Day
EMPLOY- EE_RE- CORDS	WEEKLY_BACKUP	Inc Bk	2004.06.04 17.04.20	1 H	1 D	Any
EMPLOY- EE_RE- CORDS	EMPLOYEE_BACKUP	Inc Bk	2004.06.04 17.04.20	1 H		(*)

Field descriptions: Schedules for a specific policy domain

Domain

Specifies the name of the policy domain to which the specified schedule belongs.

* (Asterisk)

Specifies whether the corresponding schedule has expired. If there is an asterisk in this column, the corresponding schedule has expired.

Schedule Name

Specifies the name of the schedule.

Action

Specifies the action that occurs when this schedule is processed.

Start Date/Time

Specifies the initial starting date and time for this schedule.

Duration

Specifies the length of the startup window for this schedule.

Period Specifies the time between startup windows (assuming `DAYOFWEEK=ANY`). The column is blank for enhanced schedules.

Day Specifies the day of the week on which the startup windows for the schedule begin. The column contains an asterisk for enhanced schedules.

Example: Display detailed client schedules

From a managed server, display detailed information about client schedules. See “Field descriptions: Detailed client schedules” on page 883 for field descriptions.

```
query schedule * type=client format=detailed
```

```

Policy Domain Name: ADMIN_RECORDS
Schedule Name: ADMIN_BACKUP
Description:
  Action: Incremental
  Subaction:
  Options:
  Objects:
  Priority: 5
Start Date/Time: 2004.06.04 17.04.20
Duration: 1 Hour(s)
Schedule Style: Classic
  Period: 1 Day(s)
  Day of Week: Any
  Month:
  Day of Month:
  Week of Month:
  Expiration:
Last Update by (administrator): $$CONFIG_MANAGER$$
Last Update Date/Time: 2004.06.04 17.51.49
Managing profile: ADMIN_INFO
Policy Domain Name: EMPLOYEE_RECORDS
Schedule Name: EMPLOYEE_BACKUP
Description:
  Action: Incremental
  Subaction:
  Options:
  Objects:
  Priority: 5
Start Date/Time: 2004.06.04 17.04.33
Duration: 1 Hour(s)
Schedule Style: Enhanced
  Period:
  Day of Week: Any
  Month: Mar,Jun,Nov
  Day of Month: -14,14,22
  Week of Month: Last
  Expiration:
Last Update by (administrator): $$CONFIG_MANAGER$$
Last Update Date/Time: 2004.06.04 17.18.30
Managing profile: EMPLOYEE

```

Field descriptions: Detailed client schedules

Policy Domain Name

Specifies the name of the policy domain.

Schedule Name

Specifies the name of the schedule.

Description

Specifies the description of the schedule.

Action

Specifies the type of action that occurs when this schedule is run. See the DEFINE SCHEDULE command for a listing of actions.

Subaction

Specifies that the type of operation identified by the **ACTION** parameter is to be scheduled. See the DEFINE SCHEDULE command for a listing of subactions.

Options

Specifies the options that are supplied to the DSMC command when the schedule is run.

Objects

Specifies the objects for which the specified action is performed.

QUERY SCHEDULE

Priority

Specifies the priority value for the schedule.

Start Date/Time

Specifies the initial starting date and time for the schedule.

Duration

Specifies the length of the startup window for the schedule.

Schedule Style

Specifies whether classic or enhanced schedule rules are used.

Period Specifies the time between startup windows (assuming DAYOFWEEK=ANY). This is not displayed for enhanced syntax schedules.

Day of Week

Specifies the day of the week on which the startup windows for the schedule begin. Using a standard format displays an asterisk in the day of week field for enhanced schedules.

Month

Specifies the months during which the schedule will run. This is not displayed for classic syntax schedules.

Day of Month

Specifies the days of the month during which the schedule will run. This is not displayed for classic syntax schedules.

Week of Month

Specifies the weeks (first, second, third, fourth, or last) of the month during which the schedule will run. This is not displayed for classic syntax schedules.

Expiration

Specifies the date and time on which this schedule expires. If this column is blank, the schedule does not expire.

Last Update by (administrator)

Specifies the name of the administrator that most recently updated the schedule. If this field contains a \$\$CONFIG_MANAGER\$\$, the schedule is associated with a domain that is managed by the configuration manager.

Last Update Date/Time

Specifies the last date and time the schedule was last updated.

Managing Profile

Specifies the profile or profiles to which the managed server subscribed to get the definition of this schedule.


QUERY SCHEDULE (Query an administrative schedule)

Use this command to display information about one or more administrative schedules.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

schedule_name

Specifies the name of the schedule to be queried. You can use a wildcard character to specify this name.

Type=Administrative (Required)

Specifies that the query displays administrative command schedules.

Format

Specifies how information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the schedules.

Detailed

Specifies that detailed information is displayed for the schedules.

The standard format displays a blank period column and an asterisk in the day column for enhanced schedules. Issue `FORMAT=DETAILED` to display complete information about an enhanced schedule.

Example: Display detailed information on administrative command schedules

From a managed server, display detailed information about administrative command schedules. See “Field descriptions” on page 886 for field descriptions.

```

query schedule * type=administrative
format=detailed
    
```

QUERY SCHEDULE

```
Schedule Name: BACKUP_ARCHIVEPOOL
Description:
  Command: backup db
  Priority: 5
Start Date/Time: 2004.06.04 16.57.15
Duration: 1 Hour(s)
Schedule Style: Classic
  Period: 1 Day(s)
  Day of Week: Any
  Month:
  Day of Month:
  Week of Month:
Expiration:
  Active: No
Last Update by (administrator): $$CONFIG MANAGERS$
Last Update Date/Time: 2004.06.04 17.51.49
Managing Profile: ADMIN_INFO
Schedule Name: MONTHLY_BACKUP
Description:
  Command: q status
  Priority: 5
Start Date/Time: 2004.06.04 16.57.14
Duration: 1 Hour(s)
Schedule Style: Enhanced
  Period:
  Day of Week: Tue,Thu,Fri
  Month: Aug,Nov
  Day of Month:
  Week of Month: Second,Third
Expiration:
  Active: No
Last Update by (administrator): $$CONFIG MANAGER
Last Update Date/Time: 2004.06.04 17.51.49
Managing Profile: ADMIN_INFO
```

Field descriptions

Schedule Name

Specifies the name of the schedule.

Description

Specifies the description of the schedule.

Command

Specifies the command that is scheduled.

Priority

Specifies the priority value for this schedule.

Start Date/Time

Specifies the initial starting date and time for this schedule.

Duration

Specifies the length of the startup window.

Schedule Style

Specifies whether classic or enhanced schedule rules are used.

Period Specifies the time between startup windows (assuming DAYOFWEEK=ANY). This is not displayed for enhanced syntax schedules.

Day of Week

Specifies the day of the week on which the startup windows begin.

Month

Specifies the months during which the schedule will run. This is not displayed for classic syntax schedules.

Day of Month

Specifies the days of the month during which the schedule will run. This is not displayed for classic syntax schedules.

Week of Month

Specifies the weeks (first, second, third, fourth, or last) of the month during which the schedule will run. This is not displayed for classic syntax schedules.

Expiration

Specifies the date after which this schedule will no longer be used. If this column is blank, the schedule does not expire.

Active Specifies whether the schedule has been processed according to the time and date set for this schedule.

Last Update by (administrator)

Specifies the name of the administrator that most recently updated the schedule. If this field contains a `$$CONFIG_MANAGER$$`, the schedule is associated with a domain that is managed by the configuration manager.

Last Update Date/Time

Specifies the last date and time the schedule was modified.

Managing Profile

Specifies the profile or profiles to which the managed server subscribed to get the definition of this schedule.

QUERY SCRIPT (Query Tivoli Storage Manager scripts)

Use this command to display information about scripts.


You can use this command in conjunction with the DEFINE SCRIPT command to create a new script using the contents from another script.

Privilege class

The privilege class required for this command depends on whether the OUTPUTFILE parameter is specified in the command.

- If the OUTPUTFILE parameter is not specified, any administrator can issue this command.
- If the OUTPUTFILE parameter is specified and the REQSYSAUTHOUTFILE server option is set to YES, the administrator must have system privilege.
- If the OUTPUTFILE parameter is specified and the REQSYSAUTHOUTFILE server option is set to NO, the administrator must have operator, policy, storage, or system privilege.

Syntax



Parameters

script_name

Specifies the name of the script for which information is to be displayed. You can include a wildcard character to specify this name.

Important: If you do not specify a script, the query displays information about all scripts. The time used to process this command and the amount of information displayed can be extensive.

Format

Specifies the output format for displaying script information. The default is STANDARD. Possible values are:

Standard

Specifies that only the script name and description in a script are displayed.

Lines

Specifies that the script name, the line number of the commands, comment lines, and the commands in the script are displayed.

Detailed

Specifies that detailed information about the script is displayed.

Raw

Specifies that commands contained in the script are written to a file named with the OUTPUTFILE parameter. This format is a way of directing output from a script to a file so that it can be copied into another script using the DEFINE SCRIPT command.

If no output file is specified, the Tivoli Storage Manager server outputs the "query script" with "format=raw" to the console.

Outputfile

Specifies the name of the file to which output is directed when you specify FORMAT=RAW. The file you specify must reside on the server running this command. If the file exists, the query output is appended to the end of the file.

Example: List the script descriptions

Display the standard information about scripts.

query script *

Name	Description
QCOLS	Display columns for a specified SQL table
QSAMPLE	Sample SQL Query
EXAMPLE	Backup the store pools and database when no sessions

Example: Display the contents of a script with line numbers

Display the lines of information for a script named Q_AUTHORITY.

query script q_authority format=lines

Name	Line Number	Command
Q_AUTHORITY	1	/* -----*/
	5	/* Script Name: Q_AUTHORITY */
	10	/* Description: Display administrators that */
	15	/* have the authority to issue */
	20	/* commands requiring a */
	25	/* specific privilege. */
	30	/* Parameter 1: privilege name - in the form */
	35	/* x_priv - EX. policy_priv */
	40	/* Example: run q_authority storage_priv */
	45	/* -----*/
	50	select admin_name from admins where -
	55	upper(system_priv) <> 'NO' or -
	60	upper(\$1) <> 'NO'

Example: Create a script from an existing script

Query the ENGDEV script and direct the output to a file named MY.SCRIPT.

query script engdev format=raw outputfile=my.script

Example: Create a script from an existing script

Create a new script by using the contents from the file MY.SCRIPT. Name the new script AGADM.

define script agadm file=my.script

Example: Display detailed script information

Display detailed information about scripts. See “Field descriptions” for field descriptions.

```
query script * format=detailed
```

```

                Name: QCOLS
                Line Number: DESCRIPTION
                Command: Display columns for a specified SQL
                        table
Last Update by (administrator): SERVER_CONSOLE
                Last Update Date/Time: 12/02/1997 16:05:29

                Name: QCOLS
                Line Number: 1
                Command: select colname from columns where
                        tabname='$1'
Last Update by (administrator): SERVER_CONSOLE
                Last Update Date/Time: 12/02/1997 16:05:29
    
```

Field descriptions

Name The name of the script.

Line Number
The line number of the script or the string DESCRIPTION.

Command
The command included on the line number displayed in the previous field.

Last Update by (administrator)
The name of the administrator that has defined or most recently updated the script.

Last Update Date/Time
The date and time that the administrator defined or updated the script.

Related commands

Table 270. Commands related to QUERY SCRIPT

Command	Description
COPY SCRIPT	Creates a copy of a script.
DEFINE SCRIPT	Defines a script to the Tivoli Storage Manager server.
DELETE SCRIPT	Deletes the script or individual lines from the script.
RENAME SCRIPT	Renames a script to a new name.
RUN	Runs a script.
UPDATE SCRIPT	Changes or adds lines to a script.


QUERY SERVER (Query a server)

Use this command to display information about a server definition.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

server_name

Specifies the name of the server to be queried. You can use wildcard characters to specify this name. This parameter is optional. The default is all server names.

Format

Specifies how the information is displayed. The parameter is optional. The default is STANDARD.

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: List all servers

Display information in standard format about all servers. See “Field descriptions” on page 892 for field descriptions.

```
query server *
```

Server Name	Comm. Method	High-level Address	Low-level Address	Days Since Last Access	Server Password Set	Virtual Volume Password Set	Allow Replacement
SERVER_A	TCPIP	9.115.35.6	1501	11	Yes	No	No
SERVER_B	TCPIP	9.115.45.24	1500	<1	Yes	No	No
ASTRO	TCPIP	9.115.32.21	1500	24	Yes	No	No

Example: Display detailed information about a specific server

From a managed server, display detailed information about SERVER_A. See “Field descriptions” on page 892 for field descriptions.

```
query server server_a format=detailed
```

QUERY SERVER

```
Server Name: SERVER_A
Comm. Method: TCPIP
High-level Address: 9.115.4.15
Low-level Address: 1500
Description:
Allow Replacement:
Node Name:
Last Access Date/Time: 06/09/1998 12:51:51
Days Since Last Access: <1
Locked?:
Compression: Client's Choice
Archive Delete Allowed?: No
URL:
Registration Date/Time: 06/09/1998 12:51:51
Registering Administrator: $$CONFIG_MANAGER$$
Bytes Received Last Session: 0
Bytes Sent Last Session: 0
Duration of Last Session (sec): 0.00
Pct. Idle Wait Last Session: 0.00
Pct. Comm. Wait Last Session: 0.00
Pct. Media Wait Last Session: 0.00
Grace Deletion Period: 5
Managing Profile:
Server Password Set Date/Time:
Days since Server Password Set:
Invalid Sign-on count for Server:
Virtual Volume Password Set: No
Virtual Volume Password Set Date/Time:(?)
Days Since Virtual Volume Password Set:(?)
Invalid Sign-on Count for Virtual Volume Node: 0
Validate Protocol: No
SSL: No
```

Field descriptions

Server Name

The name of the server.

Communication Method

The communication method used to connect to the server.

High-level Address

The IP address (in dotted decimal format) of the server.

Low-level Address

The port number of the server.

Description

The server description.

Allow Replacement

Whether a server definition on a managed server can be replaced with a definition from a configuration manager.

Node Name

The name of the client node.

Last Access Date/Time

The last date and time that the client node accessed the server.

Days Since Last Access

The number of days since the client node accessed the server.

Locked?

Whether the client node is locked out of Tivoli Storage Manager.

Compression

The type of compression performed by Tivoli Storage Manager on client files.

Archive Delete Allowed?

Whether the client node can delete its own archive files. A value of (?) denotes that this field is not set and does not apply to this definition.

URL The URL used to access this server from a web browser-based interface.

Registration Date/Time

The date and time that the client node was registered.

Registering Administrator

The name of the administrator that registered the client node.

Bytes Received Last Session

The number of bytes received by the server during the last client node session.

Bytes Sent Last Session

The number of bytes sent to the client node.

Duration of Last Session (sec)

The length of the last client node session, in seconds.

Pct. Idle Wait Last Session

The percentage of the total session time that the client was not performing any functions.

Pct. Comm. Wait Last Session

The percentage of the total session time that the client waited for a response from the server.

Pct. Media Wait Last Session

The percentage of the total session time that the client waited for a removable volume to be mounted.

Grace Deletion Period

How many days an object remains on the target server after it has been marked for deletion.

Managing Profile?

The profile from which the managed server got the definition of this server.

Server Password Set Date/Time

When the password for the server was set.

Days since Server Password Set

The number of days since the server password was set.

Invalid Sign-on count for Server

The maximum number of invalid sign-on attempts that the server can accept.

Virtual Volume Password Set

Whether the password used to sign on to the target server has been set.

Virtual Volume Password Set Date/Time

When the password for virtual volume support was last set.

Days Since Virtual Volume Password Set

The number of days since the password for virtual volume support was set.

QUERY SERVER

Invalid Sign-on Count for Virtual Volume Node

The maximum number of invalid sign-on attempts that are accepted on the target server.

Validate Protocol

Specifies whether the storage agent has data validation enabled.

SSL Specifies whether Secure Sockets Layer (SSL) communication is used.

Related commands

Table 271. Commands related to **QUERY SERVER**

Command	Description
DEFINE DEVCLASS	Defines a device class.
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE DEVCLASS	Deletes a device class.
DELETE FILESPACE	Deletes data associated with clients file spaces.
DELETE SERVER	Deletes the definition of a server.
QUERY NODE	Displays partial or complete information about one or more clients.
RECONCILE VOLUMES	Reconciles source server virtual volume definitions and target server archive objects.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
SET REPLSERVER	Specifies a target replication server.
UPDATE DEVCLASS	Changes the attributes of a device class.
UPDATE NODE	Changes the attributes associated with a client node.
UPDATE SERVER	Updates information about a server.

QUERY SERVERGROUP (Query a server group)

Use this command to display information about server groups and group members.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

group_name

Specifies the server group to query. This parameter is optional. You can use wildcard characters to specify this name.

Example: List server groups

From a managed server, query all server groups. “Field descriptions” for field descriptions.

```
query servergroup *
```

Server Group	Group Member	Description	Managing Profile
ADMIN_GROUP	SERVER_A SERVER_B SERVER_C SERVER_D	Headquarters	ADMIN_INFO

Field descriptions

Server Group

The name of the server group.

Group Member

The group members.

Description

The description of the server group.

Managing Profile

The profile or profiles to which the managed server subscribed to get the definition of the server groups.

Related commands

Table 272. Commands related to QUERY SERVERGROUP

Command	Description
COPY SERVERGROUP	Creates a copy of a server group.
DEFINE SERVERGROUP	Defines a new server group.
DELETE SERVERGROUP	Deletes a server group.
QUERY SERVER	Displays information about servers.

QUERY SERVERGROUP

Table 272. Commands related to QUERY SERVERGROUP (continued)

Command	Description
RENAME SERVERGROUP	Renames a server group.
UPDATE SERVERGROUP	Updates a server group.


QUERY SESSION (Query client sessions)

Use this command to display information about administrative and client node sessions.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

sessnum

Specifies the number of the administrative or client node session to query. This parameter is optional. If you do not specify a value for this parameter, all sessions display.

MINTIMethreshold

Specifies to display sessions that have had at least this number of minutes elapse from the time the client sent data to the server for storage. This parameter is optional. The minimum number of minutes is 1. The maximum number of minutes is 99999999.

MAXTHROUGHput

Specifies to display sessions that are transferring data at a rate less than this number of kilobytes per second. This parameter is optional. The minimum number of kilobytes per second is 0. The maximum number of kilobytes per second is 99999999.

Format

Specifies how the information displays. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information displays for the session.

Detailed

Specifies that complete information displays for the session.

Example: List active client node sessions

Display information on all administrative and client node sessions that are communicating with the server. See “Field descriptions” on page 898 for field descriptions.

```
query session
```

Sess Number	Comm. Method	Sess State	Wait Time	Bytes Sent	Bytes Recvd	Sess Type	Platform	Client Name
77	TCP/IP	Run	0 S	134	159	Admin	OS/2	ADMIN

Example: Display detailed information on active client node sessions

Display information in full detail about all administrative and client node sessions that are communicating with the server. See “Field descriptions” for field descriptions.

```
query session format=detailed
```

```

    Sess Number: 77
    Comm. Method: Tcp/Ip
    Sess State: MediaW
    Wait Time: 13 S
    Bytes Sent: 2.2 K
    Bytes Recvd: 1.8 K
    Sess Type: Node
    Platform: AIX
    Client Name: N25XY
    Media Access Status: Waiting for mount of output volume:
                        COPYVOL1, 8MMPPOOL ( 12 seconds )

                        Waiting for mount point:
                        8MMCLASS, ONSITEPOOL ( 45 seconds )

                        Waiting for output volume:
                        COPYVOL2, OFFSITEPOOL ( 30 seconds )
    Proxy By Storage Agent: mickeynode
    User Name:
    Date/Time First Data Sent: 07/08/08    18:01:21
    Actions: BKIns

```

Field descriptions

Sess Number

Specifies a unique session identification number assigned by the server.

Comm. Method

Specifies the method being used by the client to communicate with the server.

Sess State

Specifies the current communications state of the server. Possible states are:

End The session is ending (session resources are released).

IdleW Waiting for client's next request (session is idle).

MediaW

The session is waiting for access to a sequential access volume.

RecvW

Waiting to receive an expected message from the client.

Run

The server is executing a client request (and not waiting to send data).

SendW

The server is waiting to send data to the client (waiting for data to be delivered to the client node that has already been sent).

Start

The session is starting (authentication is in progress).

Wait Time

Specifies the amount of time (seconds, minutes, or hours) the server has been in the current state shown.

Bytes Sent

Specifies the number of bytes of data sent to the client node since the session was initiated.

Bytes Recvd

Specifies the number of bytes of data received from the client node since the session was initiated.

Sess Type

Specifies the type of session in process: ADMIN for an administrative session, NODE for a client node session, or SERVER. SERVER specifies the server starts a session and initiates server-to-server operations such as central configuration, library sharing, and storage agent sessions.

Platform

Specifies the type of operating system associated with the client.

Client Name

Specifies the name of the client node or the administrator.

Media Access Status

Specifies the type of media wait state. When a session is in a media wait state, this field displays a list of all mount points and sequential volumes for the session. The list of mount points specifies the device class and the associated storage pool. The list of volumes specifies the primary storage pool volumes in addition to any copy storage pool and active-data pool volumes along with their assigned storage pool.

The server allows multiple sessions to read and one session to write to a volume concurrently in a storage pool associated with the FILE device type. As a result, a volume in a storage pool with a device type of FILE can appear as the current volume for more than one session.

Proxy by Storage Agent

Specifies the storage agent that is the proxy for LAN-Free data movement for the node.

User Name

Specifies the user ID of the node, on a multi-user system, that connects to the server when it is not the same system user who originally connected to the server.

Date/Time First Data Sent

Specifies the date and time that the client first sent data to the server for storage.

Actions

Displays a list of actions that have been performed during the session. An action is listed only once, even if the action occurs multiple times during a session. Possible actions are:

BkIns One or more backup objects were stored on the server. The operation might have been an incremental backup or a selective backup.

BkUpd One or more attributes were updated for a backup object that is stored on the server.

BkDel One or more backup objects stored on the server have been deleted.

QUERY SESSION

BkRebind

One or more backup objects stored on the server have been bound to a different management class.

NoQueryRestore

A no-query restore operation was initiated from the client, to restore backed-up files from the server to the client system.

ArIns One or more archive objects were stored on the server.

ObjRtrv

One or more files were retrieved from the server. This might have been to retrieve archive files, or to restore backup data (except for backup data from a no-query restore operation).

MigIns

One or more files have been migrated and stored on the server by Tivoli Storage Manager for Space Management (HSM client).

MigDel

One or more space-managed files that were stored on the server have been deleted.

MigRebind

One or more space-managed files that are stored on the server have been bound to a different management class.

MigRecall

One or more space-managed files that are stored on the server have been recalled.

MigUpd

The attributes for one or more space-managed files that are stored on the server have been updated.

FSAdd

The client node added one or more new file spaces to server storage.

FSUpd

The client node updated attributes for one or more file spaces that are defined to the server.

DefAuth

A **SET ACCESS** command has been processed by the client node, which caused an authorization rule for access to the client node's data to be added.

Related commands

Table 273. Command related to QUERY SESSION

Command	Description
CANCEL SESSION	Cancels active sessions with the server.


QUERY SHREDSTATUS (Query shredding status)

Use this command to display information about data waiting to be shredded.

Privilege class

To issue this command you must have administrator privilege.

Syntax



Parameters

Format

Specifies how the information is displayed. This parameter is optional. The default is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed. This is the default.

Detailed

Specifies that complete information is displayed.

Example: Display summary shredding information

Show partial information about data shredding on the server. See “Field descriptions” for field descriptions.

```
query shredstatus
```

Shredding Active	Objects Awaiting Shred
-----	-----
NO	4

Example: Display detailed shredding information

Display detailed information about data shredding on the server. See “Field descriptions” for field descriptions.

```
query shredstatus format=detailed
```

Shredding Active	Objects Awaiting Shred	Occupied Space (MB)	Data Left To Shred (MB)
-----	-----	-----	-----
NO	4	182	364

Field descriptions

Shredding Active

Indicates whether or not the server is actively shredding data at this time.

Objects Awaiting Shred

The number of objects currently waiting to be shredded.

QUERY SHREDSTATUS

Occupied Space (MB)

The amount of server storage space occupied by the objects currently waiting to be shredded, in megabytes. This is the amount of space that will become available when the objects are shredded.

Data Left to Shred (MB)

The amount of data that still needs to be shredded.

Related commands

Table 274. Commands related to **QUERY SHREDSTATUS**

Command	Description
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
EXPORT NODE	Copies client node information to external media or directly to another server.
GENERATE BACKUPSET	Generates a backup set of a client's data.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
MOVE DATA	Moves data from a specified storage pool volume to another storage pool volume.
QUERY STGPOOL	Displays information about storage pools.
SETOPT	Updates a server option without stopping and restarting the server.
SHRED DATA	Manually starts the process of shredding deleted data.
UPDATE STGPOOL	Changes the attributes of a storage pool.

QUERY SPACETRIGGER

STGP00L Expansion prefix

The prefix associated with the space trigger.

STGP00L

The storage pool name associated with the query.

Last Update by (administrator)

The administrator who last updated the storage pool space trigger.

Last Update Date/Time

The date and time when the administrator last updated the storage pool space trigger.

Related commands

Table 275. Commands related to QUERY SPACETRIGGER

Command	Description
DEFINE SPACETRIGGER	Defines a space trigger to expand the space for a storage pool.
DELETE SPACETRIGGER	Deletes the storage pool space trigger.
UPDATE SPACETRIGGER	Changes attributes of storage pool space trigger.

QUERY SSLKEYRINGPW (Query SSL key database file password)

Use this command to display the Secure Sockets Layer (SSL) key database file password. The key database file password is needed in order to update the key database file with new certificates or to designate a default certificate.

Privilege class

You must have system privileges to issue this command.

Syntax

►►—Query SSLKEYRINGPW—————►►

Parameters

The **QUERY SSLKEYRINGPW** command contains no parameters.

Example: Display the key database file password

Issue the command:

```
query sslkeyringpw
```

214qX{;8TP

Related commands

Table 276. Commands related to **QUERY SSLKEYRINGPW**

Command	Description
DELETE KEYRING	Deletes password information in the certificate key database.
SET SSLKEYRINGPW	Sets or updates the key database file password.

QUERY STATUS (Query system parameters)

Use the **QUERY STATUS** command to display information about system parameters.

Use this command for the following reasons:

- To display the service level of the server
- To display information about the general server parameters, such as those defined by the **SET** commands
- To request information about client sessions, such as the availability of the server, password authentication, accounting settings, or the retention period for the information retained in the activity log
- To display information about the central scheduler, such as the central scheduling mode of the server
- To display the maximum number of repeated attempts that are allowed after a failed attempt to execute a scheduled command
- To display whether sub-files can be backed up to this server, as indicated by the **SET SUBFILE** command

Privilege class

Any administrator can issue this command.

Syntax

►►—Query Status—◄◄

Parameters

None.

Example: Query the status of a configuration manager

Display general information about server parameters. The command is issued from a configuration manager. For descriptions of displayed fields, see “Field descriptions” on page 907.

```
query status
```

```

Server Name: SERV1
Server host name or IP address: 9.115.44.121
Server TCP/IP port number: 1699
Server URL:
Crossdefine: On
Server Password Set: Yes
Server Installation Date/Time: 10/30/2008 01:27:01 PM
Server Restart Date/Time: 10/30/2008 01:27:37 PM
Authentication: On
Password Expiration Period: 90 Day(s)
Invalid Sign-on Attempt Limit: 0
Minimum Password Length: 0
Registration: Closed
Subfile Backup: No
Inbound Sessions Disabled: DRSERVER1 DRSERVER2
Outbound Sessions Disabled: DRSERVER2 DRSERVER3
Availability: Enabled
Accounting: Off
Activity Log Retention: 5 Day(s)
Activity Summary Retention Period: 30 Day(s)
Activity Log Number of Records: 15,623
Activity Log Size: 18 M
License Audit Period: 30 Day(s)
Last License Audit: 10/30/2008 01:27:03 PM
Server License Compliance: Valid
Central Scheduler: Active
Maximum Sessions: 25
Maximum Scheduled Sessions: 12
Event Record Retention Period: 10 Day(s)
Client Action Duration: 5 Day(s)
Schedule Randomization Percentage: 25
Query Schedule Period: Client
Maximum Command Retries: Client
Retry Period: Client
Client-side Deduplication Verification Level: 0%
Scheduling Modes: Any
Active Receivers: CONSOLE ACTLOG
Configuration manager: Off
Refresh interval: 60
Last refresh date/time:
Context Messaging: Off
Table of Contents(TOC) Load Retention: 120 Minute(s)
Machine Globally Unique ID:
Archive Retention Protection: Off
Database Reporting Mode: Partial
Database Directories: /space/trible/db2/cn_db2/dbp
ath1,/space/trible/db2/cn_db
2/dbpath2
Database Total Space: 3,497,600
Database Used Space: 2,856,645.5
Database Free Space: 320,477.281
Free Space Available (MB): 17,725.25
Encryption Strength: AES
Client CPU Information Refresh Interval: 180
OutBound Replication: Enabled
Target replication server: DRSERVER
Default Replication Rule for Backup: ALL_DATA
Default Replication Rule for Archive: ALL_DATA
Default Replication Rule for Space Managed: ALL_DATA
Replication Record Retention Period: 14 Day(s)
Default Authentication: LDAP
LDAP Password Set: Yes
LDAP User: uid=tsmsserver,ou=Users,cn=aixdata

```

Field descriptions

Server Name

Specifies the name of the server.

Server host name or IP address

Specifies the server TCP/IP address.

QUERY STATUS

Server TCP/IP port number

Specifies the server port address.

Server URL

Specifies the URL address of the server that issued this command.

Crossdefine

Specifies whether another server running the **DEFINE SERVER** command automatically defines itself to this server. See the **SET CROSSDEFINE** command.

Server Password Set

Specifies whether the password has been set for the server.

Server Installation Date/Time

Specifies the date and time when the server was installed.

Server Restart Date/Time

Specifies the last date and time when the server was started.

Authentication

Specifies whether password authentication is set on or off.

Password Expiration Period

Specifies the period, in days, after which the administrator or client node password expires.

Invalid Sign-on Attempt Limit

Specifies the number of invalid sign-on attempts before a node is locked.

Minimum Password Length

Specifies the minimum number of characters for the password.

Registration

Specifies whether client node registration is open or closed.

Subfile Backup

Specifies whether subfiles can be backed up to this server, as indicated by the **SET SUBFILE** command.

Inbound Sessions Disabled

Specifies the names of servers from which server-to-server communications are not allowed. To enable inbound server sessions, use the **ENABLE SESSIONS** command.

Outbound Sessions Disabled

Specifies the names of servers to which server-to-server communications are not allowed. To enable outbound server sessions, use the **ENABLE SESSIONS** command.

Availability

Specifies whether the server is enabled or disabled.

Accounting

Specifies whether an accounting record is generated at the end of each client node session.

Activity Log Retention

Specifies the number of days information is retained in the activity log, or the size of the log.

Activity Summary Retention Period

Specifies the number of days information is retained in the SQL activity summary table.

Activity Log Number of Records

Specifies the number of records in the activity log.

Activity Log Size

Specifies the size of the activity log.

License Audit Period

Specifies the period, in days, after which the license manager automatically audits the Tivoli Storage Manager license. Additional licensing information is available by using the **QUERY LICENSE** command.

Last License Audit

Specifies the date and time when the last license audit occurred. Additional licensing information is available by using the **QUERY LICENSE** command.

Server License Compliance

Specifies whether the server is in compliance (Valid) or out of compliance (Failed) with the license terms. Use the **QUERY LICENSE** command to see what factors are causing the server to fail to comply with the license terms.

Central Scheduler

Specifies whether central scheduling is running (active or inactive).

Maximum Sessions

Specifies the maximum number of client/server sessions.

Maximum Scheduled Sessions

Specifies the maximum number of client/server sessions available for processing scheduled work.

Event Record Retention Period

Specifies the number of days central scheduler event records are retained.

Client Action Duration

Specifies the duration of the period during which the client processes the schedule defined with the **DEFINE CLIENTACTION** command.

Schedule Randomization Percentage

Specifies how much of the startup window is used for running scheduled events in client-polling mode.

Query Schedule Period

Specifies the frequency with which clients poll the server to obtain scheduled work, in client-polling mode. If the value in this field is Client, the polling frequency is determined by the client node.

Maximum Command Retries

Specifies the maximum number of times that a client scheduler tries to run a scheduled command after a failed attempt. If the value in this field is Client, the client node determines the maximum number.

Retry Period

Specifies the number of minutes between failed attempts by the client scheduler to contact the server or to run a scheduled command. If the value in this field is Client, the client node determines the number of minutes.

Client-side Deduplication Verification Level

Specifies a percentage of extents to be verified by the Tivoli Storage Manager server. The extents are created during client-side data deduplication.

QUERY STATUS

Scheduling Modes

Specifies the central scheduling modes supported by the server.

Active Receivers

Specifies the receivers for which event logging has begun.

Configuration manager

Specifies whether the server is a configuration manager.

Refresh interval

Specifies the interval that elapses before the managed server requests a refresh of any changes from a configuration manager.

Last refresh date/time

If the server is a managed server, specifies the date and time of the last successful refresh of configuration information from the configuration manager.

Context Messaging

Specifies whether context messaging is enabled or disabled.

Table of Contents (TOC) Load Retention

Specifies the approximate number of minutes that unreferenced TOC data are retained in the database.

Machine Globally Unique ID

The globally unique identifier (GUID) as of the last time that the server was started. This GUID identifies the host system to which the current server belongs.

Archive Retention Protection

Specifies if archive data retention protection is activated or deactivated.

Database Directories

Specifies the locations of the database directories.

Database Total Space

Specifies the total amount of space that is available, both used and free, in megabytes.

Database Used Space

Specifies the total amount of space currently in use, in megabytes.

Database Free Space

Specifies the space remaining in the file system where the path is located.

Encryption Strength

Indicates data encryption strength: AES or DES.

Client CPU Information Refresh Interval

Specifies the number of days that elapse between client scans for CPU information that is used for PVU estimation.

Outbound replication

Specifies whether replication processing is enabled or disabled. If outbound replication is disabled, new replication processes cannot start on the server.

Target replication server

Specifies the name of the server that is the target for node replication operations. If a target replication server does not exist, this field is blank.

Default Replication Rule for Backup

Specifies the server replication rule that applies to backup data. The following values are possible:

ALL_DATA

Replicates active and inactive backup data. The data is replicated with a normal priority.

ACTIVE_DATA

Replicates only active backup data. The data is replicated with a normal priority.

ALL_DATA_HIGH_PRIORITY

Replicates active and inactive backup data. The data is replicated with a high priority.

ACTIVE_DATA_HIGH_PRIORITY

Replicates only active backup data. The data is replicated with a high priority.

NONE

Backup data is not replicated.

Default Replication Rule for Archive

Specifies the server replication rule that applies to archive data. The following values are possible:

ALL_DATA

Replicates archive data. The data is replicated with a normal priority.

ALL_DATA_HIGH_PRIORITY

Replicates archive data. The data is replicated with a high priority.

NONE

Archive data is not replicated.

Default Replication Rule for Space Managed

Specifies the server replication rule that applies to space-managed data. The following values are possible:

ALL_DATA

Replicates space-managed data. The data is replicated with a normal priority.

ALL_DATA_HIGH_PRIORITY

Replicates space-managed data. The data is replicated with a high priority.

NONE

Space-managed data is not replicated.

Replication Record Retention Period

Specifies the number of days that replication history records are retained in the database of the source replication server.

Default Authentication

Specifies the default password authentication method: LOCAL or LDAP.

Authentication Target	Authentication Method
Tivoli Storage Manager server	LOCAL
LDAP directory server	LDAP

When you issue the **SET DEFAULTAUTHENTICATION** command, you define the resulting authentication method for all **REGISTER ADMIN** and **REGISTER NODE** commands. The default is LOCAL.

LDAP Password Set

This output field shows if a password is defined for the user ID that is named in the **SET LDAPUSER** command. The values are YES and NO. If YES, the user ID that is named in the **SET LDAPUSER** command can issue administrative commands on the LDAP namespace reserved for Tivoli Storage Manager. If NO, issue the **SET LDAPPASSWORD** command to set the password for the user ID that is named in the **SET LDAPUSER** command.

LDAP User

Specifies the user ID that is named in the **SET LDAPUSER** command. This user ID can issue administrative commands on the namespace reserved for Tivoli Storage Manager on the LDAP directory server.

Related commands

Table 277. Commands related to QUERY STATUS

Command	Description
BEGIN EVENTLOGGING	Starts event logging to a specified receiver.
DISABLE REPLICATION	Prevents outbound replication processing on a server.
DISABLE SESSIONS	Prevents new sessions from accessing Tivoli Storage Manager but permits existing sessions to continue.
ENABLE REPLICATION	Allows outbound replication processing on a server.
ENABLE SESSIONS	Resumes server activity following the DISABLE command or the ACCEPT DATE command.
END EVENTLOGGING	Ends event logging to a specified receiver.
QUERY LICENSE	Displays information about licenses and audits.
SET ACCOUNTING	Specifies whether accounting records are created at the end of each client session.
SET ACTLOGRETENTION	Specifies the number of days to retain log records in the activity log.
SET AUTHENTICATION	Specifies whether clients are prompted for a password at the beginning of a session.
SET CONTEXTMESSAGING	Specifies to turn on context messaging to debug an ANR9999D message.
SET CPUINFOREFRESH	Specifies the number of days between client scans for workstation information used for PVU estimates.
SET CROSSDEFINE	Specifies whether to cross define servers.
SET DEDUPVERIFICATIONLEVEL	Specifies the percentage of extents verified by the server during client-side deduplication.
SET DEFAULTAUTHENTICATION	Specifies the default password authentication method for any REGISTER NODE or REGISTER ADMIN commands.
SET EVENTRETENTION	Specifies the number of days to retain records for scheduled operations.
SET LDAPPASSWORD	Sets the password for the LDAPUSER.

Table 277. Commands related to QUERY STATUS (continued)

Command	Description
SET LDAPUSER	Sets the user who oversees the passwords and administrators on the LDAP directory server.
SET MAXCMDRETRIES	Specifies the maximum number of retries after a failed attempt to execute a scheduled command.
SET MAXSCHEDESESSIONS	Specifies the maximum number of client/server sessions available for processing scheduled work.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.
SET QUERYSCHEDPERIOD	Specifies the frequency for clients to obtain scheduled work, in client-polling mode.
SET RANDOMIZE	Specifies the randomization of start times within a window for schedules in client-polling mode.
SET REGISTRATION	Specifies whether users can register themselves or must be registered by an administrator.
SET RETRYPERIOD	Specifies the time between retry attempts by the client scheduler.
SET CONTEXTMESSAGING	Specifies to turn on context messaging to debug an ANR9999D message.
SET SCHEDMODES	Specifies the central scheduling mode for the server.
SET SERVERHLADDRESS	Specifies the high-level address of a server.
SET SERVERLLADDRESS	Specifies the low-level address of a server.
SET SERVERNAME	Specifies the name by which the server is identified.
SET SERVERPASSWORD	Specifies the server password.
SET SUMMARYRETENTION	Specifies the number of days to retain information for the activity summary table.
SET TOCLOADRETENTION	Specifies the number of minutes to retain information for unreferenced TOC sets.

QUERY STATUSTHRESHOLD (Query status monitoring thresholds)

Use this command to display information about status monitoring thresholds.

Status monitoring thresholds compare the defined conditions to the status monitoring server queries and inserts the results in the status monitoring table.


Multiple thresholds can be defined for an activity. For example, you can create a threshold that provides a warning status if storage pool capacity utilization is greater than 80%. You can then create another threshold that provides error status if storage pool capacity utilization is greater than 90%.

Note: If a threshold is already defined for an EXISTS condition, you cannot define another threshold with one of the other condition types.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

threshold_name

Specifies the threshold name. The name cannot exceed 48 characters in length.

Format

Specifies how the information is displayed. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the specified status thresholds.

Detailed

Specifies that complete information is displayed for the specified status thresholds.

activity

Specifies the activity for which you want to display status indicators. If you do not specify a value, information is displayed for all activities. For a list of activities, see the **DEFINE STATUSTHRESHOLD** command.

Condition

Restricts the output to only those matching the specified value. Possible values are:

EXists

Displays status thresholds where the condition equals EXISTS.

GT Displays status thresholds where the condition equals GT.

GE Displays status thresholds where the condition equals GE.

LT Displays status thresholds where the condition equals LT.

LE Displays status thresholds where the condition equals LE.

EQual

Displays status thresholds where the condition equals EQUAL.

Value

Displays thresholds that have the specified value. If you do not specify a value, information is displayed for all values. You can specify an integer from 0 to 9223372036854775807.

Status

Displays status thresholds that have the specified status value. If you do not specify a value, information is displayed for all values. Possible values are:

Normal

Displays the status thresholds that have a normal status value.

Warning

Displays the status thresholds that have a warning status value.

Error

Displays the status thresholds that have an error status value.

QUERY status threshold

Query all status thresholds by issuing the following command:

query statusthreshold

Threshold Name	Activity Name	Condition Name	Value	Report Status
ACTIVELOGCHECK	ACTIVE LOG UTILIZATION (%)	>	90	ERROR
AVGSTGPLW	AVERAGE STORAGE POOL UTILIZATION (%)	>	85	WARNING
AVGSTGPLE	AVERAGE STORAGE POOL UTILIZATION (%)	>	90	ERROR

Query status thresholds and display detailed format

Query status thresholds and display the output in detailed format, by issuing the following command:

```
query statusthreshold f=d
```

```
Threshold Name: ACTIVELOGCHECK
Activity Name: ACTIVE LOG UTILIZATION (%)
Condition Name: >
Value: 90
Report Status: ERROR
Server Name: TSMAWP24

Threshold Name: AVGSTGPLW
Activity Name: AVERAGE STORAGE POOL UTILIZATION (%)
Condition Name: >
Value: 85
Report Status: WARNING
Server Name: TSMAWP24

Threshold Name: AVGSTGPLE
Activity Name: AVERAGE STORAGE POOL UTILIZATION (%)
Condition Name: >
Value: 95
Report Status: ERROR
Server Name: TSMAWP24
```

Related commands

Table 278. Commands related to QUERY STATUSTHRESHOLD

Command	Description
"DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)" on page 337	Defines a status monitoring threshold.
"DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)" on page 451	Deletes a status monitoring threshold.
"QUERY MONITORSTATUS (Query the monitoring status)" on page 799	Displays information about monitoring alerts and server status settings.
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)" on page 1142	Specifies whether to enable client at-risk activity interval evaluation.
"SET STATUSMONITOR (Specifies whether to enable status monitoring)" on page 1144	Specifies whether to enable status monitoring.
"SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)" on page 1146	Specifies the refresh interval for status monitoring.
"SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)" on page 1147	Specifies whether to use client at-risk skipped files as failure evaluation.
"UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)" on page 1375	Changes the attributes of an existing status monitoring threshold.


QUERY STGPOOL (Query storage pools)

Use this command to display information about one or more storage pools. You can also use this command to monitor migration processes for the storage pool.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

pool_name

Specifies the storage pool to query. This parameter is optional. You can use wildcard characters to specify this name. If you do not specify a value for this parameter, all storage pools are displayed.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Pooltype

Specifies the type of storage pool to query. This parameter is optional. The default value is ANY. Possible values are:

ANY

Query primary storage pools, copy storage pools, and active-data pools.

PRimary

Query only primary storage pools.

COpy

Query only copy storage pools.

ACTIVEdata

Query only active-data pools.

Example: Display detailed random-access disk storage pool information

Display details for the storage pool named DISKPOOL. See “Field descriptions” on page 922 for field descriptions.

QUERY STGPOOL

```
query stgpool diskpool format=detailed
```

```
Storage Pool Name: DISKPOOL
Storage Pool Type: Primary
Device Class Name: DISK
Estimated Capacity: 66 G
Space Trigger Util: 0.0
  Pct Util: 0.0
  Pct Migr: 3.1
  Pct Logical: 100.0
  High Mig Pct: 90
  Low Mig Pct: 70
  Migration Delay: 0
Migration Continue: Yes
Migration Processes: 1
Reclamation Processes: 1
  Next Storage Pool:
  Reclaim Storage Pool:
Maximum Size Threshold: No Limit
  Access: Read/Write
Description:
Overflow Location:
Cache Migrated Files?:
  Collocate?: Group
Reclamation Threshold: 60
Offsite Reclamation Limit:
Maximum Scratch Volumes Allowed: 32
Number of Scratch Volumes Used: 1
  Delay Period for Volume Reuse: 0 Day(s)
  Migration in Progress?: No
  Amount Migrated (MB): 0.00
Elapsed Migration Time (seconds): 0
  Reclamation in Progress?: No
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 02/24/2005 13:57:16
Storage Pool Data Format: Native
  Copy Storage Pool(s):
  Active-data Pool(s):
  Continue Copy on Error?: No
  CRC Data: Yes
  Reclamation Type: Threshold
  Overwrite Data when Deleted: 2 Time(s)
  Deduplicate Data: No
Processes for Identifying Duplicates:
  Duplicate Data Not Stored:
  Autocopy Mode: Client
Contains data deduplicated by clients?: No
```

Example: Display detailed sequential-access disk storage pool information

Display details for the storage pool named FILEPOOL. See “Field descriptions” on page 922 for field descriptions.

```
query stgpool filepool format=detailed
```

```

Storage Pool Name: FILEPOOL
Storage Pool Type: Primary
Device Class Name: FILEC
Estimated Capacity: 66 G
Space Trigger Util: 0.0
    Pct Util: 0.0
    Pct Migr: 3.1
    Pct Logical: 100.0
    High Mig Pct: 90
    Low Mig Pct: 70
    Migration Delay: 0
Migration Continue: Yes
Migration Processes: 1
Reclamation Processes: 1
Next Storage Pool:
Reclaim Storage Pool:
Maximum Size Threshold: No Limit
    Access: Read/Write
Description:
Overflow Location:
Cache Migrated Files?:
    Collocate?: Group
Reclamation Threshold: 60
Offsite Reclamation Limit:
Maximum Scratch Volumes Allowed: 32
Number of Scratch Volumes Used: 1
Delay Period for Volume Reuse: 0 Day(s)
    Migration in Progress?: No
    Amount Migrated (MB): 0.00
Elapsed Migration Time (seconds): 0
    Reclamation in Progress?: No
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 02/24/2005 13:57:16
Storage Pool Data Format: Native
Copy Storage Pool(s):
Active-data Pool(s):
Continue Copy on Error?: No
    CRC Data: Yes
    Reclamation Type: Threshold
Overwrite Data when Deleted: 2 Time(s)
Deduplicate Data: Yes
Processes for Identifying Duplicates: 2
Duplicate Data Not Stored: 1 GB
    Autocopy Mode: Client
Contains data deduplicated by clients?: Yes

```

Example: Display detailed sequential storage pool information

Display details for the active-data sequential storage pool named FILEPOOL which uses a FILE-type device class. See “Field descriptions” on page 922 for field descriptions.

query stgpool filepool format=detailed

QUERY STGPOOL

```

Storage Pool Name: FILEPOOL
Storage Pool Type: Active-data
Device Class Name: FILEC
Estimated Capacity: 0.0 M
Space Trigger Util: 0.0
    Pct Util: 0.0
    Pct Migr: 0.0
    Pct Logical: 0.0
    High Mig Pct: 90
    Low Mig Pct: 70
    Migration Delay: 0
Migration Continue: Yes
Migration Processes: 1
Reclamation Processes: 1
Next Storage Pool:
Reclaim Storage Pool:
Maximum Size Threshold: No Limit
    Access: Read/Write
Description:
Overflow Location:
Cache Migrated Files?:
    Collocate?: Group
Reclamation Threshold: 60
Offsite Reclamation Limit:
Maximum Scratch Volumes Allowed: 99
Number of Scratch Volumes Used: 0
Delay Period for Volume Reuse: 0 Day(s)
    Migration in Progress?: No
    Amount Migrated (MB): 0.00
Elapsed Migration Time (seconds): 0
    Reclamation in Progress?: No
Last Update by (administrator): SERVER_CONSOLE
Last Update Date/Time: 05/03/2004 11:37:57
Storage Pool Data Format: Native
Copy Storage Pool(s):
Active-data Pool(s):
Continue Copy on Error:
    CRC Data: Yes
    Reclamation Type: Threshold
    Deduplicate Data: No
Processes for Identifying Duplicates:
Duplicate Data Not Stored:
Contains data deduplicated by clients?: No

```

Example: Display summary information for a specific storage pool

Display information for the storage pool named POOL1. See “Field descriptions” on page 922 for field descriptions.

```
query stgpool pool1
```

Storage Pool Name	Device Class Name	Estimated Capacity	Pct Util	Pct Migr	High Mig Pct	Low Mig Pct	Next Storage Pool
POOL1	DISK	58.5 M	0.8	0.7	90	70	POOL2

Example: Display detailed 8-mm tape storage pool information

Display details for the storage pool named 8MMPOOL. See “Field descriptions” on page 922 for field descriptions.

```
query stgpool 8mmpool format=detailed
```

```

Storage Pool Name: 8MMPPOOL
Storage Pool Type: Primary
Device Class Name: 8MMTAPE
Estimated Capacity: 0.0 M
Space Trigger Util: 0.0
  Pct Util: 0.0
  Pct Migr:
    Pct Logical: 0.0
    High Mig Pct: 90
    Low Mig Pct: 70
  Migration Delay: 0
Migration Continue: Yes
Migration Processes: 1
Reclamation Processes: 1
  Next Storage Pool:
Reclaim Storage Pool:
Maximum Size Threshold: 5 M
  Access: Read/Write
  Description: Main storage pool
  Overflow Location: Room1234/Bldg31
Cache Migrated Files?:
  Collocate?: No
  Reclamation Threshold: 60
Offsite Reclamation Limit:
Maximum Scratch Volumes Allowed: 5
Number of Scratch Volumes Used: 3
  Delay Period for Volume Reuse: 0 Day(s)
  Migration in Progress?: No
  Amount Migrated (MB): 0.00
Elapsed Migration Time (seconds): 0
  Reclamation in Progress?: No
Last Update by (administrator): ADMIN
  Last Update Date/Time: 03/22/2002 06:55:45
Storage Pool Data Format: Native
  Copy Storage Pool(s): COPYPOOL1
  Active-data Pools: ACTIVEPOOL1 ACTIVEPOOL2
  Continue Copy on Error: Yes
  CRC Data: Yes
  Reclamation Type: Threshold
  Deduplicate Data: No
Processes for Identifying Duplicates:
  Duplicate Data Not Stored:
  Autocopy Mode: Client
Contains data deduplicated by clients?: No

```

Example: Display detailed NAS2CLASS storage pool information

Display details for a storage pool, NAS2LIBPOOL. When you set up this storage pool, you set the data format to NETAPPDUMP. See “Field descriptions” on page 922 for field descriptions.

```
query stgpool nas2libpool format=detailed
```

QUERY STGPOOL

```
Storage Pool Name: NAS2
Storage Pool Name: NAS2LIBPOOL
Storage Pool Type: Primary
Device Class Name: NAS2CLASS
Estimated Capacity: 0.0 M
Space Trigger Util:
  Pct Util: 0.0
  Pct Migr:
  Pct Logical: 0.0
  High Mig Pct:
  Low Mig Pct:
  Migration Delay:
  Migration Continue:
  Migration Processes:
  Reclamation Processes:
  Next Storage Pool:
  Reclaim Storage Pool:
  Maximum Size Threshold:
    Access: Read/Write
  Description:
  Overflow Location:
  Cache Migrated Files?:
  Collocate?: Group
  Reclamation Threshold:
  Offsite Reclamation Limit:
  Maximum Scratch Volumes Allowed: 50
  Number of Scratch Volumes Used: 0
  Delay Period for Volume Reuse: 0 Day(s)
  Migration in Progress?:
  Amount Migrated (MB):
  Elapsed Migration Time (seconds):
  Reclamation in Progress?:
  Last Update by (administrator): SERVER_CONSOLE
  Last Update Date/Time: 07/08/2004 16:24:43
  Storage Pool Data Format: NetApp Dump
  Copy Storage Pool(s):
  Active-data Pool(s):
  Continue Copy on Error?: No
  CRC Data: No
  Reclamation Type:
  Deduplicate Data: No
  Processes for Identifying Duplicates:
  Duplicate Data Not Stored:
  Autocopy Mode: Client
  Contains data deduplicated by clients?: No
```

Field descriptions

Storage Pool Name

The name of the storage pool.

Storage Pool Type

The type of storage pool.

Device Class Name

The name of the device class assigned to the storage pool.

Estimated Capacity

The estimated capacity of the storage pool in megabytes (M) or gigabytes (G).

For DISK devices, this is the capacity of all volumes in the storage pool, including any volumes that are varied offline.

For sequential-access storage pools, estimated capacity is the total estimated space of all the sequential-access volumes in the storage pool, regardless of their access mode. At least one volume must be used in a

sequential-access storage pool (either a scratch volume or a private volume) to calculate estimated capacity.

For tape and FILE, the estimated capacity for the storage pool includes the following factors:

- The capacity of all the scratch volumes that the storage pool already acquired or can acquire. The number of scratch volumes is defined by the **MAXSCRATCH** parameter on the **DEFINE STGPOOL** or **UPDATE STGPOOL** command.
- The capacity of all the private volumes that are defined to the storage pool using the **DEFINE VOLUME** command.

The calculations for estimated capacity depend on the availability of the storage for the device assigned to the storage pool.

Space Trigger Util

Utilization of the storage pool, as calculated by the storage pool space trigger, if any, for this storage pool. Note that you can define space triggers for storage pools associated with DISK or FILE device types only.

For sequential access devices, space trigger utilization is expressed as a percentage of the number of used bytes on each sequential access volume relative to the size of the volume, and the estimated capacity of all existing volumes in the storage pool. It does not include potential scratch volumes. Unlike the calculation for percent utilization (Pct Util), the calculation for space trigger utilization favors creation of new private file volumes by the space trigger over usage of additional scratch volumes.

For disk devices, space trigger utilization is expressed as a percentage of the estimated capacity, including cached data. However, it excludes data that resides on any volumes that are varied offline. The value for space trigger utilization can be higher than the value for percent migration (Pct Migr) if you issue **QUERY STGPOOL** while a file creation is in progress. The value for space trigger utilization is determined by the amount of space actually allocated while the transaction is in progress. The value for percent migration represents only the space occupied by committed files. At the end of the transaction, these values are synchronized.

The value for space trigger utilization includes cached data on disk volumes. Therefore, when cache is enabled and migration occurs, the value remains the same because the migrated data remains on the volume as cached data. The value decreases only when the cached data expires or when the space that cached files occupy needs to be used for no-cached files.

Pct Util

An estimate of the utilization of the storage pool, as a percentage.

For sequential access devices, this is expressed as a percentage of the number of active bytes on each sequential access volume and the estimated capacity of all volumes in the storage pool, including the number of potential scratch volumes that may be allocated.

For disk devices, this is expressed as a percentage of the estimated capacity, including cached data and data that resides on any volumes that are varied offline. The value for Pct Util can be higher than the value for Pct Migr if you issue this command while a file creation transaction is in progress. The value for Pct Util is determined by the amount of space actually allocated (while the transaction is in progress). The value for Pct

Migr represents only the space occupied by committed files. At the end of the transaction, these values become synchronized.

The Pct Util value includes cached data on disk volumes. Therefore, when cache is enabled and migration occurs, the Pct Util value remains the same because the migrated data remains on the volume as cached data. The Pct Util value decreases only when the cached data expires or when the space that cached files occupy needs to be used for noncached files.

Pct Migr (*primary storage pools only*)

An estimate of the percentage of data in the storage pool that can be migrated. The server uses this value and the high and low migration thresholds to determine when to start and stop migration.

For random-access disk devices, this value is specified as a percentage of the value for the estimated capacity, excluding cached data, but including data on any volumes varied offline.

For sequential-access disk devices, this value is specified as a percentage of the value for the estimated capacity, including the capacity of all the scratch volumes specified for the pool. For other types of sequential-access devices, this value is the percentage of the total number of volumes in the pool that contain at least one byte of active data. The total number of volumes includes the maximum number of scratch volumes.

The Pct Util value includes cached data on a volume; the Pct Migr value excludes cached data. Therefore, when cache is enabled and migration occurs, the Pct Migr value decreases but the Pct Util value remains the same because the migrated data remains on the volume as cached data. The Pct Util value decreases only when the cached data expires or when the space that cached files occupy needs to be used for noncached files.

Pct Logical

The logical occupancy of the storage pool as a percentage of the total occupancy. Logical occupancy is space occupied by client files that may or may not be part of an aggregate. A Pct Logical value less than 100% indicates that there is vacant space within aggregates in the storage pool.

High Mig Pct (*primary storage pools only*)

The high migration threshold, which specifies when the server can begin migration for the storage pool. The server starts migration processes when capacity utilization reaches this threshold.

Low Mig Pct (*primary storage pools only*)

The low migration threshold, which specifies when the server can stop migration for the storage pool. The server stops migration processes when capacity utilization reaches this threshold.

Migration Delay (*primary storage pools only*)

The minimum number of days that a file must remain in a storage pool before the server can migrate the file to the next storage pool. For a disk storage pool, the days are counted from the time that the file was stored in the storage pool or last retrieved by a client. For a sequential access storage pool, the days are counted from the time that the file was stored in the storage pool.

Migration Continue (*primary storage pools only*)

Whether the server continues to migrate files to the next storage pool even if the files have not been in the pool for the number of days specified by the migration delay.

Migration Processes

The number of parallel processes that are used for migrating files from a random or sequential access primary storage pool.

Reclamation Processes

The number of parallel processes that are used for reclaiming the volumes in a sequential access primary or copy storage pool.

Next Storage Pool (*primary storage pools only*)

The storage pool that is the destination for data that is migrated from this storage pool.

Reclaim Storage Pool (*primary, sequential access storage pools only*)

If specified, the storage pool that is the destination for data moved from volumes during reclamation processing. If no pool is specified, by default reclamation processing moves data only among volumes within the same storage pool.

Maximum Size Threshold (*primary storage pools only*)

The maximum size of files that may be stored in the storage pool.

Access

How the data in the storage pool can be accessed.

Description

The description of the storage pool.

Overflow Location (*sequential access storage pools only*)

The location where volumes in the storage pool are stored when they are ejected from an automated library with the MOVE MEDIA command.

Cache Migrated Files? (*random access storage pools only*)

Whether caching is enabled for files migrated to the next storage pool.

Collocate? (*sequential access storage pools only*)

Whether collocation is disabled or enabled. If collocation is disabled, the value of this field is No. If collocation is enabled, the possible values are Group, Node, and Filespace.

Reclamation Threshold (*sequential access storage pools only*)

The threshold that determines when volumes in a storage pool are reclaimed. The server compares the percentage of reclaimable space on a volume to this value to determine if reclamation is necessary.

Offsite Reclamation Limit

The number of offsite volumes having their space reclaimed during reclamation for this storage pool. This field applies only when POOLTYPE=COPY.

Maximum Scratch Volumes Allowed (*sequential access storage pools only*)

The maximum number of scratch volumes that the server can request for the storage pool.

Number of Scratch Volumes Used (*sequential access storage pools only*)

The number of scratch volumes used in the storage pool.

Delay Period for Volume Reuse (*sequential access storage pools only*)

The number of days that must elapse after all files have been deleted from a volume, before the server returns the volume to scratch or reuses the volume.

Migration in Progress? (*primary storage pools only*)

Whether at least one migration process is active for the storage pool.

Amount Migrated (MB) (*primary storage pools only*)

The amount of data, in megabytes, that has been migrated, if migration is in progress. If migration is not in progress, this value indicates the amount of data migrated during the last migration. When multiple, parallel migration processes are used for the storage pool, this value indicates the total amount of data migrated by all processes.

Elapsed Migration Time (seconds) (*primary storage pools only*)

The amount of time that has elapsed since migration began, if migration is active. If migration is not active, this value indicates the amount of time required to complete the last migration. When multiple, parallel migration processes are used for the storage pool, this value indicates the total time from the beginning of the first process until the completion of the last process.

Reclamation in Progress? (*sequential access storage pools only*)

Whether a reclamation process is active for the storage pool.

Last Update by (administrator)

The name of the administrator that has defined or most recently updated the storage pool.

Last Update Date/Time

The date and time that an administrator defined or most recently updated the storage pool.

Storage Pool Data Format

The type of data format used to write data to this storage pool (for example NATIVE, NETAPPDUMP, CELERRADUMP, or NDMPDUMP).

Copy Storage Pool (s)

The copy storage pools listed here will have data simultaneously written to them when data is backed up or archived to the primary storage pool queried by this command.

Active-data Pool(s)

The active-data pools listed here will have data simultaneously written to them when data is backed up to the primary storage pool queried by this command.

Continue Copy on Error?

Whether a server should continue writing data to other copy storage pools in the list or terminate the entire transaction when a write failure occurs to one of the copy pools in the list. This field applies only to primary random-access and primary sequential-access storage pools.

CRC Data

Whether data is validated by a cyclic redundancy check (CRC) when data is transferred during data storage and retrieval on a device.

Reclamation Type

Whether volumes in this storage pool are reclaimed by threshold or by SnapLock retention date.

Overwrite Data when Deleted

The number of times data will be physically overwritten after it has been deleted from the database.

Deduplicate Data

Whether data in the storage pool will be deduplicated.

Processes for Identifying Duplicates

The number of duplicate-identification processes that are specified as the default for the storage pool. The number of duplicate-identification processes that are specified in this field might not equal the number of duplicate-identification processes that are running.

Duplicate Data Not Stored

Represents the amount of data that was removed from the storage pool by reclamation processing. The value of this field also represents the amount of storage space that was saved in this storage pool as a result of server-side data deduplication.

Autocopy Mode

Indicates whether data is written simultaneously to copy storage pools or active-data pools during client store sessions, server import processes, server data migration processes, or all three operations. The value CLIENT indicates either client store or server import operations. The value ALL indicates that simultaneous-write operations occur whenever this pool is a target for any of the eligible operations.

If the storage pool is a copy storage pool or an active-data pool or if the simultaneous-write function is disabled, this field is blank.

Contains data deduplicated by clients?

Indicates whether the storage pool contains data that was deduplicated by clients. Storage pools that contain data deduplicated by clients are not accessible for LAN-free data movement by storage agents V6.1 or earlier.

Related commands

Table 279. Commands related to QUERY STGPOOL

Command	Description
COPY ACTIVE DATA	Copies active backup data.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE STGPOOL	Deletes a storage pool from server storage.
UPDATE STGPOOL	Changes the attributes of a storage pool.

QUERY SUBSCRIBER


QUERY SUBSCRIBER (Display subscriber information)

Use this command on a configuration manager to display information about subscribers and their profile subscriptions.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

server_name

Specifies the name of a managed server for which subscription information is displayed. You can use wildcard characters to specify multiple server names. This parameter is optional. The default is all managed servers.

PROFILE

Specifies a profile name for which information is displayed. You can use wildcard characters to specify multiple profile names. This parameter is optional. The default is all profiles.

Example: List a configuration manager's profile subscriptions

Display subscriber information for all profile subscriptions to this configuration manager. See “Field descriptions” for field descriptions.

```
query subscriber
```

Subscriber	Profile name	Is current?	Last update date/time
SERVER2	DEFAULT_PROFILE	Yes	Thu, May 14, 1998 01:14:42 PM
SERVER2	SETUP	Yes	Thu, May 14, 1998 01:14:42 PM

Field descriptions

Subscriber

The name of the subscriber (managed server).

Profile name

The name of the profile.

Is current?

Whether the subscription has been refreshed with the current information associated with the profile. Possible values are:

Yes The managed server is current.

No The managed server is not current. If this field is NO after the profile has been refreshed, check the server messages for error conditions that might cause the refresh to fail.

Unknown

Either the managed server has a more recent version of the profile than the configuration manager, or the profile no longer exists on the configuration manager, but the subscription is still associated with the profile.

Last update date/time

Specifies the date and time that configuration information for the subscription was successfully distributed to the subscriber.

Related commands

Table 280. Commands related to QUERY SUBSCRIBER

Command	Description
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE SUBSCRIBER	Deletes obsolete managed server subscriptions.
DELETE SUBSCRIPTION	Deletes a specified profile subscription.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
QUERY SUBSCRIPTION	Displays information about profile subscriptions.


QUERY SUBSCRIPTION (Display subscription information)

Use this command on a managed server to display profile subscription information.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

profile_name

Specifies the name of the profile for which subscription information is displayed. You can use wildcard characters to specify multiple names. This parameter is optional. The default is all profiles.

Example: Display description information

Display subscription information for all profiles.

```
query subscription
```

Configuration manager	Profile name	Last update date/time
-----	-----	-----
SERVER1	ADMIN_INFO	Thu, May 14, 1998 01:35:13 PM
SERVER1	DEFAULT_PROFILE	Thu, May 14, 1998 01:35:13 PM
SERVER1	EMPLOYEE	Thu, May 14, 1998 01:35:13 PM

Field descriptions

Configuration manager

The name of the configuration manager.

Profile name

The name of the profile.

Last update date/time

When the most recent configuration information was successfully distributed to the subscriber.

Related commands

Table 281. Commands related to QUERY SUBSCRIPTION

Command	Description
DEFINE SUBSCRIPTION	Subscribes a managed server to a profile.
DELETE SUBSCRIBER	Deletes obsolete managed server subscriptions.
DELETE SUBSCRIPTION	Deletes a specified profile subscription.

Table 281. Commands related to QUERY SUBSCRIPTION (continued)

Command	Description
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
QUERY SUBSCRIBER	Displays information about subscribers and their subscriptions to profiles.

QUERY SYSTEM (Query the system configuration and capacity)

Use this command to obtain consolidated information about the server's configuration and capacity.

This command consolidates output from the following commands:

QUERY ASSOCIATION	QUERY OPTION
QUERY COPYGROUP	QUERY PROCESS
QUERY DB	QUERY SCHEDULE
QUERY DBSPACE	QUERY SESSION
QUERY DEVCLASS	QUERY STATUS
QUERY DOMAIN	QUERY STGPOOL
QUERY LOG	QUERY VOLHISTORY
QUERY MGMTCLASS	QUERY VOLUME

Privilege class

Any administrator can issue this command.

Syntax

►►—Query SYStem—◄◄

Example: Display consolidated system information

Issue the **QUERY SYSTEM** command to obtain consolidated system information. For sample outputs for these query commands, see the individual commands.

```
query system
```

Related commands

Table 282. Commands related to **QUERY SYSTEM**

Command	Description
QUERY ASSOCIATION	Displays the clients associated with one or more schedules.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY DB	Displays allocation information about the database.
QUERY DBSPACE	Displays information about the storage space defined for the database.
QUERY DEVCLASS	Displays information about device classes.
QUERY DOMAIN	Displays information about policy domains.
QUERY LOG	Displays information about the recovery log.
QUERY MGMTCLASS	Displays information about management classes.
QUERY OPTION	Displays information about server options.
QUERY PROCESS	Displays information about background processes.

Table 282. Commands related to QUERY SYSTEM (continued)

Command	Description
QUERY SCHEDULE	Displays information about schedules.
QUERY SESSION	Displays information about all active administrator and client sessions with Tivoli Storage Manager.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
QUERY STGPOOL	Displays information about storage pools.
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
QUERY VOLUME	Displays information about storage pool volumes.

QUERY TAPEALERTMSG (Display status of SET TAPEALERTMSG command)

Use this command to display the status of the SET TAPEALERTMSG command. You can enable or disable tape alerts. When enabled, Tivoli Storage Manager can retrieve diagnostic information from a tape or library device and display it using ANR messages. When disabled, Tivoli Storage Manager will not query a device for this information.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

►►—Query TAPEAlertmsg—◄◄

Example: Display the status of the QUERY TAPEALERTMSG command

Use the **QUERY TAPEALERTMSG** command to determine if tape alerts are to be retrieved from devices and displayed in the form of ANR messages.

```
query tapealertmsg
```

```
ANR2017I Administrator SERVER_CONSOLE issued command:
QUERY TAPEALERTMSG
ANR8960I QUERY TAPEALERTMSG: The display of Tape Alerts from SCSI
devices is Enabled.
```

Related commands

Table 283. Commands related to QUERY TAPEALERTMSG

Command	Description
SET TAPEALERTMSG	Specifies whether tape and library devices report diagnostic information to the server.

QUERY TOC (Display table of contents for a backup image)


Use this command to display directory and file information contained in the table of contents (TOC) for a specified backup image. This command does not load table of contents information into the IBM Tivoli Storage Manager database. The specified table of contents are read from a storage pool each time the **QUERY TOC** command is issued.

This command cannot be issued from the server console. If the table of contents is stored on removable media, a mount point is required and output is delayed while the storage pool volume is mounted.

Privilege class

To issue this command you must have either system privilege, policy privilege for the domain to which the node is assigned, or client owner authority over the node.

Syntax



Parameters

node_name (Required)

Specifies the name of the NAS node to which the table of contents (TOC) belongs. You cannot use wildcards to specify this name.

filespace_name (Required)

Specifies the name of the file space to which the table of contents belongs. The file space name you specify cannot contain wildcard characters.

CREATIONDate

Specifies the creation date of the backup image for which the table of contents is to be displayed. This parameter is optional. If you specify **CREATIONDATE**, you must also specify **CREATIONTIME**. If you do not specify these parameters, the contents of the latest backup image for the specified node and file space will be displayed, provided that this image has a table of contents. You can only specify the creation date as the following:

Value	Description	Example
MM/DD/YYYY	A specific date	05/15/2002

This specifies that you want to display the contents of the backup image created on this date. You can obtain this date from the output of the **QUERY NASBACKUP** command.

CREATIONTime

Specifies the creation time of the backup image for which the table of contents is to be displayed. This parameter is optional. If you specify **CREATIONTIME**, you

must also specify **CREATIONDATE**. If you do not specify these parameters, the contents of the latest backup image for the specified node and file space will be displayed, provided that this image has a table of contents. You can only specify the creation time as the following:

Value	Description	Example
HH:MM:SS	A specific time on the specified creation date.	10:30:08

This specifies that you want to display the contents of the backup image created on this time for the specified date. You can obtain this time from the output of the **QUERY NASBACKUP** command.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed for the files.

Detailed

Specifies that complete information is displayed for the files, including the hexadecimal representation of each file or directory name.

Example: Display detailed table of contents information for a specific node

Use the **QUERY TOC** command to display information in the table of contents belonging to NAS node NETAPP in the file space /vol/vol1 created on 12/06/2002 at 11:22:46. Specify a detailed format.

```
query toc netapp /vol/vol1 creationdate=12/06/2002 creationtime=11:22:46
format=detailed
```

```
Objects in the image backed up on 12/06/2002 11:22:46
for filesystem /vol/vol1 in node NETAPP:

      Object Name: /.etc
Hexadecimal Object Name: 2f657463
      Object Type: Directory
      Object Size: 4,096
Last data Modification Date/Time: 07/31/2002 14:21:19

      Object Name: /.etc/oldmaps/ndmp
Hexadecimal Object Name: 2f6574632f6f6c646d6170
                          732f6e646d70
      Object Type: Directory
      Object Size: 4,096
Last data Modification Date/Time: 07/31/2002 14:21:19

      Object Name: /.etc/oldmaps/ndmp/TSM
                          /vol/vol1/3df0e8fd
Hexadecimal Object Name: 2f6574632f6f6c646d6170
                          732f6e646d702f54534d2
                          02f766f6c2f766f6c312f3
                          364663065386664
      Object Type: File
      Object Size: 36,864
Last data Modification Date/Time: 12/06/2002 11:14:22
```

Related commands

Table 284. Commands related to QUERY TOC

Command	Description
BACKUP NODE	Backs up a network-attached storage (NAS) node.
QUERY NASBACKUP	Displays information about NAS backup images.
RESTORE NODE	Restores a network-attached storage (NAS) node.


QUERY VIRTUALFSMAPPING (Query a virtual file space mapping)

Use this command to query a virtual file space mapping definition.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

node_name

Specifies the client node to which the virtual file space belongs. You can use wildcard characters to specify this name. This parameter is optional. The default is all client node names. You must specify a value for this parameter if you specify a virtual file space name.

virtual_file_space_name

Specifies the name of the virtual file space mappings to be queried. You can use wildcard characters to specify this name. This parameter is optional. If a value is not specified, all virtual file space mappings are queried. Virtual file space mapping names are case sensitive. Use the **QUERY VIRTUALFSMAPPING** command to determine the correct capitalization for the virtual file space mapping to be queried.

Example: Display virtual file spaces for a specific node

Display the currently defined virtual file spaces for node NAS1. See “Field descriptions” for field descriptions.

```
query virtualfsmapping nas1
```

Node Name	Virtual Filespace Mapping Name	Filespace Name	Path	Hexadecimal Path?
NAS1	/mikesdir	/vol/vol2	/mikes	No
NAS1	/tmpdir	/vol/vol1	/tmp	No
NAS1	/nonASCIIDir	/vol/vol3	2f73657276657231	Yes

Field descriptions

Node Name

Specifies the name of the client node.

Virtual Filespace Mapping Name

Specifies the name of the virtual file space mapping.

Filespace Name

The name of the file space that belongs to the node.

File space names can be in a different code page or locale than the server. If they are, the names in the Administration Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, Tivoli Storage Manager converts the name to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

Path Specifies the path to the client node.

Hexadecimal Path

Indicates whether the path is hexadecimal.

Related commands

Table 285. Commands related to QUERY VIRTUALFSMAPPING

Command	Description
DEFINE VIRTUALFSMAPPING	Define a virtual file space mapping.
DELETE VIRTUALFSMAPPING	Delete a virtual file space mapping.
UPDATE VIRTUALFSMAPPING	Update a virtual file space mapping.

QUERY VOLHISTORY (Display sequential volume history information)

Use this command to display sequential volume history information. To save sequential volume history information to one or more files, use the BACKUP VOLHISTORY command.


Use the VOLUMEHISTORY server option to specify one or more volume history files. After the server is restarted, Tivoli Storage Manager updates volume information in both the database and the files.

Use the QUERY BACKUPSET command to query specified backup set information.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

BEGINDate

Specifies that you want to display information beginning with records created on the specified date. This parameter is optional. The default is the earliest date for which history information exists.

You can specify the date using one of the values below:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY

Value	Description	Example
TODAY- <i>days</i> or - <i>days</i>	The current date minus days specified	TODAY-7 or -7. To display information beginning with records created a week ago, specify BEGINDATE=TODAY-7 or BEGINDATE=-7.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDDate

Specifies that you want to display information ending with records created on the specified date. This parameter is optional. The default is the current date.

You can specify the date using one of the values below:

Value	Description	Example
MM/DD/YYYY	A specific date	09/15/1998
TODAY	The current date	TODAY
TODAY- <i>days</i> or - <i>days</i>	The current date minus days specified. The maximum number of days is 9999.	TODAY-1 or -1. To display records created up to yesterday, specify ENDDATE=TODAY-1 or ENDDATE=-1.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies that you want to display information beginning with records created at the specified time. This parameter is optional. The default is midnight (00:00:00).

You can specify the time using one of the values below:

QUERY VOLHISTORY

Value	Description	Example
<i>HH:MM:SS</i>	A specific time on the specified begin date	12:33:28
NOW	The current time on the specified begin date	NOW
NOW+ <i>HH:MM</i> or + <i>HH:MM</i>	The current time plus hours and minutes on the specified begin date	NOW+03:00 or +03:00. If you issue this command at 9:00 with BEGINTIME=NOW+03:00 or BEGINTIME=+03:00, Tivoli Storage Manager displays records with a time of 12:00 or later on the begin date.
NOW- <i>HH:MM</i> or - <i>HH:MM</i>	The current time minus hours and minutes on the specified begin date	NOW-03:30 or -03:30. If you issue this command at 9:00 with BEGINTIME=NOW-03:30 or BEGINTIME=-03:30, Tivoli Storage Manager displays records with a time of 5:30 or later on the begin date.

ENDTime

Specifies that you want to display information ending with records created at the specified time on the end date. This parameter is optional. The default is the current time.

You can specify the time using one of the values below:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time on the specified end date	10:30:08
NOW	The current time on the specified end date	NOW
NOW+ <i>HH:MM</i> or + <i>HH:MM</i>	The current time plus hours and minutes on the specified end date	NOW+03:00 or +03:00. If you issue this command at 9:00 with ENDTIME=NOW+03:00 or ENDTIME=+03:00, Tivoli Storage Manager displays records with a time of 12:00 or later on the end date.
NOW- <i>HH:MM</i> or - <i>HH:MM</i>	The current time minus hours and minutes on the specified end date	NOW-03:30 or -03:30 If you issue this command at 9:00 with ENDTIME=NOW-3:30 or ENDTIME=-3:30, Tivoli Storage Manager displays records with a time of 5:30 or earlier on the end date.

Type

Specifies the type of records to display from the volume history file. This parameter is optional. The default is ALL. Possible values are:

ALL

Specifies all records.

BACKUPSET

Specifies to display only information about backup set volumes.

DBBackup

Specifies to display only records that contain information about full and incremental database backup volumes, that is with the volume types of BACKUPFULL and BACKUPINCR.

DBRpf

Specifies to display only records that contain information about full and incremental database backup volumes and recovery plan file object volumes (volume types of BACKUPFULL, BACKUPINCR, and RPFILe).

DBSnapshot

Specifies to display only records that contain information about volumes used for database snapshot backups.

EXPort

Specifies only records that contain information about export volumes.

REMOte

Specifies to display only records that contain information about volumes used by library clients.

RPFILe

Specifies to display only records that contain information about file objects of a recovery plan that are saved on a target server and that were created assuming database full and incremental backups. The parameter displays only records about recovery plan files that are saved on another Tivoli Storage Manager server by using the server-to-server virtual volume function for Tivoli Storage Manager.

RPFSnapshot

Specifies to display only records that contain information about file objects of a recovery plan that are saved on a target server and that were created assuming database snapshot backups. **RPFSnapshot** only displays records about recovery plan files that are saved on another Tivoli Storage Manager server by using the server-to-server virtual volume function for Tivoli Storage Manager.

STGDelete

Specifies only records that contain information about deleted sequential storage pool volumes.

STGNew

Specifies only records that contain information about new sequential access storage volumes.

STGReuse

Specifies only records that contain information about reused sequential storage pool volumes.

Example: Display volume history information

Display volume history information stored in the Tivoli Storage Manager database. See “Field descriptions” on page 944 for field descriptions. Issue the command:

```
query volhistory
```

QUERY VOLHISTORY

```
Date/Time: 02/25/2011 18:28:06
Volume Type: STGNEW
Backup Series: 0
Backup Operation: 0
Volume Seq: 0
Device Class: FILE
Volume Name: /adsmfct/server/prv011
Volume Location:
Command:
Database Backup ID High: 0
Database Backup ID LOW: 21,968,899
Database Backup Home Position: 0
Database Backup HLA: /NODE0000/
Database Backup LLA: FULL_BACKUP.20100121103759.1
Database Backup Total Data Bytes (MB): 12,494.88
Database Backup total Log Bytes (MB): 220.11
Database Backup Block Num High: -1
Database Backup Block Num Low: -1
Database Backup Stream Id: 1
Database Backup Volume Sequence for Stream: 2
```

Note: The volume history file will contain additional fields that do not appear in the query output. These fields are specific to database backup and restore support. They are not intended for use or modification by Tivoli Storage Manager administrators. The fields will be bracketed with a message indicating these are for Tivoli Storage Manager internal use only and not meant to be modified.

Field descriptions

Date/Time

The date and time that the volume was created.

Volume Type

The type of volume:

BACKUPFULL

Full database backup volume.

BACKUPINCR

Incremental database backup volume.

BACKUPSET

Client backup set volume.

DBSNAPSHOT

Snapshot database backup volume.

EXPORT

Export volume.

REMOTE

A volume used on the library client, which is the Tivoli Storage Manager server named in the Volume Location field. See the volume history on the server that is the library client to get details about how the volume is used.

RPFIL

Recovery plan file object volume created assuming full and incremental database backups.

RPFSSnapshot

Recovery plan file object volume created assuming snapshot database backups.

STGDELETE

Deleted sequential access storage pool volume.

STGNEW

Added sequential access storage pool volume.

STGREUSE

Reused sequential access storage pool volume.

Backup Series

The value of this field depends on the volume type:

- For BACKUPFULL or BACKUPINCR volume types: the backup series identifier.
- For the DBSNAPSHOT volume type: the identifier of the backup series that is associated with the DBSNAPSHOT entry.
- For the RPFIL volume type: the identifier of the backup series that is associated with the RPFIL entry.
- For the RPFNSNAPSHOT volume type: the identifier of the backup series that is associated with the RPFNSNAPSHOT entry.
- For BACKUPSET volume types: this field is blank.
- For all other volume types: always 0.

A backup series is a full backup and all incremental backups that apply to that full backup. Another series begins with the next full backup of the database.

Backup Operation

For BACKUPFULL or BACKUPINCR volume types: the operation number of this backup volume within the backup series. The full backup within a backup series is operation 0. The first incremental backup for that full backup is operation 1, the second incremental backup is operation 2, and so on.

For DBSNAPSHOT volume types: the operation number of this DBSNAPSHOT volume within the DBSNAPSHOT series.

For all other volume types: always 0.

This field is blank when the volume type is BACKUPSET.

Volume Seq

The sequence or position of the volume within the backup series.

- For BACKUPFULL or BACKUPINCR volume types: the sequence, or position, of the volume within the backup series. Volume sequence 1 identifies the first volume used for the first operation (a full backup), and so on. For example, if the full backup occupies three volumes, these volumes are identified as volume sequence 1, 2, and 3, respectively. The first volume of the next operation (the first incremental backup) is then volume sequence 4.
- For BACKUPSET volume types: the sequence, or position, of the volume within the BACKUPSET series.
- For DBSNAPSHOT volume types: the sequence, or position, of the volume within the DBSNAPSHOT series. Volume sequence 1 identifies the first volume used for the first DBSNAPSHOT operation, and so on.
- For EXPORT volume types: the sequence number of the volume when it was used for exporting data.

QUERY VOLHISTORY

- For RPFIL volume types: the value of this field is always one (1).
- For all other volume types: always 0.

Device Class

The name of the device class associated with this volume.

Volume Name

The name of the volume.

Volume Location

The location of the volume. This information is available only for the following volume types:

BACKUPFULL
BACKUPINCR
EXPORT
REMOTE
RPFIL

For the volume type of REMOTE, this location field is the server name of the library client that owns this volume.

For the volume type of RPFIL, this location field is the server name defined in the device class definition used by the PREPARE command when the DEVCLASS parameter is specified.

Command

When the volume type is EXPORT or BACKUPSET and the volume sequence is 1 (for example, the first volume), this field shows the command that was used to generate the volume. If the EXPORT or BACKUPSET is on more than one volume, the command is displayed with the first volume but not with any of the other volumes.

For any volume type other than EXPORT or BACKUPSET, this field is blank.

Note: The following fields are not used by Tivoli Storage Manager servers that are V6.3 or later. However, the fields are displayed for compatibility with earlier releases.

- Database Backup ID High
- Database Backup ID Low
- Database Backup Home Position
- Database Backup HLA
- Database Backup LLA
- Database Backup Total Data Bytes (MB)
- Database Backup Total Log Bytes (MB)
- Database Backup Block Num High
- Database Backup Block Num Low

Related commands

Table 286. Commands related to QUERY VOLHISTORY

Command	Description
BACKUP VOLHISTORY	Records volume history information in external files.
DELETE VOLHISTORY	Removes sequential volume history information from the volume history file.

Table 286. Commands related to QUERY VOLHISTORY (continued)

Command	Description
PREPARE	Creates a recovery plan file.
QUERY RPFIL	Displays information about recovery plan files.
QUERY BACKUPSET	Displays backup sets.
UPDATE VOLHISTORY	Adds or changes location information for a volume in the volume history file.


QUERY VOLUME (Query storage pool volumes)

Use this command to display information about one or more storage pool volumes.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

volume_name

Specifies the volume to query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a name, all storage pool volumes are included in the query.

ACcESS

Specifies that output is restricted by volume access mode. This parameter is optional. You can specify multiple access modes by separating the modes with commas and no intervening spaces. If you do not specify a value for this parameter, output is not restricted by access mode. Possible values are:

READWrite

Display volumes with an access mode of READWRITE. Client nodes and server processes can read from and write to files stored on the volumes.

READOnly

Display volumes with an access mode of READONLY. Client nodes and server processes can only read files stored on the volumes.

UNAVailable

Display volumes with an access mode of UNAVAILABLE. Neither client nodes nor server processes can access files stored on the volumes.

OFFsite

Display copy storage pool volumes with an access mode of OFFSITE. The volumes are at offsite locations from which they cannot be mounted.

DESTROYED

Display primary storage pool volumes with an access mode of DESTROYED. The volumes have been designated as permanently damaged.

Status

Specifies that output is restricted by volume status. This parameter is optional. You can specify multiple status values by separating values with commas and no intervening spaces. If you do not specify a value for this parameter, output is not restricted by volume status. Possible values are:

ONline

Display random access volumes that are available to the server.

OFFline

Display random access volumes that are not available to the server.

EMPTy

Display sequential access volumes that have no data.

PENding

Display volumes with a status of PENDING. These could be sequential-access volumes from which all files have been deleted, but for which the time specified by the REUSEDELAY parameter on the **DEFINE STGPOOL** command has not elapsed. These volumes could also be random-access disk volumes that were deleted, but that still contain discarded data that is waiting to be shredded. After the data is shredded, the volume will be physically deleted.

FILLing

Display sequential access volumes that the server has written to but has not yet filled to capacity.

FULL

Display sequential access volumes that the server has filled.

STGPool

Specifies the storage pool to include in the query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a storage pool name, all storage pools are included in the query.

DEVclass

Specifies the device class to include in the query. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a device class name, all devices are included in the query.

Format

Specifies how the information is displayed. This parameter is optional. The default value is STANDARD. Possible values are:

Standard

Specifies that partial information is displayed.

Detailed

Specifies that complete information is displayed.

Example: List all file storage pool volumes

Display information on all storage pool volumes with the device class name of FILE. See “Field descriptions” on page 952 for field descriptions.

query volume devclass=file

Volume Name	Storage Pool Name	Device Class Name	Estimated Capacity	Pct Util	Volume Status
/FCT/SERVER/COV011	COPYSTG	FILE	0.0 M	0.0	Pending
/FCT/SERVER/COV012	COPYSTG	FILE	0.0 M	0.0	Empty
/FCT/SERVER/COV013	COPYSTG	FILE	0.0 M	0.0	Empty
/FCT/SERVER/PRV011	PRIMESTG	FILE	0.0 M	0.0	Empty
/FCT/SERVER/PRV012	PRIMESTG	FILE	0.0 M	0.0	Empty

Example: Display detailed information about a specific storage pool volume

Display details about the storage pool volume named */fct/server/cov011*. See “Field descriptions” on page 952 for field descriptions.

query volume cov011 format=detailed

```

Volume Name: /FCT/SERVER/COV011
Storage Pool Name: COPYSTG
Device Class Name: DISK
Estimated Capacity: 10.0 M
Scaled Capacity Applied:
Pct Util: 6.7
Volume Status: On-line
Access: Read/Write
Pct. Reclaimable Space: 3.2
Scratch Volume?: Yes
In Error State?: No
Number of Writable Sides: 1
Number of Times Mounted: 11
Write Pass Number: 1
Approx. Date Last Written: 04/14/1998 16:17:26
Approx. Date Last Read: 04/01/1998 13:26:18
Date Became Pending:
Number of Write Errors: 0
Number of Read Errors: 0
Volume Location:
Volume is MVS Lanfree Capable: No
Last Update by (administrator): COLLIN
Last Update Date/Time: 05/01/1998 14:07:27
Begin Reclaim Period:
End Reclaim Period:
Logical Block Protection:
Drive Encryption Key Manager:
    
```

Example: Display detailed information about a storage pool volumes with a specific device class

Display details about a volume in a storage pool with a device class name of FILECLASS. See “Field descriptions” on page 952 for field descriptions.

query volume devclass=fileclass format=detailed

```

Volume Name: /WORM_FILESYS/0000000E.BFS
Storage Pool Name: FILEPOOL
Device Class Name: FILECLASS
Estimated Capacity: 2.0 G
Scaled Capacity Applied:
  Pct Util: 0.0
Volume Status: Filling
  Access: Read/Write
Pct. Reclaimable Space: 0.0
  Scratch Volume?: Yes
  In Error State?: No
Number of Writable Sides: 1
Number of Times Mounted: 1
  Write Pass Number: 1
Approx. Date Last Written: 03/22/2004 15:23:46
Approx. Date Last Read: 03/22/2004 15:23:46
Date Became Pending:
Number of Write Errors: 0
Number of Read Errors: 0
Volume Location:
Volume is MVS Lanfree Capable: No
Last Update by (administrator):
  Last Update Date/Time: 03/22/2004 15:23:46
  Begin Reclaim Period: 03/22/2005
  End Reclaim Period: 04/22/2005
Logical Block Protection:
Drive Encryption Key Manager:

```

Example: Display detailed information about a specific storage pool volume

Display details about a storage pool volume named 000642. The volume is in a storage pool associated with a 3592 device class. See “Field descriptions” on page 952 for field descriptions.

query volume 000642 format=detailed

```

Volume Name: 000642
Storage Pool Name: 3592POOL
Device Class Name: 3592CLASS
Estimated Capacity: 2.0 G
Scaled Capacity Applied:
  Pct Util: 0.0
Volume Status: Filling
  Access: Read/Write
Pct. Reclaimable Space: 0.0
  Scratch Volume?: Yes
  In Error State?: No
Number of Writable Sides: 1
Number of Times Mounted: 1
  Write Pass Number: 1
Approx. Date Last Written: 03/22/2004 15:23:46
Approx. Date Last Read: 03/22/2004 15:23:46
Date Became Pending:
Number of Write Errors: 0
Number of Read Errors: 0
Volume Location:
Volume is MVS Lanfree Capable: No
Last Update by (administrator):
  Last Update Date/Time: 03/22/2004 15:23:46
  Begin Reclaim Period: 03/22/2005
  End Reclaim Period: 04/22/2005
Logical Block Protection: Yes
Drive Encryption Key Manager: Tivoli Storage Manager

```

Field descriptions

Volume Name

The name of the storage pool volume.

Storage Pool Name

The storage pool to which the volume is defined.

Device Class Name

The device class assigned to the storage pool.

Estimated Capacity

The estimated capacity of the volume, in megabytes (M), gigabytes (G), or terabytes (T).

For DISK devices, this value is the capacity of the volume.

For sequential access devices, this value is an estimate of the total space available on the volume, based on the device class.

Scaled Capacity Applied

The percentage of capacity to which a volume is scaled. For example, a value of 20 for a volume whose maximum capacity is 300 GB indicates that the volume can only store 20 percent of 300 GB, or 60 GB. This attribute applies only to IBM 3592 devices.

Pct Util

An estimate of the utilization of the volume. The utilization includes all space occupied by both files and aggregates, including empty space within aggregates.

For DISK volumes, the utilization also includes space occupied by cached data.

Volume Status

The status of the volume.

Access

Whether the volume is available to the server.

Pct. Reclaimable Space (*sequential access volumes only*)

The amount of space on this volume that can be reclaimed because data has expired or been deleted. This value is compared to the reclamation threshold for the storage pool to determine if reclamation is necessary. Reclaimable space includes empty space within aggregates.

When determining which volumes in a storage pool to reclaim, the Tivoli Storage Manager server first determines the reclamation threshold. The reclamation threshold is indicated by the value of the **THRESHOLD** parameter on the **RECLAIM STGPOOL** command or, if that value was not specified, the value of the **RECLAIM** parameter in a storage pool definition. The server then examines the percentage of reclaimable space for each volume in the storage pool. If the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool, the volume is a candidate for reclamation.

For example, suppose storage pool FILEPOOL has a reclamation threshold of 70 percent. This value indicates that the server can reclaim any volume in the storage pool that has a percentage of reclaimable space that is greater than 70 percent. The storage pool has three volumes:

- FILEVOL1 with 65 percent reclaimable space
- FILEVOL2 with 80 percent reclaimable space

- FILEVOL3 with 95 percent reclaimable space

When reclamation begins, the server compares the percent of reclaimable space for each volume with the reclamation threshold of 70 percent. In this example, FILEVOL2 and FILEVOL3 are candidates for reclamation because their percentages of reclaimable space are greater than 70.

For volumes belonging to a SnapLock storage pool, the value is displayed but is not used.

Scratch Volume? (*sequential access volumes only*)

Whether this volume will be returned to scratch when the volume becomes empty.

In Error State?

Whether the volume is in an error state. The server cannot write to volumes in an error state.

Number of Writable Sides

This information is reserved for IBM Tivoli Storage Manager.

Number of Times Mounted

The number of times that the server has opened the volume for use. The number of times that the server has opened the volume is not always the same as the number of times that the volume has been physically mounted in a drive. After a volume is physically mounted, the server can open the same volume multiple times for different operations, for example for different client backup sessions.

Write Pass Number (*sequential access volumes only*)

The number of times the volume has been written to from the beginning to the end.

Approx. Date Last Written

The approximate date on which the volume was last written.

Approx. Date Last Read

The approximate date on which the volume was last read.

Date Became Pending

The date that the status of the volume changed to pending.

Number of Write Errors

The number of writing errors that have occurred on the volume.

Number of Read Errors

The number of reading errors that have occurred on the volume.

Volume Location

The location of the volume.

Volume is MVS Lanfree Capable

Whether the volume is LAN-free capable. A LAN-free capable volume is one that has been defined and used (at least once) by the Tivoli Storage Manager Version 5 Release 2 or later z/OS data manager server.

Last Update by (administrator)

The administrator that defined or most recently updated the volume.

Last Update Date/Time

When the volume was defined or most recently updated.

Begin Reclaim Period

Represents the date after which the Tivoli Storage Manager server will

QUERY VOLUME

begin reclaiming this volume, but not later than the date represented by the end reclaim period. If, when the reclaim period begins, there are files on the volume that have not expired, they will be moved to a new WORM volume during reclamation processing. This field displays a date only if this volume is in a storage pool for which the value of the RECLAMATIONTYPE parameter is SNAPLOCK.

If more than one archive is stored on the same volume, the start of the volume's reclamation period is based on the date of the most recent archive. For SnapLock volumes, the RETVer parameter of the **DEFINE COPYGROUP** command determines how long an archive is stored. If RETVer is set to 100 days, the volume's reclamation period will start 100 days after the first archive is stored on it. If a second archive is stored on the same volume, the reclamation start date will be adjusted to 100 days after the new archive is stored. If the RETVer value is changed after the first archive is stored, the latest reclamation date will apply for all of the archives on the volume. For example, assume RETVer is set to 100 for an initial archive, but is then changed to 50. If a second archive is stored on the volume three days after the first, the reclamation period will not start until 100 days after the first archive was stored.

End Reclaim Period

Represents the date by which the Tivoli Storage Manager must complete reclamation processing on this volume to ensure continued protection of the data. It also represents the Last Access Date physical file attribute in the NetApp Filer, which prevents the file from being deleted until after that date. This field displays a date only if this volume is in a storage pool for which the value of the RECLAMATIONTYPE parameter is SNAPLOCK.

Logical Block Protection

Specifies whether logical block protection is enabled for the volume. You can use logical block protection only with the following types of drives and media:

- IBM LTO5 and later
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later
- Oracle StorageTek T10000C drives

Drive Encryption Key Manager

The drive encryption key manager. This field applies only to volumes in a storage pool associated with a device type of 3592, LTO, or ECARTRIDGE.

Related commands

Table 287. Commands related to **QUERY VOLUME**

Command	Description
DEFINE DEVCLASS	Defines a device class.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.
DELETE VOLUME	Deletes a volume from a storage pool.
UPDATE DEVCLASS	Changes the attributes of a device class.
UPDATE VOLUME	Updates the attributes of storage pool volumes.

Table 287. Commands related to QUERY VOLUME (continued)

Command	Description
VARY	Specifies whether a disk volume is available to the server for use.

QUIT (End the interactive mode of the administrative client)

Use this command to end an administrative client session in interactive mode.

You cannot use the **QUIT** command from the `SERVER_CONSOLE` administrative ID, or the console, batch, or mount modes of the administrative client.

Privilege class

Any administrator can issue this command.

Syntax

▶▶—QUIT—▶▶

Parameters

None.

Example: End an interactive administrative client session

End an administrative client session in the interactive mode.

```
quit
```

Related commands

None.

RECLAIM STGPOOL (Reclaim volumes in a sequential-access storage pool)

Use this command to reclaim volumes in a sequential-access storage pool. Reclamation makes the fragmented space on volumes usable again by moving any remaining files from one volume to another volume. This action makes the original volume available for reuse. Reclamation does not move inactive versions of backup data from volumes in active-data pools. For storage pools defined with `RECLAMATIONTYPE=SNAPLOCK`, this command also deletes empty WORM FILE volumes that have exceeded their reclaim period.

For storage pools that use Write Once Read Many (WORM) devices, reclamation is not necessary because WORM volumes are not reusable. However, you can run reclamation to consolidate data onto fewer volumes. Volumes emptied by reclamation can be checked out of the library, freeing slots for new volumes.

This command can be used only with sequential access storage pools. The storage pool data format cannot be `NETAPPDUMP`, `CELERRADUMP`, or `NDMPDUMP`.

Use this command only if you are not going to use automatic reclamation for the storage pool. You can set the `RECLAIM` parameter of the storage pool definition to 100 to prevent automatic reclamation from running. This command honors the values of the `RECLAIMPROCESS` and `RECLAIMSTGPOOL` attributes of the storage pool definition. This command also honors the values of the `OFFSITERECLAIMLIMIT` and `RECLAIM` parameters of the storage pool definition, if not overridden by the `OFFSITERECLAIMLIMIT` and `THRESHOLD` command parameters. Use caution when setting the value of the `RECLAIM` parameter of an active-data pool to 100. Doing so prevents the removal of inactive versions of backup data.

This command creates one or more reclamation processes that can be canceled with the `CANCEL PROCESS` command. The number of processes is limited by the `RECLAIMPROCESS` attribute of the storage pool definition. To display information about background processes, use the `QUERY PROCESS` command.

Remember: Issuing this command removes duplicate data in a primary storage pool, copy storage pool, or active-data pool that has been set up for data deduplication.


Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool being reclaimed and the reclaim storage pool, if applicable.

Syntax

```

>> RECLaim STGpool pool_name [THreshold=number]
    [DURATION=minutes] [Wait=No] [Wait=No] [Yes]
  
```



Parameters

pool_name (Required)

Specifies the storage pool in which volumes are to be reclaimed.

Duration

Specifies the maximum number of minutes that the reclamation runs before being automatically canceled. You can specify a number from 1 to 9999. This parameter is optional.

After the specified number of minutes elapses, the next time the server checks the reclamation process the server stops the reclamation process. The server checks the reclamation process when the server mounts another eligible volume from the storage pool being reclaimed, and when the server begins to reclaim a new batch of files from the currently mounted volume. As a result, the reclamation can run longer than the value you specified for this parameter.

Until the server checks the reclamation process, there is no indication that the duration period has expired. When the server stops the reclamation process, the server issues message ANR4927W: Reclamation terminated for volume xxx - duration exceeded.

If you do not specify this parameter, the process stops only when no more volumes meet the threshold.

If you specify a duration value for reclamation of a copy storage pool with off-site volumes, you might cause the reclamation to terminate before any volumes are reclaimed. In most situations when initiating reclamation for a copy storage pool with off-site volumes, consider limiting the number of off-site volumes to be reclaimed rather than limiting the duration. For details, see the **OFFSITERECLAIMLIMIT** parameter.

Threshold

Specifies the percentage of reclaimable space that a volume must have in order to be eligible for reclamation. Reclaimable space is the amount of space occupied by files that are expired or deleted from the Tivoli Storage Manager database. Reclaimable space also includes unused space.

You can specify a number from 1 to 99. This parameter is optional. If not specified, the RECLAIM attribute of the storage pool definition is used.

When determining which volumes in a storage pool to reclaim, the Tivoli Storage Manager server first determines the reclamation threshold. The server then examines the percentage of reclaimable space for each volume in the storage pool. If the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool, the volume is a candidate for reclamation.

For example, suppose storage pool FILEPOOL has a reclamation threshold of 70 percent. This value indicates that the server can reclaim any volume in the storage pool that has a percentage of reclaimable space that is greater than 70 percent. The storage pool has three volumes:

- FILEVOL1 with 65 percent reclaimable space
- FILEVOL2 with 80 percent reclaimable space
- FILEVOL3 with 95 percent reclaimable space

When reclamation begins, the server compares the percent of reclaimable space for each volume with the reclamation threshold of 70 percent. In this example, FILEVOL2 and FILEVOL3 are candidates for reclamation because their percentages of reclaimable space are greater than 70. To determine the percentage of reclaimable space for a volume, issue the **QUERY VOLUME** command and specify **FORMAT=DETAILED**. The value in the field Pct. Reclaimable Space is the percentage of reclaimable space for the volume.

Specify a value of 50 percent or greater for this parameter so that files stored on two volumes can be combined into a single target volume.

OFFSITERECLAIMLimit

Specifies the maximum number of off-site storage pool volumes that the server tries to reclaim. This parameter is valid only for copy storage pools. Limiting the number of off-site volumes can prevent the server from spending more time analyzing volumes rather than reclaiming them. You can specify a number from 0 to 99999. This parameter is optional. If not specified, the **OFFSITERECLAIMLIMIT** attribute of the storage pool definition is used.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is No. Possible values are:

No Specifies that the server processes this command in the background.

You can continue with other tasks while the command is being processed. Messages created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the **CANCEL PROCESS** command. If you cancel this process, some files might have already been moved to new volumes before the cancellation.

Yes

Specifies that the server processes this command in the foreground. The operation must complete before you can continue with other tasks. Output messages are displayed to the administrative client when the operation completes. Messages are also displayed either in the activity log or the server console, or both, depending on where the messages are logged.

Note: You cannot specify WAIT=YES from the server console.

Example: Reclaim volumes in a sequential-access storage pool

Reclaim volumes in the storage pool named TAPEPOOL. Specify that the server end the reclamation as soon as possible after 60 minutes.

```
reclaim stgpool tapepool duration=60
```

Related commands

Table 288. Commands related to RECLAIM STGPOOL

Command	Description
CANCEL PROCESS	Cancels a background server process.
MIGRATE STGPOOL	Performs migration for the storage pool.
MOVE DRMEDIA	Moves DRM media on-site and off-site.

RECLAIM STGPOOL

Table 288. Commands related to RECLAIM STGPOOL (continued)

Command	Description
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY PROCESS	Displays information about background processes.
QUERY STGPOOL	Displays information about storage pools.


RECONCILE VOLUMES (Reconcile differences in the virtual volume definitions)

Issue this command from the source server to reconcile differences between virtual volume definitions on the source server and archive files on the target server. Tivoli Storage Manager finds all volumes of the specified device class on the source server and all corresponding archive files on the target server. The target server inventory is also compared to the local definition for virtual volumes to see if inconsistencies exist.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

device_class_name

Specifies the device class name of the virtual volumes. If you do not specify a name, Tivoli Storage Manager reconciles all virtual volumes. This parameter is optional.

FIX

Specifies whether or not Tivoli Storage Manager attempts to correct any identified inconsistencies. This parameter is optional. The default is NO. Possible values are:

No Specifies that Tivoli Storage Manager does not fix any inconsistencies.

Yes

Specifies that Tivoli Storage Manager makes the following corrections:

- Tivoli Storage Manager marks as unavailable storage pool volumes on the source server that cannot be located on the target server. Volumes that are only found in the volume history, such as database backups and import and export volumes, are reported as being inconsistent.
- Archive files on the target server that do not correspond to any virtual volumes on the source server are marked for deletion from the target server.

The following table shows the details of the actions taken:

RECONCILE VOLUMES

FIX=	At the Source Server	At the Target Server	Action
NO	Volumes exist	No files exist	Report error
		Files exist but are marked for deletion	
		Active files exist but attributes do not match	
	Volumes do not exist	Active files exist	Report error
Files exist but are marked for deletion		None	
YES	Volumes exist	No files exist	Report error Storage pool volumes: Marked as unavailable
		Files exist but marked for deletion	Report error Storage pool volumes: If attributes match, mark files on the target server as active again, mark volumes on the source server as unavailable, and recommend that an AUDIT VOLUME be done to verify the data. If attributes do not match, mark volumes as unavailable.
		Active files exist but attributes do not match	Report error Storage pool volumes: Mark as unavailable and recommend that an AUDIT VOLUME be done to verify the data.
	Volumes do not exist	Active files exist	Mark files for deletion on the target server.
		Files exist but marked for deletion	None

Example: Reconcile differences in the virtual volume definitions

Reconcile the differences between all virtual volumes definitions on the source server and archive files on the target server to correct any inconsistencies.

```
reconcile volumes remotel fix=yes
```

Related commands

Table 289. Commands related to RECONCILE VOLUMES

Command	Description
DEFINE DEVCLASS	Defines a device class.
DEFINE SERVER	Defines a server for server-to-server communications.
DELETE SERVER	Deletes the definition of a server.
QUERY SERVER	Displays information about servers.
UPDATE SERVER	Updates information about a server.

REGISTER commands

Use the **REGISTER** commands to define or add objects to Tivoli Storage Manager.

- “REGISTER ADMIN (Register an administrator ID)” on page 965
- “REGISTER LICENSE (Register a new license)” on page 969
- “REGISTER NODE (Register a node)” on page 971

REGISTER ADMIN (Register an administrator ID)


Use this command to add an administrator ID to the server. After registration, the administrator ID can issue a limited set of commands, including all query commands. To provide additional privileges, use the **GRANT AUTHORITY** command.

Privilege class

To issue this command, you must have system privilege.

Remember: When you register an administrator with the same name as an existing node, be aware of the administrator authentication method and the **SSLREQUIRED** setting. Any node that has the same name as the administrator that is being registered inherits those settings.

Syntax



Notes:

- 1 A password is not required if you register an administrator and select `AUTHENTICATION=LDAP`. At logon, you are prompted for a password.
- 2 The **PASSEXP** command does not apply to administrators who authenticate to an LDAP directory server.
- 3 The default value can change if you issued the **SET DEFAULTAUTHENTICATION** command and specified LDAP.

Parameters

admin_name (Required)

Specifies the name of the administrator to be registered. The maximum length of the name is 64 characters.

You cannot specify an administrator name of `NONE`.

password (Required)

Specifies the password of the administrator to be registered. The maximum length of the password is 64 characters. The password is not case-sensitive if authentication is with the Tivoli Storage Manager server. Passwords that authenticate with an LDAP directory server are case-sensitive. See “Naming Tivoli Storage Manager objects” on page 13 for the characters that are available for specifying a password.

Tip: A password is not required if you register an administrator and select AUTHENTICATION=LDAP. At logon, you are prompted for a password.

PASSExp

Specifies the number of days the password remains valid. You can set the password expiration period from 0 to 9999 days. A value of 0 means that the password never expires. This parameter is optional. If you do not specify this parameter, the password is set with the global expiration period of 90 days. This parameter does not affect passwords that authenticate with an LDAP directory server.

CONTACT

Specifies information identifying the administrator being registered. This parameter is optional. The maximum length of this string is 255 characters. The contact information must be enclosed in quotation marks if it contains any blanks.

FORCEPwreset

Specifies whether the administrator is required to change or reset the password. This parameter is optional. The default value is NO. Possible values are:

No Specifies that the administrator does not need to change or reset the password while attempting to sign on to the server.

Yes

Specifies that the administrator's password expires at the next sign-on. The client or administrator must change or reset the password then. If a password is not specified, you receive an error message.

EMAILAddress

Specifies more contact information. The information that is specified by this parameter is not acted upon by Tivoli Storage Manager.

AUTHentication

This parameter specifies the authentication method for the administrator user ID. Specify one of the following values: LDAP or LOCAL. The parameter is optional and defaults to LOCAL. The default can change to LDAP if you use the **SET DEFAULTAUTHENTICATION** command and specify LDAP.

Local

Specifies that the local Tivoli Storage Manager server database is used.

LDap

Specifies that the administrator user ID authenticates passwords with an LDAP directory server. Passwords that authenticate with an LDAP directory server are case-sensitive.

Tip: A password is not required if you register an administrator and select AUTHENTICATION=LDAP. At logon, you are prompted for a password.

SSLrequired

Specifies whether the administrator user ID must use Secure Sockets Layer

(SSL) to communicate with the backup-archive client from the Tivoli Storage Manager server. When you authenticate passwords with an LDAP directory server, you must protect the sessions by using SSL or another network security method.

Yes

SSL is required.

No SSL is not required.

Default

SSL is required for an administrator if the password that is associated with its user ID authenticates with an LDAP directory server. SSL is not required for an administrator ID that authenticates its password with the Tivoli Storage Manager server (LOCAL).

Alert

Specifies whether alerts are sent to an administrators email address.

Yes

Specifies that alerts are sent to the specified administrators email address.

No Specifies that alerts are not sent to the specified administrators email address. This is the default value.

Tip: Alert monitoring must be enabled, and email settings must be correctly defined to successfully receive alerts by email. To view the current settings, issue the **QUERY MONITORSETTINGS** command.

Example: Register an administrator ID

Define an administrator, LARRY, with the password PASSONE. You can identify LARRY as second-shift personnel by specifying this information with the CONTACT parameter. Issue the command:

```
register admin larry passone contact='second shift'
```

Example: Register an administrator ID and set the authentication method

Define an administrator ID for Harry, use the password Pa##\$twO. Using the ID and password, Harry can authenticate to the LDAP directory server. Issue the command:

```
register admin Harry Pa##$tw0 authentication=ldap
```

If the password specified does not adhere to LDAP password standards, the administrator is registered, and at log-in a new password must be entered.

Related commands

Table 290. Commands related to REGISTER ADMIN

Command	Description
GRANT AUTHORITY	Assigns privilege classes to an administrator.
LOCK ADMIN	Prevents an administrator from accessing Tivoli Storage Manager.
QUERY ADMIN	Displays information about one or more Tivoli Storage Manager administrators.

REGISTER ADMIN

Table 290. Commands related to REGISTER ADMIN (continued)

Command	Description
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
REGISTER NODE	Defines a client node to the server and sets options for that user.
REMOVE ADMIN	Removes an administrator from the list of registered administrators.
RENAME ADMIN	Changes a Tivoli Storage Manager administrator's name.
SET DEFAULTAUTHENTICATION	Specifies the default password authentication method for any REGISTER NODE or REGISTER ADMIN commands.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.
UNLOCK ADMIN	Enables a locked administrator to access Tivoli Storage Manager.
UPDATE ADMIN	Changes the password or contact information associated with any administrator.
UPDATE NODE	Changes the attributes associated with a client node.

REGISTER LICENSE (Register a new license)

Use this command to register new licenses for server components, including Tivoli Storage Manager (base), Tivoli Storage Manager Extended Edition, and System Storage Archive Manager.

Licenses are stored in enrollment certificate files. The enrollment certificate files contain licensing information for the server product. The NODELOCK file preserves the licensing information for your installation. Your license agreement determines what you are licensed to use, even if you cannot use the REGISTER LICENSE command to register all components. You are expected to comply with the license agreement and use only what you have purchased. Use of the REGISTER LICENSE command implies that you agree to and accept the license terms specified in your license agreement.

Important:

- Before upgrading from a previous version of Tivoli Storage Manager, you must delete or rename the NODELOCK file.
- To unregister licenses, you must erase the NODELOCK file in the server instance directory of your installation, and reregister any previously registered licenses.
- You cannot register licenses for Tivoli Storage Manager for Mail, Tivoli Storage Manager for Databases, Tivoli Storage Manager for ERP, and Tivoli Storage Manager for Space Management.

To generate a report that can help you understand the license requirements for your system, run the **QUERY PVUESTIMATE** command. The report contains estimates of the number of client devices and PVU totals for server devices. The estimates are not legally binding.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

▶▶—REGister LICense—FILE—=— tsmbasic.lic
                               |
                               |—tsmee.lic
                               |—dataret.lic
                               |—*.lic

```

Parameters

FILE

Specifies the name of the enrollment certificate file containing the license to be registered. The specification can contain a wildcard (*). Enter the complete file name or a wildcard in place of the file name. The file names are case-sensitive. The following values can be used:

tsmbasic.lic

To license base IBM Tivoli Storage Manager.

tsmee.lic

To license IBM Tivoli Storage Manager Extended Edition. This includes the disaster recovery manager, large libraries, and NDMP.

REGISTER LICENSE

dataret.lic

To license System Storage Archive Manager. This is required to enable Data Retention Protection as well as Expiration and Deletion Suspension (Deletion Hold).

***.lic**

To license all IBM Tivoli Storage Manager licenses for server components.

Example: Register a license

Register the base Tivoli Storage Manager license.

```
register license file=tsmbasic.lic
```

Related commands

Table 291. Commands related to REGISTER LICENSE

Command	Description
AUDIT LICENSES	Verifies compliance with defined licenses.
QUERY LICENSE	Displays information about licenses and audits.
QUERY PVUESTIMATE	Displays processor value unit estimates.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET LICENSEAUDITPERIOD	Specifies the number of days between automatic license audits.

REGISTER NODE (Register a node)

Use this command to register a node to the server.

This command also automatically creates an administrative user ID with client owner authority over the node. You can use this administrative user ID to access the web Backup-Archive Client from remote locations through a web browser. If an administrative user ID exists with the same name as the node being registered, an administrative user ID is not automatically defined. The client node is registered without an administrative user ID. This process also applies if your site uses open registration.

If a client requires a different policy domain than STANDARD, you must register the client node with this command or update the registered node.

If you do not specify an administrator name or select **USERID=NONE**, the administrator is given the same name as the node. To keep the same-named node and administrator synchronized, the authentication method and the **SSLREQUIRED** setting for the administrator are updated to match the node. Passwords that are shared between same-named nodes and administrators are kept synchronized during an authentication change.


If you register a node that has the same name as an administrator, the administrator authentication method and **SSLREQUIRED** setting change to match those of the node.

Important: When `sslrequired=serveronly` in a **REGISTER NODE** command, the admin **SSLREQUIRED** setting reverts to YES. To use a non-SSL session with a storage agent, rename the admin with the identical name by issuing the **RENAME ADMIN** command.


Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.


Syntax



REGISTER NODE



REGISTER NODE



Notes:

- 1 A password is not required if you register an administrator and select AUTHENTICATION=LDAP. At logon you are prompted for a password.

REGISTER NODE

- 2 The **PASSEXP** command does not apply to administrators who authenticate to an LDAP directory server.
- 3 You can specify the **BKREPLRULEDEFAULT**, **ARREPLRULEDEFAULT**, or **SPREPLRULEDEFAULT** parameter only if you specify the **REPLSTATE** parameter.
- 4 The default value can change if you issued the **SET DEFAULTAUTHENTICATION** command and specified LDAP.

Parameters

node_name (Required)

Specifies the name of the client node to be registered. The maximum length of the name is 64 characters.

You cannot specify a node name of NONE.

password (Required)

Specifies the client node password, which has a maximum length of 64 characters. A password that authenticates with the Tivoli Storage Manager server is considered "LOCAL" and is not case-sensitive. A password that authenticates with an LDAP directory server is considered "LDAP" and is case-sensitive. See "Naming Tivoli Storage Manager objects" on page 13 for a table that lists the characters available for specifying a password.

Tip: A password is not required if you register a node and select AUTHENTICATION=LDAP. At logon you are prompted for a password.

PASSEXP

Specifies the number of days the password remains valid. You can set the password expiration period from 0 to 9999 days. A value of 0 means that the password never expires. This parameter is optional. If you do not specify this parameter, the server common-password expiration period is used. The common password expiration period is 90 days unless changed by issuing the **SET PASSEXP** command.

You can change the password expiration period by using the **UPDATE NODE** or **SET PASSEXP** commands. You can issue the **SET PASSEXP** command to set a common expiration period for all administrators and client nodes. You can also use the command to selectively set password expiration periods. If you selectively set a password expiration period by using the **REGISTER NODE** command, the **UPDATE NODE** command, or the **SET PASSEXP** command, the expiration period is excluded from common password expiration periods that were created using the **SET PASSEXP** command.

You can use the **RESET PASSEXP** command to reset the password expiration period to the common expiration period. The **PASSEXP** command does not apply to nodes that authenticate to an LDAP directory server.

USerid

Specifies the administrative user ID with client owner authority. If you do not specify a user ID, by default an administrative user ID with client owner authority is created using the node name of the client node to be registered. This parameter is optional. When PASSWORDACCESS=GENERATE is used by the client to change the password, the administrative ID with the same name can be used to access the web Backup-Archive Client from a remote location. You can specify one of the following values:

NONE

Specifies that an administrative user ID is not automatically defined.

user_id

Specifies a node ID different than the node being registered. You can use this parameter to grant client owner authority to an existing administrative user ID.

CONtact

Specifies a text string of information identifying the node. The parameter is optional. The maximum length of the text string is 255 characters. The contact information must be enclosed in quotation marks if it contains any blanks.

DOmain

Specifies the name of the policy domain to which the node is assigned. The parameter is optional. If you do not specify a policy domain name, the node is assigned to the default policy domain (STANDARD).

When a source server is registered as a node, it is assigned to a policy domain. Data from the source server is stored in the storage pool specified in the archive copy group of the default management class of that domain.

COMPression

Specifies whether the client node compresses its files before sending them to the server for backup and archive. The parameter is optional. The default value is CLIENT.

Note: This parameter does not apply to nodes with a type of NAS or SERVER.

You can specify one of the following values:

Client

Specifies that the client determines whether to compress files.

Yes

Specifies that the client node compresses its files before sending them to the server for backup and archive.

No Specifies that the client node does not compress its files before sending them to the server for backup and archive.

ARCHDElete

Specifies whether the client node can delete its own archive files from the server. The parameter is optional. The default value is YES. You can specify one of the following values:

Yes

Specifies that the client node can delete its own archive files from the server.

No Specifies that the client node cannot delete its own archive files from the server.

BACKDElete

Specifies whether the client node can delete its own backup files from the server. The parameter is optional. The default value is NO. You can specify one of the following values:

No Specifies that the client node cannot delete its own backup files from the server.

Yes

Specifies that the client node can delete its own backup files from the server.

REGISTER NODE

CLOptset

Specifies the name of the option set to be used by the client. The parameter is optional.

FORCEPwreset

Specifies whether to force a client to change or reset the password. The parameter is optional. The default value is NO. You can specify one of the following values:

No Specifies that the password expiration period is set by the **SET PASSEXP** command. The client does not have to change or reset the password while attempting to logon to the server.

Yes

Specifies that the client node password expires at the next logon. The client must change or reset the password at that time. If a password is not specified, you receive an error message.

Type

Specifies the type of node being registered. The parameter is optional. The default value is CLIENT. You can specify one of the following values:

Client

Specifies that the client node is a Backup-Archive Client, Tivoli Storage Manager for Space Management client, or application client.

NAS

Specifies that the node is a Network Attached Storage (NAS) file server whose data is protected using NDMP operations. The node name cannot be SERVER.

Note: The name of the NAS node must be the same as the data mover. Therefore, the name cannot be changed after a corresponding data mover is defined.

Server

Specifies that the client node is a source server being registered on the target server.

URL

Specifies the URL address that is used in your web browser to administer the client node. The parameter is optional.

MAXNUMMP

Specifies the maximum number of mount points a node is allowed to use on the server or storage agent only for operations such as backup, archive, and Tivoli Storage Manager for Space Management migration. The parameter is optional and does not apply to nodes with a type of NAS or SERVER. The default value is 1. You can specify an integer from 0- 999. A value of 0 specifies that a node cannot acquire any mount point for a client data store operation. The MAXNUMMP value is not evaluated or enforced during client data read operations such as restore, retrieve, and Tivoli Storage Manager for Space Management recall. However, mount points in use for data read operations are evaluated against attempted concurrent data store operations for the same client node and might prevent the data store operations from being able to acquire mount points.

For volumes in a storage pool associated with the FILE device type, the server can have multiple sessions to read and one process to write to the same

volume concurrently. To increase concurrency and provide efficient access for nodes with data in FILE storage pools, increase the value of the **MAXNUMMP** parameter.

For nodes that store data into primary storage pools with the simultaneous-write function enabled, you might have to adjust the value of the **MAXNUMMP** parameter to specify the correct number of mount points for each client session. A client session requires one mount point for the primary storage pool and one mount point for each copy storage pool and each active-data pool.

For server-to-server backup, if one server is at a different version than the other server, set the number of mount points on the target server to a value higher than 1. Otherwise, you receive an error.

A storage agent independently tracks the number of points used during a client session. If a node has a storage agent installed it might exceed the **MAXNUMMP** value. The **MAXNUMMP** value might also be exceeded under conditions where the node does not have to wait for a mount point.

Note: The server might preempt a client operation for a higher priority operation and the client might lose a mount point if no other mount points are available. See the *Administrator's Guide* for more information.

KEEPMP

Specifies whether the client node keeps the mount point for the entire session. The parameter is optional. The default value is **NO**. You can specify one of the following values:

Yes

Specifies that the client node must retain the mount point during the entire session. If policy definitions cause data to be stored to a disk storage pool after storing data to a sequential access storage pool, any mount points held by the session will not be released.

No Specifies that the client node release the mount point during the session. If policy definitions cause data to be stored to a disk storage pool after storing data to a sequential access storage pool, any mount points held by the session will be released.

AUTOFSRename

Specify whether file spaces are automatically renamed when you upgrade the client system to support Unicode or specify if file spaces are renamed by the client, if needed. The default is **NO**. Setting the parameter to **YES** enables automatic renaming, which occurs when the client runs one of the following operations: archive, selective backup, full incremental backup, or partial incremental backup. The automatic renaming changes the names of existing backed-up file spaces that are not in Unicode in server storage. Then the file spaces are backed up in Unicode. You can use this parameter for Unicode-enabled Tivoli Storage Manager clients using Windows, Macintosh OS X, and NetWare operating systems.

After the client with support for Unicode is installed, any new file spaces that the client backs up are stored in server storage using the UTF-8 code page. UTF-8 is a byte-oriented encoding form specified by the Unicode Standard.

You can specify one of the following values:

Yes

Existing file spaces are automatically renamed when you upgrade to a client that supports Unicode and the client runs one of the following

REGISTER NODE

operations: archive, selective backup, full incremental backup, or partial incremental backup. The renaming occurs whether the client uses the graphical user interface, the command line, or the client scheduler.

For example, the server renames a drive as follows:

```
Original name: D_DRIVE  
New name: D_DRIVE_OLD
```

The new name indicates that the file space is stored on the server in a format that is not Unicode.

- No** Existing file spaces are not automatically renamed when the client system upgrades to a client that supports Unicode, and the client runs one of the following operations: archive, selective backup, full incremental backup, or partial incremental backup.

Client

The option `AUTOFSRENAME` in the client's option file determines whether file spaces are renamed.

By default, the client option is set to `PROMPT`. When the client system upgrades to a client that supports Unicode and the client runs a Tivoli Storage Manager operation with the graphical user interface or the command line, the program displays a one-time prompt to the user about whether to rename file spaces.

When the client scheduler runs an operation, the program does not prompt for a choice about renaming, and does not rename file spaces. Backups of existing file spaces are sent as before (not in Unicode).

Refer to the appropriate *Backup-Archive Clients Installation and User's Guide* book for more information on the client option.

VALIDateprotocol

Specify whether Tivoli Storage Manager performs a cyclic redundancy check (CRC) to validate the data sent between the client and server. The parameter is optional. The default is `NO`. You can specify one of the following values:

- No** Specifies that data validation not be performed on any data sent between the client and server.

Dataonly

Specifies that data validation be performed only on file data that is sent between the client and server. This does not include the file metadata. This mode impacts performance because additional overhead is required to calculate and compare CRC values between the client and the server.

All

Specifies that data validation be performed on all client file data, client file metadata, and Tivoli Storage Manager server metadata that is sent between the client and server. This mode impacts performance as additional overhead is required to calculate and compare CRC values between the client and the server.

TXNGroupmax

Specifies the number of files per transaction commit that are transferred between a client and a server. Client performance may be improved by using a larger value for this option.

The default value is 0. Specifying 0 indicates the node will use the server global value that is set in the server options file. To use a value other than the

server global value, specify a value of 4 through 65,000 for this parameter. The node value takes precedence over the server value.

Attention: Increasing the TXNGROUPMAX value will result in increased recovery log utilization. Higher recovery log utilization may increase the risk of running out of log space. Evaluate each node's performance before changing this parameter. For more information on managing the recovery log, see the *Administrator's Guide*.

DATAWritepath

Specifies the transfer path used when the client sends data to the server, storage agent, or both, during storage operations such as backup or archive. The parameter is optional. The default is ANY.

Note: If a path is unavailable, the node cannot send any data. For example, if you select the LAN-free option but a LAN-free path is not defined, the operation will not work.

You can specify one of the following values:

ANY

Specifies that data is sent to the server, storage agent, or both, by any available path. A LAN-free path will be used if one is available. If a LAN-free path is unavailable, the data will be moved using the LAN.

LAN

Specifies that data is sent using the LAN.

LANFree

Specifies that data is sent using a LAN-free path.

DATAReadpath

Specifies the transfer path used when the server, storage agent, or both read data for a client, during operations such as restore or retrieve. The parameter is optional. The default is ANY.

Note: If a path is unavailable, data cannot be read. For example, if you select the LAN-free option but a LAN-free path is not defined, the operation will not work.

You can specify one of the following values:

ANY

Specifies that the server, storage agent, or both use any available path to read data. A LAN-free path will be used if one is available. If a LAN-free path is unavailable, the data will be read using the LAN.

LAN

Specifies that data is read using the LAN.

LANFree

Specifies that data is read using a LAN-free path.

TARGETLevel

Specifies the client deployment package that is targeted for this node. You can substitute an applicable release package for V.R.M.F, which stands for Version.Release.Modification.Fix Level. For example: TARGETLevel=6.2.0.0.

You must specify each segment with a number that is applicable to a deployment package. You cannot use an asterisk in any field as a substitution for a valid number. The parameter is optional.

REGISTER NODE

Restriction: The **TARGETLEVEL** parameter does not apply to nodes with a type of NAS or SERVER.

SESSIONINITiation

Controls whether the server or the client initiates sessions. The default is that the client initiates sessions. The parameter is optional.

Clientorserver

Specifies that the client may initiate sessions with the server by communicating on the TCP/IP port defined with the server option **TCPPOINT**. Server-prompted scheduling may also be used to prompt the client to connect to the server.

SERVEROnly

Specifies that the server will not accept client requests for sessions. All sessions must be initiated by server-prompted scheduling on the port defined for the client with the **REGISTER** or **UPDATE NODE** commands. You cannot use the client acceptor (dsmcad) to start the scheduler when **SESSIONINITIATION** is set to **SERVERONLY**.

HLAddress

Specifies the client IP address that the server contacts to initiate scheduled events. This optional parameter is used only when **SESSIONINITIATION** is set to **SERVERONLY**, regardless of any addresses previously used by the client to contact the server. If **SESSIONINITIATION SERVERONLY** is not in use, this option has no effect.

The address can be specified either in numeric or host name format. If a numeric address is used, it will be saved without verification by a domain name server. If the address is not correct, it can cause failures when the server attempts to contact the client. Host name format addresses will be verified with a domain name server. Verified names will be saved and resolved with Domain Name Services when the server contacts the client.

LLAddress

Specifies the client port number on which the client listens for sessions from the server. This optional parameter is used only when **SESSIONINITIATION** is set to **SERVERONLY**, regardless of any addresses previously used by the client to contact the server. If **SESSIONINITIATION SERVERONLY** is not in use, this option has no effect.

The value for this parameter must match the value of client option **TCPCLIENTPORT**. The default value is 1501.

EMAILAddress

This parameter is used for additional contact information. The information specified by this parameter is not acted upon by Tivoli Storage Manager.

DEDUPLICATION

Specifies where data deduplication can occur for this node. You can specify one of the following values:

SERVEROnly

Specifies that data stored by this node can be deduplicated on the server only. This value is the default.

Clientorserver

Specifies that data stored by this node can be deduplicated on either the client or the server. For data deduplication to take place on the client, you must also specify a value of **YES** for the **DEDUPLICATION** client option.

You can specify this option in the client option file or in the client option set on the Tivoli Storage Manager server.

BACKUPINITiation

Specifies whether the non-root user ID on the client node can back up files to the server. The parameter is optional. The default value is ALL, indicating that non-root user IDs can back up data to the server. You can select one of the following values:

ALL

Specifies that non-root user IDs can back up files to the server. ALL is the default if BACKUPINITIATION is not specified.

ROOT

Specifies that the root user ID can back up files to the server. If you are using the V6.4 or later backup-archive client, authorized users have the same privileges as the root user ID.

For more information about authorized users, see the *Backup-Archive Clients Installation and User's Guide*.

Restriction: The attribute is ignored by the server if the backup-archive client connects from a platform other than AIX, HP-UX, Linux, Oracle Solaris, or Mac OS.

Remember: The application programming interface (API) is affected by the BACKUPINITIATION parameter on the server. By default, all API users are allowed to back up data. Setting the parameter to ROOT on an API node is not recommended.

REPLState

Specifies whether data that belongs to the client node is ready to be replicated. This parameter is optional. Specify this parameter only if you are issuing the **REGISTER NODE** command on a server that is configured to replicate data to a target replication server. If you register a client node on a source replication server and set up replication for the node, do not register the node on the target replication server. The client node is created automatically on the target server the first time that replication occurs.

You can select one of the following values:

ENabled

Specifies that the client node is configured for replication and is ready to replicate. When you specify this parameter, the replication mode in the client node definition on the source replication server is automatically set to SEND. This setting indicates that data that belongs to the client node is sent to a target server during replication.

When replication first occurs for the client node, the replication state of the node on the target replication server is automatically set to ENABLED. The replication mode on the target replication server is set to RECEIVE. This setting indicates that data that belongs to the client node is received from a source replication server. To determine the replication state and mode, issue the **QUERY NODE** command on a source or a target replication server.

DISabled

Specifies that the node is configured for replication but that replication does not occur until you enable it.

REGISTER NODE

BKREPLRuledefault, ARREPLRuledefault, and SPREPLRuledefault

Specifies the replication rule that applies to a data type if the file space rules for the data type are set to DEFAULT.

Restriction: You can specify the **BKREPLRULEDEFAULT**, **ARREPLRULEDEFAULT**, or **SPREPLRULEDEFAULT** parameter only if you specify the **REPLSTATE** parameter.

BKREPLRuledefault

Specifies the replication rule for backup data.

ARREPLRuledefault

Specifies the replication rule for archive data.

SPREPLRuledefault

Specifies the replication rule for space-managed data.

If the file space rules for the data type are set to DEFAULT and you do not specify a rule for the **BKREPLRULEDEFAULT**, **ARREPLRULEDEFAULT**, or **SPREPLRULEDEFAULT** parameter, data is replicated according to the server rule for the data type.

You can specify normal-priority replication or high-priority replication rules. In a replication process that includes both normal and high-priority data, high-priority data is replicated first. Before you specify a rule, consider the order in which you want the data to be replicated.

You can specify the following rules:

ALL_DATA

Replicates active and inactive backup data, archive data, or space-managed data. The data is replicated with a normal priority.

ACTIVE_DATA

Replicates only active backup data. The data is replicated with a normal priority. This rule is valid only for **BKREPLRULEDEFAULT**.

Attention:

- If you specify **ACTIVE_DATA**, inactive backup data that belongs to client nodes is not replicated.
- If the replication rule for backup data in any file spaces that belong to the client node is **DEFAULT**, inactive backup data in those file spaces on the target replication server is deleted.

ALL_DATA_HIGH_PRIORITY

Replicates active and inactive backup data, archive data, or space-managed data. Data is replicated with a high priority.

ACTIVE_DATA_HIGH_PRIORITY

Replicates only active backup data. The data is replicated with a high priority. This rule is valid only for **BKREPLRULEDEFAULT**.

Attention:

- If you specify `ACTIVE_DATA_HIGH_PRIORITY`, inactive backup data that belongs to client nodes is not replicated.
- If the replication rule for backup data in any file spaces that belong to the client node is `DEFAULT`, inactive backup data in those file spaces on the target replication server is deleted.

DEFAULT

Replicates data according to the server replication rule for backup data.

For example, suppose that you want to replicate the archive data in all the file spaces that belongs to a client node. Replication of the archive data is a high priority. One method to accomplish this task is to specify `ARREPLRULEDEFAULT=DEFAULT`. Ensure that the file space rules for archive data are also set to `DEFAULT` and that the server rule for archive data is set to `ALL_DATA_HIGH_PRIORITY`.

Restriction: If a node is configured for replication, the file space rules are set to `DEFAULT` after the node stores data on the source replication server.

NONE

Data of the specified type is not replicated.

For example, if you do not want to replicate space-managed data that belongs to a client node, specify `SPREPLRULEDEFAULT=NONE`

ROLEOVERRIDE

Specifies whether to override the reported role of the client for processor value unit (PVU) estimation reporting. The default is `USEREPORTED`.

The role reported by the client is either client-device (for example, a workstation) or server-device (for example, file/print server, application server, database). By default, the client reports its role based on the client type and the operating system. All clients initially report their role as server-device, except for Backup-Archive Clients running Microsoft Windows workstation distributions (Windows Vista) and Macintosh OS X.

Specify one of the following values:

Client

Specifies a client-device.

Server

Specifies a server-device.

Other

Specifies that this node is not to be used for PVU estimation reporting. This can be useful when there are multiple nodes deployed for a physical system (for example, virtual environments, test nodes, retired nodes, and nodes not in production or clustering).

Usereported

Use the reported role provided by the client.

AUTHentication

This parameter specifies the password authentication method for the node. Specify one of the following values: `LDAP` or `LOCAL`. The parameter is optional and defaults to `LOCAL`. The default can change to `LDAP` if you use the `SET DEFAULTAUTHENTICATION` command and specify `LDAP`.

REGISTER NODE

Local

Specifies that the local Tivoli Storage Manager server database is used.

LDap

Specifies that the node uses an LDAP directory server for password authentication.

Tip: A password is not required if you register a node and select AUTHENTICATION=LDAP. At logon you are prompted for a password.

SSLrequired

Specifies whether the node must use Secure Sockets Layer (SSL) to communicate with the Tivoli Storage Manager server. When you authenticate passwords with an LDAP directory server, you must protect the sessions by using SSL or another network security method.

Yes

SSL is required.

No SSL is not required.

Default

SSL is required for a node if its password authenticates with an LDAP directory server. SSL is not required for a node that authenticates its password with the Tivoli Storage Manager server (LOCAL).

SERVERonly

SSL is required between the backup-archive client and the server. SSL is not required between the backup-archive client and the storage agent.

Example: Register a client node that only the root user can back up

Register the client node `mete0rite` with password `KingK0ng` to permit only the root user to back up files to the server.

```
register node mete0rite KingK0ng
backupinit=root
```

Example: Register a client node and password and set compression on

Register the client node `J0E0S2` with the password `SECRETCODE` and assign this node to the `DOM1` policy domain. This node can delete its own backup and archive files from the server. All files are compressed by the client node before they are sent to the server. This command automatically creates a `J0E0S2` administrative user ID with password `SECRETCODE`. In addition, the administrator now has client owner authority to the `J0E0S2` node.

```
register node joeos2 secretcode domain=dom1
archdelete=yes backdelete=yes
compression=yes
```

Example: Register a client node and password and prevent the administrative ID from being automatically assigned

Register the client node `FRAN` and prevent an administrative user ID from being automatically defined

```
register node fran xy34z userid=none
```

Example: Grant client owner authority for an existing administrative user

Grant client owner authority to an existing administrative user ID, *HELPADMIN*, when registering the client node JAN. This would not automatically create an administrator ID named JAN, but would grant client owner authority for this node to the *HELPADMIN* administrator.

```
register node jan pwdsafe userid=helpadmin
```

Example: Register a NAS file server node that uses NDMP operations

Register a node name of NAS1 for a NAS file server that is using NDMP operations. Assign this node to a special NAS domain.

```
register node nas1 pw4pw domain=nasdom type=nas
```

Example: Register a node and specify data validation

Register a node name of ED whose data is backed up to a primary storage pool that also has a list of 2 copy storage pools defined. The primary storage pool is of disk device type, and the 2 copy storage pools are of sequential device type. Specify 2 as the maximum number of mount points. The client's network has been unstable during the last few weeks. The client wants assurance that his data is not corrupted as it is sent over the network. Specify that Tivoli Storage Manager should validate all data that is sent by this user to the server.

```
register node ed pw45twx maxnummp=2 validateprotocol=all
```

Example: Register a node and specify the maximum number of files per transaction commit

Register a node name of ED and set the TXNGroupmax to 1,000.

```
register node ed pw45twx txngroupmax=1000
```

Example: Register a node and allow it to deduplicate data on the client system

Register a node name of JIM and allow it to deduplicate data on the client system.

```
register node jim jim deduplication=clientorserver
```

Example: Register a node name of ED and set the role as a server-device for PVU estimation reporting

Register a node name of ED and set the role as a server-device for PVU estimation reporting.

```
register node ed pw45twx roleoverride=server
```

Example: Register a node on a source replication server

Define NODE1 to a source replication server. Specify a replication rule for the backup data that belongs to NODE1 so that active backup data is replicated with a high priority. Enable replication for the node.

```
register node node1 bkreplruledefault=active_data_high_priority replstate=enabled
```

Example: Register a node that uses a password that authenticates with an LDAP directory server

Register a node name of node17 that must authenticate to an LDAP directory server.

```
register node node17 n0de17n0de#seventeen authentication=ldap
```

You must enter a new password the next time you log on if the new password does not adhere to requirements for a “complex” password.

Related commands

Table 292. Commands related to REGISTER NODE

Command	Description
DEFINE ASSOCIATION	Associates clients with a schedule.
DEFINE DATAMOVER	Defines a data mover to the Tivoli Storage Manager server.
DEFINE MACHNODEASSOCIATION	Associates a Tivoli Storage Manager node with a machine.
DELETE FILESPACE	Deletes data associated with clients file spaces.
LOCK NODE	Prevents a client from accessing the server.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY PVUESTIMATE	Displays an estimate of the client-devices and server-devices being managed.
QUERY REPLNODE	Displays information about the replication status of a client node.
REGISTER ADMIN	Defines a new administrator without granting administrative authority.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
REMOVE REPLNODE	Removes a node from replication.
RENAME NODE	Changes the name for a client node.
RESET PASSEXP	Resets the password expiration for nodes or administrators.
SET DEFAULTAUTHENTICATION	Specifies the default password authentication method for any REGISTER NODE or REGISTER ADMIN commands.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.
SET CPUINFOREFRESH	Specifies the number of days between client scans for workstation information used for PVU estimates.
SET DEDUPVERIFICATIONLEVEL	Specifies the percentage of extents verified by the server during client-side deduplication.
SET REGISTRATION	Specifies whether users can register themselves or must be registered by an administrator.

Table 292. Commands related to REGISTER NODE (continued)

Command	Description
UNLOCK NODE	Enables a locked user in a specific policy domain to access the server.
UPDATE ADMIN	Changes the password or contact information associated with any administrator.
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE NODE	Changes the attributes associated with a client node.

REMOVE commands

Use the **REMOVE** commands to remove an object from Tivoli Storage Manager.

- “REMOVE ADMIN (Delete an administrator user ID)” on page 989
- “REMOVE NODE (Delete a node or an associated machine node)” on page 991
- “REMOVE REPLNODE (Remove a client node from replication)” on page 993

REMOVE ADMIN (Delete an administrator user ID)

Use this command to remove an administrator user ID from the system.

You cannot remove the last system administrator user ID or the SERVER_CONSOLE administrative ID from the system.

Using SYNCLDAPDELETE=YES with this command removes the administrator user ID from the LDAP directory server. Do not specify SYNCLDAPDELETE=YES if the LDAP namespace is shared by multiple Tivoli Storage Manager servers. If you want to remove an administrator user ID from all Tivoli Storage Manager servers and all LDAP directory servers, specify SYNCLDAPDELETE=YES.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

admin_name (Required)

Specifies the name of the administrator to be removed.

SYNCLDAPDELETE

If the administrative user ID is stored on the LDAP directory server, you can delete the entry on that server by specifying a value of YES. Do not delete the entry if it is used by another server that shares the same administrator namespace on the LDAP directory server.

Example: Remove an administrator user ID

Remove the administrator user ID larry. Issue the following command:

```
remove admin larry
```

Example: Remove any corresponding entries for administrator user IDs that authenticate to multiple LDAP directory servers

To remove the entry in the LDAP directory server for administrator user ID harry, issue the following command:

```
remove admin harry syncldapdelete=yes
```

Related commands

Table 293. Commands related to REMOVE ADMIN

Command	Description
LOCK ADMIN	Prevents an administrator from accessing Tivoli Storage Manager.
QUERY ADMIN	Displays information about one or more Tivoli Storage Manager administrators.

REMOVE ADMIN

Table 293. Commands related to REMOVE ADMIN (continued)

Command	Description
REGISTER ADMIN	Defines a new administrator without granting administrative authority.
RENAME ADMIN	Changes a Tivoli Storage Manager administrator's name.

REMOVE NODE (Delete a node or an associated machine node)

Use this command to remove a node from the server. If you are using disaster recovery manager and the node to be removed is associated with a machine, the association between the node and the machine is also deleted. If a node is a member of a collocation group and you remove the node from the server, the node is also removed from the collocation group.

Using `SYNCLDAPDELETE = YES` with this command removes the node's entry from the LDAP directory server. Do not specify `SYNCLDAPDELETE = YES` if the LDAP namespace is shared by multiple IBM Tivoli Storage Manager servers. If you want to remove a node from all Tivoli Storage Manager servers and all LDAP directory servers, specify `SYNCLDAPDELETE = YES`.

When a node is removed, the corresponding administrative ID is removed only if:

- The administrator name is identical to the node name.
- The administrator has client owner or client access authority *only* to the node being removed.
- The administrator is not a managed object.


Important: Before you can remove a node, you must delete all backup and archive file spaces that belong to that node.

Before you can remove a NAS node that has a corresponding data mover defined, you must first delete any paths from the data mover, delete the data mover, delete all virtual file space definitions for the node, then remove the NAS node.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax



Parameters

node_name (Required)

Specifies the name of the node to be removed.

`SYNCLdapdelete`

If the node ID is stored on the LDAP directory server, you can delete the entry on that server by specifying a value of YES. Do not delete the entry if it is used by another server that shares the same namespace on the LDAP directory server.

REMOVE NODE

Example: Remove a client node

Remove the client node LARRY.

```
remove node larry
```

Example: Remove any corresponding entries for a node ID that authenticates to an LDAP directory server

To remove the client node elroy that has corresponding entries that share namespace on an LDAP directory server, issue the following command:

```
remove node elroy syncldapdelete=yes
```

Related commands

Table 294. Commands related to REMOVE NODE

Command	Description
DELETE DATAMOVER	Deletes a data mover.
DELETE FILESPACE	Deletes data associated with clients file spaces.
DELETE PATH	Deletes a path from a source to a destination.
DELETE VIRTUALFSMAPPING	Delete a virtual file space mapping.
LOCK NODE	Prevents a client from accessing the server.
QUERY COLLOGGROUP	Displays information about collocation groups.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY SESSION	Displays information about all active administrator and client sessions with Tivoli Storage Manager.
REGISTER NODE	Defines a client node to the server and sets options for that user.
RENAME NODE	Changes the name for a client node.

REMOVE REPLNODE (Remove a client node from replication)

Use this command to remove a node from replication if you no longer want to replicate the data that belongs to the node.

To completely remove a client node from replication, you must issue this command on the source and target replication servers that have this node configured for replication. You can only issue this command from an administrative command-line client. You cannot issue this command from the server console.

If you issue the **REMOVE REPLNODE** command for a client node whose replication mode is set to **SEND** or **RECEIVE**, the mode is set to **NONE**. The replication state is also set to **NONE**. After you remove a client node from replication, the target replication server can accept backup, archive, and space-managed data directly from the node.

Remember: If a client node is removed from replication, information in the database about replication for the node is deleted. If the client node is enabled for replication later, the replication process replicates all the data that is specified by replication rules and settings.


Issuing the **REMOVE REPLNODE** command does not delete data that belongs to the client node. To delete file space data that belongs to the client node, issue the **DELETE FILESPACE** command for each of the file spaces that belong to the node. If you do not want to keep the client node definition, issue the **REMOVE NODE** command. To delete file space data and the client node definition, issue **DELETE FILESPACE** and **REMOVE NODE** on the target replication server.

If a node replication process is running for a client node that is specified by this command, the command fails and the replication information for the node is not removed.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

node_name or *node_group_name* (Required)

Specifies the name of the client node or defined group of client nodes that you want to remove from replication. To specify multiple client node names and client-node group names, separate the names with commas with no intervening spaces. You can use wildcard characters to specify client node names, but not to specify client-node group names.

REMOVE REPLNODE

Example: Remove three client nodes and a client node group from replication

The names of the client nodes are NODE1, NODE2, and NODE3. The name of the client node group is PAYROLL. Issue the following command on the source and target replication servers:

```
remove replnode node*,payroll
```

Related commands

Table 295. Commands related to REMOVE REPLNODE

Command	Description
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.

RENAME commands

Use the **RENAME** commands to change the name of an existing object.

- “RENAME ADMIN (Rename an administrator)” on page 996
- “RENAME FILESPACE (Rename a client file space on the server)” on page 998
- “RENAME NODE (Rename a node)” on page 1001
- “RENAME SCRIPT (Rename a Tivoli Storage Manager script)” on page 1003
- “RENAME SERVERGROUP (Rename a server group)” on page 1004
- “RENAME STGPOOL (Change the name of a storage pool)” on page 1005

RENAME ADMIN (Rename an administrator)

Use this command to change the name of an existing administrator. Existing information for this administrator such as password, contact information, and privilege classes is not altered.

If you assign an existing administrative user ID to another person, it is recommended that you use the **UPDATE ADMIN** command to change the password.


When an administrator and a node share a name and you change the administrator authentication method, the node authentication method also changes. If you rename an administrator to the same name as an existing node, the authentication method and the **SSLREQUIRED** setting for the node can change. If those settings are different, after the renaming, both administrator and node will have the same authentication method and **SSLREQUIRED** setting.

You cannot rename the SERVER_CONSOLE administrative ID.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

current_admin_name (Required)

Specifies the name of the administrator user ID to be renamed.

new_admin_name (Required)

Specifies the new administrator name. The maximum length of the name is 64 characters.

SYNC1dapdelete

If the administrative user ID is stored on the LDAP directory server, you can rename the entry on that server by specifying a value of YES. Do not rename the entry if it is used by another server that shares the same administrator namespace on the LDAP directory server. The default is NO.

Yes

Select YES to delete the corresponding LDAP entry in order to rename it. If you want to rename all of the same-named administrators to keep them in synch with the LDAP server, select YES.

No

Does not delete the corresponding entries. Select NO if the corresponding entries are being used by another server that shares administrative namespace in the LDAP directory server. Select NO if you want to rename only one administrator and have the administrators on the shared LDAP server keep their same name. After submitting the command, issue the `UPDATE ADMIN AUTHENTICATION=LDAP` command to create an LDAP entry for

the new administrator name. Updating the administrator this way guarantees that the administrator can authenticate against the LDAP directory server. The default is NO.

Example: Rename an administrator

Rename the Tivoli Storage Manager administrator CLAUDIA to BILL.

```
rename admin claudia bill
```

Example: Rename an administrator and all same-named administrators that authenticate to an LDAP server

```
rename admin claudia bill syncldapdelete=yes
```

Related commands

Table 296. Commands related to RENAME ADMIN

Command	Description
QUERY ADMIN	Displays information about one or more Tivoli Storage Manager administrators.
UPDATE ADMIN	Changes the password or contact information associated with any administrator.

RENAME FILESPACE (Rename a client file space on the server)

Use this command to rename an existing client file space on the server to a new file space name or to rename imported file spaces.


You might want to rename a file space that has been imported or to cause the creation of new Unicode-enabled file spaces for Unicode-enabled clients.

Privilege class

Any administrator with unrestricted policy authority or with restricted policy authority over the client's policy domain can issue this command.

Syntax

►► `REName Filespace` *node_name* *current_file_space_name* *new_file_space_name* ►►



Notes:

- 1 This parameter is the default when you specify `NAMETYPE=Unicode`.

Parameters

node_name (Required)

Specifies the name of the client node to which the file space to be renamed belongs.

current_file_space_name (Required)

Specifies the name of the file space to be renamed. A file space name is case-sensitive and must be specified exactly as defined to the server. Virtual file space mapping names are allowed.

new_file_space_name (Required)

Specifies the new name for the file space. A client file space name is case-sensitive and must be specified exactly as it is to be defined to the server. This parameter cannot be an existing virtual file space mapping name. If the *current_file_space_name* is a virtual file space, the *new_file_space_name* must follow all the rules for defining a virtual file space name. See the **DEFINE VIRTUALFSMAPPING** command for more information.

Attention: If the new name type is hexadecimal, be careful to specify valid UTF-8 hexadecimal values so the server's code page displays the file space name as intended. For example, do not specify a value that can be interpreted as a backspace.

NAMETYPE

Specify how you want the server to interpret the current file space name that you enter. This parameter is useful when the server has clients with support for Unicode. You can use this parameter for Unicode-enabled Tivoli Storage Manager clients using Windows, Macintosh OS X and NetWare operating systems.

The default value is SERVER. If a virtual file space mapping name is specified, you must use SERVER. Possible values are:

SERVER

The server uses the server's code page to interpret the file space name.

UNICODE

The server converts the file space name entered from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server's code page. Conversion can fail if the string includes characters that are not available in the server code page, or if the server has a problem accessing system conversion routines.

FSID

The server interprets the file space name as the file space ID (FSID).

NEWNAMETYPE

Specify how you want the server to interpret the new file space name that you enter. The default is SERVER if you specified the NAMETYPE as SERVER, or if the file space to be renamed is not Unicode. The default is UNICODE if you specified the NAMETYPE as UNICODE, or if the file space to be renamed is Unicode. If a virtual file space mapping name is specified, you must use SERVER. Possible values are:

SERVER

The server uses the server's code page to interpret the file space name.

UNICODE

The server converts the file space name entered from the server code page to the UTF-8 code page. The success of the conversion depends on the actual characters in the name and the server's code page. If the conversion is not successful, you may want to specify the HEXADECIMAL parameter.

HEXADecimal

The server interprets the file space name that you enter as the hexadecimal representation of a name in Unicode. Using hexadecimal ensures that the server is able to correctly rename the file space, regardless of the server's code page.

To view the hexadecimal representation of a file space name, you can use the **QUERY FILESPACE** command with **FORMAT=DETAILED**.

Restriction: You cannot specify a new name of a type that is different from the original name. You can rename a file space that is Unicode to another name in Unicode. You can rename a file space that is not Unicode using a new name in the server's code page. You cannot mix the two types.

Example: Rename an imported file space to prevent overwriting

An AIX client node named LARRY backed up file space /r033 to the Tivoli Storage Manager server. The file space was exported to tape and later re-imported to the server. When this file space was imported, a system-generated name, /r031, was created for it because /r033 already existed for client node LARRY.

Client node LARRY, however, already had a file space named /r031 that was not backed up, therefore, was unknown to the server. Unless the imported file space is renamed, it will overlay file space /r031 because the file space name generated by the IMPORT function is the same as a file space on client node LARRY that is unknown to the server.

RENAME FILESPACE

Use the following command to rename imported file space /r031. The new name, /imported-r033, identifies that the new file space is an imported image of file space /r033.

```
rename filesystem larry /r031 /imported-r033
```

Example: Rename file space to create a Unicode-enabled file space

Client JOE is using an English Unicode-enabled Tivoli Storage Manager client. JOE has backed up several large file spaces that are not Unicode enabled in server storage. File space \\joe\c\$ contains some files with Japanese file names that cannot be backed up to a file space that is not Unicode enabled. Because the file spaces are large, the administrator does not want to convert all of JOE's file spaces to Unicode-enabled file spaces at this time. The administrator wants to only rename the non-Unicode file space, \\joe\c\$, so that the next backup of the file space causes the creation of a new Unicode-enabled file space. The new Unicode-enabled file space will allow the Japanese files to be successfully backed up.

Use the following command to rename \\joe\c\$:

```
rename filesystem joe \\joe\c$ \\joe\c$_old
```

Related commands

Table 297. Commands related to RENAME FILESPACE

Command	Description
DEFINE VIRTUALFSMAPPING	Define a virtual file space mapping.
DELETE FILESPACE	Deletes data associated with clients file spaces.
EXPORT NODE	Copies client node information to external media or directly to another server.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY OCCUPANCY	Displays file space information by storage pool.

RENAME NODE (Rename a node)

Use this command to rename a node.

If you are assigning an existing node ID to another person, it is recommended that you use the **UPDATE NODE** command to change the password.

Restrictions:


- You cannot rename a NAS node name that has a corresponding data mover defined. If the data mover has defined paths, the paths must first be deleted. For details on the steps that are required when you want to rename a NAS node, refer to the *Administrator's Guide*.
- If a node is configured for replication, it cannot be renamed.

If you rename a node to the same name as an existing administrator, the administrator authentication method and **SSLREQUIRED** setting are updated to match the node. When a node and an administrator share a name and you change the node authentication method or the node **SSLREQUIRED** setting, the administrator settings also change. You must have system level authority to update the node authentication method or the node **SSLREQUIRED** setting and also update a same-named administrator.

Privilege class

You must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax



Parameters

current_node_name (Required)

Specifies the name of the node to be renamed.

new_node_name (Required)

Specifies the new name of the node. The maximum length of the name is 64 characters.

SYNCldapdelete

If the node ID is stored on the LDAP directory server, you can rename the entry on that server by specifying a value of YES. Do not rename the entry if it is used by another server that shares the same namespace on the LDAP directory server. This parameter applies to nodes that authenticate passwords with an LDAP directory server. The default is NO.

Yes

Select YES to delete the corresponding LDAP entry in order to rename it. If you want to rename all of the same-named nodes to keep them in synch with the LDAP server, select YES.

RENAME NODE

No Does not delete the corresponding entries. Select NO if the corresponding entries are being used by another server that shares administrative namespace in the LDAP directory server. Select NO if you want to rename only one node and keep the same-named nodes on the shared LDAP server, the same. After submitting the command, issue the UPDATE NODE AUTHENTICATION=LDAP command to create an LDAP entry for the new node name. Updating the node this way guarantees that the node can authenticate against the LDAP directory server. The default is NO.

Example: Rename a node

Rename the node JOE to JOYCE.

```
rename node joe joyce
```

Example: Rename a node that shares a namespace with other servers

Rename the node JOYCE to JOE and do not delete the previous name from corresponding LDAP directory servers. Specify **SYNCLDAPDELETE=NO** to ensure that other servers that share a namespace on the LDAP directory server can still connect to the LDAP directory server.

```
rename node joyce joe syncldapdelete=no
```

Related commands

Table 298. Commands related to RENAME NODE

Command	Description
QUERY NODE	Displays partial or complete information about one or more clients.
UPDATE NODE	Changes the attributes associated with a client node.

RENAME SCRIPT (Rename a Tivoli Storage Manager script)

Use this command to rename a Tivoli Storage Manager script.

Privilege class

To issue this command, you must have operator, policy, system, storage, or system privilege.

Syntax

```
►►—REName SCRIpt—current_script_name—new_script_name—◄◄
```

Parameters

current_script_name **(Required)**

Specifies the name of the script to rename.

new_script_name **(Required)**

Specifies the new name for the script. The name can contain as many as 30 characters.

Example: Rename a script

Rename SCRIPT1 to a new script named SCRIPT2.

```
rename script script1 script2
```

Related commands

Table 299. Commands related to RENAME SCRIPT

Command	Description
COPY SCRIPT	Creates a copy of a script.
DEFINE SCRIPT	Defines a script to the Tivoli Storage Manager server.
DELETE SCRIPT	Deletes the script or individual lines from the script.
QUERY SCRIPT	Displays information about scripts.
RUN	Runs a script.
UPDATE SCRIPT	Changes or adds lines to a script.

RENAME SERVERGROUP

RENAME SERVERGROUP (Rename a server group)

Use this command to rename a server group.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—RENAME SERVERGRoup—current_group_name—new_group_name—◀◀
```

Parameters

current_group_name **(Required)**

Specifies the server group to rename.

new_group_name **(Required)**

Specifies the new name of the server group. The maximum length of the name is 64 characters.

Example: Rename a server group

Rename server group WEST_COMPLEX to BIG_WEST.

```
rename servergroup west_complex big_west
```

Related commands

Table 300. Commands related to RENAME SERVERGROUP

Command	Description
COPY SERVERGROUP	Creates a copy of a server group.
DEFINE SERVERGROUP	Defines a new server group.
DELETE SERVERGROUP	Deletes a server group.
QUERY SERVERGROUP	Displays information about server groups.
UPDATE SERVERGROUP	Updates a server group.

RENAME STGPOOL (Change the name of a storage pool)

Use this command to change the name of a storage pool. You can change storage pool names so that the same names can be used on a configuration manager and its managed servers.

When you change a storage pool name, any administrators with restricted storage privilege for the old storage pool are automatically given restricted storage privilege for the renamed storage pool. However, if a management class or copy group specifies the existing storage pool as the destination for files, the destination is not changed to the new storage pool name.

If processes are active when a storage pool is renamed, the old name may still be displayed in messages or queries for those processes.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—REName STGpool—*current_pool_name*—*new_pool_name*—◄◄

Parameters

current_pool_name **(Required)**

Specifies the storage pool to rename.

new_pool_name **(Required)**

Specifies the new name of the storage pool. The maximum length of the name is 30 characters.

Example: Change the name of a storage pool

Rename storage pool STGPOOLA to STGPOOLB:

```
rename stgpool stgpoola stgpoolb
```

Related commands

Table 301. Commands related to RENAME STGPOOL

Command	Description
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE STGPOOL	Deletes a storage pool from server storage.
QUERY STGPOOL	Displays information about storage pools.
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.
UPDATE STGPOOL	Changes the attributes of a storage pool.

REPLICATE NODE (Replicate data in file spaces that belong to a client node)

Use this command to replicate data in file spaces that belong to one or more client nodes or defined groups of client nodes.

When you issue this command, a process is started in which data that belongs to the specified client nodes is replicated according to replication rules. Files that are no longer stored on the source replication server, but that exist on the target replication server, are deleted during this process.

If a node replication process is already running for a client node specified by this command, the node is skipped, and replication begins for other nodes that are enabled for replication.

Tip: When the **QUERY PROCESS** command is issued during node replication, the output can show unexpected results for the number of completed replications. The reason is that, for node replication purposes, each file space is considered to contain three logical file spaces:

- One for backup objects
- One for archive objects
- One for space-managed objects

By default, the **QUERY PROCESS** command generates results for each logical file space. Other factors also affect the output of the **QUERY PROCESS** command:


- If a file space has a replication rule that is set to **NONE**, the file space is not included in the count of file spaces that are being processed.
- If you specify data types in the **REPLICATE NODE** command, only those data types are included in the count of file spaces that are being processed, minus any file spaces that are specifically excluded.


Issue this command on the server that acts as a source for replicated data.

Privilege class

To issue this command, you must have system privilege.

Syntax





Notes:

- 1 Do not mix file space identifiers (FSIDs) and file space names in the same command.
- 2 Do not specify FSID if you use wildcard characters for the client node name.

Parameters

node_name or *node_group_name* (Required)

Specifies the name of the client node or defined group of client nodes whose data is to be replicated. You can also specify a combination of client node names and client-node group names. To specify multiple client node names or client-node group names, separate the names with commas with no intervening spaces. You can use wildcard characters with client node names but not with client-node group names. The replication rules for all file spaces in the specified client nodes are checked.

file_space_name or *FSID*

Specifies the name of the file space or the file space identifier (FSID) to be replicated. A name or FSID is optional. If you do not specify a name or an FSID, all the data in all the file spaces for the specified client nodes is eligible for replication.

file_space_name

Specifies the name of the file space that has data to be replicated. File space names are case sensitive. To determine the correct capitalization for the file space, issue the **QUERY FILESPACE** command. Separate multiple names with commas with no intervening spaces. When you specify a name, you can use wildcard characters.

A server that has clients with Unicode-enabled file spaces might have to convert the file space name. For example, the server might have to convert

REPLICATE NODE

a name from the server code page to Unicode. For details, see the **NAMETYPE** parameter. If you do not specify a file space name, or if you specify only a single wildcard character for the name, you can use the **CODETYPE** parameter to limit the operation to Unicode file spaces or to non-Unicode file spaces.

FSID

Specifies the file space identifier for the file space to be replicated. The server uses FSIDs to find the file spaces to replicate. To determine the FSID for a file space, issue the **QUERY FILESPACE** command. Separate multiple FSIDs with commas with no intervening spaces. If you specify an FSID, the value of the **NAMETYPE** parameter must be FSID.

NAMETYPE

Specifies how you want the server to interpret the file space names that you enter. You can use this parameter for Tivoli Storage Manager clients that are Unicode-enabled and that have Windows, Macintosh OS X, or NetWare operating systems.

Use this parameter only when you enter a partly-qualified or fully-qualified file space name. The default value is SERVER. You can specify one of the following values:

SERVER

The server uses the server code page to interpret file space names.

UNICODE

The server converts file space names from the server code page to the UTF-8 code page. The success of the conversion depends on the characters in the name and the server code page. Conversion can fail if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines.

FSID

The server interprets file space names using their file space identifiers.

CODETYPE

Specifies the type of file spaces to be included in node replication processing. Use this parameter only when you enter a single wildcard character for the file space name. The default value is BOTH, which specifies that file spaces are included regardless of code page type. You can specify one of the following values:

UNICODE

Specifies only file spaces that are in Unicode.

NONUNICODE

Specifies only file spaces that are not in Unicode.

BOTH

Specifies all file spaces regardless of code page type.

DATATYPE

Specifies the type of data to be replicated. Data is replicated according to the replication rule that applies to the data type. This parameter is optional. You can specify one or more data types. If you do not specify a data type, all backup, archive, and space-managed data is replicated. Separate multiple data types with commas with no intervening spaces. You cannot use wildcard characters. You can specify one of the following values:

ALL

Replicates all backup, archive, and space-managed data in a file space

according to the rule that is assigned to the data type. For example, suppose NODE1 has a single file space. The following replication rules apply:

- The file space rules for backup and archive data in the file space are set to ALL_DATA.
- The file space rule for space-managed data is set to DEFAULT.
- The client node rule for space-managed data is set to NONE.

If you issue REPLICATE NODE NODE1 DATATYPE=ALL, only backup data and archive data are replicated.

BACKUP

Replicates active and inactive backup data in a file space if the controlling replication rule is ALL_DATA, ACTIVE_DATA, ALL_DATA_HIGH_PRIORITY, or ACTIVE_DATA_HIGH_PRIORITY.

BACKUPActive

Replicates only active backup data in a file space if the controlling replication rule is ACTIVE_DATA or ACTIVE_DATA_HIGH_PRIORITY.

ARCHive

Replicates only archive data in a file space if the controlling replication rule is ALL_DATA or ALL_DATA_HIGH_PRIORITY.

SPACEManaged

Replicates only space-managed data in a file space if the controlling replication rule is ALL_DATA or ALL_DATA_HIGH_PRIORITY.

PRIority

Specifies the data to replicate based upon the priority of the replication rule. You can specify one of the following values:

All

Replicates all data in a file space if the controlling replication rule is ALL_DATA, ACTIVE_DATA, ALL_DATA_HIGH_PRIORITY, or ACTIVE_DATA_HIGH_PRIORITY.

High

Replicates only data in a file space that has a controlling replication rule of ALL_DATA_HIGH_PRIORITY or ACTIVE_DATA_HIGH_PRIORITY.

Normal

Replicates only data in a file space that has a controlling replication rule of ALL_DATA or ACTIVE_DATA.

MAXSESSions

Specifies the maximum allowable number of data sessions to use for sending data to a target replication server. This parameter is optional. The value can be 1 - 99. The default value is 10. Increasing the number of sessions can improve node replication throughput.

When setting this value, consider the number of logical and physical drives that can be dedicated to the replication process. To access a sequential-access volume, Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on the following factors:

- Other Tivoli Storage Manager and system activity
- The mount limits of the device classes for the sequential access storage pools that are involved

REPLICATE NODE

Ensure that there are sufficient mount points and drives available to allow node replication processes to complete. Each replication session might need a mount point on the source and target replication servers for storage pool volumes. If the device type is not FILE, each session might also need a drive on both the source and target replication servers.

When setting a value for **MAXSESSIONS**, also consider the available bandwidth and the processor capacity of the source and target replication servers.

Tip:

- The value specified by the **MAXSESSIONS** parameter applies only to data sessions. Data sessions are sessions during which data is sent to a target replication server. However, if you issue a **QUERY SESSION** command, the total number of sessions might exceed the number of data sessions. The difference is due to short control sessions that are used for querying and setting up replication operations.
- Remember that the value of the **MAXSESSIONS** parameter represents the maximum allowable number of sessions and that the number of sessions that are used for replication depends on the amount of data to be replicated. If you are replicating only a small amount of data, you do not achieve any benefit by increasing the number of sessions. The total number of sessions might be less than the value that is specified by the **MAXSESSIONS** parameter.

Preview

Specifies whether to preview data. This parameter is optional. The default value is NO. You can specify one of the following values:

No Specifies that the data is replicated to the target server but that the data is not previewed.

Yes

Specifies that data is previewed but not replicated. If you specify **PREVIEW=YES**, only volumes that must be physically mounted, such as tape and optical volumes, are displayed. Volumes that are assigned to storage pools that have a device class of FILE are not displayed.

The following information is displayed in the output:

- The names of client nodes whose data would be replicated.
- The number of files that would be replicated or deleted.
- The estimated amount of time it would take to complete the node replication process.
- A list of volumes that would be mounted

If the data belonging to a client node that is specified by the **REPLICATE NODE** command was never replicated and you specify **PREVIEW=YES**, the node and its file spaces are automatically defined on the target replication server.

LISTfiles

Specifies whether to list the names of files that would be replicated. This parameter is optional. The default is NO. Specifying this parameter signifies that the **WAIT** parameter is set to YES and that you cannot issue the **WAIT** parameter from the server console.

You can specify one of the following values:

No Specifies that the names of files that would be replicated are not displayed.

Yes

Specifies that the names of files that would be replicated are displayed.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default value is NO. You can specify one of the following values:

No Specifies that the command processes in the background. To monitor the background processing of the **REPLICATE NODE** command, issue the **QUERY PROCESS** command.

Yes

Specifies that the command processes in the foreground. Messages are not displayed until the command completes processing. You cannot specify **WAIT=YES** from the server console.

Example: Replicate data by data type and priority

Replicate high-priority active backup data and high-priority archive data that belongs to all the client nodes in group PAYROLL.

```
replicate node payroll datatype=backupactive,archive priority=high
```

Example: Replicate all the data that belongs to a node according to the assigned replication rules

NODE1 has a single file space. The following replication rules apply:

- File space rules:
 - Backup data: ACTIVE_DATA
 - Archive data: DEFAULT
 - Space-managed data: DEFAULT
- Client node rules:
 - Backup data: DEFAULT
 - Archive data: ALL_DATA_HIGH_PRIORITY
 - Space-managed data: DEFAULT
- Server rules:
 - Backup data: ALL_DATA
 - Archive data: ALL_DATA
 - Space-managed data: NONE

```
replicate node node1 priority=all
```

Active backup data in the file space is replicated with normal priority. Archive data is replicated with high priority. Space-managed data is not replicated.

Related commands

Table 302. Commands related to REPLICATE NODE

Command	Description
CANCEL REPLICATION	Cancels node replication processes.
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.

REPLICATE NODE

Table 302. Commands related to REPLICATE NODE (continued)

Command	Description
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLNODE	Displays information about the replication status of a client node.
QUERY REPLRULE	Displays information about node replication rules.
QUERY SERVER	Displays information about servers.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REMOVE REPLNODE	Removes a node from replication.
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE REPLRULE	Enables or disables replication rules.
VALIDATE REPLICATION	Verifies replication for file spaces and data types.

REPLY (Allow a request to continue processing)

Use this command and an identification number to inform the server that you have completed a requested operation. Not all server requests require a reply. This command is required only if the request message specifically indicates that a reply is needed.

Privilege class

To issue this command, you must have system privilege or operator privilege.

Syntax

```
►►—REPLY—request_number—┐
                           └LABEL—==—volume_label—┘
```

Parameters

request_number (Required)

Specifies the identification number of the request.

LABEL

Specifies the label to be written on a volume when you reply to a message from a LABEL LIBVOLUME command process. This parameter is optional.

Example: Reply to a request

Respond to a reply request using 3 as the request number.

```
reply 3
```

Related commands

Table 303. Commands related to REPLY

Command	Description
CANCEL REQUEST	Cancels pending volume mount requests.
QUERY REQUEST	Displays information about all pending mount requests.

RESET PASSEXP (Reset password expiration)

Use the **RESET PASSEXP** command to reset the password expiration period to the common expiration period for administrator and client node passwords. The **RESET PASSEXP** command does not apply to passwords that are stored on an LDAP directory server.

Restriction: You cannot reset the password expiration period to the common expiration period with the **SET PASSEXP** command.


Use the **QUERY STATUS** command to display the common password expiration period.

Restriction: If you do not specify either the **NODE** or **ADMIN** parameters, the password expiration period for all client nodes and administrators will be reset.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

Node

Specifies the name of the node whose password expiration period you would like to reset. To specify a list of nodes, separate the names with commas and no intervening spaces. This parameter is optional.

Admin

Specifies the name of the administrator whose password expiration period you would like to reset. To specify a list of administrators, separate the names with commas and no intervening spaces. This parameter is optional.

Example: Reset the password expiration for specific client nodes

Reset the password expiration period for client nodes bj and katie.

```
reset passexp node=bj,katie
```

Example: Reset the password expiration for all users

Reset the password expiration period for all users to the common expiration period.

```
reset passexp
```

Related commands*Table 304. Commands related to RESET PASSEXP*

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.
UPDATE ADMIN	Changes the password or contact information associated with any administrator.
UPDATE NODE	Changes the attributes associated with a client node.

RESTART EXPORT (Restart a suspended export operation)

Use this command to restart a suspended export operation.

An export operation is suspended when any of the following conditions is detected:

- A **SUSPEND EXPORT** command is issued for the running export operation
- Segment preemption - the file being read for export is deleted by some other process
- Communication errors on a server-to-server export
- No available mount points
- Necessary volumes are unavailable
- I/O errors encountered

Important: Nodes or file spaces (on the exporting server) in the original export operation that are subsequently renamed are not included in the resumed operation. Any remaining data for nodes or file spaces on the target server that are deleted prior to resumption are discarded.

Privilege class

You must have system privilege to issue this command.

Syntax



Parameters

export_identifier

This optional parameter is the unique identifier for the suspended server-to-server export operation. You can use the wildcard character to specify this name. The export identifier name can be found by issuing the **QUERY EXPORT** command to list all the currently suspended server-to-server export operations.

Example: Restart a suspended export

Restart the suspended export operation identified by the export identifier EXPORTALLACCTNODES.

```
restart export exportallacctnodes
```

Related commands

Table 305. Commands related to **RESTART EXPORT**

Command	Description
CANCEL EXPORT	Deletes a suspended export operation.
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.

Table 305. Commands related to RESTART EXPORT (continued)

Command	Description
QUERY EXPORT	Displays the export operations that are currently running or suspended.
SUSPEND EXPORT	Suspends a running export operation.

RESTORE commands

Use the **RESTORE** commands to restore Tivoli Storage Manager storage pools or volumes.

- “RESTORE NODE (Restore a NAS node)” on page 1019
- “RESTORE STGPOOL (Restore storage pool data from a copy pool or an active-data pool)” on page 1024
- “RESTORE VOLUME (Restore primary volume data from a copy pool or an active-data pool)” on page 1028

RESTORE NODE (Restore a NAS node)


Use this command to initiate a restore operation for a network-attached storage (NAS) node.

You can use the **RESTORE NODE** command to restore backups that were created by using either the client's **BACKUP NAS** command or the server's **BACKUP NODE** command. NAS data may be restored from primary or copy native Tivoli Storage Manager pools; primary or copy NAS pools; or any combination needed to achieve the restore.

Privilege class

To issue this command, you must have system privilege, policy privilege for the domain to which the node is assigned, or client owner authority over the node.

Syntax



Parameters

node_name (Required)

Specifies the name of the node to restore. You cannot use wildcard characters or specify a list of names.

source_file_system (Required)

Specifies the name of the file system to restore. You cannot use wildcard characters for this name. You cannot specify more than one file system to restore. Virtual file space names are allowed.

destination_file_system

Specifies that the file server restores the data to an existing, mounted file system on the file server. This parameter is optional. The default is the original location of the file system on the file server. Virtual file space names are allowed.

RESTORE NODE

FILELIST

Specifies the list of file or directory names to be restored. This parameter is optional. The default is to restore the entire file system. If this value is specified, the server attempts to restore the objects from the appropriate image. If the **PITDATE** and **PITTIME** parameters are specified, then the file is restored from the last backup image prior to the specified time. If no **PITDATE** and **PITTIME** parameters are specified, the file is restored from the latest backup image of the file system.

If the image is a differential backup, objects are first restored from the corresponding full backup and then from the differential backup. The restore is done by scanning the appropriate images for the specified objects and restoring any that are found. The TOCs for these images is not accessed, so the IBM Tivoli Storage Manager server does not check whether the objects are actually contained within the images.

The folder path and file name must be entered using forward slash (/) symbols. No ending forward slash (/) is needed at the end of the file name. All arguments that contain a space must have double quotation marks ("argument with spaces") surrounding the entire argument.

```
FILELIST="/path/to/filename1 with blanks",/path/to/filename2_no_blanks
```

Any file names that contain commas must have double quotation marks surrounding the entire argument, surrounded by single quotation marks ("argument with commas").

```
FILELIST='"/path/to/filename1,with,commas"',/path/to/filename2_no_commas
```

To restore a complete directory, specify a directory name instead of a file name. All files in the directory and its subdirectories are restored. An ending forward slash (/) is not needed at the end of the directory name:

```
FILELIST=/path/to/mydir
```

file_name

Specifies one or more file or directory names to be restored. The names you specify cannot contain wildcards. Multiple names must be separated with commas and no intervening blanks. File names are case-sensitive.

FILE:*file_list*

Specifies the name of a file that contains a list of the file or directory names to be restored. In the specified file, each file or directory name must be on a separate line. Blank lines and comment lines that begin with an asterisk are ignored. For example:

To restore files FILE01, FILE02, and FILE03, create a file named RESTORELIST that contains a line for each file:

```
FILE01  
FILE02  
FILE03
```

You can specify the files to be restored with the command as follows:

```
FILELIST=FILE:RESTORELIST
```

NAMETYPE

Specifies how you want the server to interpret the names specified as **FILELIST**=*file_name* or the names listed in the file specified with **FILELIST**=*file_list*. This parameter is useful when the names may contain Unicode characters. It has no effect if the FILELIST parameter is not specified. The default value is SERVER. Possible values are:

SERVER

The server uses the server's code page to interpret the names.

HEXadecimal

The server interprets the names that you enter as the hexadecimal representation of a name in Unicode. To view the hexadecimal representation of a file or directory name, you can use the **QUERY TOC** command with **FORMAT=DETAILED**.

UNICODE

The server interprets the names as being UTF-8 encoded. This option only applies when you have specified a list with **FILELIST=FILE:file_list**.

Restriction: Network Data Management Protocol (NDMP) has limitations that prevent Tivoli Storage Manager from reporting whether or not individual files and directories are successfully restored.

PITDate

Specifies the point-in-time date. When used with the **PITTIME** parameter, **PITDATE** establishes the point in time from which you want to select the data to restore. The latest data that was backed up on or before the date and time that you specify will be restored. This parameter is optional. The default is **TODAY**.

You can specify the date by using one of the following values:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	06/25/2001
TODAY	The current date	TODAY
<i>TODAY-days</i> or <i>-days</i>	The current date minus days specified	TODAY-7 or -7 . To restore data that was backed up a week ago, specify PITDATE=TODAY-7 or PITDATE=-7 .
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
<i>EOLM-days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
<i>BOTM+days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

PITTime

Specifies the point-in-time time. When used with the **PITDATE** parameter, **PITTIME** establishes the point in time from which you want to select the data to restore. The latest data that was backed up on or before the date and time that you specify will be restored. This parameter is optional. The default is the current time.

You can specify the time by using one of the following values:

RESTORE NODE

Value	Description	Example
<i>HH:MM:SS</i>	A specific time on the specified date	12:33:28
NOW	The current time on the specified date	NOW
NOW- <i>HH:MM</i> or - <i>HH:MM</i>	The current time minus hours and minutes on the specified begin date	NOW-03:30 or -03:30. If you issue this command at 9:00 with PITTIME=NOW-03:30 or PITTIME=-03:30, the server restores backup records with a time of 5:30 or later on the point-in-time date.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. The default is NO. Possible values are:

No Specifies that the server processes this command in the background. Use the **QUERY PROCESS** command to monitor the background processing of this command.

Yes

Specifies that the server processes this command in the foreground. You wait for the command to complete before continuing with other tasks. The server then displays the output messages to the administrative client when the command completes.

Restriction: You cannot specify WAIT=YES from the server console.

TYPE

Specifies the type of image to restore. The default value for this parameter is BACKUPIMAGE and it is used to restore data from standard NDMP base or differential backups. Other image types represent backup methods that might be specific to a particular file server. Possible values are:

BACKUPImage

Specifies that the file system should be restored from the appropriate standard NDMP backup images. This is the default method for performing an NDMP restore operation. Using the BACKUPIMAGE type, you can restore data from base and differential backups, and data at the file level.

SNAPMirror

Specifies that the file system should be retrieved from a NetApp SnapMirror image. SnapMirror images are block-level full-backup images of a NetApp file system. A SnapMirror image can only be restored to a file system that has been prepared as a SnapMirror target volume. Refer to the documentation that came with your NetApp file server for details.

After a SnapMirror image is retrieved and copied to a target file system, Tivoli Storage Manager breaks the SnapMirror relationship that was created by the file server during the operation. After the restore is complete, the target file system returns to the same state as that of the original file system at the point-in-time of the backup.

When setting the **TYPE** parameter to SNAPMIRROR, note the following restrictions:

Restrictions:

- You cannot specify the FILELIST parameter.
- Neither the *source_file_system_name* nor the *destination_file_system_name* can be a virtual filesystem name.
- This parameter is valid for NetApp and IBM N-Series file servers only.

Example: Restore a complete directory

Restore all of the files and subdirectories in the directory /mydir.

```
restore node nasnode /myfs /dest filelist=/path/to/mydir
```

Example: Restore data from a file system

Restore the data from the /vol/vol10 file system on NAS node NAS1.

```
restore node nas1 /vol/vol10
```

Example: Restore a directory-level backup to the same location

Restore the directory-level backup to the original location. The source is the virtual file space name /MIKESDIR and no destination is specified.

```
restore node nas1 /mikesdir
```

For this example and the next example, assume the following virtual file space definitions exist on the server for the node NAS1.

VFS Name	Filesystem	Path
/mikesdir	/vol/vol2	/mikes
/TargetDirVol2	/vol/vol2	/tmp
/TargetDirVol1	/vol/vol1	/tmp

Example: Restore a directory-level backup to a different file system

Restore the directory-level backup to a different file system but preserve the path.

```
restore node nas1 /mikesdir /vol/vol0
```

Related commands

Table 306. Commands related to RESTORE NODE

Command	Description
BACKUP NODE	Backs up a network-attached storage (NAS) node.
DEFINE VIRTUALFSMAPPING	Define a virtual file space mapping.
QUERY NASBACKUP	Displays information about NAS backup images.
QUERY TOC	Displays details about the table of contents for a specified backup image.

RESTORE STGPOOL (Restore storage pool data from a copy pool or an active-data pool)

Use this command to restore files from one or more copy storage pools or active-data pools to a primary storage pool.

Tivoli Storage Manager restores all the primary storage pool files that:

- Have been identified as having errors
- Reside on a volume with an access mode of DESTROYED

You can also use this command to identify volumes that contain damaged, primary files. During restore processing, a message is issued for every volume in the restored storage pool that contains damaged, non-cached files. Use the **QUERY CONTENT** command to identify damaged, primary files on a specific volume.

In addition to restoring data to primary storage pools that have NATIVE or NONBLOCK data formats, this command also lets you restore data to primary storage pools that have NDMP data formats (NETAPPDUMP, CELERRADUMP, or NDMPDUMP). The primary storage pool must have the same data format as the copy storage pool from which data is to be restored. Tivoli Storage Manager supports backend data movement for NDMP images. For details, see the *Administrator's Guide*.

Note: To restore NAS client-node data to NAS storage pools, you must manually change the access mode of the volumes to DESTROYED using the **UPDATE VOLUME** command. However, if you are using disaster recovery manager, the plan file will contain the information the server needs to automatically mark the volumes as DESTROYED.

Restoration of files might be incomplete if backup file copies in copy storage pools or active-data pools were moved or deleted by other Tivoli Storage Manager processes during restore processing. To prevent this problem, do not issue the following commands for copy storage pool or active-data pool volumes while restore processing is in progress:

- **MOVE DATA**
- **DELETE VOLUME (DISCARDATA=YES)**
- **AUDIT VOLUME (FIX=YES)**


Also, you can prevent reclamation processing for your copy storage pools by setting the RECLAIM percentage to 100 with the **UPDATE STGPOOL** command.

Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the primary storage pool for which files are to be restored. If you are a restricted storage administrator and you want to restore files to a new primary storage pool, you must also have authority for the new storage pool.

Syntax

```
▶▶—RESTORE STGpool—primary_pool_name—┬──COPYstgpool—==—copy_pool_name—┘▶▶
```

**A (Yes):**

|—ACTIVEDATAPool—=*active-data_pool_name*—|

Parameters***primary_pool_name* (Required)**

Specifies the name of the primary storage pool that is being restored.

COPYstgpool

Specifies the name of the copy storage pool from which the files are to be restored. This parameter is optional. If this parameter is not specified, files are restored from any copy pool in which copies can be located. Do not use this parameter with the `ACTIVEDATAONLY` or `ACTIVEDATAPPOOL` parameters.

ACTIVEDATAOnly

Specifies that active versions of backup files are to be restored from active-data pools only. This parameter is optional. The default is `NO`. If this parameter is not specified, files are restored from copy-storage pools. Do not use this parameter with the `COPYSTGPOOL` parameter. Possible values are:

No Specifies that the storage pool will not be restored from active-data pools.

Yes

Specifies that the storage pool will be restored from active-pool or pools that you specify using the `ACTIVEDATAPPOOL` parameter. If you specify `YES` as a value for `ACTIVEDATAONLY`, but do not specify a value for `ACTIVEDATAPPOOL`, files are restored from any active-data pool in which active versions of backup files can be located.

Attention: Restoring a primary storage pool from an active-data pool might cause some or all inactive files to be deleted from the database if the server determines that an inactive file needs to be replaced but cannot find it in the active-data pool.

ACTIVEDATAPool

Specifies the name of the active-data pool from which the active versions of backup files are to be restored. This parameter is optional. If this parameter is not specified, files are restored from any active-data pool in which active versions of backup files can be located.

NEWstgpool

Specifies the name of the new storage pool to which to restore the files. This parameter is optional. If this parameter is not specified, files are restored to the original primary storage pool (the pool being restored).

RESTORE STGPOOL

MAXProcess

Specifies the maximum number of parallel processes that are used for restoring files. Using multiple, parallel processes may improve throughput for the restore. This parameter is optional. You can specify a value from 1 to 999. The default is 1.

When determining this value, consider the number of mount points (logical drives) and physical drives that can be dedicated to this operation. To access a sequential access volume, Tivoli Storage Manager uses a mount point, and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the restore.

Each process needs a mount point for copy storage pool volumes, and, if the device type is not FILE, each process also needs a drive. If you are restoring files in a sequential storage pool, each process needs an additional mount point for primary storage pool volumes and, if the device class is not FILE, an additional drive. For example, suppose you specify a maximum of 3 processes to restore a primary sequential storage pool from a copy storage pool of the same device class. Each process requires two mount points and two drives. To run all three processes, the device class must have a mount limit of at least 6, and at least 6 mount points and 6 drives must be available.

To preview a restore, only one process is used and no mount points or drives are needed.

Preview

Specifies if you want to preview but not perform the restore. The preview lets you identify volumes required to restore the storage pool. The preview displays:

- A list of primary storage pool volumes that contain damaged files.
- The number of files and the number of bytes to be restored, assuming that the access mode of the required copy storage pool volumes is READWRITE or READONLY when the restore operation is performed.
- A list of copy storage pool volumes containing files to be restored. These volumes must be mounted if you perform the restore.
- A list of any volumes containing files that cannot be restored.

Note: For only a list of offsite copy storage pool volumes to be mounted during a restore, change the access mode of the copy pool volumes to UNAVAILABLE. This prevents reclamation and move data processing of the volumes until they are moved onsite for the restore.

This parameter is optional. The default is NO. Possible values are:

No Specifies that the restore is done.

Yes

Specifies that you want to preview the restore but not do the restore.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is NO. Possible values are:

No Specifies that the server processes this command in the background.

You can continue with other tasks while the command is being processed.

Messages created from the background process are displayed either in the activity log or the server console, depending on where messages are logged. To cancel a background process, use the **CANCEL PROCESS** command. If you cancel this process, some files may have already been restored prior to the cancellation.

Yes

Specifies that the server performs this operation in the foreground. The operation must complete before you can continue with other tasks. The server then displays the output messages to the administrative client when the operation completes.

Note: You cannot specify WAIT=YES from the server console.

Example: Restore files from a copy storage pool to the primary storage pool

Restore files from any copy storage pool to the primary storage pool, PRIMARY_POOL.

```
restore stgpool primary_pool
```

Example: Restore files from a specific active-data pool to the primary storage pool

Restore files from active-data pool ADP1 to the primary storage pool PRIMARY_POOL.

```
restore stgpool primary_pool activedataonly=yes activedatapool=adp1
```

Related commands

Table 307. Commands related to RESTORE STGPOOL

Command	Description
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
CANCEL PROCESS	Cancels a background server process.
COPY ACTIVE DATA	Copies active backup data.
QUERY CONTENT	Displays information about files in a storage pool volume.
QUERY PROCESS	Displays information about background processes.
RESTORE VOLUME	Restores files stored on specified volumes in a primary storage pool from copy storage pools.
UPDATE STGPOOL	Changes the attributes of a storage pool.
UPDATE VOLUME	Updates the attributes of storage pool volumes.

RESTORE VOLUME (Restore primary volume data from a copy pool or an active-data pool)

Use this command to restore all files on damaged volumes in a primary storage pool that was backed up to a copy storage pool or copied to an active-data pool. Tivoli Storage Manager does not restore cached copies of files and removes those cached files from the database during restore processing.

In addition to restoring data to volumes in storage pools that have NATIVE or NONBLOCK data formats, this command also lets you restore data to volumes in storage pools that have NDMP data formats (NETAPPDUMP, CELERRADUMP, or NDMPDUMP). The volumes to be restored must have the same data format as the volumes in the copy storage pool. Tivoli Storage Manager supports backend data movement for NDMP images. For details, see the *Administrator's Guide*.

This command changes the access mode of the specified volumes to DESTROYED. When all files on a volume are restored to other locations, the destroyed volume is empty and is deleted from the database.

The restoration may be incomplete for one or more of the following reasons:

- Files were either never backed up or the backup copies are marked as damaged. Use the **QUERY CONTENT** command to get more information on the remaining files on the volume.
- A copy storage pool was specified on the **RESTORE** command, but files were backed up to a different copy storage pool. Use the **PREVIEW** parameter when you issue the **RESTORE** command again to determine if this is the problem.
- Volumes in the copy storage pool needed to perform the restore operation are offsite or unavailable. Check the activity log for messages that occurred during restore processing.
- Backup file copies in copy storage pools were moved or deleted by other processes during a restore. See note 3.
- An active-data pool was specified for the restore, and inactive files were not available to be copied.

Important:

1. You cannot restore volumes in storage pools defined with a CENTERA device class.
2. Before you restore a random-access volume, issue the **VARY** command to vary the volume offline.
3. To prevent copy storage pools files from being moved or deleted by other processes, do not issue the following commands for copy storage pool volumes during a restore:
 - MOVE DATA
 - DELETE VOLUME (DISCARDDATA=YES)
 - AUDIT VOLUME (FIX=YES)


To prevent reclamation processing of copy storage pools, issue the **UPDATE STGPOOL** command with the **RECLAIM** parameter set to 100.

Privilege class

To issue this command you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the primary storage pool. If you have restricted privilege and want to restore files to a new primary storage pool, you

must also have authority for the new storage pool.

Syntax



A (Yes):



Parameters

`volume_name` (Required)

Specifies the name of the primary storage pool volume to be restored. To specify a list of volumes that belong to the same primary storage pool, separate the names with commas and no intervening spaces.

`COPYstgpool`

Specifies the name of the copy storage pool from which the files are to be restored. This parameter is optional. If you do not specify this parameter, files are restored from any copy pool in which copies can be located. Do not use this parameter with the `ACTIVEDATAONLY` or `ACTIVEDATAPool` parameters.

`ACTIVEDATAonly`

Specifies that active versions of backup files are to be restored from active-data pools only. This parameter is optional. The default is `NO`. If this parameter is not specified, files are restored from copy-storage pools. Do not use this parameter with the `COPYSTGPOOL` parameter. Possible values are:

No Specifies that the storage pool will not be restored from active-data pools.

Yes

Specifies that the storage pool will be restored from active-pool or pools that you specify using the `ACTIVEDATAPool` parameter. If you specify `YES` as a value for `ACTIVEDATAONLY`, but do not specify a value for `ACTIVEDATAPool`, files are restored from any active-data pool in which active versions of backup files can be located.

RESTORE VOLUME

Attention: Restoring a volume from an active-data pool might cause some or all inactive files to be deleted from the database if the server determines that an inactive file needs to be replaced but cannot find it in the active-data pool.

ACTIVEDATAPool

Specifies the name of the active-data pool from which the active versions of backup files are to be restored. This parameter is optional. If this parameter is not specified, files are restored from any active-data pool in which active versions of backup files can be located.

NEWstgpool

Specifies the name of the new storage pool to which to restore the files. This parameter is optional. If you do not specify this parameter, files are restored to the original primary storage pool.

MAXProcess

Specifies the maximum number of parallel processes to use for restoring files. Using parallel processes may improve throughput. This parameter is optional. You can specify a value from 1 to 999. The default is 1.

When determining this value, consider the number of mount points (logical drives) and physical drives that can be dedicated to this operation. To access a sequential access volume, Tivoli Storage Manager uses a mount point, and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the restore.

Each process needs a mount point for copy storage pool volumes. If the device type is not FILE, each process also needs a drive. If you are restoring a sequential storage pool, each process needs an additional mount point for primary storage pool volumes and, if the device type is not FILE, an additional drive. For example, suppose you specify a maximum of three processes to back up a primary sequential storage pool to a copy storage pool of the same device class. Each process requires two mount points and two drives. To run all three processes, the device class must have a mount limit of at least 6, and at least 6 mount points and 6 drives must be available.

To preview a backup, only one process is used and no mount points or drives are needed.

Preview

Specifies if you want to preview but not perform the restore. You can use this option to identify the offsite volumes required to restore a storage pool. This parameter is optional. The default is NO. Possible values are:

No Specifies that you want to perform the restore operation.

Yes

Specifies that you want to preview the restore operation but restore the data.

Tip: If you preview a restore to see a list of offsite copy pool volumes to be mounted, you should change the access mode of the identified volumes to UNAVAILABLE. This prevents reclamation and MOVE DATA processing for these volumes until they are transported to the onsite location for use in restore processing.

The preview displays the following:

- The number of files and bytes to be restored, if the access mode of the copy storage pool volumes is READWRITE or READONLY when the restoration is performed.
- A list of copy storage pool volumes containing files to be restored. These volumes must be mounted if you perform the restore.
- A list of volumes containing files that cannot be restored.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. This default is NO. Possible values are:

No Specifies that the server processes this command in the background.

You can continue with other tasks while the command is being processed. Messages created from the background process are displayed either in the activity log or the server console, depending on where messages are logged.

To cancel a background process, use the **CANCEL PROCESS** command. If you cancel this process, some files may have already been backed up prior to the cancellation.

Yes

Specifies that the server processes this command in the foreground. The operation must complete before you can continue with other tasks. The server then displays the output messages to the administrative client when the command completes.

Remember: You cannot specify WAIT=YES from the server console.

Example: Restore primary volume data files

Restore files stored on volume PVOL2 in primary storage pool PRIMARY_POOL.

```
restore volume pv012
```

Example: Restore primary volume data files from an active-data pool

Restore files stored on volume VOL001 in primary pool PRIMARY_POOL from active-data pool ADP1.

```
restore volume vol001 activedataonly=yes activedatapool=adp1
```

Related commands

Table 308. Commands related to RESTORE VOLUME

Command	Description
BACKUP STGPOOL	Backs up a primary storage pool to a copy storage pool.
COPY ACTIVATEDATA	Copies active backup data.
CANCEL PROCESS	Cancels a background server process.
QUERY PROCESS	Displays information about background processes.
RESTORE STGPOOL	Restores files to a primary storage pool from copy storage pools.

REVOKE commands

Use the **REVOKE** commands to revoke privileges or access.

- “REVOKE AUTHORITY (Remove administrator authority)” on page 1033
- “REVOKE PROXYNODE (Revoke proxy authority for a client node)” on page 1037

REVOKE AUTHORITY (Remove administrator authority)

Use this command to revoke one or more privilege classes from an administrator.

You can also use this command to reduce the number of policy domains to which a restricted policy administrator has authority and the number of storage pools to which a restricted storage administrator has authority.


If you use the **REVOKE AUTHORITY** command without the **CLASSES**, **DOMAINS**, and **STGPOLS** parameters, you will revoke all privileges for the specified administrator.

At least one administrator must have system privilege; therefore, if the administrator is the only one with system privilege, you cannot revoke the authority.


Privilege class

To issue this command, you must have system privilege.

Syntax



A:



Notes:

- 1 If all these parameters are omitted, all administrator privileges will be revoked for this administrator.

Parameters

admin_name (Required)

Specifies the name of the administrator whose administrative privilege is to be revoked or reduced.

Classes

Specifies one or more administrative privilege classes to be revoked. You can specify more than one class by separating each with a comma.

REVOKE AUTHORITY

System

Indicates that system authority is to be revoked for this administrator. If CLASSES=SYSTEM is specified, no other classes can be specified, and the DOMAINS and STGPOOLS parameters cannot be specified.

Policy

Indicates that policy privilege is to be revoked for this administrator. To revoke all policy privilege, specify CLASSES=POLICY and do not specify the DOMAINS parameter.

Storage

Indicates that storage privilege is to be revoked for this administrator. To revoke all storage privilege, specify CLASSES=STORAGE and do not specify the STGPOOLS parameter.

Operator

Indicates that operator privilege is to be revoked for this administrator.

Node

Indicates that node privilege is to be revoked for this user.

AUTHORITY

Indicates the authority level to revoke for a user with node privilege. This parameter is optional.

If an administrator already has system or policy privilege to the policy domain to which the node belongs, this command will not change the administrator's privilege. Possible authority levels are:

Access

Indicates that client access authority is revoked. This is the default when CLASSES=NODE is specified.

Note: A client node can set the REVOKEREMOTEACCESS option to prevent access by a user with node privilege and client access authority. If a user with node privilege has client owner authority, or has system or policy privileges to the policy domain to which the node belongs, that administrator can still access the web backup-archive client.

Owner

Indicates that client owner authority is revoked.

Domains

Indicates that you want to revoke an administrator's client access or client owner authority to all clients in the specified policy domain. This parameter cannot be used together with the NODE parameter.

Node

Indicates that you want to revoke an administrator's client access or client owner authority to the node. This parameter cannot be used together with the DOMAIN parameter.

Domains

When used with CLASSES=POLICY, specifies a list of policy domains that can no longer be managed by a restricted policy administrator. (The administrator was authorized to manage these domains until the **REVOKE** command was issued.) This parameter is optional. The items in the list are separated by commas, with no intervening spaces. You can use wildcard characters to specify a name. Authority for all matching domains is revoked. If DOMAINS is specified, the parameter CLASSES=POLICY is optional.

STGpools

Specifies a list of storage pools that can no longer be managed by a restricted policy administrator. (The administrator had been authorized to manage these storage pools until the **REVOKE** command was issued.) This parameter is optional. The items in the list are separated by commas, with no intervening spaces. You can use wildcard characters to specify a name. Authority for all matching storage pools will be revoked. If STGPOOLS is specified then the parameter CLASSES=STORAGE is optional.

Usage notes

1. To change an unrestricted storage administrator to a restricted storage administrator, you must first use this command to revoke the unrestricted privilege. Then, use the **GRANT AUTHORITY** command to grant the administrator restricted storage privilege and to identify the storage pools to which the administrator has authority.

To revoke unrestricted storage privilege from an administrator, specify the CLASSES=STORAGE parameter. You cannot use the STGPOOLS parameter to revoke authority for selected storage pools from an unrestricted storage administrator.

2. To change an unrestricted policy administrator to a restricted policy administrator, you must first use this command to revoke the unrestricted privilege. Then, use the **GRANT AUTHORITY** command to grant the administrator restricted policy privilege and to identify the policy domains to which the administrator has authority.

To revoke unrestricted policy privilege from an administrator, specify the CLASSES=POLICY parameter. You cannot use the DOMAINS parameter to revoke authority for selected domains from an unrestricted administrator.

Example: Revoke certain administrative privileges

Revoke part of administrator CLAUDIA's privileges. CLAUDIA has restricted policy privilege for the policy domains EMPLOYEE_RECORDS and PROG1. Restrict CLAUDIA's policy privilege to the EMPLOYEE_RECORDS policy domain.

```
revoke authority claudia classes=policy
domains=employee_records
```

Example: Revoke all administrative privileges

Administrator LARRY currently has operator and restricted policy privilege. Revoke all administrative privileges for administrator LARRY. To revoke all administrative privileges for an administrator, identify the administrator, but do not specify CLASSES, DOMAINS, or STGPOOLS. LARRY remains an administrator but he can only use those commands that can be issued by any administrator.

```
revoke authority larry
```

Example: Revoke node privilege

Help desk personnel user CONNIE currently has node privilege with client owner authority for client node WARD3. Revoke her node privilege with client owner authority.

```
revoke authority connie classes=node
authority=owner node=ward3
```

REVOKE AUTHORITY

Related commands

Table 309. Commands related to REVOKE AUTHORITY

Command	Description
GRANT AUTHORITY	Assigns privilege classes to an administrator.
QUERY ADMIN	Displays information about one or more Tivoli Storage Manager administrators.

REVOKE PROXYNODE (Revoke proxy authority for a client node)

Use this command to revoke authority for an agent client node to perform backup and restore operations for a target node on the Tivoli Storage Manager server.

Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege

Syntax

```
►►—REVOke PROXynode TArget—==—target_node_name—AGent—==—agent_node_name—►►
```

Parameters

TArget (Required)

Specifies the target node to which an agent node has been granted proxy authority. Wildcard characters and comma-separated lists of node names are allowed.

AGent (Required)

Specifies which node has authority to act as proxy to the target node. Wildcard characters and comma-separated lists of node names are allowed.

Example: Revoke a node's proxy authority

To revoke authority from target node NASCLUSTER to act as proxy for all agent nodes which start with the letter M, issue the following command.

```
revoke proxynode target=nascluster agent=m*
```

Related commands

Table 310. Commands related to REVOKE PROXYNODE

Command	Description
GRANT PROXYNODE	Grant proxy authority to an agent node.
QUERY PROXYNODE	Display nodes with authority to act as proxy nodes.

ROLLBACK (Rollback uncommitted changes in a macro)

Use this command within a macro to undo any processing changes made by commands run by the server but not yet committed to the database. A committed change is permanent and cannot be rolled back. The **ROLLBACK** command is useful for testing macros.

Ensure that your administrative client session is not running with the ITEMCOMMIT option when using this command.

Important: **SETOPT** commands inside a macro cannot be rolled back.

Privilege class

Any administrator can issue this command.

Syntax

▶▶—ROLLBACK—▶▶

Parameters

None.

Example: Rollback changes in a macro

Run the REGN macro with the **ROLLBACK** command to verify that the macro works without committing any changes.

Macro Contents:

```
/* Macro to register policy
administrators and grant authority */
REGister Admin sara hobby
GRant AUTHority sara CLasses=Policy
REGister Admin ken plane
GRant AUTHority ken CLasses=Policy
ROLLBACK /* prevents any changes from being committed */
```

Related commands

Table 311. Commands related to ROLLBACK

Command	Description
COMMIT	Makes changes to the database permanent.
MACRO	Runs a specified macro file.

RUN (Run a Tivoli Storage Manager script)

Use this command to run a Tivoli Storage Manager script. To issue this command on another server, the script being run must be defined on that server.


You can include **RUN** commands in scripts as long as they do not create loops. For example, you should avoid including **RUN** commands where SCRIPT_A runs SCRIPT_B and SCRIPT_B runs SCRIPT_A.

Important: Tivoli Storage Manager does not have a command that can cancel a script after it starts. To stop a script, you must halt the server.

Privilege class

To issue this command, you must have operator, policy, system, storage, or system privilege.

Syntax



Parameters

script_name (Required)

Specifies the name of the script you want processed. The name you specify cannot be a substitution variable, such as \$1.

substitution_value

Specifies one or more values to substitute for variables when the script is run. In a script, a substitution variable consists of a '\$' character, followed by a number. When you run the script, Tivoli Storage Manager replaces the substitution variables defined in a script with the values you supply with this command. You must specify values for each substitution variable defined in the script or the script will fail. This parameter is optional.

Preview

Specifies whether to preview the command lines of a script without actually processing the script. The default is NO.

Possible values are:

Yes

Specifies that the command lines included in a script are displayed, but the script is not processed.

No

Specifies that the command lines included in a script are displayed and the script is processed.

Verbose

Specifies whether command lines, variable substitution, and conditional logic

RUN

testing used in a script are displayed as the script is being processed. This parameter is ignored if PREVIEW=YES is specified. The default is NO.

Possible values are:

Yes

Specifies that the command lines, variable substitution, and conditional logic testing are displayed as the script is being processed.

No

Specifies that the command lines, variable substitution, and conditional logic testing do not display as the script is being processed.

Example: View the commands generated by a script with a table name substitution variable

To run the following example script, called QSAMPLE, you issue a **RUN** command that specifies the table name ACTLOG as the value for the substitution variable, \$1. Use the output to preview the commands generated by the script before running the commands.

```
001 /* This is a sample SQL Query in wide format */
005 SET SQLDISPLAYMODE WIDE
010 SELECT colname FROM -
015 COLUMNS WHERE TABNAME='$1'
run qsample actlog preview=yes
```

```
ANR1461I RUN: Executing command script QSAMPLE.
ANR1466I RUN: Command script QSAMPLE, Line 5 :
           set sqldisplaymode wide.
ANR1466I RUN: Command script QSAMPLE, Line 15 :
           select colname from columns where tabname='ACTLOG'.
ANR1470I RUN: Command script QSAMPLE completed successfully
           (PREVIEW mode)
```

Example: Run a script to display and run the commands generated by the script

Run the same script as show in the prior example to display both the generated commands and the results of the commands.

```
run qsample actlog verbose=yes
```

```
ANR1461I RUN: Executing command script QSAMPLE.
ANR1466I RUN: Command script QSAMPLE, Line 5 :
           set sqldisplaymode wide.
ANR1466I RUN: Command script QSAMPLE, Line 5 : RC=RC_OK
ANR1466I RUN: Command script QSAMPLE, Line 15 :
           select colname from columns where tabname='ACTLOG'.
```

```
COLNAME
-----
DATE_TIME
MSGNO
SEVERITY
MESSAGE
ORIGINATOR
NODENAME
OWNERNAME
SCHEDNAME
DOMAINNAME
SESSID
```

```
ANR1462I RUN: Command script QSAMPLE, Line 15 : RC=RC_OK
ANR1462I RUN: Command script QSAMPLE completed successfully.
```

Example: Run a script to display just the results of the commands in the script

Run the previous example script, without displaying just the results of the generated commands in the script.

```
run qsampl actlog verbose=no
```

```
COLNAME
-----
DATE_TIME
MSGNO
SEVERITY
MESSAGE
ORIGINATOR
NODENAME
OWNERNAME
SCHEDNAME
DOMAINNAME
SESSID

ANR1462I RUN: Command script QSAMPLE completed successfully.
```

Related commands

Table 312. Commands related to RUN

Command	Description
COPY SCRIPT	Creates a copy of a script.
DEFINE SCRIPT	Defines a script to the Tivoli Storage Manager server.
DELETE SCRIPT	Deletes the script or individual lines from the script.
QUERY SCRIPT	Displays information about scripts.
RENAME SCRIPT	Renames a script to a new name.
UPDATE SCRIPT	Changes or adds lines to a script.

SELECT (Perform an SQL query of the Tivoli Storage Manager database)

Use the **SELECT** command to create and format a customized query of the Tivoli Storage Manager database.

Tivoli Storage Manager provides an SQL interface to a DB2 program. Restrictions and guidelines for handling SQL queries are handled directly by DB2.

To help you find what information is available, Tivoli Storage Manager provides three system catalog tables:

SYSCAT.TABLES

Contains information about all tables that can be queried with the **SELECT** command.

SYSCAT.COLUMNS

Describes the columns in each table.

SYSCAT.ENUMTYPES

Defines the valid values for each type and the ordering of those values for columns that have an enumerated data type (SQL93).

You can issue the **SELECT** command to query these tables to determine the location of the information that you want.

Usage notes

You cannot issue the **SELECT** command from a server console.

Because the select command does not lock and unlock records, contention for a record can cause the server to erroneously issue message ANR2034E: **SELECT: No match found using this criteria**. Check your selection criteria, and if you believe that it is correct, try the command again.

To stop the processing of a **SELECT** command after it starts, cancel the administrative session from which the command was issued. Cancel the session from either the server console or another administrative session.

Temporary table spaces are used to process SQL queries within DB2. Inadequate temporary space can cause SQL queries to fail.

To export output to a comma-separated file for import into a spreadsheet, use **-comma** and **>** command-line options on the **dsmdmc** command.

Privilege class

Any administrator can issue this command.

Syntax

For **SELECT** statement syntax and guidelines, see the DB2 Information Center: <http://pic.dhe.ibm.com/infocenter/db2luw/v9r7>. Search on *Select-statement* for usage guidelines.

Important: The appropriate syntax for the timestamp Select statement is:

```
SELECT * FROM SUMMARY WHERE ACTIVITY='EXPIRATION' AND START_TIME
>'2009-05-10 00:00:00' AND START_TIME <'2009-05-11 23:23:23'
```

List of examples

The **SELECT** command is used to customize a wide variety of queries. To give you an idea of what you can do with the command, this section includes many examples. There are, however, many more possibilities. Query output is shown only for the more complex commands to illustrate formatting.

The following list summarizes the example **SELECT** commands:

- List administrator user ID passwords that are authenticated with an external LDAP directory server
- List available tables
- List client nodes and administrative clients that are currently locked from server access
- List client nodes and administrative clients that have not specified the correct password lately
- List nodes in the standard policy domain that are not associated with the daily backup schedule DAILYBACKUP
- List the administrators that have policy authority
- List type E (ERROR) or W (WARNING) messages that have been issued in the time period for which activity log records have been maintained
- List the administrative schedules that have been defined or altered by administrator JAKE
- List the relative administrative schedule priorities
- List the management classes that have an archive copy group with a retention period greater than 365 days
- List the client nodes that are in each policy domain
- Count how many files have been archived from each node
- List the clients that are using space management
- Determine how many volumes would be reclaimed if the reclamation threshold is changed to 50 percent for storage pool TAPE
- Determine how many backup files would be affected for each node if the DAILY management class in the STANDARD policy domain is changed or deleted
- For all active client sessions, determine how long have they been connected and their effective throughput in bytes per second
- Determine how long the current background processes have been running and determine their effective throughput in time and files per second
- Count the number of client nodes are there for each platform type
- Count the number of file spaces each client node has and list the client nodes ascending order
- Obtain statistical information for calculating the number of off-site volumes that have their space reclaimed during reclamation of a storage pool
- Obtain PVU estimate detail records
- Obtain information about the node roles
- Obtain information about status

SELECT

Example: List administrator user IDs that authenticate to the Tivoli Storage Manager server

List all the administrator user IDs whose passwords authenticate with the Tivoli Storage Manager server:

```
select admin_name from admins where
authentication=local
```

Example: List available tables

List all the tables available for querying the IBM Tivoli Storage Manager database.

```
select * from syscat.tables
```

```
ABSCHEMA: SERVER1
TABNAME: ACTLOG
CREATE_TIME: 1999-05-01 07:39:06
COLCOUNT: 10
INDEX_COLCOUNT: 1
UNIQUE_INDEX: FALSE
REMARKS: Server activity log

TABSCHEMA: SERVER1
TABNAME: ADMIN_SCHEDULES
CREATE_TIME: 1995-05-01 07:39:06
COLCOUNT: 14
INDEX_COLCOUNT: 1
UNIQUE_INDEX: TRUE
REMARKS: Administrative command schedules

TABSCHEMA: SERVER1
TABNAME: ADMINS
CREATE_TIME: 1995-05-01 07:39:06
COLCOUNT: 15
INDEX_COLCOUNT: 1
UNIQUE_INDEX: TRUE
REMARKS: Server administrators

TABSCHEMA: SERVER1
TABNAME: ARCHIVES
CREATE_TIME: 1995-05-01 07:39:06
COLCOUNT: 10
INDEX_COLCOUNT: 5
UNIQUE_INDEX: FALSE
REMARKS: Client archive files
```

Example: List client nodes and administrative clients that are currently locked from server access

```
select node_name from nodes where locked='YES'
```

```
select admin_name from admins where locked='YES'
```

Example: List client nodes and administrative clients that have not specified the correct password lately

```
select node_name from nodes where invalid_pw_count <>0
```

```
select admin_name from admins where invalid_pw_count <>0
```

Example: List nodes in the standard policy domain that are not associated with the daily backup schedule DAILYBACKUP

```
select node_name from nodes where domain_name='STANDARD' and
node_name not in (select node_name from associations
where domain_name='STANDARD' and
schedule_name='DAILYBACKUP')
```

Example: List the administrators who have policy authority

```
select admin_name from admins where
upper(system_priv) <>'NO'
or upper(policy_priv) <>'NO'
```

Example: List type E (ERROR) or W (WARNING) messages that have been issued in the time period for which activity log records have been maintained

```
select date_time,msgno,message from actlog
where severity='E' or severity='W'
```

Example: List the administrative schedules that have been defined or altered by administrator JAKE

```
select schedule_name from admin_schedules
where chg_admin='JAKE'
```

Example: List the relative administrative schedule priorities

```
select schedule_name,priority from admin_schedules order
by priority
```

Example: List the management classes that have an archive copy group with a retention period greater than 365 days

```
select domain_name,set_name,class_name from ar_copygroups
where retver='NOLIMIT' or cast(retver as integer) >365
```

Example: List the management classes that specify more than five backup versions

```
select domain_name,set_name,class_name from bu_copygroups
where verexists ='NOLIMIT' or
cast(verexists as integer)>5
```

Example: List the client nodes that are using the client option set named SECURE

```
select node_name from nodes where option_set='SECURE'
```

Example: List the client nodes that are in each policy domain

```
select domain_name,num_nodes from domains
```

Example: Count how many files have been archived from each node

Attention: This command might take a long time to complete.

```
select node_name,count(*) from archives
group by node_name
```

Example: List the clients that are using space management

```
select node_name from auditocc where spacemg_mb <>0
```

Example: Determine how many volumes would be reclaimed if the reclamation threshold is changed to 50 percent for storage pool TAPE

```
select count(*) from volumes where stgpool_name='TAPE'
and upper(status)='FULL' and pct_utilized < 50
```

SELECT

Example: Determine how many backup files would be affected for each node if the DAILY management class in the STANDARD policy domain is changed or deleted

Note: This command takes significant time and resources to complete.

```
select node_name, count(*) as "Files" from backups
where class_name='DAILY' and node_name in
(select node_name from nodes where domain_name='STANDARD')
group by node_name
```

Example: For all active client sessions, determine how long have they been connected and their effective throughput in bytes per second

```
select session_id as "Session",
client_name as "Client",
state as "State",
current_timestamp-start_time as "Elapsed Time",
(cast(bytes_sent as decimal(18,0)) /
cast(second(current_timestamp-start_time) as decimal(18,0)))
as "Bytes sent/second",
(cast(bytes_received as decimal(18,0)) /
cast(second(current_timestamp-start_time) as decimal(18,0)))
as "Bytes received/second"
from sessions
```

```
Session: 24
Client: ALBERT
State: Run
Elapsed Time: 0 01:14:05.000000
Bytes sent/second: 564321.9302768451
Bytes received/second: 0.0026748857944

Session: 26
Client: MILTON
State: Run
Elapsed Time: 0 00:06:13.000000
Bytes sent/second: 1638.5284210992221
Bytes received/second: 675821.6888561849
```

Example: Determine how long the current background processes have been running and determine their effective throughput in time and files per second

Note: Expiration does not report the number of bytes processed.

```
select process_num as "Number",
process,
current_timestamp-start_time as "Elapsed Time",
(cast(files_processed as decimal(18,0)) /
cast(second(current_timestamp-start_time) as decimal(18,0)))
as "Files/second",
(cast(bytes_processed as decimal(18,0)) /
cast(second(current_timestamp-start_time) as decimal(18,0)))
as "Bytes/second"
from processes
```

```
Number: 1
PROCESS: Expiration
Elapsed Time: 0 00:24:36.000000
Files/second: 6.3216755870092
Bytes/second: 0.00000000000000
```

Example: Count the number of client nodes for each platform type

```
select platform_name,count(*) as "Number of Nodes"
from nodes group by platform_name
```

PLATFORM_NAME	Number of Nodes
AIX	6
SunOS	27
Win32	14
Linux	20

Example: Count the number of file spaces each client node has and list the client nodes ascending order

```
select node_name, count(*) as "number of filespaces"
from filespaces group by node_name order by 2
```

NODE_NAME	number of filespaces
ALBERT	2
MILTON	2
BARNEY	3
SEBASTIAN	3
MAILHOST	4
FALCON	4
WILBER	4
NEWTON	4
JEREMY	4
WATSON	5
RUSSELL	5

Example: Obtain statistical information for calculating the number of off-site volumes that have their space reclaimed during reclamation of a storage pool.

```
select * from summary where activity='OFFSITE RECLAMATION'
```

```
START_TIME: 2004-06-16 13:47:31.000000
END_TIME: 2004-06-16 13:47:34.000000
ACTIVITY: OFFSITE RECLAMATION
NUMBER: 4
ENTITY: COPYPOOL
COMMMETH:
ADDRESS:
SCHEDULE_NAME:
EXAMINED: 170
AFFECTED: 170
FAILED: 0
BYTES: 17821251
IDLE: 0
MEDIAS: 0
PROCESSES: 2
SUCCESSFUL: YES
VOLUME_NAME:
DRIVE_NAME:
LIBRARY_NAME:
LAST_USE:
COMM_WAIT:
NUM_OFFSITE_VOL: 2
```

SELECT

Example: Identify which storage pools contain data that was deduplicated by clients

```
select stgpool_name,has_client_dedup_data from stgpools
```

STGPOOL_NAME	HAS_CLIENT_DEDUP_DATA
ADPOOL	NO
ARCHIVEPOOL	NO
BACKUPPOOL	NO
COPYDEDUP	NO
COPYNODEDUP	NO
FILEPOOL	YES
FILEPOOL2	NO
LANFREEFILEPOOL	YES
SPACEMGPPOOL	NO

Example: Obtain information about the database

```
select * from db
```

```
DATABASE_NAME: TSMDB1
TOT_FILE_SYSTEM_MB: 2048000
USED_DB_SPACE_MB: 12576
FREE_SPACE_MB: 1576871
TOTAL_PAGES: 983044
USABLE_PAGES: 982908
USED_PAGES: 977736
FREE_PAGES: 5172
BUFF_HIT_RATIO: 96.2
TOTAL_BUFF_REQ: 53967
SORT_OVERFLOW: 0
LOCK_ESCALATION: 0
PKG_HIT_RATIO: 70.0
LAST_REORG: 2010-07-15 17:32:55.000000
FULL_DEV_CLASS: OUTFILE
NUM_BACKUP_INCR: 0
LAST_BACKUP_DATE: 2010-01-21 10:37:59.000000
PHYSICAL_VOLUMES: 0
PAGE_SIZE:
NUM_BACKUP_STREAMS: 4
```

Example: Obtain PVU estimate detail records

Generate the PVU estimate for a node named ACCTSRECSRV, which is used by the IBM Tivoli Storage Manager Extended Edition product.

```
select * from pvuestimate_details where node_name='ACCTSRECSRV'
```

```

PRODUCT: PRODTSMEE
LICENSE_NAME: MGSYSLAN
NODE_NAME: ACCTSRECSRV
LAST_USED: 2008-01-20 16:12:24.000000
TRYBUY: FALSE
PROC_VENDOR: IBM
PROC_BRAND: POWER5+ QCM
PROC_TYPE: 4
PROC_MODEL:
PROC_COUNT: 2
      ROLE: SERVER
ROLE_OVERRIDE: USERREPORTED
ROLE_EFFECTIVE: SERVER
VALUE_UNITS: 50
VALUE_FROM_TABLE: YES
      PVU: 100
SCAN_ERROR : NO
API_CLIENT: NO
PVU_AGNOSTIC: NO
HYPERVISOR: VMWARE
      GUID: 01.2e.1c.80.e5.04-
           .11.da.aa.ab.00.-
           15.58.0b.d9.47
VERSION: 6
RELEASE: 3
LEVEL: 1
VENDOR_D: IBM(R)
BRAND_D: POWER5(TM) QCM
TYPE_D: Quad-core Module
MODEL_D: All Existing

```

Field descriptions

PRODUCT

Rollup of license types into products at the level presented in the **QUERY PVUESTIMATE** command. Possible values are PRODTSMEE, PRODTSM, PRODSAM, PRODMAIL, PRODDDB, PRODSYSB, PRODSpace, PRODSAN, PRODERP, or blank.

LICENSE_NAME

The license assigned to this node.

NODE_NAME

The node name.

LAST_USED

Date and time the identified node last connected to the system under this license.

TRYBUY

Indicates if running under try and buy mode. Possible values are TRUE or FALSE.

PROC_VENDOR

Processor vendor name as reported by the client.

PROC_BRAND

Processor brand name as reported by the client.

PROC_TYPE

Processor type as reported by the client. This value also reflects the number of cores. Example values are 1=SINGLE CORE, 2=DUO CORE, and 4=QUAD CORE.

PROC_MODEL

Processor model as reported by the client.

SELECT

PROC_COUNT

Processor quantity.

ROLE Node role. Possible values are CLIENT, SERVER, or OTHER.

ROLE_OVERRIDE

Override value specified in the **UPDATE NODE** command.

ROLE_EFFECTIVE

Actual role based on the values in the **ROLE** and **ROLE_OVERRIDE** fields.

VALUE_UNITS

Assigned processor value unit (PVU) for the processor.

PVU Calculated PVU value.

$PVU \text{ per node} = \text{number of processors per node} * \text{processor type} * \text{pvu value}$

where the processor type represents the number of cores, and the pvu value is the value defined for the processor type in the IBM PVU table.

VALUE_FROM_TABLE

Flag that indicates whether the PVU was calculated based on the IBM PVU table. Possible values are YES or NO. If NO, a value of 100 is applied for each node defined as a server. If no role is defined for a node, the role of server is assumed for purposes of PVU calculation.

SCAN_ERROR

Flag that indicates whether license information was reported by client. Possible values are YES or NO.

API_CLIENT

Flag that indicates an API application. Possible values are YES or NO.

PVU_AGNOSTIC

Flag indicating that the client version release level is earlier than Tivoli Storage Manager V6.3. If the version is earlier than 6.3, valid PVU metrics are not expected. Possible values are YES or NO.

HYPERVISOR

Name of the virtual machine software as reported by the client.

GUID Globally Unique Identifier (GUID) of the computer where the node is located. The GUID is obtained from the node table.

VERSION

Version of client.

RELEASE

Release of client.

LEVEL

Level of client.

VENDOR_D

Processor vendor display value from the PVU table.

BRAND_D

Processor brand display value from the PVU table.

TYPE_D

Processor type display value from the PVU table.

MODEL_D

Processor model display value from the PVU table.

Example: Obtain role and PVU-related information

The following example shows partial results for a selected node, including PVU-related information and role information. Possible roles are CLIENT, SERVER, or OTHER. PVU is calculated only for nodes defined as servers.

```
select * from nodes
```

```
ROLE: CLIENT  
ROLE_0: USEREPORTED  
PVENDOR: INTEL  
PBRAND: INTEL  
PTYPE: 4  
PMODEL:  
PCOUNT: 1  
HYPERVISOR:  
PAPI: NO  
SCANERROR: NO
```

SET commands

Use the **SET** commands to specify values that affect many different Tivoli Storage Manager operations.

- “SET ACCOUNTING (Set accounting records on or off)” on page 1055
- “SET ACTLOGRETENTION (Set the retention period or the size of the activity log)” on page 1056
- “SET ALERTACTIVEDURATION (Set the duration of an active alert)” on page 1058
- “SET ALERTCLOSEDDURATION (Set the duration of a closed alert)” on page 1059
- “SET ALERTEMAIL (Set the alert monitor to email alerts to administrators)” on page 1060
- “SET ALERTEMAILFROMADDR (Set the email address of the sender)” on page 1062
- “SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)” on page 1063
- “SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)” on page 1064
- “SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)” on page 1066
- “SET ALERTMONITOR (Set the alert monitor to on or off)” on page 1067
- “SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)” on page 1065
- “SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)” on page 1068
- “SET ARCHIVERETENTIONPROTECTION (Activate data retention protection)” on page 1069
- “SET ARREPLRULEDEFAULT (Set the server replication rule for archive data)” on page 1071
- “SET AUTHENTICATION (Set password authentication)” on page 1073
- “SET BKREPLRULEDEFAULT (Set the server replication rule for backup data)” on page 1074
- “SET CLIENTACTDURATION (Set the duration period for the client action)” on page 1077
- “SET CONFIGMANAGER (Specify a configuration manager)” on page 1078
- “SET CONFIGREFRESH (Set managed server configuration refresh)” on page 1079
- “SET CONTEXTMESSAGING (Set message context reporting on or off)” on page 1080
- “SET CPUINFOREFRESH (Refresh interval for the client workstation information scan)” on page 1081
- “SET CROSSDEFINE (Specifies whether to cross-define servers)” on page 1082
- “SET DBRECOVERY (Set the device class for automatic backups)” on page 1083
- “SET DBREPORTMODE (Set the level of database reporting)” on page 1085
- “SET DEFAULTAUTHENTICATION (Set the default authentication method for **REGISTER NODE** and **REGISTER ADMIN** commands)” on page 1088
- “SET DRMACTIVEDATASTGPOOL (Specify the active-data pools to be managed by DRM)” on page 1090

- “SET DRMCHECKLABEL (Specify label checking)” on page 1091
- “SET DRMCMDFILENAME (Specify the name of a file to contain commands)” on page 1092
- “SET DRMCOPYSTGPOOL (Specify the copy storage pools to be managed by DRM)” on page 1093
- “SET DRMCOURIERNAME (Specify the courier name)” on page 1094
- “SET DRMDBBACKUPEXPIREDDAYS (Specify DB backup series expiration)” on page 1095
- “SET DRMFILEPROCESS (Specify file processing)” on page 1097
- “SET DRMINSTRPREFIX (Specify the prefix for recovery instructions file names)” on page 1098
- “SET DRMNOTMOUNTABLENAME (Specify the not mountable location name)” on page 1100
- “SET DRMPLANPREFIX (Specify a prefix for recovery plan file names)” on page 1101
- “SET DRMPLANVPOSTFIX (Specify replacement volume names)” on page 1103
- “SET DRMPRIMSTGPOOL (Specify the primary storage pools to be managed by DRM)” on page 1104
- “SET DRMRPFEXPIREDDAYS (Set criteria for recovery plan file expiration)” on page 1105
- “SET DRMVAULTNAME (Specify the vault name)” on page 1107
- “SET EVENTRETENTION (Set the retention period for event records)” on page 1108
- “SET INVALIDPWLIMIT (Set the number of invalid logon attempts)” on page 1109
- “SET DEFAULTAUTHENTICATION (Set the default authentication method for **REGISTER NODE** and **REGISTER ADMIN** commands)” on page 1088
- “SET LDAPPASSWORD (Set the LDAP password for the server)” on page 1111
- “SET LDAPUSER (Define the server administrator for the LDAP directory server)” on page 1112
- “SET LICENSEAUDITPERIOD (Set license audit period)” on page 1114
- “SET MAXCMDRETRIES (Set the maximum number of command retries)” on page 1115
- “SET MAXSCHEDESESSIONS (Set maximum scheduled sessions)” on page 1116
- “SET MINPWLENGTH (Set minimum password length)” on page 1118
- “SET PASSEXP (Set password expiration date)” on page 1121
- “SET QUERYSCHEDPERIOD (Set query period for polling client nodes)” on page 1123
- “SET RANDOMIZE (Set randomization of scheduled start times)” on page 1124
- “SET REGISTRATION (Set open or closed registration)” on page 1126
- “SET REPLRETENTION (Set the retention period for replication records)” on page 1128
- “SET REPLSERVER (Set the target replication server)” on page 1130
- “SET RETRYPERIOD (Set time between retry attempts)” on page 1132
- “SET SCHEDMODES (Select a central scheduling mode)” on page 1133
- “SET SERVERHLADDRESS (Set the high-level address of a server)” on page 1135
- “SET SERVERLLADDRESS (Set the low-level address of a server)” on page 1136

SET commands

- “SET SERVERNAME (Specify the server name)” on page 1137
- “SET SERVERPASSWORD (Set password for server)” on page 1138
- “SET SPREPLRULEDEFAULT (Set the server replication rule for space-managed data)” on page 1139
- “SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)” on page 1142
- “SET STATUSMONITOR (Specifies whether to enable status monitoring)” on page 1144
- “SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)” on page 1146
- “SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)” on page 1147
- “SET SUBFILE (Set subfile backup for client nodes)” on page 1149
- “SET SUMMARYRETENTION (Set number of days to keep data in activity summary table)” on page 1150
- “SET TAPEALERTMSG (Set tape alert messages on or off)” on page 1151
- “SET TOCLOADRETENTION (Set load retention period for table of contents)” on page 1152

SET ACCOUNTING (Set accounting records on or off)

Use this command to determine whether an accounting record is created every time a client node session ends. An accounting record tracks the amount of storage used by a client node session.

Use the `QUERY STATUS` command to determine whether accounting records are generated. At installation, this value is set to `OFF`.

The accounting records are stored in an accounting file named `dsmacnt.log`.

The environment variable, `DSMSERV_ACCOUNTING_DIR`, specifies the directory where the accounting file is located.

For more information about accounting, see the *Administrator's Guide*.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>>—Set ACCounting—[ON]
                    [OFF]

```

Parameters

ON Specifies that the server creates an accounting record every time a client node session ends.

OFF Specifies that the server does not create accounting records.

Example: Create accounting records

To create an accounting record at the end of each client node session issue the command:

```
set accounting on
```

Related commands

Table 313. Commands related to SET ACCOUNTING

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

Date

Specifies that activity log management is retention-based.

Size

Specifies that activity log management is size-based.

Example: Set the activity log retention period

Set the server to retain activity log records for 60 days. Issue the command:

```
set actlogretention 60
```

Example: Set the activity log size

Set the server to limit the size of the activity log to 300 MB. Issue the command:

```
set actlogretention 300 mgmtstyle=size
```

Related commands

Table 314. Command related to SET ACTLOGRETENTION

Command	Description
QUERY ACTLOG	Displays messages from the server activity log.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET ALERTACTIVEDURATION (Set the duration of an active alert)

Use this command to specify how long an alert remains active before it becomes inactive. If an active alert is triggered again, the duration is restarted.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set ALERTACTiveduration—*number_mins*—◄◄

Parameters

number_mins (Required)

Specifies the number of minutes that an alert remains active before it becomes inactive. Specify a value from 1 to 99999. The initial server default value is 480 minutes.

Set the duration of an active alert to one day

Issue the following command to specify that alerts remain active for 1440 minutes before they change to inactive status:

```
set alertactiveduration 1440
```

Related commands

Table 315. Commands related to SET ALERTACTIVEDURATION

Command	Description
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)" on page 1066	Specifies how long an alert remains inactive before it is closed.
"SET ALERTCLOSEDDURATION (Set the duration of a closed alert)" on page 1059	Specifies how long an alert remains closed before it is deleted.
"SET ALERTMONITOR (Set the alert monitor to on or off)" on page 1067	Specifies whether alert monitoring is set to on or off.
"SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)" on page 1068	Specifies how often the alert monitor updates and prunes alerts from the database.

SET ALERTCLOSEDDURATION (Set the duration of a closed alert)

Use this command to specify how long an alert remains closed before it is deleted.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—Set ALERTClosedduration—number_mins—◄◄
```

Parameters

number_mins (Required)

Specifies the number of minutes that an alert remains closed before it is deleted. Setting the value to 0 causes alerts to be deleted immediately after they are closed. Specify a value from 0 to 99999. The default value is set to 60 minutes when the Tivoli Storage Manager server database is initially formatted.

Delete alerts two hours after they are closed

Specify that alerts remain closed for 120 minutes before they are deleted:

```
set alertclosedduration 120
```

Related commands

Table 316. Commands related to SET ALERTCLOSEDDURATION

Command	Description
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"SET ALERTACTIVEDURATION (Set the duration of an active alert)" on page 1058	Specifies how long an alert remains active before it is moved to inactive status.
"SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)" on page 1066	Specifies how long an alert remains inactive before it is closed.
"SET ALERTMONITOR (Set the alert monitor to on or off)" on page 1067	Specifies whether alert monitoring is set to on or off.
"SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)" on page 1068	Specifies how often the alert monitor updates and prunes alerts from the database.

SET ALERTEMAIL (Set the alert monitor to email alerts to administrators)

Use this command to enable alerts to be sent to specified administrators by email.

Alert monitoring must be started, and the monitoring settings must be configured to send alerts by email.

You can use the server to email alerts to an administrator using the following set of commands:

- Verify that alert monitoring is set to on by issuing the **QUERY MONITORSETTINGS** command.
- Start the alert monitor on the server by issuing the **SET ALERTMONITOR** command.
- Add administrator names to one or more alert triggers by issuing the **DEFINE ALERTTRIGGER** or **UPDATE ALERTTRIGGER** command.
- Set up the administrator name to receive alerts by email. Administrators must have **ALERT=YES** and an **EMAILADDRESS** defined. Issue the **REGISTER ADMIN** or **UPDATE ADMIN** command.
- Define the SMTP host server that is used to send email, by issuing the **SET ALERTEMAILSMTPHOST** command.
- Define the SMTP port by issuing the **SET ALERTEMAILSMTPPORT** command.
- Define the email address of the sender of the alerts by issuing the **SET ALERTEMAILFROMADDR** command.

Privilege class

To issue this command, you must have system privilege.

Syntax

►► Set ALERTEMail ON OFF ◀◀

Parameters

ON

Specifies that alerts can be sent to specified administrators by email.

OFF

Specifies that alerts cannot be sent to specified administrators by email.

Enable alerts to be sent to the administrator when it occurs

Enable alerts to be sent by email by issuing the following command:

```
set alertemail on
```

Related commands

Table 317. Commands related to SET ALERTEMAIL

Command	Description
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.

Table 317. Commands related to SET ALERTEMAIL (continued)

Command	Description
"SET ALERTEMAILFROMADDR (Set the email address of the sender)" on page 1062	Specifies the email address of the alert sender.
"SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)" on page 1063	Specifies the SMTP mail server host name that is used to send alerts by email.
"SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)" on page 1064	Specifies the SMTP mail server port that is used to send alerts by email.
"SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)" on page 1065	Specifies the administrators that want to receive alert summaries by email.

SET ALERTEMAILFROMADDR (Set the email address of the sender)

Use this command to specify the email address of the alert sender.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set ALERTEMAILFRomaddr—*email_address*—◄◄

Parameters

email_address (Required)

Specifies the email address of the sender. Email addresses are in the form of *name@domain*. Email names, including the address, cannot exceed 64 characters in length, and the domain name cannot exceed 255 characters in length.

Specify the email address of the alert sender

Specify the email address of the sender by issuing the following command:

```
set alertemailfromaddr djadmin@mydomain.com
```

Related commands

Table 318. Commands related to SET ALERTEMAILFROMADDR

Command	Description
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"SET ALERTEMAIL (Set the alert monitor to email alerts to administrators)" on page 1060	Enables alerts to be sent by email to specified administrators.
"SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)" on page 1063	Specifies the SMTP mail server host name that is used to send alerts by email.
"SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)" on page 1064	Specifies the SMTP mail server port that is used to send alerts by email.
"SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)" on page 1065	Specifies the administrators that want to receive alert summaries by email.

SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)

Use this command to specify the Simple Mail Transfer Protocol (SMTP) mail server host name that is used to send the alert email.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set ALERTEMAILSMTPHost—*host_name*—◄◄

Parameters

host_name (Required)

Specifies the SMTP mail server host name.

Specify the host name for the SMTP mail server as mail.domain.com

Specify mail.domain.com as the SMTP mail server, by issuing the following command:

```
set alertemailsmtp host mail.domain.com
```

Related commands

Table 319. Commands related to SET ALERTEMAILSMTPHOST

Command	Description
“SET ALERTEMAIL (Set the alert monitor to email alerts to administrators” on page 1060	Enables alerts to be sent by email to specified administrators.
“SET ALERTEMAILFROMADDR (Set the email address of the sender)” on page 1062	Specifies the email address of the alert sender.
“SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)” on page 1064	Specifies the SMTP mail server port that is used to send alerts by email.
“SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)” on page 1065	Specifies the administrators that want to receive alert summaries by email.

SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)

Use this command to specify the port number for the SMTP mail server. This mail server is used to send the alerts by email.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set ALERTEMAILSMTPPort—*TCP_port*—————◄◄

Parameters

TCP_port (Required)

Specifies the port number of the SMTP mail server. Specify a value of 1 through 32767. The default port number is 25.

Specify the port number of the SMTP mail server

Specify port number 450 as your SMTP mail server by issuing the following command:

```
set alertemailsmtpport 450
```

Related commands

Table 320. Commands related to SET ALERTEMAILSMTPPORT

Command	Description
"SET ALERTEMAIL (Set the alert monitor to email alerts to administrators" on page 1060	Enables alerts to be sent by email to specified administrators.
"SET ALERTEMAILFROMADDR (Set the email address of the sender)" on page 1062	Specifies the email address of the alert sender.
"SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)" on page 1063	Specifies the SMTP mail server host name that is used to send alerts by email.
"SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)" on page 1065	Specifies the administrators that want to receive alert summaries by email.

SET ALERTSUMMARYTOADMINS (Set the list of administrators to receive alert summaries by email)

Use this command to specify the administrators that want to receive alert summaries by email, every hour.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set ALERTSUMMARYToadmins *admin_name* ◀◀

Parameters

admin_name (Required)

Specifies the administrator name that wants to receive alert summaries by email. You can specify up to three administrator names by separating them with commas and no intervening spaces.

Specify two administrators to receive alert summaries

Specify that administrators HARRY and COLIN want to receive alert summaries, by issuing the following command:

```
set alertsummarytoadmins HARRY,COLIN
```

Related commands

Table 321. Commands related to SET ALERTSUMMARYTOADMINS

Command	Description
“SET ALERTEMAIL (Set the alert monitor to email alerts to administrators” on page 1060	Enables alerts to be sent by email to specified administrators.
“SET ALERTEMAILFROMADDR (Set the email address of the sender)” on page 1062	Specifies the email address of the alert sender.
“SET ALERTEMAILSMTPHOST (Set the SMTP mail server host name)” on page 1063	Specifies the SMTP mail server host name that is used to send alerts by email.
“SET ALERTEMAILSMTPPORT (Set the SMTP mail server host port)” on page 1064	Specifies the SMTP mail server port that is used to send alerts by email.

SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)

Use this command to specify how long an alert remains inactive. After the inactive duration is past, the alert is closed.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set ALERTINactiveduration—*number_mins*—————►►

Parameters

number_mins (Required)

Specifies the number of minutes that an alert remains inactive before it becomes closed. Specify a value from 1 to 99999. The initial server default value is 480 minutes.

Change alert status from inactive to closed after 60 minutes

Issue the following command to specify that an alert remains in inactive status for 60 minutes before it changes to closed status:

```
set alertinactiveduration 60
```

Related commands

Table 322. Commands related to SET ALERTINACTIVEDURATION

Command	Description
"SET ALERTACTIVEDURATION (Set the duration of an active alert)" on page 1058	Specifies how long an alert remains active before it is moved to inactive status.
"SET ALERTCLOSEDDURATION (Set the duration of a closed alert)" on page 1059	Specifies how long an alert remains closed before it is deleted.
"SET ALERTMONITOR (Set the alert monitor to on or off)" on page 1067	Specifies whether alert monitoring is set to on or off.
"SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)" on page 1068	Specifies how often the alert monitor updates and prunes alerts from the database.

SET ALERTMONITOR (Set the alert monitor to on or off)

Use this command to turn the alert monitor on or off.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>> Set ALERTMONITOR {OFF|ON} <<<
  
```

Parameters

ON

Specifies that the Tivoli Storage Manager server monitors alerts.

OFF

Specifies that the Tivoli Storage Manager server does not monitor alerts. When the Tivoli Storage Manager server database is initially formatted, the alert monitoring setting is set to OFF.

Turn on alert monitoring

Turn on alert monitoring by issuing the following command:

```
set alertmonitor on
```

Related commands

Table 323. Commands related to SET ALERTMONITOR

Command	Description
"SET ALERTACTIVEDURATION (Set the duration of an active alert)" on page 1058	Specifies how long an alert remains inactive before it is closed.
"SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)" on page 1066	Specifies how long an alert remains inactive before it is closed.
"SET ALERTCLOSEDDURATION (Set the duration of a closed alert)" on page 1059	Specifies how long an alert remains closed before it is deleted.
"SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)" on page 1068	Specifies how often the alert monitor updates and prunes alerts from the database.

SET ALERTUPDATEINTERVAL (Set how often the alert monitor updates and prunes alerts)

Use this command to specify how often the alert monitor updates and prunes alerts that are stored in the Tivoli Storage Manager server database.

During this check interval, the alert monitor examines each alert on the server and completes the following actions:

- The alert monitor determines whether the active or inactive durations elapsed. If the specified duration elapses, the alert status is updated to the next state. For example:
 - Active to Inactive
 - Inactive to Closed
- If an alert is closed for the duration that is specified by the **SET ALERTCLOSEDDURATION** command, the alert is deleted.

You can use the **QUERY MONITORSETTINGS** command to determine whether alert monitoring is on. Use the **SET ALERTMONITOR** command to turn on alert monitoring.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
➤—Set ALERTUPDateinterval—number_mins—➤
```

Parameters

number_mins (Required)

Specifies the length of time, in minutes, that the monitor waits before alerts are updated and pruned on the server. Specify a value from 1 to 9999. The server has an initial default value of 10 minutes.

Set alert update interval to 60 minutes

Specify that alerts are updated every hour by issuing the following command:

```
set alertupdateinterval 60
```

Related commands

Table 324. Commands related to SET ALERTUPDATEINTERVAL

Command	Description
"SET ALERTACTIVEDURATION (Set the duration of an active alert)" on page 1058	Specifies how long an alert remains active before it is moved to inactive status.
"SET ALERTINACTIVEDURATION (Set the duration of an inactive alert)" on page 1066	Specifies how long an alert remains inactive before it is closed.
"SET ALERTCLOSEDDURATION (Set the duration of a closed alert)" on page 1059	Specifies how long an alert remains closed before it is deleted.
"SET ALERTMONITOR (Set the alert monitor to on or off)" on page 1067	Specifies whether alert monitoring is set to on or off.

SET ARCHIVERETENTIONPROTECTION (Activate data retention protection)

Use this command to activate and deactivate archive data retention protection. The server cannot contain any data in order for this command to work. At installation, the value is set to OFF.

When archive data retention protection is active:

- Only archive copies can be stored on the server.
- No archive copy can be deleted until the **RETVER** parameter in the **DEFINE COPYGROUP** (archive) command is satisfied.

Defining storage pools of type RECLAMATIONTYPE=SNAPLOCK is only supported on servers with data retention protection enabled.

Use the **QUERY STATUS** command to display the status of archive data retention protection.

Privilege class

To issue this command you must have system privilege or unrestricted storage privilege.

Syntax

```

>>—Set ARCHIVERETENTIONPROTECTION—OFF
ON

```

Parameters

OFF

Specifies that archive data retention protection is not active.

ON Specifies the archive data retention protection is active.

Example: Activate data retention protection

Activate archive data retention protection by issuing the following command:

```
set archiveretentionprotection on
```

Related commands

Table 325. Commands related to SET ARCHIVERETENTIONPROTECTION

Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
AUDIT VOLUME	Compares database and storage pool information, and optionally, resolves any inconsistencies.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE VOLUME	Assigns a volume to be used for storage within a specified storage pool.

SET ARCHIVERETENTIONPROTECTION

Table 325. Commands related to SET ARCHIVERETENTIONPROTECTION (continued)

Command	Description
DELETE FILESPACE	Deletes data associated with clients file spaces.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
UPDATE COPYGROUP	Changes one or more attributes of a copy group.

SET ARREPLRULEDEFAULT (Set the server replication rule for archive data)

Use this command to set the server replication rule for archive data.

Restriction: The replication rule that you set with this command is applied only if file space rules and client node rules for archive data are set to DEFAULT.

Issue this command on the server that acts as a source for replicated data.

You can specify a normal-priority replication rule or a high-priority replication rule. In a replication process that includes both normal-priority and high-priority data, high-priority data is replicated first. Before you specify a rule, consider the order in which you want the data to be replicated.

For example, suppose that your client nodes contain archive data and backup data. Replication of the archive data is a higher priority than the backup data. To prioritize the archive data, issue the **SET ARREPLRULEDEFAULT** command and specify the ALL_DATA_HIGH_PRIORITY replication rule. To prioritize the backup data, issue the **SET BKREPLRULEDEFAULT** command and specify the ALL_DATA replication rule for backup data. The ALL_DATA rule for backup data replicates backup data with a normal priority.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>> Set ARREPLRuledefault ALL_DATA
                          ALL_DATA_HIGH_PRIORITY
                          NONE
  
```

Parameters

ALL_DATA

Replicates archive data with a normal priority.

ALL_DATA_HIGH_PRIORITY

Replicates archive data with a high priority.

NONE

Archive data is not replicated.

Example: Set the server replication rule for archive data

Set up the default rule for archive data to replicate with a high priority.

```
set arreplruledefault all_data_high_priority
```

Related commands

Table 326. Commands related to SET ARREPLRULEDEFAULT

Command	Description
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.

SET ARREPLRULEDEFAULT

Table 326. Commands related to SET ARREPLRULEDEFAULT (continued)

Command	Description
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLRULE	Displays information about node replication rules.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
SET BKREPLRULEDEFAULT	Specifies the server node-replication rule for backup data.
SET SPREPLRULEDEFAULT	Specifies the server node-replication rule for space-managed data.
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE REPLRULE	Enables or disables replication rules.
VALIDATE REPLICATION	Verifies replication for file spaces and data types.

SET AUTHENTICATION (Set password authentication)

Use this command to specify whether administrators and client nodes need a password to access the server. The value is set to ON at installation.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—Set Authentication—ON  
Off—◄◄
```

Parameters

ON Specifies that administrators and client nodes need a password to access the server.

Off

Specifies that administrators and client nodes that authenticate to the IBM Tivoli Storage Manager server do not require a password to access the server. Administrators and client nodes that authenticate to an LDAP directory server are not affected.

Example: Set authentication on

Set authentication to ON to require administrators and client nodes to enter a password when accessing the server.

```
set authentication on
```

Related commands

Table 327. Command related to SET AUTHENTICATION

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET BKREPLRULEDEFAULT (Set the server replication rule for backup data)

Use this command to set the server replication rule for backup data.

Restriction: The replication rule that you set with this command is applied only if file space rules and client node rules for backup data are set to DEFAULT.

Issue this command on the server that acts as a source for replicated data.


You can specify normal-priority replication rules or high-priority replication rules. In a replication process that includes both normal and high-priority data, high-priority data is replicated first. Before you specify a rule, consider the order in which you want the data to be replicated.

For example, suppose that your client nodes contain archive data and active backup data. Replication of the active backup data is a higher priority than the archive data. To prioritize the backup data, issue the **SET BKREPLRULEDEFAULT** command and specify the **ACTIVE_DATA_HIGH_PRIORITY** replication rule. To prioritize the archive data, issue the **SET ARREPLRULEDEFAULT** command and specify the **ALL_DATA** replication rule for archive data. The **ALL_DATA** rule for archive data replicates archive data with a normal priority.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

ALL_DATA

Replicates active and inactive backup data. The data is replicated with normal priority.

ACTIVE_DATA

Replicates active backup data. The data is replicated with normal priority.

Attention:

- If you specify `ACTIVE_DATA`, inactive backup data that belongs to client nodes is not replicated.
- If the replication rules for backup data in any file spaces and any client nodes is `DEFAULT`, inactive backup data in those file spaces on the target replication server is deleted. For example, suppose the rules for backup data in file space `/a` in `NODE1` and file space `/c` in `NODE2` are `DEFAULT`. The rules for backup data in `NODE1` and `NODE2` are also `DEFAULT`. If you specify `ACTIVE_DATA` as the server rule, inactive data in file spaces `/a` and `/c` on the target server is deleted when replication occurs.

ALL_DATA_HIGH_PRIORITY

Replicates active and inactive backup data. Data is replicated with a high priority.

ACTIVE_DATA_HIGH_PRIORITY

Replicates active backup data. Data is replicated with a high priority.

Attention:

- If you specify `ACTIVE_DATA`, inactive backup data that belongs to client nodes is not replicated.
- If the replication rules for backup data in any file spaces and any client nodes is `DEFAULT`, inactive backup data in those file spaces on the target replication server is deleted. For example, suppose the rules for backup data in file space `/a` in `NODE1` and file space `/c` in `NODE2` are `DEFAULT`. The rules for backup data in `NODE1` and `NODE2` are also `DEFAULT`. If you specify `ACTIVE_DATA_HIGH_PRIORITY` as the server rule, inactive data in file spaces `/a` and `/c` is deleted.

NONE

Backup data is not replicated.

Example: Set the server replication rule for backup data

Set up the default rule for backup data to replicate only active data and to replicate the data with a high priority.

```
set bkreplruledefault active_data_high_priority
```

Related commands

Table 328. Commands related to SET BKREPLRULEDEFAULT

Command	Description
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLRULE	Displays information about node replication rules.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.

SET BKREPLRULEDEFAULT

Table 328. Commands related to SET BKREPLRULEDEFAULT (continued)

Command	Description
SET ARREPLRULEDEFAULT	Specifies the server node-replication rule for archive data.
SET REPLRETENTION	Specifies the retention period for replication history records.
SET SPREPLRULEDEFAULT	Specifies the server node-replication rule for space-managed data.
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE REPLRULE	Enables or disables replication rules.
VALIDATE REPLICATION	Verifies replication for file spaces and data types.

SET CLIENTACTDURATION (Set the duration period for the client action)

Use this command to specify the duration for the schedule that was defined with the DEFINE CLIENTACTION command. A client action defines a schedule that runs one time on a client.

The program deletes these event records whether or not the client has processed the schedule. However, the schedules are not deleted until after the first event records are deleted. The retention period for events defaults to 10 days at installation.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—SET CLIENTACTDuration—days—◄◄
```

Parameters

days (Required)

Specifies the number of days during which the schedule for the client action is active. You can specify an integer from 0 to 999. The default is 5 days.

The number of days you specify determines how long the database retains the schedule before deletion. A value of 0 indicates that the schedule duration is indefinite, and the schedule and associations are not deleted from the database.

Example: Set a 15–day duration period for the client action

To specify that the schedule for the client action be active for 15 days issue the following command.

```
set clientactduration 15
```

Related commands

Table 329. Commands related to SET CLIENTACTDURATION

Command	Description
DEFINE CLIENTACTION	Defines a command to be performed at a client node.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET CONFIGMANAGER (Specify a configuration manager)

Use this command to specify whether a server is a configuration manager. On a configuration manager, you can define configuration profiles to which other servers can subscribe.

You cannot designate a server as a configuration manager if the server subscribes to one or more profiles on another configuration manager.

If a server is a configuration manager, you cannot change this designation until you delete all profiles, including the default profile.

Issue the **QUERY STATUS** command to determine if a server is a configuration manager. When a server is installed, it is not designated as a configuration manager.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

ON Specifies that the server is a configuration manager.

When you designate a server as a configuration manager, IBM Tivoli Storage Manager creates a default profile named `DEFAULT_PROFILE` and associates with the profile all servers and server groups defined on the configuration manager. You can modify or delete the default profile.

Off

Specifies that the server is not a configuration manager.

Example: Specify a configuration manager

Designate a server as a configuration manager.

```
set configmanager on
```

Related commands

Table 330. Commands related to SET CONFIGMANAGER

Command	Description
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET CONFIGREFRESH	Specifies a time interval for managed servers to contact configuration managers.

SET CONFIGREFRESH (Set managed server configuration refresh)

Use this command on a managed server to specify how often that server contacts its configuration manager for updated configuration information.

To display the current setting, issue the **QUERY STATUS** command. At installation, the interval is set to 60 minutes.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set CONFIGRefresh—*minutes*—◀◀

Parameters

minutes (Required)

Specifies the interval, in minutes, before a managed server contacts its configuration manager for configuration updates. Specify an integer from 0 to 10000.

- If the value is greater than 0, the managed server immediately contacts the configuration manager. The next contact occurs when the specified interval is reached.
- If the value is 0, the managed server does not contact the configuration manager.

This value is ignored if the server does not subscribe to at least one profile on a configuration manager.

Example: Set a 45-minute refresh interval

Specify that a managed server contacts its configuration manager every 45 minutes.

```
set configrefresh 45
```

Related commands

Table 331. Commands related to SET CONFIGREFRESH

Command	Description
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
NOTIFY SUBSCRIBERS	Notifies servers to refresh their configuration information.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UPDATE PROFILE	Changes the description of a profile.

SET CONTEXTMESSAGING (Set message context reporting on or off)

Use this command to get additional information when ANR9999D messages occur. Tivoli Storage Manager polls the server components for information that includes process name, thread name, session ID, transaction data, locks that are held, and database tables that are in use.

Note: When consecutive messages are issued from the same code area by the same thread, only the first of these messages will report the context information.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>> Set CONTEXTmessaging { ON | Off }
    
```

Parameters

ON Specifies to enable message context reporting.

Off Specifies to disable message context reporting.

Example: Set message context reporting on or off

Turn on context messaging to receive additional information that could help determine the cause of ANR9999D messages.

```
set contextmessaging on
```

Related commands

Table 332. Commands related to SET CONTEXTMESSAGING

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET CPUINFOREFRESH (Refresh interval for the client workstation information scan)

Use this command to specify the number of days between client scans of workstation information that is used to estimate the processor value unit (PVU).

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set CPUINFOREFRESH—*days*—————►►

Parameters

days (Required)

Specifies the number of days between scans for client devices. To retrieve the current setting, issue the **QUERY STATUS** command. The possible values are 1 - 9999. The default is 180.

Example: Set the amount of time before the next refresh to 90 days

```
SET CPUINFOREFRESH 90
```

Related commands

Table 333. Commands related to SET CPUINFOREFRESH

Command	Description
QUERY PVUESTIMATE	Displays an estimate of the client-devices and server-devices being managed.

SET CROSSDEFINE

SET CROSSDEFINE (Specifies whether to cross-define servers)

Use this command to specify whether a server is automatically defined to another server.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►► Set CROSSDefine [ON | Off] ►►
```

Parameters

ON Specifies that a server may be cross-defined to another server. To automatically define one server to another, you must also permit cross defining in the server definition.

Off

Specifies that a server may not be cross-defined to another server.

Example: Specifies whether to cross-define servers

Set cross define on to allow a server to be cross-defined to another server.

```
set crossdefine on
```

Related commands

Table 334. Command related to SET CROSSDEFINE

Command	Description
DEFINE SERVER	Defines a server for server-to-server communications.
SET SERVERHLADDRESS	Specifies the high-level address of a server.
SET SERVERLLADDRESS	Specifies the low-level address of a server.
SET SERVERPASSWORD	Specifies the server password.

SET DBRECOVERY (Set the device class for automatic backups)

Use this command to specify the device class and number of data streams to be used for automatic backups.

If you issue the BACKUP DB command, and the device class is not the one that is specified in the SET DBRECOVERY command, a warning message is issued. However, the backup operation continues and is not affected.

Privilege class

To issue this command, you must have system or unrestricted storage privilege.

Syntax

```

▶▶ SET DBRECOVERY device_class_name [NUMStreams=1] [NUMStreams=number]

```

Parameters

device_class_name (Required)

Specifies the device class to use for database backups.

NUMStreams

Specifies the number of parallel data movement streams to use when backing up the database. The default value is 1, and the maximum number is 4. Increasing this value causes a corresponding increase in the number of database backup sessions to be used and in the number of drives to be used for the device class. A **NUMSTREAMS** value specified in the **BACKUP DB** command overrides any value set in the **SET DBRECOVERY** command. The **NUMSTREAMS** value is used for all types of database backups.

If a value is specified that is greater than the number of drives available for the device class, the number of available drives are used. The available drives are those defined to the device class by the MOUNTLIMIT parameter or by the number of online drives for the specified device class. The session is displayed in the **QUERY SESSION** output.

If you increase the number of streams, more volumes are used from the corresponding device class for this operation. Using more volumes might improve the speed of the database backups, but at the cost of more volumes that are not fully used.

For details about using multiple streams for database backup and restore operations, see the *Administrator's Guide*.

Example: Specify a device class for database backups

Specify the DBBACK device class for database backups. Issue the command:
 set dbrecovery dback

Example: Specify a device class and number of streams for database backups

Specify the DBBACK device class for database backups, and specify that the backup is to use two data movement streams. Issue the command:

SET DBRECOVERY

```
set dbrecovery dbback numstreams=2
```

Related commands

Table 335. Commands related to SET DBRECOVERY

Command	Description
BACKUP DB	Backs up the Tivoli Storage Manager database to sequential access volumes.
QUERY DB	Displays allocation information about the database.
QUERY DBSPACE	Displays information about the storage space defined for the database.

SET DBREPORTMODE (Set the level of database reporting)

Use this command to set the level of database diagnostic reporting to be presented on the service console and activity log.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

None

Specifies that no database diagnostic reporting is performed. This setting should not normally be used.

Partial

Specifies that the only events that are reported are those that are exceptions and that might represent errors. This is the default.

Full

Specifies that all database diagnostic information is reported. This includes information that is normal and expected.

Example: Report only events that are likely to be errors

Issue the following command:

```
set dbreportmode partial
```

Related commands

Table 336. Commands related to SET DBREPORTMODE

Command	Description
QUERY ACTLOG	Displays messages from the server activity log.

SET DEDUPVERIFICATIONLEVEL (Set the percentage of extents to verify)

Use this command to verify extents sent to the server during client-side data deduplication.

A rogue application that resides on a client system and that imitates the client, API, or GUI application can initiate an attack on the server. To reduce server vulnerability to such attacks, you can specify a percentage of client extents for the server to verify.

If the server detects that a security attack is in progress, the current session is canceled. In addition, the setting of the **DEDUPLICATION** parameter on the **REGISTER NODE** command is changed. The setting is changed from **CLIENTORSERVER** to **SERVERONLY**. The **SERVERONLY** setting disables client-side data deduplication for that node.


The server also issues a message that a potential security attack was detected and that client-side data deduplication was disabled for the node. If client-side data deduplication is disabled, all other client operations (for example, backup operations) continue. Only client-side data deduplication is disabled. If client-side data deduplication is disabled for a node because a potential attack was detected, the server deduplicates the data that is eligible for client-side data deduplication.

For more information about data deduplication, see the *Administrator's Guide*.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

percent_value (Required)

Specify an integer value 0 - 100 to indicate the percentage of client extents to be verified. A value of 0 indicates that no client extents are verified. The default for this command is 0.

Tips:

- Verifying extents consumes processing power and adversely affects server performance. For optimal performance, do not specify values greater than 10 for this command.
- To display the current value for **SET DEDUPVERIFICATIONLEVEL**, issue the **QUERY STATUS** command.

Example: Specify a minimum level of data deduplication verification

To specify that 1% of extents created during client-side data deduplication are verified, issue the following command:

```
set dedupverificationlevel 1
```

Example: Turn off data deduplication verification

To specify that none of the extents created during client-side data deduplication are verified, issue the following command:

```
set dedupverificationlevel 0
```

Related commands

Table 337. Commands related to SET DEDUPVERIFICATIONLEVEL

Command	Description
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
QUERY CONTENT	Displays information about files in a storage pool volume.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER NODE	Defines a client node to the server and sets options for that user.
UPDATE NODE	Changes the attributes associated with a client node.
UPDATE STGPOOL	Changes the attributes of a storage pool.

SET DEFAULTAUTHENTICATION

Table 338. Commands related to SET DEFAULTAUTHENTICATION (continued)

Command	Description
SET LDAPUSER	Sets the user who oversees the passwords and administrators on the LDAP directory server.
REGISTER ADMIN	Defines a new administrator without granting administrative authority.
REGISTER NODE	Defines a client node to the server and sets options for that user.

SET DRMACTIVEDATASTGPOOL (Specify the active-data pools to be managed by DRM)

Use this command to specify names of the active-data pools to be recovered after a disaster. Tivoli Storage Manager uses these names if the **PREPARE**, **MOVE DRMEDIA**, or **QUERY DRMEDIA** command does not include the **ACTIVEDATASTGPOOL** parameter.


By default, volumes in active-data pools are not eligible for processing by disaster recovery manager. To process active-data pool volumes, you must issue the **SET DRMACTIVEDATASTGPOOL** command, or you must use the **ACTIVEDATASTGPOOL** command-line parameter on the **MOVE DRMEDIA**, **QUERY DRMEDIA**, or **PREPARE** command.

Use the **QUERY DRMSTATUS** command to display the current settings.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

active-data_pool_name (Required)

Specifies the active-data pool names. Separate multiple names with commas with no intervening spaces. You can use wildcard characters. The specified names will overwrite any previous settings. If you enter a null string (""), all current names are removed, and no active-data pool volumes in **MOUNTABLE** state are processed if they were not explicitly entered as **MOVE DRMEDIA**, **QUERY DRMEDIA**, or **PREPARE** command parameters.

Example: Set an eligible active-data pool

Set **ACTIVEDATAPOOL1** as the eligible active-data pool.

```
set drmactivedatapool activedatastgpool1
```

Related commands

Table 339. Commands related to **SET DRMACTIVEDATASTGPOOL**

Command	Description
MOVE DRMEDIA	Moves DRM media on-site and off-site.
PREPARE	Creates a recovery plan file.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.
SET DRMCOPYSTGPOOL	Specifies that copy storage pools are managed by DRM.
SET DRMPRIMSTGPOOL	Specifies that primary storage pools are managed by DRM.

SET DRMCHECKLABEL (Specify label checking)

Use this command to specify whether Tivoli Storage Manager reads the labels of sequential media checked out by the **MOVE DRMEDIA** command. At installation, the value of the DRMCHECKLABEL is set to YES.

Use the **QUERY DRMSTATUS** command to check the current setting.

This command does not apply to 349X device types.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

Yes

Specifies that Tivoli Storage Manager reads the labels of sequential media checked out by the **MOVE DRMEDIA** command.

No Specifies that Tivoli Storage Manager does not read the labels of sequential media checked out by the **MOVE DRMEDIA** command.

Example: Specify no label checking

Specify that Tivoli Storage Manager does not perform label checking.

```
set drmchecklabel no
```

Related commands

Table 340. Commands related to SET DRMCHECKLABEL

Command	Description
MOVE DRMEDIA	Moves DRM media on-site and off-site.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMCMDFILENAME (Specify the name of a file to contain commands)

Use this command to name a file that can contain the commands created when the **MOVE DRMEDIA** or **QUERY DRMEDIA** commands are issued. If the **SET DRMCMDFILENAME** is not issued, the **MOVE DRMEDIA** or **QUERY DRMEDIA** command generates a file name.

Use the **QUERY DRMSTATUS** command to display the current command file name.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—Set DRMCMDFilename—file_name—►►
```

Parameters

file_name (Required)

Specifies a full path name for a file to contain the commands created by the **MOVE DRMEDIA** or **QUERY DRMEDIA** command.

Attention: If a file of the same name already exists, **MOVE DRMEDIA** or **QUERY DRMEDIA** command tries to use it, and the existing data is overwritten.

Example: Specify a file name to contain DRMEDIA commands

Specify a file name of `/adsm/drm/orm/exec.cmds`.

```
set drmcmdfilename /adsm/drm/orm/exec.cmds
```

Related commands

Table 341. Commands related to SET DRMCMDFILENAME

Command	Description
MOVE DRMEDIA	Moves DRM media on-site and off-site.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMCOPYSTGPOOL (Specify the copy storage pools to be managed by DRM)

Use this command to specify names of the copy storage pools to be recovered after a disaster. Tivoli Storage Manager uses these names if the **PREPARE** command does not include the **COPYSTGPOOL** parameter.

If the **MOVE DRMEDIA** or **QUERY DRMEDIA** command does not include the **COPYSTGPOOL** parameter, the command processes the volumes in the **MOUNTABLE** state that are in the copy storage pool named by the **SET DRMCOPYSTGPOOL** command. At installation, all copy storage pools are eligible for DRM processing.

Use the **QUERY DRMSTATUS** command to display the current settings.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

▶▶—Set DRMCOPYstgpool—copy_pool_name—▶▶

```

Parameters

copy_pool_name (Required)

Specifies the copy storage pool names. Separate multiple names with commas and no intervening spaces. You can use wildcard characters. The specified names replace any previous setting. If you enter a null string (""), all current names are removed, and all copy storage pools are eligible for processing.

Example: Set an eligible copy storage pool

Set **COPYSTGPOOL1** as the eligible copy storage pool.

```
set drmcopystgpool copystgpool1
```

Related commands

Table 342. Commands related to SET DRMCOPYSTGPOOL

Command	Description
MOVE DRMEDIA	Moves DRM media on-site and off-site.
PREPARE	Creates a recovery plan file.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.
SET DRMPRIMSTGPOOL	Specifies that primary storage pools are managed by DRM.

SET DRMCOURIERNAME (Specify the courier name)

Use this command to specify the courier name. At installation, this name is set to COURIER. The **MOVE DRMEDIA** command uses the courier name to set the location of volumes that are moving to the COURIER state.

You can use the **QUERY DRMSTATUS** to see the name of the courier.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set DRMCOURiername—*courier_name*—◄◄

Parameters

courier_name (Required)

Specifies the name of the courier. The name can be up to 255 characters. Enclose the name in quotation marks if it contains any blank characters.

Example: Set the courier name

Set the name of the courier to Joe's Courier Service.

```
set drmcouriername "Joe's Courier Service"
```

Related commands

Table 343. Commands related to SET DRMCOURIERNAME

Command	Description
MOVE DRMEDIA	Moves DRM media on-site and off-site.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMDBBACKUPEXPIREDAYS (Specify DB backup series expiration)

Use this command to specify when a database backup series is eligible to be expired.

The value set by this command applies to both a snapshot and a full plus incremental database backup series. Any type of database backup series is eligible for expiration if all of the following are true:

- The age of the last volume of the series has exceeded the expiration value set by this command.
- For volumes that are not virtual volumes, all volumes in the series are in the VAULT state.
- The volume is not part of the most recent database backup series.

Remember: The most recent backup series of either type is not deleted. See the **MOVE DRMEDIA** command for more information on the expiration of database backup volumes that are not virtual volumes. See the **EXPIRE INVENTORY** command for more information on expiration of database backup volumes that are virtual volumes.

Use the **QUERY DRMSTATUS** to see the number of days specified.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set DRMDBBackupexpiredays—*days*—◄◄

Parameters

days (Required)

Specifies the number of days that must elapse since a database series was created before it is eligible to be expired. The number of days must match the volume reuse delay period for copy storage pools managed by disaster recovery manager. Specify an integer from 0 to 9999.

Example: Set the database backup series expiration

Set the database backup series expiration value to 60.

```
set drmdbbackupexpiredays 60
```

Related commands

Table 344. Commands related to SET DRMDBBACKUPEXPIREDAYS

Command	Description
DSMSERV RESTORE DB	Restores a Tivoli Storage Manager database.
MOVE DRMEDIA	Moves DRM media on-site and off-site.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMDBBACKUPEXPIREDAYS

Table 344. Commands related to SET DRMDBBACKUPEXPIREDAYS (continued)

Command	Description
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.

SET DRMINSTRPREFIX (Specify the prefix for recovery instructions file names)

Use this command to specify a prefix to the recovery instructions file name. If you issue this command, IBM Tivoli Storage Manager uses the specified prefix if the **PREPARE** command is issued without the INSTRPREFIX parameter.

Use the **QUERY DRMSTATUS** command to display the current value for the prefix.

the prefix is the current Tivoli Storage Manager server working directory.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set DRMINSTRPrefix—*prefix*—————►►

Parameters

prefix (Required)

Specifies a path name prefix for the files that contain the recovery instructions. When processing the **PREPARE** command, Tivoli Storage Manager appends the name of the appropriate recovery plan file stanza to find the file. The maximum length is 250 characters.

The prefix can be one of the following:

- **Directory path:** End the prefix with a forward slash (/). For example:

```
/admsrv/recinstr/
```

For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name would be:

```
/admsrv/recinstr/RECOVERY.INSTRUCTIONS.GENERAL
```

- **Directory path followed by a string:** Tivoli Storage Manager treats the string as part of the file name. For example:

```
/admsrv/recinstr/accounts
```

For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name would be:

```
/admsrv/recinstr/accounts.RECOVERY.INSTRUCTIONS.GENERAL
```

- **String only:** Tivoli Storage Manager specifies the directory path and appends the appropriate recovery plan file stanza name.
 - Tivoli Storage Manager uses the name of the current working directory. For example, the current working directory is `/opt/tivoli/tsm/server/bin`. You specify the following:

```
shipping
```

For the RECOVERY.INSTRUCTIONS.GENERAL file, the resulting file name would look like this:

```
/opt/tivoli/tsm/server/bin/shipping.RECOVERY.INSTRUCTIONS.GENERAL
```

Example: Specify the recovery plan prefix

Specify reading the recovery plan instructions from directory /drmp1an/primesrv.
set drminstrprefix /drmp1an/primesrv/

Related commands

Table 346. Commands related to SET DRMINSTRPREFIX

Command	Description
PREPARE	Creates a recovery plan file.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMNOTMOUNTABLENAME (Specify the not mountable location name)

Use this command to specify the name of the onsite location for storing the media. At installation, the name is set to NOTMOUNTABLE. Use the **QUERY DRMSTATUS** command to see the location name.

The location name is used by the **MOVE DRMEDIA** command to set the location of volumes that are moving to the NOTMOUNTABLE state.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set DRMNOTMOUNTABLENAME—*location*—◄◄

Parameters

location (Required)

Specifies the name of the onsite location for storing the media. The name can be up to 255 characters. Enclose the name in quotation marks if it contains any blank characters.

Example: Specify the name of the onsite location

Set the name of the location to room 123/31.

```
set drmnountable "room 123/31"
```

Related commands

Table 347. Commands related to SET DRMNOTMOUNTABLENAME

Command	Description
MOVE DRMEDIA	Moves DRM media on-site and off-site.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMPLANPREFIX (Specify a prefix for recovery plan file names)

Use this command to specify a prefix for a recovery plan file name.

If you issue this command, Tivoli Storage Manager uses the specified prefix if the **PREPARE** command does not include the PLANPREFIX parameter.

Use the **QUERY DRMSTATUS** command to display the current value for the recovery plan prefix.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set DRMPLANPrefix—*prefix*—►►

Parameters

prefix (Required)

Specifies the prefix for a recovery plan file name. The maximum length of the prefix is 250 characters. If you enter a null string (""), the current prefix is removed, and the server uses the algorithm described in the PLANPREFIX parameter in the **PREPARE** command.

For the prefix, you can specify:

- **A directory path followed by a forward slash (/):** Tivoli Storage Manager appends to the prefix the date and time in the `yyyymmdd.hhmmss` format. For example, the **SET DRMPLANPREFIX** is set to the following:

```
/admsrv/recplans/
```

The resulting recovery plan file name is:

```
/admsrv/recplans/19971115.051421
```

- **A directory path followed by a string:** Tivoli Storage Manager uses the string as part of the file name. Tivoli Storage Manager appends to the prefix the date and time in the `.yyyymmdd.hhmmss` format (note the initial period). For example, the **SET DRMPLANPREFIX** is set to the following:

```
/admsrv/recplans/accounting
```

The resulting recovery plan filename is:

```
/admsrv/recplans/accounting.19971115.051421
```

- **A string that is not preceded by a directory path:** Tivoli Storage Manager appends to the prefix the date and time information in the `.yyyymmdd.hhmmss` format (note the initial period). Tivoli Storage Manager determines the directory path as follows:
 - Tivoli Storage Manager uses the directory path name of the current working directory of the Tivoli Storage Manager server. For example, the current Tivoli Storage Manager working directory is `/opt/tivoli/tsm/server/bin`. The **SET DRMPLANPREFIX** command is set to the following:


```
shipping
```

The resulting recovery plan file name is:

SET DRMPLANPREFIX

/opt/tivoli/tsm/server/bin/shipping.19971115.051421

Example: Specify a prefix for recovery plan file names

Specify a prefix so that the generated recovery plan files are stored in the following directory:

/drmp1an/primsrv

Issue the command:

```
set drmp1anprefix /drmp1an/primsrv/
```

Related commands

Table 348. Commands related to SET DRMPLANPREFIX

Command	Description
PREPARE	Creates a recovery plan file.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMPLANVPOSTFIX (Specify replacement volume names)

Use this command to specify the character to be appended to replacement volume names in the recovery plan file. The character can help you find or generate replacement volume names when you use the recovery plan file.

At installation, the character is set to @. Tivoli Storage Manager generates replacement names for primary storage pool volumes that were added by the **DEFINE VOLUME** command. Use the appended character to:

- Find replacement volume names in the recovery plan stanzas so that you can change the names at recovery time. For example, you may not know the names of the available tape volumes at the recovery site.
- Generate replacement volume names. You need a naming convention that works for any device type in your primary storage pools. Consider the following:
 - The generated length of replacement volume name
 - Legal characters in the replacement volume name
 - Conflicts with existing volume names
 - A replacement volume name must be different from any destroyed, existing, or new volume name.

Use the **QUERY DRMSTATUS** command to see the character added to the end of the replacement volume names.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set DRMPLANVpostfix—*character*—————►►

Parameters

character (Required)

Specifies the character appended to the replacement volume names in the recovery plan file. Specify an alphanumeric or special character.

Example: Specify the appended character for replacement volume names

Set the character appended to the replace volume names to R.

```
set drmplnvpostfix R
```

Related commands

Table 349. Commands related to SET DRMPLANVPOSTFIX

Command	Description
PREPARE	Creates a recovery plan file.
QUERY DRMSTATUS	Displays DRM system parameters.

SET DRMPRIMSTGPOOL (Specify the primary storage pools to be managed by DRM)

Use this command to specify the names of primary storage pools that you want to recover. If the **PREPARE** command does not include the PRIMSTGPOOL parameter, DRM processes the names specified in this command.

Use the **QUERY DRMSTATUS** command to display the current settings. At installation, all primary storage pools defined to the server are eligible for DRM processing.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

▶▶—Set DRMPRIMstgpool—primary_pool_name—▶▶

```

Parameters

primary_pool_name (Required)

Specifies the names of the primary storage pool names you want to recover. Separate multiple names with commas and no intervening spaces. You can use wildcard characters to specify names. The names that you specify replace any previous setting. If you enter a null string (""), all current names are removed, and all primary storage pools are eligible for DRM processing.

Example: Set a primary storage pool to be managed by DRM

Set the primary storage pool to be managed by DRM to PRIMSTGPOOL1.

```
set drmprimstgpool primstgpool1
```

Related commands

Table 350. Commands related to SET DRMPRIMSTGPOOL

Command	Description
PREPARE	Creates a recovery plan file.
QUERY DRMSTATUS	Displays DRM system parameters.
SET DRMCOPYSTGPOOL	Specifies that copy storage pools are managed by DRM.

SET DRMRPFEXPIREDDAYS (Set criteria for recovery plan file expiration)

Use this command to specify when recovery plan files are eligible for expiration. This command and expiration processing apply only to recovery plan files that were created with the DEVCLASS parameter specified on the **PREPARE** command (that is, virtual volumes of type RPFIL and RPSNAPSHOT). Expiration processing on the source server expires plan files stored on the target server. Locally created recovery plan files are not expired.

An RPFIL file is associated with a full plus incremental database backup series. An RPSNAPSHOT file is associated with a database snapshot backup series.

Attention: The latest RPFIL and RPSNAPSHOT files are never deleted.

A recovery plan file is eligible for expiration if both of the following are true:

- The last recovery plan file of the series has exceeded the expiration value specified with the **SET DRMRPFEXPIREDDAYS** command.
- The latest recovery plan file is not associated with the most recent database backup series.

See the **EXPIRE INVENTORY** command for more information about expiration processing.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—Set DRMRPFExpiredays—days—————►►
```

Parameters

days (Required)

Specifies the number of days that must elapse before a recovery plan file expires. You can specify a number from 0 to 9999. At installation, this value is set to 60.

Example: Set the recovery plan expiration

Set the recovery plan file expiration value to 30.

```
set drmrpfexpireddays 30
```

Related commands

Table 351. Commands related to SET DRMRPFEXPIREDDAYS

Command	Description
PREPARE	Creates a recovery plan file.
QUERY DRMSTATUS	Displays DRM system parameters.
QUERY RPFCONTENT	Displays the contents of a recovery plan file.
QUERY RPFIL	Displays information about recovery plan files.

SET DRMRPFEXPIREDAYS

Table 351. Commands related to SET DRMRPFEXPIREDAYS (continued)

Command	Description
QUERY VOLHISTORY	Displays sequential volume history information that has been collected by the server.
SET DRMDBBACKUPEXPIREDAYS	Specifies criteria for database backup series expiration.

SET DRMVAULTNAME (Specify the vault name)

Use this command to specify the vault name. At installation the name is set to VAULT. Use the **QUERY DRMSTATUS** command to see the name of the vault.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—SET DRMAu1tname—vault_name—◄◄
```

Parameters

vault_name (Required)

Specifies the name of the vault. The name can be up to 255 characters. Enclose the name in quotation marks if it contains any blank characters.

Example: Specify a vault name

Specify ironmountain as the vault name.

```
set drmvaultname ironmountain
```

Related commands

Table 352. Commands related to SET DRMVAULTNAME

Command	Description
MOVE DRMEDIA	Moves DRM media on-site and off-site.
QUERY DRMEDIA	Displays information about disaster recovery volumes.
QUERY DRMSTATUS	Displays DRM system parameters.

SET EVENTRETENTION (Set the retention period for event records)

Use this command to set the retention period for event records in the server database that will allow you to monitor completed schedules. An event record is created whenever processing of a scheduled command is started or missed.

You can adjust the length of time that the server maintains event information to avoid insufficient or outdated data. The server automatically removes the event records from the database after the retention period passes and the startup window for the event has elapsed.

You can issue the **QUERY EVENT** command to display information about scheduled and completed events.

You can issue the **DELETE EVENT** command to delete event records regardless of whether their retention period has passed.

You can issue the **QUERY STATUS** command to display the value for the event retention period. At installation, this value is set to 10 days.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set Eventretention—*days*—◄◄

Parameters

days (Required)

The number of days that the database retains event records. You can specify an integer from 0 to 9999. A value of 0 indicates that only event records for the current day are retained.

Example: Set the retention period for event records

Set the retention period to 15 days.

```
set eventretention 15
```

Related commands

Table 353. Commands related to SET EVENTRETENTION

Command	Description
DELETE EVENT	Deletes event records before a specified date and time.
QUERY EVENT	Displays information about scheduled and completed events for selected clients.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET INVALIDPWLIMIT (Set the number of invalid logon attempts)

Use this command to set the number of invalid logon attempts that are allowed before a node is locked.

The **SET INVALIDPWLIMIT** command also applies to LDAP directory servers that store complex node passwords. LDAP directory servers can limit the number of invalid password attempts independent of the IBM Tivoli Storage Manager server. You might not want to set up the LDAP directory server for invalid attempts for the Tivoli Storage Manager namespace if you use the **SET INVALIDPWLIMIT** command.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
➤—Set—INVALIDPwlimit—number—➤
```

Parameters

number (Required)

Specifies the number of invalid logon attempts allowed before a node is locked.

You can specify an integer from 0 to 9999. A value of 0 means that invalid logon attempts are not checked. A value of 1 means that if a user issues an invalid password one time, the node is locked by the server. The default is 0.

Important: If your password is authenticated with an LDAP directory server, it can be managed by the LDAP server and the Tivoli Storage Manager server. Not all Tivoli Storage Manager server commands affect passwords that authenticate with an LDAP server. For example, the **SET PASSEXP** and **RESET PASSEXP** commands do not affect passwords that authenticate with an LDAP directory server. You can manage your password features through the Tivoli Storage Manager server. If you issued the **SET INVALIDPWLIMIT** command, all Tivoli Storage Manager passwords are controlled by the limit that you set. If you configure the LDAP directory server to limit the number of invalid password attempts, a conflict might occur.

Example: Define the number of allowed invalid login attempts

Set the number of invalid logon attempts allowed.

```
set invalidpwlimit 6
```

Related commands

Table 354. Commands related to SET INVALIDPWLIMIT

Command	Description
QUERY ADMIN	Displays information about one or more Tivoli Storage Manager administrators.
QUERY NODE	Displays partial or complete information about one or more clients.

SET INVALIDPWLIMIT

Table 354. Commands related to SET INVALIDPWLIMIT (continued)

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET MINPWLENGTH	Sets the minimum length for client passwords.

SET LDAPPASSWORD (Set the LDAP password for the server)

Use this command to define a password for the user name or ID that you specify in the **SET LDAPUSER** command.

This user can then gain access to the Tivoli Storage Manager space on the LDAP directory server to issue administrative operations, such as resetting passwords.

Requirement: You must define the **LDAPURL** option and issue the **SET LDAPUSER** command before submitting the **SET LDAPPASSWORD** command. If the **LDAPURL** option is not defined when you set the **LDAPPASSWORD**, you must restart the Tivoli Storage Manager server after defining the **LDAPURL** option.

Privilege class

To issue this command you must have system privilege.

Syntax

```
►►—Set LDAPPASSWORD—ldap_user_password—◄◄
```

Parameters

ldap_user_password

Specifies the password that the server uses when it authenticates to the LDAP directory server. The maximum length of the password is 64 characters. If you have equal signs within your password, you must contain the whole password within quotation marks in the **SET LDAPPASSWORD** command. You can use the following characters:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
0 1 2 3 4 5 6 7 8 9
~ ! @ # $ % ^ & * _ - + = ` | ( ) { } [ ] : ; < > , . ? /
```

Example: Set the LDAP password

```
set ldappassword LdAp20&12PaSsWoRd
```

Example: Set an LDAP password that includes an equal sign

```
set ldappassword "LdAp=LastWoRd"
```

Related commands

Table 355. Commands related to SET LDAPPASSWORD

Command	Description
AUDIT LDAPDIRECTORY	Audit an IBM Tivoli Storage Manager-controlled namespace on an LDAP directory server.
SET DEFAULTAUTHENTICATION	Specifies the default password authentication method for any REGISTER NODE or REGISTER ADMIN commands.
SET LDAPUSER	Sets the user who oversees the passwords and administrators on the LDAP directory server.

SET LDAPUSER (Define the server administrator for the LDAP directory server)

Use this command to specify the ID of a user or account that can conduct Tivoli Storage Manager administrative operations when accessing the LDAP directory server. The user ID must have access to the base distinguished name (Base DN) on the LDAP directory server.

This command determines which ID issues administrative operations, such as resetting node passwords or entering user accounts, on the LDAP directory server. You must grant Base DN access to the **LDAPUSER**. The ID that is defined in the **SET LDAPUSER** command must be able to create, delete, and modify objects under the Base DN.

See the *Administrator's Guide* for information about access control lists (ACLs) and accessing an LDAP directory server to authenticate passwords.

To change the **LDAPUSER**, reissue the **SET LDAPUSER** command.

Privilege class

To issue this command you must have system privilege.

Syntax

```
▶▶—Set LDAPUser—ldap_user_dn—————▶▶
```

Parameters

ldap_user_dn

Specifies the user DN that is used when the Tivoli Storage Manager server needs to issue administrative commands to the Tivoli Storage Manager namespace on the LDAP directory server. If the user DN contains spaces or equal signs (=), surround the entire DN with quotation marks. The maximum length of the DN is 256 characters.

If you do not know the DN when you enter the **SET LDAPUSER** command, contact your LDAP administrator. If you use Microsoft Active Directory, you can determine the DN by using one of the following methods:

Base your query on the name attribute

Issue the **dsquery user -name tsmuser** command.

Base your query on the user principal name (UPN)

Issue the **dsquery user -upn tsmuser@tsmdev.storage.tucson.ibm.com** command.

Base your query on the SAM account name

dsquery user -samid tsmuser

For more information about Tivoli Directory Server and defining the DN, see Distinguished names (DNs).

Example: Define the LDAPUSER gnuthium

```
set ldapuser "cn=gnuthium,ou=manufacturing,o=ibm,c=us"
```

This command sets the user DN with a common name (cn) of gnuthium, with organizational unit (ou) as manufacturing, organization (o) as IBM, and country (c) as the United States.

Important: Contain your user DN within single or double quotation marks.

Related commands

Table 356. Commands related to SET LDAPUSER

Command	Description
AUDIT LDAPDIRECTORY	Audit an IBM Tivoli Storage Manager-controlled namespace on an LDAP directory server.
SET DEFAULTAUTHENTICATION	Specifies the default password authentication method for any REGISTER NODE or REGISTER ADMIN commands.
SET LDAPPASSWORD	Sets the password for the LDAPUSER.

SET LICENSEAUDITPERIOD (Set license audit period)

Use this command to specify the period, in days, between automatic license audits performed by IBM Tivoli Storage Manager.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

days

Specifies the number of days between automatic server license audits. This parameter is optional. The default value is 30. You can specify an integer from 1 to 30, inclusive.

Example: Specify a 14 day server license audit

Specify that the server audits licenses every 14 days.

```
set licenseauditperiod 14
```

Related commands

Table 357. Commands related to SET LICENSEAUDITPERIOD

Command	Description
AUDIT LICENSES	Verifies compliance with defined licenses.
QUERY AUDITOCCUPANCY	Displays the server storage utilization for a client node.
QUERY LICENSE	Displays information about licenses and audits.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER LICENSE	Registers a license with the Tivoli Storage Manager server.

SET MAXCMDRETRIES (Set the maximum number of command retries)

Use this command to set the maximum number of times that a scheduler on a client node can retry a failed, scheduled command.

You can use the command to override the maximum number of retries that are specified by the client node. A client's value is overridden only if the client is able to connect with the server.

This command is used with the **SET RETRYPERIOD** command to regulate the time and the number of retry attempts to rerun failed command.

You can issue the **QUERY STATUS** command to display the current retry value. At installation, Tivoli Storage Manager is configured so that each client determines its own retry value.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>> Set MAXCMDRetries number
  
```

Parameters

number

Specifies the maximum number of times the scheduler on a client node can retry a failed scheduled command. This parameter is optional.

The default is that each client determines its own value for this parameter. You can specify an integer from 0 to 9999. See the appropriate client documentation for more information on setting the maximum command retries from the client.

Example: Set the maximum number of command retries to 2

Retry, only twice, a failed attempt to process a scheduled command.

```
set maxcmdretries 2
```

Related commands

Table 358. Command related to SET MAXCMDRETRIES

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET RETRYPERIOD	Specifies the time between retry attempts by the client scheduler.

SET MAXSCHEDESESSIONS (Set maximum scheduled sessions)

Use this command to set the number of sessions that the server can use to process scheduled operations. This command specifies the maximum number of scheduled sessions as a percentage of the total number of available server sessions.

Limiting the number of sessions ensures that some are available for unscheduled operations, such as backup or archive. You can increase either the total number of sessions (with the `MAXSESSIONS` parameter) or the maximum percentage of scheduled sessions. Increasing the total number of sessions available, however, can affect server performance. Increasing the maximum percentage of scheduled sessions can reduce the sessions available for unscheduled operations.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
➤—Set MAXSCHedsessions—percent—➤
```

Parameters

percent (Required)

Specifies the percentage of total server sessions that can be used for scheduled operations. You can specify an integer from 0 to 100. The `MAXSESSIONS` parameter in the server options file determines the maximum number of total available server sessions.

If you set the maximum percentage of scheduled sessions to 0, no scheduled events can begin. If you set the maximum percentage of scheduled sessions to 100, the maximum number of scheduled sessions is the value of the `MAXSESSIONS` option.

Tip: If the maximum number of scheduled sessions do not coincide with the percentage that you set in the `SET MAXSCHEDESESSIONS` command, run the `SET MAXSCHEDESESSIONS` command again. Look in the `MAXSESSIONS` option and determine the number that is specified there. If the `MAXSESSIONS` option number changed and you did not issue the `SET MAXSCHEDESESSIONS` command since the change, the maximum number of scheduled sessions can change.

Set a maximum of 20 sessions for scheduled activities

The `MAXSESSIONS` option has a value of 80. If you want no more than 20 sessions to be available for scheduled activity, set the percentage to 25.

```
set maxschedsessions 25
```

Related commands

Table 359. Commands related to SET MAXSCHEDESESSIONS

Command	Description
QUERY OPTION	Displays information about server options.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET MINPWLENGTH

SET MINPWLENGTH (Set minimum password length)

Use this command to set the minimum length of a password.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set—MINPwlength—*length*—◄◄

Parameters

length (Required)

Specifies the minimum length of a password. You can specify an integer from 0 to 64. A value of 0 means that the password length is not checked. The default value for minimum password length is set to 0.

Example: Set the minimum password length

Set the minimum password length to 5 characters.

```
set minpwlenght 5
```

Related commands

Table 360. Commands related to SET MINPWLENGTH

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET INVALIDPWLIMIT	Sets the number of invalid logon attempts before a node is locked.

SET MONITORINGADMIN (Set the name of the monitoring administrator)

Use this command to set the name of the monitoring administrator that is used to connect to the servers in the monitored server group.

To display the name of the monitored server group, issue the **QUERY MONITORSETTINGS** command.

The administrator name that you specify must match the name of an existing administrator, otherwise the command fails.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>> Set MONITORINGADMIN admin_name

```

Parameters

admin_name

Specifies administrator names. You can remove names by issuing the command without specifying a value, or by specifying an empty value ("").

Set the monitoring administrator name

Set the name of the monitoring administrator to MONADMIN, by issuing the following command:

```
set monitoringadmin monadmin
```

Remove the monitoring administrator name

Remove the monitoring administrator, by issuing the following command:

```
set monitoringadmin ""
```

Related commands

Table 361. Commands related to SET MONITORINGADMIN

Command	Description
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"SET MONITOREDSEVERGROUP (Set the group of monitored servers)" on page 1120	Set the group of monitored servers.

SET MONITOREDSEVERGROUP (Set the group of monitored servers)

Use this command to set the group of servers that are being monitored for alerts and status. You can also use this command to change or remove the group of monitored servers.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>>—Set MONITOREDSEVERGroup group_name

```

Parameters

group_name

Specifies the Tivoli Storage Manager server group name that contains all monitored servers. You can remove a monitored server group name by issuing the command without specifying a value, or by specifying an empty value (""). Any existing monitoring for alerts and status from remote servers is ended.

Set the name of a monitored server group

Set the name of a monitored server group SUBS, by issuing the following command:

```
set monitoredservergroup subs
```

Remove the name of a monitored server group

Remove the monitored server group, by issuing the following command:

```
set monitoredservergroup
```

Related commands

Table 362. Commands related to SET MONITOREDSEVERGROUP

Command	Description
"DEFINE SERVERGROUP (Define a server group)" on page 333	Defines a new server group.
"DEFINE GRPMEMBER (Add a server to a server group)" on page 246	Defines a server as a member of a server group.
"DELETE GRPMEMBER (Delete a server from a server group)" on page 426	Deletes a server from a server group.
"QUERY SERVERGROUP (Query a server group)" on page 895	Displays information about server groups.
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"SET MONITORINGADMIN (Set the name of the monitoring administrator)" on page 1119	Set the name of the monitoring administrator.

SET PASSEXP (Set password expiration date)

Use this command to set the expiration period for administrator and client node passwords. You can set a common password expiration period for all administrators and client node passwords or selectively set password expiration periods.

Restriction: The **SET PASSEXP** command does not apply to passwords that authenticate with an LDAP directory server.

You can override this setting for one or more nodes by using the **REGISTER NODE** or **UPDATE NODE** command with the **PASSEXP** parameter.

Important: If you do not specify the **NODE** or **ADMIN** parameters, *all* client node and administrator passwords will use the new password expiration period. If you selectively set a password expiration period for a client node or administrator that does not already have a set password expiration period, it is not modified if you later set a password expiration for all users. The **NODE** or **ADMIN** parameters must be specified to change the password expiration period for client nodes or administrators that have selectively set password expiration periods.


You can use the **RESET PASSEXP** command to reset the password expiration period to the common expiration period.

Use the **QUERY STATUS** command to display the common password expiration period, which at installation is set to 90 days.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

days (Required)

Specifies the number of days that a password remains valid.

You can specify from 1 to 9999 if you do not specify the **NODE** or the **ADMIN** parameter. If you specify the **NODE** or the **ADMIN** parameter, you can specify from 0 to 9999. A value of 0 means that the password never expires. If a password expires, the server prompts for a new password when the administrator or client node contacts the server.

Node

Specifies the name of the node whose password expiration period you would like to set. To specify a list of nodes, separate the names with commas and no intervening spaces. This parameter is optional.

SET PASSEXP

Admin

Specifies the name of the administrator whose password expiration period you would like to set. To specify a list of administrators, separate the names with commas and no intervening spaces. This parameter is optional.

Example: Set the administrator and client node password expiration

Set the administrator and client node password expiration period to 45 days.

```
set passexp 45
```

Example: Set an administrator's password expiration

Set the administrator LARRY's password expiration period to 120 days.

```
set passexp 120 admin=larry
```

Related commands

Table 363. Commands related to SET PASSEXP

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REGISTER NODE	Defines a client node to the server and sets options for that user.
RESET PASSEXP	Resets the password expiration for nodes or administrators.
UPDATE ADMIN	Changes the password or contact information associated with any administrator.
UPDATE NODE	Changes the attributes associated with a client node.

SET QUERYSCHEDPERIOD (Set query period for polling client nodes)

Use this command to regulate how often client nodes contact the server to obtain scheduled work when it is running in the client-polling scheduling mode.

Each client can set its own retry period at the time its scheduler is started. You can use this command to override the value specified by all clients that can connect with the server.

If client nodes poll more frequently for schedules, the nodes receive changes to schedules more quickly. However, increased polling by the client nodes also increases network traffic.

You can issue the **QUERY STATUS** command to display the value for the period between schedule queries. At installation, Tivoli Storage Manager is configured so that each client node determines its own value for this setting.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>>—Set QUERYSChedperiod—┐
                           └hours┘
  
```

Parameters

hours

Specifies the maximum number of hours the scheduler on a client node waits between attempts to contact the server to obtain a schedule. This parameter is optional. You can specify an integer from 1 to 9999. If you do not specify a value for this parameter, each client determines its own value for this parameter.

Example: Set the polling period for all client nodes

Have all clients using the polling scheduling mode contact the server every 24 hours.

```
set queryschedperiod 24
```

Related commands

Table 364. Commands related to SET QUERYSCHEDPERIOD

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET SCHEDMODES	Specifies the central scheduling mode for the server.

SET RANDOMIZE (Set randomization of scheduled start times)

Use this command to set randomized start times within the startup window of each schedule for clients by using the client-polling scheduling mode. A startup window is the start time and duration during which a schedule must be initiated. A client-polling scheduling mode is a client/server communication technique where the client queries the server for work.

Each schedule has a window during which it can be run. To balance network and server load, the start times for clients can be scattered across that window. Use this command to specify the fraction of the window over which start times for clients are distributed.

The randomization occurs at the beginning of the window to allow time for retries, if necessary. When the scheduling mode is not set to polling, randomization does not occur if the client's first contact with the server is after the start time for the event.

You can issue the **QUERY STATUS** command to display the value for the schedule randomization percentage. At installation, the value is 25 percent.

Set the randomization percentage to a value greater than 0 to prevent communication errors. Communication errors can result from a large group of clients contacting the server simultaneously. If you do experience communication errors, you can increase the randomization percentage so that client contact is spread out. This decreases the chance for communication overload and failure.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set **RAN**Domize—*percent*—————►►

Parameters

percent (Required)

Specifies the percentage of the startup window over which the start times for individual clients are distributed. You can specify an integer from 0 to 50.

A value of 0 indicates that no randomization occurs and that all clients run schedules at the beginning of the startup windows.

A value of 50 indicates that clients are assigned start times that are randomly scattered across the first half of each startup window.

At installation, this value is 25, indicating that the first 25 percent of the window is used for randomization.

If you have specified **DURUNITS=INDEFINITE** in the **DEFINE SCHEDULE** command, the percentage is applied to a 24 hour period. For example, a value of 25 percent would result in a 6 hour window.

Example: Set randomization of scheduled start times

Set randomization to 50 percent.

```
set randomize 50
```

Related commands*Table 365. Commands related to SET RANDOMIZE*

Command	Description
DEFINE SCHEDULE	Defines a schedule for a client operation or an administrative command.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET SCHEDMODES	Specifies the central scheduling mode for the server.

SET REGISTRATION (Set open or closed registration)

Use this command to require a system or policy administrator to register client nodes (closed registration) or to permit each user to register his own workstation as a client node with the server (open registration).

With *closed* registration, an administrator defines the following:

- The node name and password for each workstation
- The policy domain to which the client node belongs
- Whether the user can choose to compress files before sending them to server storage
- Whether the user can delete backup or archive files from server storage

With *open* registration, when a user accesses the server from an unregistered client node, the server prompts the user for a node name, password, and contact information, and registers the workstation. The server sets the following defaults:

- Each client node is assigned to the policy domain named STANDARD
- Each user defines whether data compression is used before files are sent to server storage
- Each user is allowed to delete archived files from server storage. The user cannot delete backup files
- The authentication method is based on the **DEFAULTAUTHENTICATION** server setting.
- SSL (Secure Sockets Layer) is required for registering client nodes that authenticate with an LDAP directory server.
- The administrator can reassign domains or change node attributes using the **UPDATE NODE** command

Existing registered client nodes are not affected by changes in the registration process.

Use the **QUERY STATUS** command to display the status of registration. At installation, this value is set to CLOSED.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node belongs.

Syntax

►► Set REGISTRATION Closed Open ►►

Parameters

Closed

Specifies that client nodes must be registered by a system or policy administrator.

Open

Specifies that users can register their workstations as client nodes with the server.

Example: Set the policy so only administrators can register client nodes

Issue the following command to limit client node registration to a policy or system administrator.

```
set registration closed
```

Example: Set the policy so users can register client nodes

Issue the following command to allow users register as client nodes to the server.

```
set registration open
```

Related commands

Table 366. Command related to SET REGISTRATION

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET REPLRETENTION (Set the retention period for replication records)

Use this command to specify the retention period for client-node replication records in the source replication-server database. You can use client node replication records to monitor running and completed processes.

A replication record is created when **REPLICATE NODE** command processing is started. By default, Tivoli Storage Manager retains client-node replication records for 30 calendar days. A *calendar day* consists of 24 hours, from midnight to midnight.


The server automatically removes replication records from the database as part of inventory expiration processing after the retention period passes. As a result, the amount of time that records are retained can exceed the specified retention period. If a replication process runs longer than the retention period, the record of the process is not deleted until the process ends, the retention period passes, and expiration runs.

Issue this command on the server that acts as a source for replicated data.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

number_of_days (Required)

The number of days that the source replication server retains replication records. You can specify an integer 0 - 9999. If you specify 0, replication records are deleted during the next inventory-expiration process. The default value is 30.

Example: Set a retention period for client-node replication records

You want to retain client-node replication records for ten days.

```
set replretention 10
```

Related commands

Table 367. Commands related to SET REPLRETENTION

Command	Description
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLNODE	Displays information about the replication status of a client node.

Table 367. Commands related to SET REPLRETENTION (continued)

Command	Description
QUERY REPLRULE	Displays information about node replication rules.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.

SET REPLSERVER (Set the target replication server)

Use this command to set the name of a target replication server. You can also use this command to change or remove a target replication server.

Issue this command on the server that acts as a source for replicated data.

To display the name of a target replication server, issue the **QUERY STATUS** command on a source replication server.

Important:

- The server name that you specify with this command must match the name of an existing server definition, and it must be the name of the server to be used as the target replication server. If the server name specified by this command does not match the server name of an existing server definition, the command fails.
- Use care when changing or removing a target replication server. If you change a target replication server, replicated client-node data is sent to a different target replication server. If you remove a target replication server, client node data is not replicated.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set REPLSERVER—*target_server_name*—◄◄

Parameters

target_server_name (Required)

Specifies the name of the target replication server. The name that you specify must match the name of an existing server. The maximum length of a name is 64 characters.

To remove a target replication server, issue the command without specifying a value.

Example: Set a target replication server

The name of the server that you want to set as the target replication server is SERVER1.

```
set replserver server1
```

Example: Remove a target replication server

The name of the server that you want to remove is SERVER1.

```
set replserver
```

Issuing this command removes SERVER1 as the target replication server for the source replication server where you issued the command. Other source replication servers might have SERVER1 as a target replication server.

Related commands*Table 368. Commands related to SET REPLSERVER*

Command	Description
DEFINE SERVER	Defines a server for server-to-server communications.
QUERY SERVER	Displays information about servers.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
UPDATE SERVER	Updates information about a server.

SET RETRYPERIOD (Set time between retry attempts)

Use this command to set the number of minutes the scheduler on a client node waits between retry attempts after a failed attempt to contact the server or after a scheduled command fails to process.

Each client can set its own retry period at the time its scheduler program is started. You can use this command to override the values specified by all clients that can connect with the server.

This command is used in conjunction with the **SET MAXCMDRETRIES** command to regulate the period of time and the number of retry attempts to run a failed command.

You can issue the **QUERY STATUS** command to display the value for the period between retries. At installation, Tivoli Storage Manager allows each client to determine its own retry period.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—Set RETRYPeriod—┐
                    └minutes┘
```

Parameters

minutes

Specifies the number of minutes the scheduler on a client node waits between retry attempts after a failed attempt to contact the server or after a scheduled command fails to process. When setting the retry period, set a time period that permits more than one retry attempt within a typical startup window. You can specify an integer from 1 to 9999.

Example: Set a fifteen minute time period between retry attempts

Have the client scheduler retry failed attempts to contact the server or to process scheduled commands every fifteen minutes.

```
set retryperiod 15
```

Related commands

Table 369. Commands related to SET RETRYPERIOD

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET MAXCMDRETRIES	Specifies the maximum number of retries after a failed attempt to execute a scheduled command.

SET SCHEDMODES (Select a central scheduling mode)

Use this command to determine how the clients communicate with the server to begin scheduled work. You must configure each client to select the scheduling mode in which it operates.

Use this command with the **SET RETRYPERIOD** command to regulate the time and the number of retry attempts to process a failed command.

You can issue the **QUERY STATUS** command to display the value for the scheduling mode supported. At installation, this value is ANY.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

ANY

Specifies that clients can run in either the client-polling or the server-prompted scheduling mode.

P011ing

Specifies that only the client-polling mode can be used. Client nodes poll the server at prescribed time intervals to obtain scheduled work.

PRompted

Specifies that only the server-prompted mode can be used. This mode is only available for clients that communicate with TCP/IP. Client nodes wait to be contacted by the server when scheduled work needs to be performed and a session is available.

Example: Restrict scheduled operations to clients using client-polling

Clients can run under both server-prompted and client-polling central scheduling. You want to temporarily restrict the scheduled operations to clients that use the client-polling mode. If you set the schedule mode to POLLING, the server discontinues prompting clients to run scheduled commands. This means that any client scheduler using the server-prompted mode waits until you set the schedule mode to ANY or PROMPTED.

```
set schedmodes polling
```

Related commands

Table 370. Command related to SET SCHEDMODES

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET SCHEDMODES

Table 370. Command related to SET SCHEDMODES (continued)

Command	Description
SET RETRYPERIOD	Specifies the time between retry attempts by the client scheduler.

SET SERVERHLADDRESS (Set the high-level address of a server)

Use this command to set the high-level address (IP) of a server. IBM Tivoli Storage Manager uses the address when you issue a **DEFINE SERVER** command with **CROSSDEFINE=YES**. You must use the **SET SERVERHLADDRESS** command for all automatic client deployments.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set SERVERHladdress—*ip_address*—————►►

Parameters

ip_address (Required)

Specifies a server high-level address as a numeric dotted decimal name or a host name. If a host name is specified, a server that can resolve the name to the dotted decimal form must be available.

Example: Set the high-level address of a server

Set the high-level address of HQ_SERVER to 9.230.99.66.

```
set serverhladdress 9.230.99.66
```

Related commands

Table 371. Command related to SET SERVERHLADDRESS

Command	Description
SET CROSSDEFINE	Specifies whether to cross define servers.
SET SERVERLLADDRESS	Specifies the low-level address of a server.
SET SERVERPASSWORD	Specifies the server password.

SET SERVERLLADDRESS (Set the low-level address of a server)

Use this command to set the low-level address of a server. Tivoli Storage Manager uses the address when you issue a **DEFINE SERVER** command with **CROSSDEFINE=YES**.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set SERVERLLaddress—*tcp_port*—◄◄

Parameters

tcp_port (Required)

Specifies the low-level address of the server. Generally, this address is identical to the TCPPOINT option in the server option file of the server.

Example: Set the low-level address of a server

Set the low-level address of HQ_SERVER to 1500.

```
set serverlladdress 1500
```

Related commands

Table 372. Command related to SET SERVERLLADDRESS

Command	Description
SET CROSSDEFINE	Specifies whether to cross define servers.
SET SERVERHLADDRESS	Specifies the high-level address of a server.
SET SERVERPASSWORD	Specifies the server password.

SET SERVERNAME (Specify the server name)

Use this command to change the server name. When you install the Tivoli Storage Manager server, the name is set at installation to SERVER1.

Use the **QUERY STATUS** command to display the server name.

If you migrate from ADSM to Tivoli Storage Manager, the name is set to ADSM or the name last specified to ADSM with a **SET SERVERNAME** command.

Important:

- If this is a source server for a virtual volume operation, changing its name can impact its ability to access and manage the data it has stored on the corresponding target server.
- To prevent problems related to volume ownership, do not change the name of a server if it is a library client.

When changing the name of a server, be aware of the following additional restrictions:

- Windows clients use the server name to identify which passwords belong to which servers. Changing the server name after the clients are connected forces the clients to reenter the passwords.
- You must set unique names on servers that communicate with each other. For details, see the *Administrator's Guide*. On a network where clients connect to multiple servers, it is recommended that all of the servers have unique names.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set SERVERname—*server_name*—◄◄

Parameters

server_name (Required)

Specifies the new server name. The name must be unique across a server network for enterprise event logging, enterprise configuration, command routing, or virtual volumes. The maximum length of the name is 64 characters.

Example: Name the server

Name the server WELLS_DESIGN_DEPT.

```
set servername wells_design_dept
```

Related commands

Table 373. Command related to SET SERVERNAME

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET SERVERPASSWORD (Set password for server)

Use this command to set the password for communication between servers to support enterprise administration and enterprise event logging and monitoring.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set SERVERPAssword—*password*—◄◄

Parameters

password (Required)

Specifies a password for the server. Other servers must have the same password in their definitions of this server.

Example: Set a server's password

Set the password for HQ_SERVER to agave.

```
set serverpassword agave
```

Related commands

Table 374. Command related to SET SERVERPASSWORD

Command	Description
SET CROSSDEFINE	Specifies whether to cross define servers.
SET SERVERHLADDRESS	Specifies the high-level address of a server.
SET SERVERLLADDRESS	Specifies the low-level address of a server.

SET SPREPLRULEDEFAULT (Set the server replication rule for space-managed data)

Use this command to set the server replication rule for space-managed data.

Restriction: The replication rule that you set with this command is applied only if file space rules and client node rules for space-managed data are set to DEFAULT.

Issue this command on the server that acts as a source for replicated data.

You can specify a normal-priority replication rule or a high-priority replication rule. In a replication process that includes both normal and high-priority data, high-priority data is replicated first. Before you specify a rule, consider the order in which you want the data to be replicated.

For example, suppose that your client nodes contain space-managed data and backup data. Replication of the space-managed data is a higher priority than the backup data. To prioritize the space-managed data, issue the **SET SPREPLRULEDEFAULT** command and specify the ALL_DATA_HIGH_PRIORITY replication rule. To prioritize the backup data, issue the **SET BKREPLRULEDEFAULT** command and specify the ALL_DATA replication rule for backup data. The ALL_DATA rule for backup data replicates backup data with a normal priority.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

>> Set SPREPLRuledefault { ALL_DATA | ALL_DATA_HIGH_PRIORITY | NONE }

```

Parameters

ALL_DATA

Replicates space-managed data with a normal priority.

ALL_DATA_HIGH_PRIORITY

Replicates space-managed data with a high priority.

NONE

Space-managed data is not replicated.

Example: Set the server replication rule for space-managed data

Set up the default rule for space-managed data to replicate with a high priority.

```
set spreplruledefault all_data_high_priority
```

Related commands

Table 375. Commands related to SET BKREPLRULEDEFAULT

Command	Description
QUERY FILESPACE	Displays information about data in file spaces that belong to a client.

SET SPREPLRULEDEFAULT

Table 375. Commands related to SET BKREPLRULEDEFAULT (continued)

Command	Description
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY REPLICATION	Displays information about node replication processes.
QUERY REPLRULE	Displays information about node replication rules.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
REPLICATE NODE	Replicates data in file spaces that belong to a client node.
SET ARREPLRULEDEFAULT	Specifies the server node-replication rule for archive data.
SET BKREPLRULEDEFAULT	Specifies the server node-replication rule for backup data.
UPDATE FILESPACE	Changes file-space node-replication rules.
UPDATE REPLRULE	Enables or disables replication rules.
VALIDATE REPLICATION	Verifies replication for file spaces and data types.


SET SSLKEYRINGPW (Set the SSL key ring password)

Use this command to provide the key database file password to the server. You can also use it to update the key database file password.

Privilege class

You must have system privileges to issue this command.

Syntax



Parameters

password

Specifies the password to use to access the key database file (`cert.kdb`). This parameter is required and is limited to 64 characters.

UPDate

Specifies whether to allow the key database file password to be updated. This parameter is optional. The default is **NO**. Possible values are:

No When specified, notifies Tivoli Storage Manager that the key database file password was changed outside of the server. The server validates the new password before recording it for use on a subsequent startup.

Yes

Specifies to use the password that you defined in the **password** parameter, along with the stored password to change the key database file password. The password is changed immediately in the key database file and Tivoli Storage Manager uses it on a subsequent startup.

Example: Query all background processes

Update the existing stored password with the new password. Restart the server to use the new password.

```
set sslkeyring newpassword update=yes
```

Related commands

Table 376. Commands related to SET SSLKEYRINGPW

Command	Description
DELETE KEYRING	Deletes password information in the certificate key database.
QUERY SSLKEYRINGPW	Displays the Secure Sockets Layer (SSL) key database file password.


SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)

Use this command to adjust the backup activity interval that is used when the status monitor assesses whether clients are at risk.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

TYPE (Required)

Specifies the type of client that should be evaluated. Specify one of the following values:

ALI

Specify this setting for all client types.

APplications

Specify this setting for only application client types.

VM

Specify this setting for virtual system clients types.

SYstems

Specify this setting for systems client types.

Interval (Required)

Specifies the amount of time, in hours, between client activity before the status monitor considers the client at-risk. You can specify an integer from 6 to 720. The interval value for all client types is set to 24 at server installation.

Set systems to use a two-week at-risk interval

Set the at-risk interval check for systems client types to 2 weeks.

```
set statusriskinterval type=systems interval=336
```

Related commands

Table 377. Commands related to

Command	Description
"DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)" on page 337	Defines a status monitoring threshold.
"DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)" on page 451	Deletes a status monitoring threshold.
"QUERY MONITORSTATUS (Query the monitoring status)" on page 799	Displays information about monitoring alerts and server status settings.
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.

Table 377. Commands related to (continued)

Command	Description
"QUERY STATUSTHRESHOLD (Query status monitoring thresholds)" on page 914	Displays information about a status monitoring thresholds.
"SET STATUSMONITOR (Specifies whether to enable status monitoring)" on page 1144	Specifies whether to enable status monitoring.
"SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)" on page 1146	Specifies the refresh interval for status monitoring.
"SET STATUSSKIPFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)" on page 1147	Specifies whether to use client at-risk skipped files as failure evaluation.
"UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)" on page 1375	Changes the attributes of an existing status monitoring threshold.


SET STATUSMONITOR (Specifies whether to enable status monitoring)

Use this command to enable and disable status monitoring.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

ON Specifies that the status monitoring is turned on.

OFF

Specifies that the status monitoring is turned off. Off is the default value.

Enable status monitoring

Set status monitoring to on to enable status monitoring.

```
set statusmonitor on
```

Related commands

Table 378. Commands related to SET STATUSMONITOR

Command	Description
"DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)" on page 337	Defines a status monitoring threshold.
"DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)" on page 451	Deletes a status monitoring threshold.
"QUERY MONITORSTATUS (Query the monitoring status)" on page 799	Displays information about monitoring alerts and server status settings.
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"QUERY STATUSTHRESHOLD (Query status monitoring thresholds)" on page 914	Displays information about a status monitoring thresholds.
"SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)" on page 1142	Specifies whether to enable client at-risk activity interval evaluation.
"SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)" on page 1146	Specifies the refresh interval for status monitoring.
"SET STATUSSKIPFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)" on page 1147	Specifies whether to use client at-risk skipped files as failure evaluation.
"UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)" on page 1375	Changes the attributes of an existing status monitoring threshold.

|
|

SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)

Use this command to specify the number of minutes between status monitoring server queries.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—Set STATUSREFreshinterval—*minutes*—◄◄

Parameters

minutes (Required)

Specifies the approximate number of minutes between status monitoring server queries. Specify an integer from 1 to 2440. The default value is 5.

Set the refresh interval for status monitoring

Specify that the server status is queried every 2 minutes, by issuing the following command:

```
set statusrefreshinterval 2
```

Related commands

Table 379. Commands related to SET STATUSREFRESHINTERVAL

Command	Description
"DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)" on page 337	Defines a status monitoring threshold.
"DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)" on page 451	Deletes a status monitoring threshold.
"QUERY MONITORSTATUS (Query the monitoring status)" on page 799	Displays information about monitoring alerts and server status settings.
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"QUERY STATUSTHRESHOLD (Query status monitoring thresholds)" on page 914	Displays information about a status monitoring thresholds.
"SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)" on page 1142	Specifies whether to enable client at-risk activity interval evaluation.
"SET STATUSMONITOR (Specifies whether to enable status monitoring)" on page 1144	Specifies whether to enable status monitoring.
"SET STATUSSKIPFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)" on page 1147	Specifies whether to use client at-risk skipped files as failure evaluation.
"UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)" on page 1375	Changes the attributes of an existing status monitoring threshold.


SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)

Use this command to enable the status monitor to consider clients as at risk when evaluating the status for each client.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

State (Required)

Specifies whether to enable the check for skipped files during the last backup. This check signifies that the client is at-risk if any files were skipped. Client data that is skipped or not backed up properly is considered at risk.

ON Specifies that the server evaluates whether a client is at risk.

OFF

Specifies that the server does not evaluate whether a client is at risk.

TYPE (Required)

Specifies the type of client that should be evaluated. Specify one of the following values:

ALL

Specify this setting for all client types.

Applications

Specify this setting for only application client types.

VM Specify this setting for virtual system clients types.

Systems

Specify this setting for systems client types.

Disable at-risk evaluation for virtual system client types

Disable the at-risk evaluation for virtual systems client types by issuing the following command:

```
set statusskipasfailure off type=vm
```

Related commands

Table 380. Commands related to SET STATUSSKIPASFAILURE

Command	Description
"DEFINE STATUSTHRESHOLD (Define a status monitoring threshold)" on page 337	Defines a status monitoring threshold.
"DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)" on page 451	Deletes a status monitoring threshold.

SET STATUSSKIPASFAILURE

Table 380. Commands related to SET STATUSSKIPASFAILURE (continued)

Command	Description
"QUERY MONITORSTATUS (Query the monitoring status)" on page 799	Displays information about monitoring alerts and server status settings.
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"QUERY STATUSTHRESHOLD (Query status monitoring thresholds)" on page 914	Displays information about a status monitoring thresholds.
"SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)" on page 1142	Specifies whether to enable client at-risk activity interval evaluation.
"SET STATUSMONITOR (Specifies whether to enable status monitoring)" on page 1144	Specifies whether to enable status monitoring.
"SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)" on page 1146	Specifies the refresh interval for status monitoring.
"UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)" on page 1375	Changes the attributes of an existing status monitoring threshold.

SET SUBFILE (Set subfile backup for client nodes)

Use this command to set up the server to allow clients to back up subfiles. On the client's workstation, the SUBFILECACHEPATH and SUBFILECACHESIZE options must be specified in the client's options file (dsm.opt). If you are using a Windows client, you must also specify the SUBFILEBACKUP option.

With subfile backups, when a client's file has been previously backed up, any subsequent backups are typically made to the portion (a subfile) of the client's file that has changed, rather than the entire file.

Use the **QUERY STATUS** command to determine whether subfiles can be backed up to the server running this command.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—Set SUBFILE—┌ Client ───────────────────────────────────────────────────────────►►
                  │ No   ───────────────────────────────────────────────────────────►►
```

Parameters

Client

Specifies that the client node can determine whether to use subfile backup.

No Specifies that the subfile backups are not to be used. At installation, this value is set to No.

Example: Set subfile backup for client nodes

Allow the client node to backup subfiles on the server.

```
set subfile client
```

Related commands

Table 381. Command related to SET SUBFILE

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SET SUMMARYRETENTION (Set number of days to keep data in activity summary table)

Use this command to specify the number of days to keep information in the SQL activity summary table.

The SQL activity summary table contains statistics about each client session and server processes. For a description of the information in the SQL activity summary table, issue the following command:

```
select colname, remarks from columns where tabname='SUMMARY'
```

Issue the **QUERY STATUS** command to display the number of days the information is kept. At installation, Tivoli Storage Manager allows each server to determine its own number of days for keeping information in the SQL activity summary table.

Privilege class

To issue this command, you must have system privilege.

Syntax

```
►►—Set Summaryretention days—►►
```

Parameters

days

Specifies the number of days to keep information in the activity summary table. You can specify a number from 0 to 9999. A value of 0 means that information in the activity summary table is not kept. A value of 1 specifies to keep the activity summary table for the current day.

Example: Specify the number of days to keep information in the SQL activity summary table

Set the server to retain the activity summary table information for 15 days.

```
set summaryretention 15
```

Related commands

Table 382. Commands related to SET SUMMARYRETENTION

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
SET ACTLOGRETENTION	Specifies the number of days to retain log records in the activity log.
QUERY ACTLOG	Displays messages from the server activity log.
SELECT	Allows customized queries of the Tivoli Storage Manager database.

SET TAPEALERTMSG (Set tape alert messages on or off)

Use this command to allow the Tivoli Storage Manager server to log notification of diagnostic information from library and drive devices. At installation, this value is set to OFF. When enabled, Tivoli Storage Manager can retrieve diagnostic information from a tape or library device and display it using ANR messages. When disabled, Tivoli Storage Manager will not query a device for this information.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax

```
▶▶—Set TAPEAlertmsg ON | OFF▶▶
```

Parameters

ON Specifies that diagnostic information will be reported to the Tivoli Storage Manager server.

OFF Specifies that diagnostic information will not be reported to the Tivoli Storage Manager server.

Example: Set tape alert messages on

Allow the Tivoli Storage Manager server to receive diagnostic information messages.

```
set tapealertmsg on
```

Related commands

Table 383. Command related to SET TAPEALERTMSG

Command	Description
QUERY TAPEALERTMSG	Displays whether the server logs hardware diagnostic information.

SET TOCLOADRETENTION (Set load retention period for table of contents)

Use this command to specify the approximate number of minutes that unreferenced table of contents data will remain loaded in the server database.

During NDMP-controlled backup operations of NAS file systems, the server can optionally collect information about files and directories in the image and store this information in a table of contents within a storage pool. The Web client can be used to examine files and directories in one or more file-system images by displaying entries from the table of contents data. The server loads the necessary table of contents data into a temporary database table.

Once the data have been loaded, the user can then select those files and directories to be restored. Because this database table is temporary, the data will only remain loaded for a specified time since the last reference to that data. At installation, the retention time is set to 120 minutes. Use the **QUERY STATUS** command to see the table of contents load retention time.

Privilege class

To issue this command you must have system privilege.

Syntax

```
►►—Set TOCLOADRetention—minutes—◄◄
```

Parameters

minutes (Required)

Specifies the approximate number of minutes that an unreferenced table of contents data is retained in the database. You can specify an integer from 30 to 1000.

Example: Define the load retention period for the table of contents

Use the command, **SET TOCLOADRETENTION**, to specify that unreferenced table of contents data is to be retained in the database for 45 minutes.

```
set toclloadretention 45
```

Related commands

Table 384. Commands related to SET TOCLOADRETENTION

Command	Description
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.

SETOPT (Set a server option for dynamic update)

Use the **SETOPT** command to update a server option dynamically without stopping and restarting the server. A SETOPT command contained in a macro or a script cannot be rolled back.

Privilege class

To issue this command, you must have system privilege.

Syntax

▶—SETOPT—*option_name*—*option_value*—▶

Parameters

option_name (Required)

Specifies a text string of information identifying the server option to be updated. The maximum length of the text string is 255 characters. The following options are available:

ADMINCOMMTIMEOUT
 ADMINIDLETIMEOUT
 ALLOWREORGINDEX
 ALLOWREORGTABLE
 CHECKTAPEPOS
 CLIENTDEDUPTXNLIMIT
 COMMTIMEOUT
 DBDIAGPATHFSTHRESHOLD
 DEDUPTIER2FILESIZE
 DEDUPTIER3FILESIZE
 DEDUPREQUIRESBACKUP
 DNSLOOKUP
 EXPINTERVAL
 EXPQUIET
 FSUSEDTHRESHOLD
 IDLETIMEOUT
 LDAPCACHEDURATION
 MAXSESSIONS
 MOVEBATCHSIZE
 MOVESIZETHRESH
 NDMPENABLEKEEPALIVE
 NDMPKEEPIDLEMINUTES
 NDMPREFDATAINTERFACE
 NUMOPENVOLSALLOWED
 RECLAIMDELAY
 RECLAIMPERIOD
 REORGBEGINTIME
 REORGDURATION
 RESOURCETIMEOUT
 RESTOREINTERVAL
 RETENTIONEXTENSION
 SANDISCOVERY
 SANREFRESHTIME
 SERVERDEDUPTXNLIMIT
 SHREDDING

SETOPT

THROUGHPUTDATATHRESHOLD
THROUGHPUTTIMETHRESHOLD
TXNGROUPMAX

option_value **(Required)**

Specifies the value for the server option.

Example: Set the maximum number of client session

Update the server option for the maximum number of client sessions to a value of 40.

```
setopt maxsessions 40
```

Related commands

Table 385. Commands related to SETOPT

Command	Description
QUERY OPTION	Displays information about server options.
QUERY SYSTEM	Displays details about the Tivoli Storage Manager server system.

SHRED DATA (Shred data)

Use this command to manually start the process of shredding deleted sensitive data. Manual shredding is possible only if automatic shredding is disabled.

You can control automatic shred processing with the `SHREDDING` server option.


This command creates a background process that can be cancelled with the **CANCEL PROCESS** command. To display information on background processes, use the **QUERY PROCESS** command.

If data from a storage pool that enforces shredding is deleted while a manual shredding process is running, it will be added to the running process.

Privilege class

To issue this command you must have system privilege.

Syntax



Parameters

DURATION

Specifies the maximum number of minutes the shredding process runs before being automatically cancelled. When the specified number of minutes elapses, the server cancels the shredding process. As soon as the process recognizes the cancellation, it ends. Because of this, the process may run longer than the value you specified for this parameter. You can specify a number from 1 to 9999. This parameter is optional. If not specified, the server will stop only after all deleted sensitive data has been shredded.

Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is No. Possible values are:

No Specifies that the server processes this command in the background. You can continue with other tasks while the command is being processed. Messages created from the background process are displayed either in the activity log or the server console, or both, depending on where messages are logged. To cancel a background process, use the **CANCEL PROCESS** command. If you cancel this process, some files might already have been shredded before the cancellation. This is the default.

Yes

Specifies that the server processes this command in the foreground. You must wait for the operation to complete before continuing with other tasks.

SHRED DATA

The server displays the output messages to the administrative client when the operation completes. Messages are also displayed either in the activity log or the server console, or both, depending on where messages are logged.

Note: You cannot specify WAIT=YES from the server console.

IOERROR

Specifies whether an I/O error encountered while shredding the data is to be considered a successful shred. This parameter is optional. The default is SHREDFailure. Possible values are:

SHREDFailure

Specifies that if the server encounters an I/O error while shredding, the data will not be considered successfully shredded and the owning file will be marked as damaged. The server will attempt to shred the data again the next time the shredding process runs, giving you a chance to correct the error and ensure the data can be properly shredded.

SHREDSuccess

Specifies that if the server encounters an I/O error while shredding and the owning file had been previously marked as damaged, the data will be considered successfully shredded. You should use this option only after the server has reported I/O errors while shredding and you are unable to correct the error.

Example: Shred data

Manually start the shredding of all deleted sensitive data. Continue the process for up to six hours before automatically cancelling it.

```
shred data duration=360
```

Related commands

Table 386. Commands related to SHRED DATA

Command	Description
CANCEL PROCESS	Cancels a background server process.
QUERY PROCESS	Displays information about background processes.
QUERY SHREDSTATUS	Displays information about data waiting to be shredded.

SUSPEND EXPORT (Suspend a currently running export operation)

Use this command to suspend a currently running server-to-server export operation which has a FILEDATA value that is not NONE. The export operation that you want to suspend must be past the initialization phase to be eligible for suspension. The state of the export operation is saved. The operation can be restarted by issuing the **RESTART EXPORT** command.

Privilege class

You must have system privilege to issue this command.

Syntax

```

▶▶—SUSPend EXPOrt—*
export_identifier

```

Parameters

EXPORTIDentifier

This optional parameter specifies the name of the export operation. You can find a name by issuing the **QUERY EXPORT** command to list all the currently running server-to-server export operations that can be suspended. You can also use the wildcard character to specify the name.

Example: Suspend a specific export operation

Suspend the running export operation EXPORTALLACCTNODES. No output is generated when you issue the **SUSPEND EXPORT** command. You must issue the **QUERY EXPORT** command to verify that the EXPORTALLACCTNODES operation is suspended.

```
suspend export exportallacctnodes
```

Example: Suspend all running export operations

Suspend all the export operations with a state of RUNNING.

```
suspend export *
```

Related commands

Table 387. Commands related to **SUSPEND EXPORT**

Command	Description
CANCEL EXPORT	Deletes a suspended export operation.
EXPORT NODE	Copies client node information to external media or directly to another server.
EXPORT SERVER	Copies all or part of the server to external media or directly to another server.
QUERY EXPORT	Displays the export operations that are currently running or suspended.
RESTART EXPORT	Restarts a suspended export operation.

UNLOCK commands

Use the **UNLOCK** commands to reestablish access after an object was locked.

- “UNLOCK ADMIN (Unlock an administrator)” on page 1159
- “UNLOCK NODE (Unlock a client node)” on page 1161
- “UNLOCK PROFILE (Unlock a profile)” on page 1163

UNLOCK ADMIN (Unlock an administrator)

Use the **UNLOCK ADMIN** command to allow a locked administrator to access the server again. You can also unlock multiple administrators that authenticate with the same method.

Privilege class

To issue this command, you must have system privilege.

Syntax

```

▶▶ UNLOCK Admin [admin_name] AUTHentication=[Local|LDap]

```

Parameters

admin_name (Required)

Specifies the name of the administrator to unlock. You can use wildcard characters to specify the administrator name. You do not have to enter an administrator name if you want to unlock all of the administrators according to their method of authentication. Use the wildcard with an authentication method to unlock multiple administrators. The parameter is required (no default wildcard).

AUTHentication

Specifies the method of password authentication that is needed for an administrator to log on.

Local

Specifies that you want to unlock administrator user IDs that authenticate passwords with the IBM Tivoli Storage Manager server.

LDap

Specifies that you want to unlock administrator user IDs that authenticate passwords with an LDAP directory server.

Example: Unlock an administrator user ID

The administrator user ID JOE is locked out of Tivoli Storage Manager. Allow JOE to access the server. Issue the following command:

```
unlock admin joe
```

Example: Unlock all administrator user IDs that authenticate passwords with an LDAP directory server

The administrator user ID that use passwords that authenticate with an LDAP directory server must be unlocked so the IDs can communicate with the Tivoli Storage Manager server.

```
unlock admin * authentication=ldap
```

UNLOCK ADMIN

Related commands

Table 388. Commands related to UNLOCK ADMIN

Command	Description
LOCK ADMIN	Prevents an administrator from accessing Tivoli Storage Manager.
QUERY ADMIN	Displays information about one or more Tivoli Storage Manager administrators.


UNLOCK NODE (Unlock a client node)

Use this command to allow a locked client node to access the server again. You can also unlock multiple nodes that use the same method of authentication.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax



Parameters

node_name (Required)

Specifies the name of the client node to unlock. You can use wildcard characters to specify the node name. You do not have to enter a node name if you want to unlock all of the nodes according to their method of authentication. Use the wildcard with an authentication method to unlock groups of nodes. The parameter is required. There is no default wildcard character available.

AUTHentication

Specifies the node password authentication method. This parameter is optional.

L0cal

Specifies that you want to unlock nodes that authenticate passwords with the IBM Tivoli Storage Manager server.

LDap

Specifies that you want to unlock nodes that authenticate passwords with an LDAP directory server.

Example: Unlock a node

The client node SMITH is locked out of IBM Tivoli Storage Manager. Allow SMITH to access the server.

```
unlock node smith
```

Example: Unlock all nodes that authenticate with the Tivoli Storage Manager server

The nodes that are not authenticating passwords with LDAP directory servers must be unlocked.

```
unlock node * authentication=local
```

Related commands

Table 389. Commands related to UNLOCK NODE

Command	Description
LOCK NODE	Prevents a client from accessing the server.

UNLOCK NODE

Table 389. Commands related to UNLOCK NODE (continued)

Command	Description
QUERY NODE	Displays partial or complete information about one or more clients.

UNLOCK PROFILE (Unlock a profile)

Use this command on a configuration manager to unlock a configuration profile so it can be distributed to subscribing managed servers.

Privilege class

To issue this command, you must have system privilege.

Syntax

►►—UNLOCK PROFILE—*profile_name*—◄◄

Parameters

profile_name (Required)

Specifies the profile to unlock. You can use wildcard characters to indicate multiple names.

Example: Unlock a profile

Unlock a profile named TOM.

```
unlock profile tom
```

Related commands

Table 390. Commands related to UNLOCK PROFILE

Command	Description
COPY PROFILE	Creates a copy of a profile.
DEFINE PROFASSOCIATION	Associates objects with a profile.
DEFINE PROFILE	Defines a profile for distributing information to managed servers.
DELETE PROFASSOCIATION	Deletes the association of an object with a profile.
DELETE PROFILE	Deletes a profile from a configuration manager.
LOCK PROFILE	Prevents distribution of a configuration profile.
QUERY PROFILE	Displays information about configuration profiles.
SET CONFIGMANAGER	Specifies whether a server is a configuration manager.
UPDATE PROFILE	Changes the description of a profile.

UPDATE commands

Use the **UPDATE** command to modify one or more attributes of an existing Tivoli Storage Manager object.

- “UPDATE ADMIN (Update an administrator)” on page 1170
- “UPDATE ALERTTRIGGER (Update a defined alert trigger)” on page 1165
- “UPDATE ALERTSTATUS (Update the status of an alert)” on page 1168
- “UPDATE BACKUPSET (Update a retention value assigned to a backup set)” on page 1174
- “UPDATE CLIENTOPT (Update a client option sequence number)” on page 1179
- “UPDATE CLOPTSET (Update a client option set description)” on page 1180
- “UPDATE COLLOGROUP (Update a collocation group)” on page 1181
- “UPDATE COPYGROUP (Update a copy group)” on page 1182
- “UPDATE DATAMOVER (Update a data mover)” on page 1190
- “UPDATE DEVCLASS (Update the attributes of a device class)” on page 1192
- “UPDATE DOMAIN (Update a policy domain)” on page 1275
- “UPDATE DRIVE (Update a drive)” on page 1277
- “UPDATE FILESPACE (Update file-space node-replication rules)” on page 1281
- “UPDATE LIBRARY (Update a library)” on page 1286
- “UPDATE LIBVOLUME (Change the status of a storage volume)” on page 1302
- “UPDATE MACHINE (Update machine information)” on page 1304
- “UPDATE MGMTCLASS (Update a management class)” on page 1306
- “UPDATE NODE (Update node attributes)” on page 1309
- “UPDATE NODEGROUP (Update a node group)” on page 1325
- “UPDATE PATH (Change a path)” on page 1326
- “UPDATE POLICYSET (Update a policy set description)” on page 1335
- “UPDATE PROFILE (Update a profile description)” on page 1337
- “UPDATE RECOVERYMEDIA (Update recovery media)” on page 1338
- “UPDATE REPLRULE (Update replication rules)” on page 1340
- “UPDATE SCHEDULE (Update a schedule)” on page 1342
- “UPDATE SCRIPT (Update a Tivoli Storage Manager script)” on page 1365
- “UPDATE SERVER (Update a server defined for server-to-server communications)” on page 1368
- “UPDATE SERVERGROUP (Update a server group description)” on page 1372
- “UPDATE SPACETRIGGER (Update the space triggers)” on page 1373
- “UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)” on page 1375
- “UPDATE STGPOOL (Update a storage pool)” on page 1379
- “UPDATE VIRTUALFSMAPPING (Update a virtual file space mapping)” on page 1419
- “UPDATE VOLHISTORY (Update sequential volume history information)” on page 1421
- “UPDATE VOLUME (Change a storage pool volume)” on page 1423


UPDATE ALERTTRIGGER (Update a defined alert trigger)

Use this command to update the attributes of one or more alert triggers.

Privilege class

To issue this command, you must have system privilege.

Syntax



Parameters

message_number (Required)

Specifies the message number that you want to associate with the alert trigger. Specify multiple message numbers, which are separated by commas, and no intervening spaces. Message numbers are a maximum of eight characters in length. Wildcard characters can be used to specify message numbers.

CATegory

Specifies the category type for the alert, which is determined by the message types. The default value is SERVER. Specify one of the following values:

APplication

Alert is classified as application category. For example, you can specify this category for messages that are associated with application (TDP) clients.

CAtalog

Alert is classified as catalog category. For example, you can specify this category for messages that are associated with the database, active log file, or archive log file.

Note: Changing the category of an alert trigger does not change the category of existing alerts on the server. New alerts are categorized with the new category.

UPDATE ALERTRIGGER

Client

Alert is classified as client category. For example, you can specify this category for messages that are associated with general client activities.

Device

Alert is classified as device category. For example, you can specify this category for messages that are associated with device classes, libraries, drives, or paths.

Server

Alert is classified as general server category. For example, you can specify this category for messages that are associated with general server activities or events.

Storage

Alert is classified as storage category. For example, you can specify this category for messages that are associated with storage pools.

Systems

Alert is classified under system clients category. For example, you can specify this category for messages that are associated with system backup and archive or hierarchical storage management (HSM) backup-archive clients.

VMclient

Alert is classified under VMclient category. For example, you can specify this category for messages that are associated with virtual machine clients.

Admin

This optional parameter specifies the name of the administrator who receives email notification of this alert. The alert trigger is defined successfully even if no administrator names are specified.

ADDadmin

Specifies the administrator name that you want to add to the list of administrators that receive email alerts. Specify multiple administrator names, which are separated by commas, and no intervening spaces.

DELadmin

Specifies the administrator name that you want to delete from the list of administrators that receive email alerts. Specify multiple administrator names, which are separated by commas, and no intervening spaces.

Update alert trigger

Add the names of the administrators that want to be notified when ANR1073E, ANR1074E alerts occur, and also delete the name of an administrator that no longer wants to be notified, by issuing the following command:

```
update alertrigger ANR1073E,ANR1074E ADDadmin=djee,cdawson,mhaye deladmin=harryh
```

Related commands

Table 391. Commands related to UPDATE ALERTRIGGER

Command	Description
"DEFINE ALERTRIGGER (Define an alert trigger)" on page 112	Associates specified messages to an alert trigger.
"DELETE ALERTRIGGER (Remove a message from an alert trigger)" on page 400	Removes a message number that can trigger an alert.

Table 391. Commands related to UPDATE ALERTTRIGGER (continued)

Command	Description
"QUERY ALERTSTATUS (Query the status of an alert)" on page 674	Displays information about alerts that have been issued on the server.
"QUERY ALERTTRIGGER (Query the list of defined alert triggers)" on page 672	Displays message numbers that trigger an alert.
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"UPDATE ALERTSTATUS (Update the status of an alert)" on page 1168	Updates the status of a reported alert.

UPDATE ALERTSTATUS


UPDATE ALERTSTATUS (Update the status of an alert)

Use this command to update the status of a reported alert.

Privilege class

Any administrator can issue this command.

Syntax



Parameters

alert_id (Required)

Specifies the alert that you want to update.

STATUS

Specifies the status type that you want to update. Alerts can be changed from active to inactive or closed, or from inactive to closed. Possible values are:

Inactive

Active alerts can be changed to inactive status.

Closed

Active and inactive alerts can be changed to closed status.

ASSIGNED

Specifies the administrator name that is assigned the alert that you want to query.

RESOLVEDBY

Specifies the administrator name that resolved the alert that you want to query.

REMARK

This parameter specifies comment text. The comment text cannot exceed 255 characters. If the description contains any blank spaces, enclose the entire text in quotation marks (""). Remove previously defined text by specifying a null string ("") for this value.

Update the comment text in an alert

Issue the following command to update the comment text for alert ID number 25 and indicate that *DJADMIN* is working on the alert:

```
update alertstatus 25 assigned=DJADMIN
```

Update alert status

Issue the following command to change alert ID number 72 to the closed status, and add a remark about how the alert was resolved:

```
update alertstatus 72 status=closed remark="Increased the file system size for the active log"
```

Related commands

Table 392. Commands related to UPDATE ALERTSTATUS

Command	Description
"DEFINE ALERTTRIGGER (Define an alert trigger)" on page 112	Associates specified messages to an alert trigger.
"DELETE ALERTTRIGGER (Remove a message from an alert trigger)" on page 400	Removes a message number that can trigger an alert.
"QUERY ALERTSTATUS (Query the status of an alert)" on page 674	Displays information about alerts that have been issued on the server.
"QUERY ALERTTRIGGER (Query the list of defined alert triggers)" on page 672	Displays message numbers that trigger an alert.
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
"UPDATE ALERTTRIGGER (Update a defined alert trigger)" on page 1165	Updates the attributes of one or more alert triggers.

UPDATE ADMIN (Update an administrator)

Use this command to change the password or contact information for an administrator. However, you cannot update the SERVER_CONSOLE administrator name.

Passwords for administrators must be changed after a length of time determined by the **SET PASSEXP** command. The **SET PASSEXP** command does not affect passwords that authenticate with an LDAP directory server.

Restriction: You cannot update the authentication method for your own user ID. If necessary, another administrator must make that change. Also, when updating a password with the **UPDATE ADMIN** command, you cannot use a wildcard with the `admin_name` parameter.


Administrators with the same name as a node can be created during a **REGISTER NODE** command. If you do not specify a name, the administrator is given the same name as the node. To keep the node and administrator with the same name synchronized, the authentication method and the **SSLREQUIRED** setting for the node are updated to match the administrator. If the administrator authentication method is changed from **LOCAL** to **LDAP** and a password is not provided, the node is put in "LDAP pending" status. A password is then requested at the next logon. Passwords between same-named nodes and administrators are kept in sync through any authentication change.


You must use the **RENAME ADMIN** command to change the name of a registered administrator.

Privilege class

To issue this command to change another administrator ID password or contact information, you must have system privilege. Any administrator can issue this command to update his or her own password or contact information.

Syntax



**Notes:**

- 1 You must specify at least one optional parameter on this command.
- 2 Passwords are optional for this command, except when changing the authentication method from LDAP to LOCAL.
- 3 The **SYNCLDAPDELETE** parameter applies only if an administrator authenticating to an LDAP directory server reverts to local authentication. If other servers use the same namespace, specify **SYNCLDAPDELETE=no**. If the LDAP entry is to be deleted to synchronize with the IBM Tivoli Storage Manager server database, select **YES**.

Parameters***admin_name* (Required)**

Specifies the name of the administrator to be updated.

password

Specifies the administrator's password. This parameter is optional for most cases. If the administrator authentication method is changed from LDAP to LOCAL, a password is required. Passwords remain current for a period determined by the password expiration period.

PASSExp

Specifies the number of days the password remains valid. You can set the password expiration period from 0 to 9999 days. A value of 0 means that the password never expires. This parameter is optional. If you do not specify this parameter, the password expiration period is unchanged. This parameter does not apply to passwords that are stored on an LDAP directory server.

CONTACT

Specifies a text string that identifies the administrator. This parameter is optional. Enclose the text string in quotation marks if it contains any blanks. To remove previously defined contact information, specify a null string ("").

FORCEPwreset

Specifies whether the administrator is required to change or reset the password. This parameter is optional. Possible values are:

No Specifies that the administrator does not need to change or reset the password while attempting to sign on to the server. The password expiration period is set by the **SET PASSEXP** command.

Yes

Specifies that the administrator's password will expire at the next sign on. The administrator must change or reset the password at that time. If a password is not specified, you will receive a syntax error.

UPDATE ADMIN

EMAILAddress

This parameter is used for additional contact information. The information specified by this parameter is not acted upon by Tivoli Storage Manager.

AUTHentication

This parameter determines the password authentication method that the administrator ID uses; either LDAP or LOCAL.

Local

Specifies that the administrator uses the local Tivoli Storage Manager server database to store passwords for authentication.

Ldap

Specifies that the administrator uses an LDAP directory server for password authentication.

SSLrequired

Specifies whether the administrator user ID must use Secure Sockets Layer (SSL) to communicate with the backup-archive client from the Tivoli Storage Manager server. When you authenticate passwords with an LDAP directory server, you must protect the sessions by using SSL or another network security method.

Yes

SSL is required.

No SSL is not required.

Default

SSL is required for an administrator if the password that is associated with its user ID authenticates with an LDAP directory server. SSL is not required for an administrator ID that authenticates its password with the Tivoli Storage Manager server (LOCAL).

SYNCLdapdelete

This parameter applies only if an administrator who authenticates to an LDAP directory server wants to revert to local authentication. If an entry is being used by another server and sharing a namespace in the LDAP directory server, select NO. If you want to synchronize an entry to a Tivoli Storage Manager server, and the entry is not used by another server or sharing a namespace, select YES. The default is NO.

ALert

Specifies whether alerts are sent to an administrators email address.

Yes

Specifies that alerts are sent to the specified administrators email address.

No Specifies that alerts are not sent to the specified administrators email address. This is the default value.

Tip: Alert monitoring must be enabled, and email settings must be correctly defined to successfully receive alerts by email. To view the current settings, issue the **QUERY MONITORSETTINGS** command.

Example: Update a password and password expiration period

Update the administrator LARRY to have the password SECRETWORD and a password expiration period of 120 days. The administrator in this example is authenticated to the Tivoli Storage Manager server. The password is not case-sensitive.

```
update admin larry secretword passexp=120
```

Example: Change the administrator's authentication method and update the password

Update the administrator LARRY to authenticate to an LDAP directory server and have the password SeCre#Tw0rd.

```
update admin larry SeCre#Tw0rd authentication=ldap
```

Password expiration does not apply to LDAP-authenticated passwords.

Related commands

Table 393. Commands related to UPDATE ADMIN

Command	Description
QUERY ADMIN	Displays information about one or more Tivoli Storage Manager administrators.
QUERY STATUS	Displays the settings of server parameters, such as those selected by the SET commands.
"QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796	Displays information about monitoring alerts and server status settings.
REGISTER ADMIN	Defines a new administrator without granting administrative authority.
REGISTER NODE	Defines a client node to the server and sets options for that user.
RENAME ADMIN	Changes a Tivoli Storage Manager administrator's name.
SET PASSEXP	Specifies the number of days after which a password is expired and must be changed.
UPDATE NODE	Changes the attributes associated with a client node.

UPDATE BACKUPSET


UPDATE BACKUPSET (Update a retention value assigned to a backup set)

Use this command to update the retention value associated with a client's backup set.

Privilege class

To issue this command, you must have system privilege or policy privilege for the domain to which the client node is assigned.

Syntax



Parameters

node_name or *node_group_name* (Required)

Specifies the names of the client nodes or node groups whose data is contained in the specified backup set to be updated. To specify multiple node and node group names, separate the names with commas and no intervening spaces. The node names that you specify can contain wildcard characters, but node group names cannot contain wildcard characters.

backup_set_name (Required)

Specifies the name of the backup set to update. The backup set name you

specify can contain wildcard characters. You can specify more than one backup set name by separating the names with commas and no intervening spaces.

RETention (Required)

Specifies the updated number of days to retain the backup set on the server. You can specify an integer from 0 to 30000. The values are:

days

Specifies the updated number of days to retain the backup set.

NOLimit

Specifies that the backup set is retained on the server indefinitely. If you specify NOLIMIT, the server retains the volumes containing the backup set forever, unless a user or administrator deletes the volumes from server storage.

Attention: Updating the retention period of a backup set may cause it to expire at a different time from other backup sets that might be stored on the same output media. In either case, the media will not be made available for other uses until all of its backup sets have expired.

BEGINDate

Specifies the beginning date in which the backup set to update was created. This parameter is optional. The default is the current date. You can use this parameter with the **BEGINTIME** parameter to specify a range for the date and time. If you specify a begin date without a begin time, the time will be at 12:00 a.m. (midnight) on the date you specify.

You can specify the date by using one of the following values:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1999
TODAY	The current date	TODAY
TODAY+ <i>days</i> or <i>+days</i>	The current date plus days specified.	TODAY +3 or +3.
TODAY- <i>days</i> or <i>-days</i>	The current date minus days specified.	TODAY-3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

BEGINTime

Specifies the beginning time in which the backup set to update was created. This parameter is optional. The default is the current time. You can use this

UPDATE BACKUPSET

parameter with the **BEGINDATE** parameter to specify a range for the date and time. If you specify a begin time without a begin date, the date will be the current date at the time you specify.

You can specify the time by using one of the following values:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08
NOW	The current time	NOW
NOW+ <i>HH:MM</i> or + <i>HH:MM</i>	The current time plus hours and minutes on the specified end date	NOW+02:00 or +02:00.
NOW- <i>HH:MM</i> or - <i>HH:MM</i>	The current time minus hours and minutes on the specified end date	NOW-02:00 or -02:00.

ENDDate

Specifies the ending date in which the backup set to update was created. This parameter is optional. You can use this parameter with the **ENDTIME** parameter to specify a range for the date and time. If you specify an end date without an ending time, the time will be at 11:59:59 p.m. on the specified end date.

You can specify the date by using one of the following values:

Value	Description	Example
<i>MM/DD/YYYY</i>	A specific date	09/15/1999
TODAY	The current date	TODAY
TODAY+ <i>days</i> or + <i>days</i>	The current date plus days specified.	TODAY +3 or +3.
TODAY- <i>days</i> or - <i>days</i>	The current date minus days specified.	TODAY -3 or -3.
EOLM (End Of Last Month)	The last day of the previous month.	EOLM
EOLM- <i>days</i>	The last day of the previous month minus days specified.	EOLM-1 To include files that were active a day before the last day of the previous month.
BOTM (Beginning Of This Month)	The first day of the current month.	BOTM
BOTM+ <i>days</i>	The first day of the current month, plus days specified.	BOTM+9 To include files that were active on the 10th day of the current month.

ENDTime

Specifies the ending time in which the backup set to update was created. This parameter is optional. You can use this parameter with the **ENDDATE** parameter to specify a range for the date and time. If you specify an end time without an end date, the date will be the current date at the time you specify.

You can specify the time by using one of the following values:

Value	Description	Example
<i>HH:MM:SS</i>	A specific time	10:30:08
NOW	The current time	NOW
NOW+ <i>HH:MM</i> or + <i>HH:MM</i>	The current time plus hours and minutes specified	NOW+02:00 or +02:00.
NOW- <i>HH:MM</i> or - <i>HH:MM</i>	The current time minus hours and minutes specified	NOW-02:00 or -02:00.

WHERERetention

Specifies the retention value, specified in days, that is associated with the backup set to update. The values are:

days

Specifies that the backup set that is retained this number of days is updated.

NOLimit

Specifies that the backup set retained indefinitely is updated.

WHEREDESCRIPTION

Specifies the description that is associated with the backup set to update. This parameter is optional. You can specify wildcard characters for the description. Enclose the description in quotation marks if it contains any blank characters.

WHEREDATATYPE

Specifies the backup sets containing the specified types of data are to be updated. This parameter is optional. The default is that backup sets for all types of data (file level, image, and application) are to be updated. To specify multiple data types, separate each data type with a comma and no intervening spaces. Possible values are:

ALL

Specifies that backup sets for all types of data (file level, image, and application) are to be updated. This is the default.

FILE

Specifies that a file level backup set is to be updated. File level backup sets contain files and directories backup up by the backup-archive client.

IMAGE

Specifies that an image backup set is to be updated. Image backup sets contain images created by the backup-archive client **BACKUP IMAGE** command.

Preview

Specifies whether to preview the list of backup sets to update, without actually updating the backup sets. This parameter is optional. The default is No. The values are:

No Specifies that the backup sets are updated.

Yes

Specifies that the server displays the backup sets to update, without actually updating the backup sets.

VERSION

Specifies the version of the backup set to update. Backup sets with the same prefix name are considered to be different versions of the same backup set. This parameter is optional. The default is to update any version that matches the criteria specified on the command. The values are:

UPDATE BACKUPSET

Any

Specifies that any version that matches the criteria specified on the command should be updated.

Latest

Specifies that only the most recent version of the backup set should be updated. If other criteria specified on the command (for example, ENDDATE or WHERERETENTION) exclude the most recent version of the backup set, then no backup set will be updated.

Example: Update a retention period

Update the retention period where the description is Healthy Computers. The retention period is assigned to backup set PERS_DATA.3099 that contains data from client node JANE. Change the retention period to 70 days.

```
update backupset jane pers_data.3099
retention=70 wheredescription="healthy computers"
```

Related commands

Table 394. Commands related to UPDATE BACKUPSET

Command	Description
DEFINE BACKUPSET	Defines a previously generated backup set to a server.
DEFINE NODEGROUP	Defines a group of nodes.
DEFINE NODEGROUPMEMBER	Adds a client node to a node group.
DELETE BACKUPSET	Updates a retention value associated with a backup set.
DELETE NODEGROUP	Deletes a node group.
DELETE NODEGROUPMEMBER	Deletes a client node from a node group.
GENERATE BACKUPSET	Generates a backup set of a client's data.
GENERATE BACKUPSETTOC	Generates a table of contents for a backup set.
QUERY BACKUPSET	Displays backup sets.
QUERY BACKUPSETCONTENTS	Displays contents contained in backup sets.
QUERY NODEGROUP	Displays information about node groups.
UPDATE NODEGROUP	Updates the description of a node group.

UPDATE CLIENTOPT (Update a client option sequence number)

Use this command to update the sequence number of a client option in a client option set.

Privilege class

To issue this command, you must have system privilege or unrestricted policy privilege.

Syntax

```
►►—UPDate CLIENTOpt—option_set_name—option_name—————►
►—current_sequence_number—new_sequence_number—————►◄
```

Parameters

option_set_name **(Required)**

Specifies the name of the option set.

option_name **(Required)**

Specifies a valid client option.

current_sequence_number **(Required)**

Specifies the current sequence number of the option.

new_sequence_number **(Required)**

Specifies the new sequence number of the option.

Example: Update a client option sequence number

To update the current client option sequence number issue the following command:

```
update clientopt eng dateformat 0 9
```

Related commands

Table 395. Commands related to UPDATE CLIENTOPT

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLIENTOPT	Adds a client option to a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
DELETE CLOPTSET	Deletes a client option set.
QUERY CLOPTSET	Displays information about a client option set.

UPDATE CLOPTSET (Update a client option set description)

Use this command to update the description for a client option set.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node is assigned.

Syntax

```
►►—UPDate CLOptset—option_set_name—DESCRiption—==—description—◄◄
```

Parameters

option_set_name (Required)

Specifies the name of the option set.

DESCRiption (Required)

Specifies a description of the client option set. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains blank characters.

Example: Update a client option set description

Update the description for a client option set named ENG.

```
update cloptset eng description="unix"
```

Related commands

Table 396. Commands related to UPDATE CLOPTSET

Command	Description
COPY CLOPTSET	Copies a client option set.
DEFINE CLIENTOPT	Adds a client option to a client option set.
DEFINE CLOPTSET	Defines a client option set.
DELETE CLIENTOPT	Deletes a client option from a client option set.
DELETE CLOPTSET	Deletes a client option set.
QUERY CLOPTSET	Displays information about a client option set.
UPDATE CLIENTOPT	Updates the sequence number of a client option in a client option set.

UPDATE COLLOGROUP (Update a collocation group)

Use this command to modify the description of a collocation group.

Privilege class

To issue this command, you must have system or unrestricted storage privilege.

Syntax

```
►►—UPDate COLLOGGroup—group_name—DESCRiption—=—description—◄◄
```

Parameters

group_name

Specifies the name of the collocation group whose description you want to update.

DESCRiption

Specifies a description of the collocation group. This parameter is required. The maximum length of the description is 255 characters. If the description contains any blanks, enclose the entire description in quotation marks.

Example: Update a collocation group

Update the collocation group, GROUP1, with a new description.

```
update colloggroup group1 "Human Resources"
```

Related commands

Table 397. Commands related to UPDATE COLLOGROUP

Command	Description
DEFINE COLLOGGROUP	Defines a collocation group.
DEFINE COLLOCMEMBER	Adds a client node to a collocation group.
DEFINE STGPOOL	Defines a storage pool as a named collection of server storage media.
DELETE COLLOGGROUP	Deletes a collocation group.
DELETE COLLOCMEMBER	Deletes a client node from a collocation group.
MOVE NODEDATA	Moves data for one or more nodes, or a single node with selected file spaces.
QUERY COLLOGGROUP	Displays information about collocation groups.
QUERY NODE	Displays partial or complete information about one or more clients.
QUERY NODEDATA	Displays information about the location and size of data for a client node.
QUERY STGPOOL	Displays information about storage pools.
REMOVE NODE	Removes a client from the list of registered nodes for a specific policy domain.
UPDATE STGPOOL	Changes the attributes of a storage pool.

UPDATE COPYGROUP (Update a copy group)

Use this command to update a backup or archive copy group. To allow clients to use the updated copy group, you must activate the policy set that contains the copy group.

Tip: The **UPDATE COPYGROUP** command fails if you specify a copy storage pool as a destination.

The **UPDATE COPYGROUP** command takes two forms, depending upon whether the update is for a backup copy group or for an archive copy group. The syntax and parameters for each form are defined separately.

- “UPDATE COPYGROUP (Update a backup copy group)” on page 1183
- “UPDATE COPYGROUP (Update a defined archive copy group)” on page 1187

Table 398. Commands related to UPDATE COPYGROUP


Command	Description
ACTIVATE POLICYSET	Validates and activates a policy set.
ASSIGN DEFMGMTCLASS	Assigns a management class as the default for a specified policy set.
COPY MGMTCLASS	Creates a copy of a management class.
DEFINE COPYGROUP	Defines a copy group for backup or archive processing within a specified management class.
DEFINE MGMTCLASS	Defines a management class.
DELETE COPYGROUP	Deletes a backup or archive copy group from a policy domain and policy set.
DELETE MGMTCLASS	Deletes a management class and its copy groups from a policy domain and policy set.
EXPIRE INVENTORY	Manually starts inventory expiration processing.
QUERY COPYGROUP	Displays the attributes of a copy group.
QUERY MGMTCLASS	Displays information about management classes.

UPDATE COPYGROUP (Update a backup copy group)

Use this command to update a defined backup copy group.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the copy group belongs.

Syntax**Parameters*****domain_name* (Required)**

Specifies the policy domain to which the copy group belongs.

***policy_set_name* (Required)**

Specifies the policy set to which the copy group belongs. You cannot update a copy group in the ACTIVE policy set.

***class_name* (Required)**

Specifies the management class to which the copy group belongs.

STANDARD

Specifies the copy group, which must be STANDARD. This parameter is optional.

Type=Backup

Specifies that you want to update a backup copy group. This parameter is optional.

DEStination

Specifies the primary storage pool where the server initially stores backup data. This parameter is optional. You cannot specify a copy storage pool as the destination.

UPDATE COPYGROUP — backup

FREquency

Specifies how frequently the server can back up a file. This parameter is optional. The server backs up a file only when the specified number of days has elapsed since the last backup. The FREQUENCY value is used only during a full incremental backup operation. This value is ignored during selective backup or partial incremental backup. You can specify an integer from 0 to 9999. The value 0 means that the server can back up a file regardless of when the file was last backed up.

VERExists

Specifies the maximum number of backup versions to retain for files that are currently on the client file system. This parameter is optional.

If an incremental backup causes the limit to be exceeded, the server expires the oldest backup version that exists in server storage. Possible values are:

number

Specifies the number of backup versions to retain for files that are currently on the client file system. You can specify an integer from 1 to 9999.

NOLimit

Specifies that you want the server to retain all backup versions.

The number of backup versions to retain is controlled by this parameter until versions exceed the retention time specified by the RETEXTRA parameter.

VERDeleted

Specifies the maximum number of backup versions to retain for files that have been deleted from the client file system after being backed up using the server. This parameter is optional.

If a user deletes a file from the client file system, the next incremental backup causes the server to change the active backup version of the file to inactive and expire the oldest versions in excess of this number. The expiration date for the remaining versions is determined by the retention time specified by the RETEXTRA or RETONLY parameter. Possible values are:

number

Specifies the number of backup versions to retain for files that are deleted from the client file system after being backed up. You can specify a value from 0 to 9999.

NOLimit

Specifies that you want the server to retain all backup versions for files that are deleted from the client file system after being backed up.

RETEExtra

Specifies the number of days that the server retains a backup version after that version becomes inactive. A version of a file becomes inactive when the client stores a more recent backup version, or when the client deletes the file from the workstation and then runs a full incremental backup. The server deletes inactive versions based on retention time even if the number of inactive versions does not exceed the number allowed by the VEREXISTS or VERDELETED parameters. This parameter is optional. Possible values are:

days

Specifies the number of days to retain inactive backup versions. You can specify an integer from 0 to 9999.

NOLimit

Specifies that you want to retain inactive backup versions indefinitely.

If you specify `NOLIMIT`, the server deletes extra backup versions based on the `VEREXISTS` parameter (when the file still exists on the client file system) or the `VERDELETED` parameter (when the file no longer exists on the client file system).

RETOnly

Specifies the number of days to retain the last backup version of a file that has been deleted from the client file system. This parameter is optional. Possible values are:

days

Specifies the number of days to retain the last remaining inactive copy of a file. You can specify an integer from 0 to 9999.

NOLimit

Specifies that you want to keep the last remaining inactive version of a file indefinitely.

If you specify `NOLIMIT`, the server retains the last remaining backup version forever, unless a user or administrator deletes the file from server storage.

MODE

Specifies whether the server backs up a file only if the file has changed since the last backup, or whenever a client requests a backup. This parameter is optional. Possible values are:

MODified

Specifies that the file is backed up only if it has changed since the last backup. A file is considered changed if any of the following is true:

- The date last modified is different
- The file size is different
- The file owner is different
- The file permissions are different

ABSolute

Specifies that the file is backed up regardless of whether it has been changed.

The `MODE` value is used only for full incremental backup. This value is ignored during partial incremental backup or selective backup.

SERialization

Specifies how the server processes files or directories when they are modified during backup processing. This parameter is optional. Possible values are:

SHRStatic

Specifies that the server backs up a file or directory only if it is not being modified during backup. The server attempts to perform a backup as many as four times, depending on the value specified for the `CHANGINGRETRIES` client option. If the file or directory is modified during each backup attempt, the server does not back it up.

Static

Specifies that the server backs up a file or directory only if it is not being modified during backup. The server attempts to perform the backup only once.

Platforms that do not support the `STATIC` option default to `SHRSTATIC`.

UPDATE COPYGROUP — backup

SHRDYnamic

Specifies that if the file or directory is being modified during a backup attempt, the server backs up the file or directory during the last attempt even though the file or directory is being modified. The server attempts to perform a backup as many as four times, depending on the value specified for the CHANGINGRETRIES client option.

DYnamic

Specifies that the server backs up a file or directory on the first attempt, regardless of whether the file or directory is being modified during backup processing.

Important: Be careful about using the SHR DYNAMIC and DYNAMIC values. Tivoli Storage Manager uses these values to determine if it backs up a file or directory while modifications are occurring. As a result, the backup version might be a fuzzy backup. A fuzzy backup does not accurately reflect what is currently in the file or directory because it contains some, but not all, modifications. If a file that contains a fuzzy backup is restored, the file may or may not be usable, depending on the application that uses the file. If a fuzzy backup is not acceptable, set SERIALIZATION to SHRSTATIC or STATIC so that Tivoli Storage Manager creates a backup version only if the file or directory is not being modified.

TOCDestination

Specifies the primary storage pool in which a table of contents (TOC) will initially be stored for any NDMP backup or backup set operation for which a TOC is generated. This parameter is optional. You cannot specify a copy storage pool as the destination. The storage pool specified for the destination must have NATIVE or NONBLOCK data format. To avoid mount delays, it is recommended that the storage pool have a device class of DISK or DEVTYPE=FILE. TOC generation is an option for NDMP backup operations, but is not supported for other image-backup operations.

To remove an existing TOC destination from the copy group, specify a null string ("") for this value.

If TOC creation is requested for a backup operation that uses NDMP and the image is bound to a management class whose backup copy group does not specify a TOC destination, the outcome will depend on the TOC parameter for the backup operation.

- If TOC=PREFERRED (the default), the backup proceeds without creation of a TOC.
- If TOC=YES, the entire backup fails because no TOC can be created.

Example: Update a backup copy group

Update the backup copy group (STANDARD) in the EMPLOYEE_RECORDS policy domain, VACATION policy set, ACTIVEFILES management class. Change the destination to DISKPOOL, with a minimum interval of seven days between backups, regardless of whether the files have been modified. Retain up to three backup versions while a file still exists on a client file system.


```
update copygroup employee_records vacation
activefiles type=backup destination=diskpool
frequency=7 verexists=3 mode=absolute
```

UPDATE COPYGROUP (Update a defined archive copy group)

Use this command to update a defined archive copy group.

Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the copy group belongs.

Syntax**Parameters*****domain_name* (Required)**

Specifies the policy domain to which the copy group belongs.

***policy_set_name* (Required)**

Specifies the policy set to which the copy group belongs. You cannot update a copy group in the ACTIVE policy set.

***class_name* (Required)**

Specifies the management class to which the copy group belongs.

STANDARD

Specifies the copy group, which must be **STANDARD**. This parameter is optional.

Type=Archive

Specifies that you want to update an archive copy group. This parameter is required.

DESTINATION

Specifies the primary storage pool where the server initially stores the archive copy. This parameter is optional. You cannot specify a copy storage pool as the destination.

FREQUENCY=Cmd

Specifies the copy frequency, which must be **CMD**. This parameter is optional.

RETVER

Specifies the number of days to keep an archive copy. This parameter is optional. Possible values are:

UPDATE COPYGROUP — archive

days

Specifies the number of days to keep an archive copy. You can specify an integer from 0 to 30000.

NOLimit

Specifies that you want to keep an archive copy indefinitely.

If you specify **NOLIMIT**, the server retains archive copies forever, unless a user or administrator deletes the file from server storage.

The value of the **RETVER** parameter can affect the management class to which the server binds an archived directory. If the client does not use the ARCHMC option, the server binds directories that are archived to the default management class. If the default management class has no archive copy group, the server binds directories that are archived to the management class with the shortest retention period.

MODE=ABSolute

Specifies that a file is always archived when the client requests it. The MODE must be ABSOLUTE. This parameter is optional.

REMin

Specifies the minimum number of days to keep an archive copy after it has been archived. This parameter is optional. The default value is 365.

SERialization

Specifies how the server processes files that are modified during archive. This parameter is optional. Possible values are:

SHRStatic

Specifies that the server does not archive a file that is being modified. The server attempts to perform an archive as many as four times, depending on the value specified for the CHANGINGRETRIES client option. If the file is modified during the archive attempt, the server does not archive the file.

Static

Specifies that the server does not archive a file that is being modified. If a file is modified during the archive attempt, the server does not archive the file.

Platforms that do not support the STATIC option default to SHRSTATIC.

SHRDynamic

Specifies that if the file is being modified during an archive attempt, the server archives the file during its last attempt even though the file is being modified. The server attempts to archive the file as many as four times, depending on the value specified for the CHANGINGRETRIES client option.

Dynamic

Specifies that the server archives a file on the first attempt, regardless of whether the file is being modified during archive processing.

Important: Be careful about using the SHRDYNAMIC and DYNAMIC values. Tivoli Storage Manager uses them to determine if it archives a file while modifications are occurring. As a result, the archive copy might be a fuzzy backup. A fuzzy backup does not accurately reflect what is currently in the file because it contains some, but not all, modifications. If a file that contains a fuzzy backup is retrieved, the file may or may not be usable, depending on the application that uses the file. If a fuzzy backup is not acceptable, set SERIALIZATION to SHRSTATIC or STATIC so that Tivoli Storage Manager creates an archive copy only if the file is not being modified.

Tip: Be cautious when selecting retention values for primary storage pools that are of type RECLAMATIONTYPE=SNAPLOCK. Volumes in these types of storage pools cannot be deleted until after their retention dates have passed.

Example: Update multiple elements of a copy group

Update the archive copy group (STANDARD) in the EMPLOYEE_RECORDS policy domain, VACATION policy set, ACTIVEFILES management class. Change the destination to TAPEPOOL. Keep archive copies for 190 days.

```
update copygroup employee_records vacation
activefiles standard type=archive
destination=tapepool retver=190
```


UPDATE DATAMOVER (Update a data mover)

Use this command to update the definition for a data mover or set a data mover off-line when the hardware is being maintained.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

data_mover_name (Required)

Specifies the name of the data mover.

HLAddress

Specifies either the new numerical IP address or the new domain name, which is used to access the NAS file server. This parameter is optional.

LLAddress

Specifies the new TCP port number to access the NAS file server for Network Data Management Protocol (NDMP) sessions. This parameter is optional.

USERid

Specifies the user ID for a user that is authorized to initiate an NDMP session with the NAS file server. For example, enter the administrative ID for a NetApp file server. This parameter is optional.

PASsword

Specifies the new password for the user ID to log onto the NAS file server. This parameter is optional.

ONLine

Specifies whether the data mover is available for use. This parameter is optional.

Yes

Specifies that the data mover is available for use.

No

Specifies that the data mover is not available for use.

Attention: If a library is controlled using a path from a data mover to the library, and the data mover is offline, the server will not be able to access the library. If the server is halted and restarted while the data mover is offline, the library will not be initialized.

Example: Update a data mover IP address

Update the data mover for the node named NAS1. Change the numerical IP address from 9.67.97.103 to 9.67.97.109.

```
update datamover nas1 haddress=9.67.97.109
```

Example: Update a data mover domain name

Update the data mover for the node named NAS1. Change the numerical IP address from 9.67.97.109 to the domain name of NETAPP2.TUCSON.IBM.COM.

```
update datamover nas1 haddress=netapp2.tucson.ibm.com
```

Related commands

Table 399. Commands related to UPDATE DATAMOVER

Command	Description
DEFINE DATAMOVER	Defines a data mover to the Tivoli Storage Manager server.
DEFINE PATH	Defines a path from a source to a destination.
DELETE DATAMOVER	Deletes a data mover.
QUERY DATAMOVER	Displays data mover definitions.
REGISTER NODE	Defines a client node to the server and sets options for that user.
UPDATE NODE	Changes the attributes associated with a client node.

UPDATE DEVCLASS (Update the attributes of a device class)

Use this command to update a defined device class.

Note: The DISK device class is predefined by IBM Tivoli Storage Manager and cannot be modified with the UPDATE DEVCLASS command.

If you are updating a device class for devices that are to be accessed through a z/OS media server, see “UPDATE DEVCLASS - z/OS media server (Update device class for z/OS media server)” on page 1249.

The syntax and parameter descriptions are provided according to the device type. The syntax and parameter information is presented in the following order.

- “UPDATE DEVCLASS (Update a 3590 device class)” on page 1193
- “UPDATE DEVCLASS (Update a 3592 device class)” on page 1197
- “UPDATE DEVCLASS (Update a 4MM device class)” on page 1202
- “UPDATE DEVCLASS (Update an 8MM device class)” on page 1206
- “UPDATE DEVCLASS (Update a DLT device class)” on page 1211
- “UPDATE DEVCLASS (Update an ECARTRIDGE device class)” on page 1216
- “UPDATE DEVCLASS (Update a FILE device class)” on page 1223
- “UPDATE DEVCLASS (Update an LTO device class)” on page 1227
- “UPDATE DEVCLASS (Update a NAS device class)” on page 1233
- “UPDATE DEVCLASS (Update OPTICAL and WORM device classes)” on page 1235
- “UPDATE DEVCLASS (Update a QIC device class)” on page 1238
- “UPDATE DEVCLASS (Update a REMOVABLEFILE device class)” on page 1242
- “UPDATE DEVCLASS (Update a SERVER device class)” on page 1244
- “UPDATE DEVCLASS (Update a VOLSAFE device class)” on page 1246

Table 400. Commands related to UPDATE DEVCLASS

Command	Description
BACKUP DEVCONFIG	Backs up Tivoli Storage Manager device information to a file.
DEFINE DEVCLASS	Defines a device class.
DEFINE LIBRARY	Defines an automated or manual library.
DELETE DEVCLASS	Deletes a device class.
QUERY DEVCLASS	Displays information about device classes.
QUERY DIRSPACE	Displays information about FILE directories.
UPDATE LIBRARY	Changes the attributes of a library.

UPDATE DEVCLASS (Update a 3590 device class)


Use the 3590 device class when you are using 3590 tape devices.

If you are defining a device class for devices that are to be accessed through a z/OS media server, see “UPDATE DEVCLASS (Update a 3590 device class for z/OS media server)” on page 1250.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined.

LIBRARY

Specifies the name of the defined library object that contains the tape drives that can be used by this device class. For information about defining a library object, see the **DEFINE LIBRARY** command.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional.

The following tables list the recording formats, estimated capacities and recording format options for 3590 devices:

Table 401. Recording formats and default estimated capacities for 3590

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
3590B	10.0 GB	Uncompressed (basic) format
3590C	See note 20.0 GB	Compressed format
3590E-B	10.0 GB	Uncompressed (basic) format, similar to the 3590B format
3590E-C	See note 20.0 GB	Compressed format, similar to the 3590C format
3590H-B	30.0 GB (J cartridge-standard length) 60.0 GB (K cartridge-extended length)	Uncompressed (basic) format, similar to the 3590B format
3590H-C	See note 60.0 GB (J cartridge-standard length) 120.0 GB (K cartridge-extended length)	Compressed format, similar to the 3590C format

Note: If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity may be greater than the listed value.

Table 402. 3590 device recording format selections

Device	Format					
	3590B	3590C	3590E-B	3590E-C	3590H-B	3590H-C
3590	Read/Write	Read/Write	–	–	–	–
Ultra SCSI	Read/Write	Read/Write	–	–	–	–
3590E	Read	Read	Read/Write	Read/Write	–	–
3590H	Read	Read	Read	Read	Read/Write	Read/Write

ESTCAPacity

Specifies the estimated capacity for the sequential access volumes categorized by this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by **K** (kilobytes), **M** (megabytes), or **G** (gigabytes). For example, `ESTCAPACITY=5M` specifies that the estimated capacity for a volume in this device class is 5 MB.

The smallest value allowed is 100 KB (that is, `ESTCAPACITY=100K`). To force the IBM Tivoli Storage Manager server to determine the estimated capacity for the volumes assigned to this device class, specify `ESTCAPACITY=""`.

PREFIX

Specifies the high level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, consider using a volume prefix that conforms to your naming conventions.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

`AB.CD2.E`

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a tape volume data set name using the default prefix is `ADSM.BFS`.

MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

However, for `EXTERNAL` library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transaction will continue and complete, but new transactions will be terminated.


UPDATE DEVCLASS (Update a 3592 device class)

If you are defining a device class for devices that are to be accessed through a z/OS media server, see “UPDATE DEVCLASS (Update a 3592 device class for z/OS media server)” on page 1255.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

device_class_name (Required)

Specifies the name of the device class to be updated. The maximum length of the device class name is 30 characters.

LIBRARY

Specifies the name of the defined library object that contains the tape drives that can be used by this device class. For information about defining a library object, see the **DEFINE LIBRARY** command.

LBProtect

Specifies whether logical block protection is used to ensure the integrity of data stored on tape. When **LBPROTECT** is set to **READWRITE** or to **WRITEONLY**, the server uses this feature of the tape drive for logical block protection and generates cyclic redundancy check (CRC) protection information for each data block written on tape. The server also validates the CRC protection information when data is read from the tape.

The following values are possible:

READWrite

Specifies that logical block protection is enabled in the server and the tape drive for both read and write operations. Data is stored with CRC information in each block. This mode affects performance because additional processor usage is required for Tivoli Storage Manager and the tape drive to calculate and compare CRC values. The **READWRITE** value does not affect backup sets and data that is generated by the **BACKUP DB** command.

When the **LBPROTECT** parameter is set to **READWRITE**, you do not have to specify the **CRCDATA** parameter in a storage pool definition because logical block protection provides better protection against data corruption.

WRITEOnly

Specifies that logical block protection is enabled in the server and the tape drive for write operations only. Data is stored containing CRC information in each block. For read operations, the server and the tape drive do not validate the CRC. This mode affects performance because additional processor usage is required for Tivoli Storage Manager to generate the CRC and for the tape drive to calculate and compare CRC values for write operations. The **WRITEONLY** value does not affect backup sets and data that are generated by the **BACKUP DB** command.

No Specifies that logical block protection is not enabled in the server and the tape drive for read and write operations. However, the server enables logical block protection on write operations for a filling volume that already has data with logical block protection.

Restriction: Logical block protection is supported only on IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.

SCALECapacity

Specifies the percentage of the media capacity that can be used to store data. This parameter is optional. Possible values are 20, 90, or 100.

Setting the scale capacity percentage to 100 provides maximum storage capacity. Setting it to 20 provides fastest access time.

Note: The scale capacity value will only take effect when data is first written to a volume. Any updates to the device class for scale capacity will not affect volumes that already have data written to them until the volume is returned to scratch status.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional.

If the drives are in a library that includes drives of different tape technology, do not use the **DRIVE** value. Use the specific format that the drives use.

Refer to the *Administrator's Guide* for more information.

The following table lists the recording formats, estimated capacities and recording format options for 3592 devices:

Table 403. Recording formats and default estimated capacities for 3592

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
3592	300 GB	Uncompressed (basic) format
3592C	See note 900 GB	Compressed format
3592-2	500 GB 700 GB	Uncompressed (basic) format JA tapes Uncompressed (basic) format JB tapes
3592-2C	1.5 TB 2.1 TB	Compressed format JA tapes Compressed format JB tapes
3592-3	640 GB 1 TB	Uncompressed (basic) format JA tapes Uncompressed (basic) format JB tapes
3592-3C	1.9 TB 3 TB	Compressed format JA tapes Compressed format JB tapes
3592-4	400 GB 1.5 TB 3.1 TB	Uncompressed (basic) format JK tapes Uncompressed (basic) format JB tapes Uncompressed (basic) format JC tape
3592-4C	1.2 TB 4.4 TB 9.4 TB	Compressed format JK tapes Compressed format JB tapes Compressed format JC tapes

Note: If this format uses the tape-drive hardware-compression feature, depending on the effectiveness of compression, the actual capacity might be different than the listed value.

Important: For optimal performance, avoid mixing different generations of drives in a single SCSI library. If you must mix drive generations in an SCSI library, use one of the special configurations that are described in the *Administrator's Guide* to prevent or minimize media problems.

Special configurations are also required for mixing different generations of 3592 drives in 349x and ACSLS libraries. For details, see the *Administrator's Guide*.

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by **K** (kilobytes), **M** (megabytes), or **G** (gigabytes). For example, `ESTCAPACITY=5M` specifies that the estimated capacity for a volume in this device class is 5 MB.

The smallest value allowed is 100 KB (that is, `ESTCAPACITY=100K`). To force the IBM Tivoli Storage Manager server to determine the estimated capacity for the volumes assigned to this device class, specify `ESTCAPACITY=""`.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, use a volume prefix that conforms to your naming conventions.

An example of a tape volume data set name using the default prefix is `ADSM.BFS`.

MOUNTRetention

Specifies the amount of time, in minutes, to retain an idle sequential access volume before dismounting it. This parameter is optional. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for `EXTERNAL` library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the `MOUNTLIMIT` parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

Note: For EXTERNAL library types, do not use the DRIVES default for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never be allowed to exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool.

DRIVEEncryption

Specifies whether drive encryption is permitted. This parameter is optional.

Updating this parameter will affect empty volumes only. If a filling volume was previously encrypted or is currently unencrypted, and you update the DRIVEENCRYPTION parameter, the volume maintains its original encrypted or unencrypted status. The filling volume also maintains its original key-management status.

ON Specifies that Tivoli Storage Manager is the key manager for drive encryption and will permit drive encryption for empty storage pool volumes only if the application method is enabled. (Other types of volumes—for example, backup sets, export volumes, and database backup volumes—will not be encrypted.) If you specify ON and you enable either the library or system method of encryption, drive encryption will not be permitted and backup operations will fail.

ALLOW

Specifies that Tivoli Storage Manager does not manage the keys for drive encryption. However, drive encryption for empty volumes is permitted if either the library or system method of encryption is enabled.

OFF

Specifies that drive encryption will not be permitted. If you enable either the library or system method of encryption, backups will fail. If you enable the application method, Tivoli Storage Manager will disable encryption and backups will be attempted.

UPDATE DEVCLASS (Update a 4MM device class)

Use the 4MM device class when you are using 4 mm tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

device_class_name (Required)

Specifies the name of the device class to be defined.

LIBRARY

Specifies the name of the defined library object that contains the 4 mm tape drives used by this device class. This parameter is optional. For information about defining a library object, see the **DEFINE LIBRARY** command.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional.

The following table lists the recording formats and estimated capacities for 4 mm devices:

Table 404. Recording formats and default estimated capacities for 4 mm tapes

Format	Estimated Capacity	Description
DRIVE	—	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
DDS1	1.3 GB (60-meter) 2.0 GB (90-meter)	Uncompressed format, applies only to 60-meter and 90-meter tapes
DDS1C	See note 1.3 GB (60-meter) 2.0 GB (90-meter)	Compressed format, applies only to 60-meter and 90-meter tapes
DDS2	4.0 GB	Uncompressed format, only applies to 120-meter tapes
DDS2C	See note 8.0 GB	Compressed format, only applies to 120-meter tapes
DDS3	12.0 GB	Uncompressed format, only applies to 125-meter tapes
DDS3C	See note 24.0 GB	Compressed format, only applies to 125-meter tapes
DDS4	20.0 GB	Uncompressed format, only applies to 150-meter tapes
DDS4C	See note 40.0 GB	Compressed format, only applies to 150-meter tapes
DDS5	36 GB	Uncompressed format, when using DAT 72 media
DDS5C	See note 72 GB	Compressed format, when using DAT 72 media
DDS6	80 GB	Uncompressed format, when using DAT 160 media
DDS6C	See note 160 GB	Compressed format, when using DAT 160 media

Note: If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity may be greater than the listed value.

ESTCAPacity

Specifies the estimated capacity for the sequential access volumes categorized by this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by **K** (kilobytes), **M** (megabytes), or **G** (gigabytes). For example, `ESTCAPACITY=5M` specifies that the estimated capacity for a volume in this device class is 5 MB.

The smallest value allowed is 100 KB (that is, `ESTCAPACITY=100K`). To force the IBM Tivoli Storage Manager server to determine the estimated capacity for the volumes assigned to this device class, specify `ESTCAPACITY=""`.

For more information on the default estimated capacity for 4 mm tapes, see Table 404 on page 1203.

PREFIX

Specifies the high level qualifier of the file name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have already established a tape label naming convention that supports your current tape management system, consider using a tape volume prefix that conforms to your naming conventions.

MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for **EXTERNAL** library types (that is, a library managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true mount limit value (including online status).

Note: For **EXTERNAL** library types, do not use the **DRIVES** default for the **MOUNTLIMIT** value. Specify the number of drives for the library as the **MOUNTLIMIT** value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transaction will continue and complete but new transactions will be terminated.


UPDATE DEVCLASS (Update an 8MM device class)

Use the 8MM device class when you are using 8 mm tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

device_class_name (Required)

Specifies the name of the device class to be updated.

LIBRARY

Specifies the name of the defined library object that contains the 8 mm tape drives that can be used by this device class. For information about defining a library object see the **DEFINE LIBRARY** command.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional.

The following table lists the recording formats and estimated capacities for 8 mm devices:

Table 405. Recording format and default estimated capacity for 8 mm tape

Format		
Medium Type	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
8200	2.3 GB	Uncompressed (standard) format, using standard 112-meter tape cartridges
8200C	See note 3.5 GB 4.6 GB	Compressed format, using standard 112-meter tape cartridges
8500	See note	Drives (Read Write)
15m	600 MB	Eliaint 820 (RW)
15m	600 MB	Exabyte 8500/8500C (RW)
15m	600 MB	Exabyte 8505 (RW)
54m	2.35 GB	Eliaint 820 (RW)
54m	2.35 GB	Exabyte 8500/8500C (RW)
54m	2.35 GB	Exabyte 8505 (RW)
112m	5 GB or 10.0 GB	Eliaint 820 (RW)
112m	5 GB or 10.0 GB	Exabyte 8500/8500C (RW)
112m	5 GB or 10.0 GB	Exabyte 8505 (RW)
160m XL	7 GB	Eliaint 820 (RW)
8500C	See note	Drives (Read Write)
15m	1.2 GB	Eliaint 820 (RW)
15m	1.2 GB	Exabyte 8500/8500C (RW)
15m	1.2 GB	Exabyte 8505 (RW)
54m	4.7 GB	Eliaint 820 (RW)
54m	4.7 GB	Exabyte 8500/8500C (RW)
54m	4.7 GB	Exabyte 8505 (RW)
112m	5 GB or 10.0 GB	Eliaint 820 (RW)
112m	5 GB or 10.0 GB	Exabyte 8500/8500C (RW)
112m	5 GB or 10.0 GB	Exabyte 8505 (RW)
160m XL	7 GB	Eliaint 820 (RW)
8900	See note	Drive (Read Write)
15m	–	Mammoth 8900 (R)
54m	–	Mammoth 8900 (R)
112m	–	Mammoth 8900 (R)
160m XL	–	Mammoth 8900 (R)
22m	2.5 GB	Mammoth 8900 (RW)
125m	–	Mammoth 8900 (RW with upgrade)
170m	40 GB	Mammoth 8900 (RW)

Table 405. Recording format and default estimated capacity for 8 mm tape (continued)

Format		
Medium Type	Estimated Capacity	Description
AIT	See note	Drive
SDX1-25C	25 GB	AIT, AIT2 and AIT3 drives
SDX1-35C	35 GB	AIT, AIT2 and AIT3 drives
SDX2-36C	36 GB	AIT2 and AIT3 drives
SDX2-50C	50 GB	AIT2 and AIT3 drives
SDX3-100C	100 GB	AIT3, AIT4, and AIT5 drives
SDX3X-150C	150 GB	AIT3-Ex, AIT4, and AIT5 drives
SDX4-200C	200 GB	AIT4 and AIT5 drives
SDX5-400C	400 GB	AIT5 drive
AITC	See note	Drive
SDX1-25C	50 GB	AIT, AIT2 and AIT3 drives
SDX1-35C	91 GB	AIT, AIT2 and AIT3 drives
SDX2-36C	72 GB	AIT2 and AIT3 drives
SDX2-50C	130 GB	AIT2 and AIT3 drives
SDX3-100C	260 GB	AIT3, AIT4, and AIT5 drives
SDX3X-150C	390 GB	AIT3-Ex, AIT4, and AIT5 drives
SDX4-200C	520 GB	AIT4 and AIT5 drives
SDX5-400C	1040 GB	AIT5 drive
M2	See note	Drive (Read Write)
75m	20.0 GB	Mammoth II (RW)
150m	40.0 GB	Mammoth II (RW)
225m	60.0 GB	Mammoth II (RW)
M2C	See note	Drive (Read Write)
75m	50.0 GB	Mammoth II (RW)
150m	100.0 GB	Mammoth II (RW)
225m	150.0 GB	Mammoth II (RW)
SAIT	See note	Drive (Read Write)
	500 GB	Sony SAIT1-500(RW)
SAITC	See note	Drive (Read Write)
	1300 GB (1.3 TB)	Sony SAIT1-500(RW)
VXA2	See note	Drive (Read Write)
V6 (62m)	20 GB	VXA-2
V10 (124m)	40 GB	
V17 (170m)	60 GB	
VXA2C	See note	Drive (Read Write)
V6 (62m)	40 GB	VXA-2
V10 (124m)	80 GB	
V17 (170m)	120 GB	
VXA3	See note	Drive (Read Write)
X6 (62m)	40 GB	VXA-3
X10 (124m)	86 GB	
X23 (230m)	160 GB	

Table 405. Recording format and default estimated capacity for 8 mm tape (continued)

Format		
Medium Type	Estimated Capacity	Description
VXA3C	See note	Drive (Read Write)
X6 (62m)	80 GB	VXA-3
X10 (124m)	172 GB	
X23 (230m)	320 GB	

Note: The actual capacities may vary depending on which cartridges and drives are used.

- For the AITC and SAITC formats, the normal compression ratio is 2.6:1.
- For the M2C format, the normal compression ratio is 2.5:1.

ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

You must specify this value as an integer followed by **K** (kilobytes), **M** (megabytes), or **G** (gigabytes). For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB.

The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K). To force the IBM Tivoli Storage Manager server to determine the estimated capacity for the volumes assigned to this device class, specify ESTCAPACITY="".

For more information on the default estimated capacity for 8 mm tapes, see Table 405 on page 1207.

PREFIX

Specifies the high level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have already established a tape label naming convention that supports your current tape management system, consider using a tape volume prefix that conforms to your naming conventions.

MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

For EXTERNAL library types (that is, a library managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true MOUNTLIMIT value (including online status).

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transaction will continue and complete, but new transactions will be terminated.

Example: Update the mount limit and capacity of an 8 mm device class

Update a device class named 8MMTAPE. Change the mount limit to 3 and the estimated capacity to 10 GB.

```
update devclass 8mmtape mountlimit=3 estcapacity=10G
```

Example: Update the mount retention period of an 8 mm device class

Update an 8 mm device class named 8MMTAPE to a 15 minute mount retention.


```
update devclass 8mmtape mountretention=15
```

UPDATE DEVCLASS (Update a DLT device class)

Use the DLT device class when you are using DLT tape devices.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax**Parameters*****device_class_name* (Required)**

Specifies the name of the device class to be updated.

LIBRARY

Specifies the name of the defined library object that contains the DLT tape drives that can be used by this device class. For information about defining a library object see the **DEFINE LIBRARY** command.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional.

The following table lists the recording formats and estimated capacities for DLT devices:

Table 406. Recording format and default estimated capacity for DLT

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
DLT1	40.0 GB	Uncompressed format, using only CompacTape III or CompacTape IV cartridges
DLT1C	See note 1 on page 1214. 80.0 GB	Compressed format, using only CompacTape III and CompacTape IV cartridges
DLT10	10.0 GB	Uncompressed format, using only CompacTape III or CompacTape IV cartridges
DLT10C	See note 1 on page 1214. 20.0 GB	Compressed format, using only CompacTape III and CompacTape IV cartridges
DLT15	15.0 GB	Uncompressed format, using only CompacTape IIIxt or CompacTape IV cartridges (not CompacTape III) Note: Valid with DLT2000XT, DLT4000, and DLT7000 drives
DLT15C	See note 1 on page 1214. 30.0 GB	Compressed format, using only CompacTape IIIxt or CompacTape IV cartridges (not CompacTape III) Valid with DLT2000XT, DLT4000, and DLT7000 drives
DLT20	20.0 GB	Uncompressed format, using only CompacTape IV cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT20C	See note 1 on page 1214. 40.0 GB	Compressed format, using only CompacTape IV cartridges Valid with DLT4000, DLT7000, and DLT8000 drives
DLT35	35.0 GB	Uncompressed format, using only CompacTape IV cartridges Valid with DLT7000 and DLT8000 drives
DLT35C	See note 1 on page 1214. 70.0 GB	Compressed format, using only CompacTape IV cartridges Valid with DLT7000 and DLT8000 drives

Table 406. Recording format and default estimated capacity for DLT (continued)

Format	Estimated Capacity	Description
DLT40	40.0 GB	Uncompressed format, using CompacTape IV cartridges Valid with a DLT8000 drive
DLT40C	See note 1 on page 1214. 80.0 GB	Compressed format, using CompacTape IV cartridges Valid with a DLT8000 drive
DLT2	80.0 GB	Uncompressed format, using Quantum DLT tape VS1 media
DLT2C	See note 1 on page 1214. 160.0 GB	Compressed format, using Quantum DLT tape VS1 media
DLT4	160.0 GB	Uncompressed format, using Quantum DLTtape VS1 cartridges. Valid with Quantum DLT-V4 drive
DLT4C	See note 1 on page 1214. 320.0 GB	Compressed format, using Quantum DLTtape VS1 cartridges. Valid with Quantum DLT-V4 drive
SDLT See note 2 on page 172.	100.0 GB	Uncompressed format, using Super DLT Tape 1 cartridges Valid with a Super DLT drive
SDLTC See note 2 on page 172.	See note 1 on page 172. 200.0 GB	Compressed format, using Super DLT Tape 1 cartridges Valid with a Super DLT drive
SDLT320 See note 2 on page 172.	160.0 GB	Uncompressed format, using Quantum SDLT I media Valid with a Super DLT drive
SDLT320C See note 2 on page 172.	See note 1 on page 172. 320.0 GB	Compressed format, using Quantum SDLT I media Valid with a Super DLT drive
SDLT600	300.0 GB	Uncompressed format, using SuperDLTtape-II media Valid with a Super DLT drive
SDLT600C	See note 1 on page 1214. 600.0 GB	Compressed format, using SuperDLTtape-II media Valid with a Super DLT drive
DLTS4	800 GB	Uncompressed format, using Quantum DLT S4 media. Valid with a DLT-S4 drive
DLTS4C	See note 1 on page 1214. 1.6 TB	Compressed format, using Quantum DLT S4 media. Valid with a DLT-S4 drive

Table 406. Recording format and default estimated capacity for DLT (continued)

Format	Estimated Capacity	Description
Note:		
1. Depending on the effectiveness of compression, the actual capacity may be greater than the listed value.		
2. IBM Tivoli Storage Manager does not support a library that contains both Backward Read Compatible (BRC) SDLT and Non-Backward Read Compatible (NBRC) SDLT drives.		

ESTCAPacity

Specifies the estimated capacity for the sequential access volumes categorized by this device class. This parameter is optional. You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by **K** (kilobytes), **M** (megabytes), or **G** (gigabytes). For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB.

The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K). To force the IBM Tivoli Storage Manager server to determine the estimated capacity for the volumes assigned to this device class, specify ESTCAPACITY="".

For more information on estimated capacities, see Table 406 on page 1212.

PREFIX

Specifies the high level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, consider using a volume prefix that conforms to your naming conventions.

MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types (that is, a library managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true MOUNTLIMIT value (including online status).

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transaction will continue and complete, but new transactions will be terminated.

UPDATE DEVCLASS — ECARTRIDGE

UPDATE DEVCLASS (Update an ECARTRIDGE device class)


Use the ECARTRIDGE device class when you are using StorageTek drives such as the StorageTek SD-3, 9490, 9840 or 9940.


If you are defining a device class for devices that are to be accessed through a z/OS media server, see “UPDATE DEVCLASS (Update an ECARTRIDGE device class for z/OS media server)” on page 1265.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax





Notes:

- 1 You can use drive encryption only for Oracle StorageTek T10000B drives with a format value of DRIVE, T10000B or T10000B-C, and for Oracle StorageTek T10000C drives with a format value of DRIVE, T10000C or T10000C-C.
- 2 You cannot specify both WORM=YES and DRIVEENCRYPTION=ON.

Parameters

device_class_name **(Required)**

Specifies the name of the device class to be updated.

LIBRARY

Specifies the name of the defined library object with the ECARTRIDGE tape drives that can be used by this device class. For information about defining a library object, see the **DEFINE LIBRARY** command.

LBProtect

Specifies whether logical block protection is used to ensure the integrity of data stored on tape. When **LBPROTECT** is set to READWRITE or to WRITEONLY, the server uses this feature of the tape drive for logical block protection and generates cyclic redundancy check (CRC) protection information for each data block written on tape. The server also validates the CRC protection information when data is read from the tape.

The following values are possible:

READWrite

Specifies that logical block protection is enabled in the server and the tape drive for both read and write operations. Data is stored with CRC information in each block. This mode affects performance because additional processor usage is required for Tivoli Storage Manager and the tape drive to calculate and compare CRC values. The READWRITE value does not affect backup sets and data that is generated by the **BACKUP DB** command.

When the **LBPROTECT** parameter is set to READWRITE, you do not have to specify the **CRCDATA** parameter in a storage pool definition because logical block protection provides better protection against data corruption.

WRITEOnly

Specifies that logical block protection is enabled in the server and the tape drive for write operations only. Data is stored containing CRC information in each block. For read operations, the server and the tape drive do not validate the CRC. This mode affects performance because additional processor usage is required for Tivoli Storage Manager to generate the CRC and for the tape drive to calculate and compare CRC values for write

UPDATE DEVCLASS — ECARTRIDGE

operations. The `WRITEONLY` value does not affect backup sets and data that are generated by the `BACKUP DB` command.

No Specifies that logical block protection is not enabled in the server and the tape drive for read and write operations. However, the server enables logical block protection on write operations for a filling volume that already has data with logical block protection.

Restriction: Logical block protection is supported only on Oracle StorageTek T10000C drives.

FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional.

Attention: If you specify `DRIVE` for a device class that has non-compatible sequential access devices, then you must mount volumes on devices that are capable of reading or writing the format established when the volume was first mounted. This can cause delays if the only sequential access device that can access the volume is already in use.

The following table lists the recording formats and estimated capacities for ECARTRIDGE devices:

Table 407. Recording formats and default estimated capacities for ECARTRIDGE tapes

Format	Estimated Capacity	Description
DRIVE	–	The server selects the highest format that is supported by the drive on which a volume is mounted. Attention: Avoid specifying <code>DRIVE</code> when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives.
18T	360 MB	18-track uncompressed (standard) (read-only) format.
18TC	1.44 GB	18-track extended (read-only), compressed format.
36T	720 MB	36-track extended (read and write) format.
36TC	2.88 GB	36-track extended (read and write), compressed format.
SD3A	10 GB	Uncompressed (standard) format, using a 10 GB 'A' cartridge with 91 meters (298 feet) of tape.
SD3AC	40 GB	Compressed format, using a 10 GB 'A' cartridge with 91 meters (298 feet) of tape.
SD3B	25 GB	Uncompressed (standard) format, using a 25 GB 'B' cartridge with 204 meters (668 feet) of tape.
SD3BC	100 GB	Compressed format, using a 25 GB 'B' cartridge with 204 meters (668 feet) of tape.
SD3C	50 GB	Uncompressed (standard) format, using a 50 GB 'C' cartridge with 392 meters (1286 feet) of tape.
SD3CC	200 GB	Compressed format, using a 50 GB 'C' cartridge with 392 meters (1286 feet) of tape.

Table 407. Recording formats and default estimated capacities for ECARTRIDGE tapes (continued)

Format	Estimated Capacity	Description
9840	20 GB	Uncompressed 9840 format, using a StorageTek 9840 cartridge
9840-C	80 GB	LZ-1 Enhanced (4:1) compressed 9840 format, using a StorageTek 9840 cartridge
9940	60 GB	Uncompressed 9940 format, using a StorageTek 9940 cartridge
9940-C	120 GB	Compressed 9940 format, using a StorageTek 9940 cartridge
9940B	200 GB	Uncompressed 9940B format, using a StorageTek 9940 cartridge
9940B-C	400 GB	Compressed 9940B format, using a StorageTek 9940 cartridge
M8100	10 GB	Uncompressed (standard) format, using a 10 GB cartridge.
M8100C	10 GB	Compressed format, using a 10 GB cartridge. Because there is no mechanism to determine the type of cartridge in an M8100 drive, the server assumes a 10 GB estimated uncompressed capacity.
T9840C	40 GB	Uncompressed T9840C format, using a StorageTek 9840 cartridge
T9840C-C	80 GB	Compressed T9840C format, using a StorageTek 9840 cartridge
T9840D	75 GB	Uncompressed T9840D format, using a StorageTek 9840 cartridge
T9840D-C	150 GB	Compressed T9840D format, using a StorageTek 9840 cartridge
T10000A	500 GB	Uncompressed T10000A format, using a StorageTek T10000 cartridge
T10000A-C	1 TB	Compressed T10000A format, using a StorageTek T10000 cartridge
T10000B	1 TB	Uncompressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000B-C	2 TB	Compressed T10000B format, using an Oracle StorageTek T10000 cartridge
T10000C	5 TB	Uncompressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge
T10000C-C	10 TB	Compressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge

Notes:

- Some formats use a tape drive hardware compression feature. Depending on the effectiveness of compression, the actual capacity may be double or more than the listed value.
- T10000A drives can read and write the T10000A format only. T10000B drives can read, but cannot write, the T10000A format. T10000C drives can read, but cannot write, the T10000A and T10000B formats.

ESTCAPacity

Specifies the estimated capacity for the sequential access volumes categorized by this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by **K** (kilobytes), **M** (megabytes), or **G** (gigabytes). For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB.

The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K). To force the IBM Tivoli Storage Manager server to determine the estimated capacity for the volumes assigned to this device class, specify ESTCAPACITY="".

For more information on the default estimated capacity for cartridge tapes, see Table 407 on page 1218.

PREFIX

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional.

If you have already established a tape label naming convention that supports your current tape management system, consider using a tape volume prefix that conforms to your naming conventions.

MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types (that is, a library managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number from 0 to 9999.

MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

DRIVES

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true MOUNTLIMIT value (including online status).

Note: For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

number

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

0 (zero)

Specifies that no new transactions can gain access to the storage pool. Any current transaction will continue and complete but new transactions will be terminated.

DRIVEEncryption

Specifies whether drive encryption is permitted. This parameter is optional.

Restriction:

1. You can use drive encryption only for the following drives:
 - Oracle StorageTek T10000B drives that have a format value of DRIVE, T10000B, or T10000B-C
 - Oracle StorageTek T10000C drives that have a format value of DRIVE, T10000C, or T10000C-C
 2. You cannot specify Tivoli Storage Manager as the key manager for drive encryption of WORM (write once, read many) media. (Specifying both WORM=YES and DRIVEENCRYPTION=ON is not supported.)
 3. If encryption is enabled for a device class, and the device class is associated with a storage pool, the storage pool should not share a scratch pool with other device classes that cannot be encrypted. If a tape is encrypted, and you plan to use it on a drive that cannot be encrypted, you must manually relabel the tape before it can be used on that drive.
- ON** Specifies that Tivoli Storage Manager is the key manager for drive encryption and permits drive encryption for empty storage pool volumes only if the application method is enabled. (Other types of volumes are not encrypted. For example, backup sets, export volumes, and database backup volumes are not encrypted.) If you specify ON and you enable another method of encryption, drive encryption is not permitted and backup operations will fail.

ALLOW

Specifies that Tivoli Storage Manager does not manage the keys for drive encryption. However, drive encryption for empty volumes is permitted if another method of encryption is enabled.

EXTERNAL

Specifies that Tivoli Storage Manager does not manage the keys for drive encryption. Use this setting with an encryption methodology that is provided by another vendor and that is used with Application Method Encryption (AME) enabled on the drive. When you specify EXTERNAL and Tivoli Storage Manager detects that AME encryption is enabled, Tivoli Storage Manager does not turn off encryption. By contrast, when you

UPDATE DEVCLASS — ECARTRIDGE

specify `ALLOW` and Tivoli Storage Manager detects that AME encryption is enabled, Tivoli Storage Manager turns off encryption.

OFF

Specifies that drive encryption is not permitted. If you enable another method of encryption, backups fail. If you enable the application method, Tivoli Storage Manager disables encryption and backups are not attempted.

UPDATE DEVCLASS (Update a FILE device class)

Use the FILE device class when you are using files on magnetic disk storage as volumes that store data sequentially (as on tape).

For information about disk subsystem requirements for FILE-type disk storage, see the *Administrator's Guide*.


The FILE device class does not support EXTERNAL libraries.

If you are defining a device class for devices that are to be accessed through a z/OS media server, see "UPDATE DEVCLASS (Update a FILE device class for z/OS media server)" on page 1272.

Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

Syntax



Parameters

device_class_name (Required)

Specifies the name of the device class to be updated.

MOUNTLimit

Specifies the maximum number of files that can be simultaneously open for input and output. This parameter is optional. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

MAXCapacity

Specifies the maximum size of any data storage files categorized by this device class. This parameter is optional.

Specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes). The minimum size is 1 MB (**MAXCAPACITY=1M**). If you are defining a FILE device class for database-backup volumes, specify a value for **MAXCAPACITY** that is appropriate for the size of the database and that minimizes the number of database volumes.

UPDATE DEVCLASS — FILE

For example, `MAXCAPACITY=5G` specifies that the maximum capacity for a volume in this device class is 5 gigabytes. The value specified should be less than or equal to the maximum supported size of a file on the target file system.

Do not define a `MAXCAPACITY` value greater than 640M when this file is for `REMOVABLEFILE` CD support. A value less than a CD's usable space (650 MB) allows for a one-to-one match between files from the `FILE` device class and copies that are on CD.

DIRectory

Specifies the directory location or locations of the files used in this device class. Enclose the entire list of directories within quotation marks, using commas to separate individual directory names. Special characters (for example, blank spaces) are permitted within directory names. For example, the directory list `"abc def,xyz"` contains two directories: `abc def` and `xyz`. This parameter is optional.

By specifying a directory name or names, you identify the locations where the server places the files that represent storage volumes for this device class.

While processing the command, the server expands the specified directory name or names into their fully qualified forms, starting from the root directory.

Important: If you are using storage agents for shared access to `FILE` volumes, you need to use the `DEFINE PATH` command to define a path for each storage agent. The path definition includes the directory names used by the storage agent to access each directory.

Later, if the server needs to allocate a scratch volume, it creates a new file in one of these directories. (The server can choose any of the directories in which to create new scratch volumes.) For scratch volumes used to store client data, the file created by the server has a file name extension of `.bfs`. For scratch volumes used to store export data, a file name extension of `.exp` is used.

For example, if you define a device class with a directory of `tsmstor` and the server needs a scratch volume in this device class to store export data, the file that the server creates might be named `/tsmstor/00566497.exp`.

Tip: If you specify multiple directories for a device class, ensure that the directories are associated with separate file systems. Space trigger functions and storage pool space calculations take into account the space remaining in each directory. If you specify multiple directories for a device class and the directories reside in the same file system, the server will calculate space by adding values representing the space remaining in each directory. These space calculations will be inaccurate. Rather than choosing a storage pool with sufficient space for an operation, the server might choose the wrong storage pool and run out of space prematurely. For space triggers, an inaccurate calculation might result in a failure to expand the space available in a storage pool. Failure to expand space in a storage pool is one of the conditions that can cause a trigger to become disabled. If a trigger is disabled because the space in a storage pool could not be expanded, you can re-enable the trigger by issuing the following command: `update spacetrigger stg`. No further changes are required to the space trigger.

Restriction: To modify a list of directories, you must replace the entire list.

SHAREd

Specifies that this `FILE` device class will be shared between the server and one

or more storage agents. To prepare for sharing, a library will be automatically defined along with a number of drives corresponding to the MOUNTLIMIT associated with the device class. If the library and drives already exist and the MOUNTLIMIT is changed, drives will either be created to reach a new higher MOUNTLIMIT value or deleted to reach a new lower value.

Storage agents using FILE volumes

You must ensure that storage agents can access newly created FILE volumes. To access FILE volumes, storage agents replace names from the directory list in the device-class definition with the names in the directory list for the associated path definition. The following illustrates the importance of matching device classes and paths to ensure that storage agents can access newly created FILE volumes.

Suppose you want to use these three directories for a FILE library:

```
/opt/tivoli1
/opt/tivoli2
/opt/tivoli3
```

1. You use the following command to set up a FILE library named CLASSA with one drive named CLASSA1 on SERVER1:

```
define devclass classa devtype=file
directory="/opt/tivoli1,/opt/tivoli2,/opt/tivoli3"
shared=yes mountlimit=1
```

2. You want the storage agent STA1 to be able to use the FILE library, so you define the following path for storage agent STA1:

- ```
define path server1 sta1 srctype=server desttype=drive device=file
directory="/opt/ibm1,/opt/ibm2,/opt/ibm3" library=classa
```

In this scenario, the storage agent, STA1, will replace the directory name /opt/tivoli1 with the directory name /opt/ibm1/ to access FILE volumes that are in the /opt/tivoli1 directory on the server.

The following results occur:

- If file volume /opt/tivoli1/file1.dsm is created on SERVER1, and if the following command is issued,

```
update devclass classa directory="/opt/otherdir,/opt/tivoli2,
/opt/tivoli3"
```

SERVER1 will still be able to access file volume /opt/tivoli1/file1.dsm, but the storage agent STA1 will not be able to access it because a matching directory name in the PATH directory list no longer exists. If a directory name is not available in the directory list associated with the device class, the storage agent can lose access to a FILE volume in that directory. Although the volume will still be accessible from the Tivoli Storage Manager server for reading, failure of the storage agent to access the FILE volume can cause operations to be retried on a LAN-only path or to fail.

### Example: Update a FILE device class for sharing

Prepare a FILE device class (named PLAINFILES) for sharing with a Tivoli Storage Manager storage agent.

```
update devclass plainfiles shared=yes
```

## UPDATE DEVCLASS — FILE

### **Example: Update the capacity of a FILE device class**

Update a file device class named STORFILES to a maximum capacity of 25 MB.

```
update devclass storfiles maxcap=25m
```

### **Example: Add a directory to a FILE device class**

Update the FILE device class, CLASSA, by adding a directory, /usr/otherdir, to the directory list. The directories /usr/tivoli2 and /usr/tivoli3 were specified when the device class was first defined.

```
update devclass classa
directory="/usr/tivoli2,/usr/tivoli3,/usr/otherdir"
```


## UPDATE DEVCLASS (Update an LTO device class)

Use the LTO device class when you are using LTO tape devices.

### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

### Syntax



### Notes:

- 1 You cannot specify DRIVEENCRYPTION=ON if your drives are using WORM (write once, read many) media.
- 2 Drive encryption is supported only for Ultrium 4, Ultrium 5, and Ultrium 6 drives and media.

### Parameters

#### *device\_class\_name* (Required)

Specifies the name of the device class to be updated. The maximum length of the device class name is 30 characters.

### **LIBRARY**

Specifies the name of the defined library object that contains the LTO tape drives used by this device class. For information about defining a library object, see the **DEFINE LIBRARY** command.

### **LBProtect**

Specifies whether logical block protection is used to ensure the integrity of data stored on tape. When **LBPROTECT** is set to **READWRITE** or to **WRITEONLY**, the server uses this feature of the tape drive for logical block protection and generates cyclic redundancy check (CRC) protection information for each data block written on tape. The server also validates the CRC protection information when data is read from the tape.

The following values are possible:

#### **READWrite**

Specifies that logical block protection is enabled in the server and the tape drive for both read and write operations. Data is stored with CRC information in each block. This mode affects performance because additional processor usage is required for Tivoli Storage Manager and the tape drive to calculate and compare CRC values. The **READWRITE** value does not affect backup sets and data that is generated by the **BACKUP DB** command.

When the **LBPROTECT** parameter is set to **READWRITE**, you do not have to specify the **CRCDATA** parameter in a storage pool definition because logical block protection provides better protection against data corruption.

#### **WRITEOnly**

Specifies that logical block protection is enabled in the server and the tape drive for write operations only. Data is stored containing CRC information in each block. For read operations, the server and the tape drive do not validate the CRC. This mode affects performance because additional processor usage is required for Tivoli Storage Manager to generate the CRC and for the tape drive to calculate and compare CRC values for write operations. The **WRITEONLY** value does not affect backup sets and data that are generated by the **BACKUP DB** command.

**No** Specifies that logical block protection is not enabled in the server and the tape drive for read and write operations. However, the server enables logical block protection on write operations for a filling volume that already has data with logical block protection.

**Restriction:** Logical block protection is supported only on IBM LTO5 and supported LTO6 drives.

### **FORMAT**

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional.

When migrating all drives from Ultrium to Ultrium 2 devices:

- Delete all existing Ultrium drive definitions and the paths associated with them.
- Define the new Ultrium 2 drives and paths.

If you are considering mixing different generations of LTO media and drives, be aware of the following restrictions. For more information about mixing drives and media, refer to the *Administrator's Guide*.

Table 408. Read - write capabilities for different generations of LTO drives

| Drives                    | Generation 1 media | Generation 2 media | Generation 3 media | Generation 4 media | Generation 5 media | Generation 6 media |
|---------------------------|--------------------|--------------------|--------------------|--------------------|--------------------|--------------------|
| Generation 1              | Read and write     | n/a                | n/a                | n/a                | n/a                | n/a                |
| Generation 2              | Read and write     | Read and write     | n/a                | n/a                | n/a                | n/a                |
| Generation 3 <sup>1</sup> | Read only          | Read and write     | Read and write     | n/a                | n/a                | n/a                |
| Generation 4 <sup>2</sup> | n/a                | Read only          | Read and write     | Read and write     | n/a                | n/a                |
| Generation 5 <sup>3</sup> | n/a                | n/a                | Read only          | Read and write     | Read and write     | n/a                |
| Generation 6 <sup>4</sup> | n/a                | n/a                | n/a                | Read only          | Read and write     | Read and write     |

<sup>1</sup> In a library with a Generation 3 drive, all Generation 1 scratch volumes need to be checked out, and all Generation 1 storage pool volumes need to be updated to read-only.

<sup>2</sup> In a library with a Generation 4 drive, all Generation 2 scratch volumes need to be checked out, and all Generation 2 storage pool volumes need to be updated to read-only.

<sup>3</sup> In a library with a Generation 5 drive, all Generation 3 scratch volumes need to be checked out, and all Generation 3 storage pool volumes need to be updated to read-only.

<sup>4</sup> In a library with a Generation 6 drive, all Generation 4 scratch volumes need to be checked out, and all Generation 4 storage pool volumes need to be updated to read-only.

The following table lists the recording formats and estimated capacities for LTO devices:

Table 409. Recording format and default estimated capacity for LTO

| Format    | Estimated capacity | Description                                                                                                                                                                                                                                                                                                                                    |
|-----------|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DRIVE     | –                  | The server selects the highest format that is supported by the drive on which a volume is mounted.<br><br><b>Attention:</b> Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives. |
| ULTRIUM   | 100 GB             | Uncompressed format, using Ultrium cartridges                                                                                                                                                                                                                                                                                                  |
| ULTRIUMC  | See note<br>200 GB | Compressed format, using Ultrium cartridges                                                                                                                                                                                                                                                                                                    |
| ULTRIUM2  | 200 GB             | Uncompressed (standard) format, using Ultrium 2 cartridges                                                                                                                                                                                                                                                                                     |
| ULTRIUM2C | See note<br>400 GB | Compressed format, using Ultrium 2 cartridges                                                                                                                                                                                                                                                                                                  |
| ULTRIUM3  | 400 GB             | Uncompressed (standard) format, using Ultrium 3 cartridges                                                                                                                                                                                                                                                                                     |

Table 409. Recording format and default estimated capacity for LTO (continued)

| Format    | Estimated capacity  | Description                                                |
|-----------|---------------------|------------------------------------------------------------|
| ULTRIUM3C | See note<br>800 GB  | Compressed format, using Ultrium 3 cartridges              |
| ULTRIUM4  | 800 GB              | Uncompressed (standard) format, using Ultrium 4 cartridges |
| ULTRIUM4C | See note<br>1.6 TB  | Compressed format, using Ultrium 4 cartridges              |
| ULTRIUM5  | 1.5 TB              | Uncompressed (standard) format, using Ultrium 5 cartridges |
| ULTRIUM5C | See note<br>3.0 TB  | Compressed format, using Ultrium 5 cartridges              |
| ULTRIUM6  | 2.5 TB              | Uncompressed (standard) format, using Ultrium 6 cartridges |
| ULTRIUM6C | See note<br>6.25 TB | Compressed format, using Ultrium 6 cartridges              |

**Note:** If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity may be greater than the listed value.

#### ESTCAPacity

Specifies the estimated capacity for the sequential access volumes categorized by this device class. This parameter is optional. You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), G (gigabytes), or T (terabytes). For example, ESTCAPACITY=100G specifies that the estimated capacity for a volume in this device class is 100 GB.

The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K). To force the IBM Tivoli Storage Manager server to determine the estimated capacity for the volumes assigned to this device class, specify ESTCAPACITY="".

For more information on estimated capacities, see Table 409 on page 1229.

#### PREFIX

Specifies the high level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, use a volume prefix that conforms to your naming conventions.

#### MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types (that is, for a library managed by an external media management

program), you can enhance device sharing between applications by setting this parameter to a low value (for example, two minutes).

**MOUNTWait**

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number from 0 to 9999.

**MOUNTLimit**

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

**DRIVES**

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

**Note:** For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

*number*

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

**0 (zero)**

Specifies that no new transactions can gain access to the storage pool.

**DRIVEEncryption**

Specifies whether drive encryption will be permitted. This parameter is optional. Drive encryption is supported only for Ultrium 4, Ultrium 5, and Ultrium 6 drives and media.

**Restriction:** If encryption is enabled for a device class, and the device class is associated with a storage pool, the storage pool should not share a scratch pool with other device classes that cannot be encrypted. If a tape is encrypted, and you plan to use it on a drive that cannot be encrypted, you must manually relabel the tape before it can be used on that drive.

**ON** Specifies that Tivoli Storage Manager is the key manager for drive encryption and will permit drive encryption for empty storage pool volumes only if the application method is enabled. (Other types of volumes will not be encrypted. For example, backup sets, export volumes, and database backup volumes are not encrypted.) If you specify ON and you enable another method of encryption, drive encryption will not be permitted and backup operations will fail.

**Note:** You cannot specify Tivoli Storage Manager as the key manager for drive encryption of WORM (write once, read many) media. (If you are using WORM media, you cannot specify `DRIVEENCRYPTION=ON`.)

### **ALLOW**

Specifies that Tivoli Storage Manager does not manage the keys for drive encryption. However, drive encryption for empty volumes is permitted if another method of encryption is enabled.

### **EXTERNAL**

Specifies that Tivoli Storage Manager does not manage the keys for drive encryption. Use this setting with an encryption methodology that is provided by another vendor and that is used with Application Method Encryption (AME) enabled on the drive. When you specify `EXTERNAL` and Tivoli Storage Manager detects that AME encryption is enabled, Tivoli Storage Manager does not turn encryption off. By contrast, when you specify `ALLOW` and Tivoli Storage Manager detects that AME encryption is enabled, Tivoli Storage Manager turns encryption off.

### **OFF**

Specifies that drive encryption will not be permitted. If you enable another method of encryption, backups will fail. If you enable the application method, Tivoli Storage Manager will disable encryption and backups will be attempted.

### **Example: Update the mount limit for an LTO device class**

Update a device class named `LTOTAPE`. Change the mount limit to 2.

```
update devclass ltotape mountlimit=2
```

## UPDATE DEVCLASS (Update a NAS device class)


Use the NAS device class when you are using NDMP (Network Data Management Protocol) operations to back up network-attached storage (NAS) file servers. The device class is for drives supported by the NAS file server for backups.

The NAS device class does not support EXTERNAL libraries.

### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

### Syntax



### Parameters

#### *device\_class\_name* (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

#### LIBRARY

Specifies the name of the defined library object that contains the SCSI tape drives used by this device class. For information about defining a library object, see the **DEFINE LIBRARY** command.

#### MOUNTRetention=0

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. Zero (0) is the only supported value for device classes with DEVType=NAS.

#### MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number from 0 to 9999.

#### MOUNTLimit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a

simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

### **DRIVES**

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true value (including online status).

**Note:** For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

### *number*

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

### **0 (zero)**

Specifies that no new transactions can gain access to the storage pool.

### **ESTCAPacity**

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

This value must be an integer followed by K (kilobytes), M (megabytes), or G (gigabytes). For example, ESTCAPACITY=100G specifies that the estimated capacity for a volume in this device class is 100 GB.

The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K). To force the IBM Tivoli Storage Manager server to determine the estimated capacity for the volumes assigned to this device class, specify ESTCAPACITY="".

### **PREFIX**

Specifies the high level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, use a volume prefix that conforms to your naming conventions.

### **Example: Update the estimated capacity for a NAS device class**

Update a device class named NASTAPE. Change the estimated capacity to 200 GB.

```
update devclass nastape library=naslib estcapacity=200G
```


## UPDATE DEVCLASS (Update OPTICAL and WORM device classes)

Use the OPTICAL device class when you are using optical devices. Use the WORM device class when you are using optical devices that have WORM capability.

### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

### Syntax



### Parameters

#### device\_class\_name (Required)

Specifies the name of the device class to be defined.

#### LIBRARY

Specifies the name of the defined library object that contains the Optical or WORM drives that can be used by this device class. For information about defining a library object see the **DEFINE LIBRARY** command.

#### FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional.

The following table lists the recording formats and estimated capacities for OPTICAL and WORM devices:

## UPDATE DEVCLASS — OPTICAL and WORM types

Table 410. Estimated capacities for OPTICAL and WORM drives

| Format | Estimated capacity | Description                                                                                                                                                                                                                                                                                                                                    |
|--------|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DRIVE  | –                  | The server selects the highest format that is supported by the drive on which a volume is mounted.<br><br><b>Attention:</b> Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives. |
| 650MB  | 650 MB             | Using a 650 MB 5.25-inch optical drive                                                                                                                                                                                                                                                                                                         |
| 1300MB | 1300 MB            | Using a 1300 MB 5.25-inch optical drive                                                                                                                                                                                                                                                                                                        |
| 2600MB | 2600 MB            | Using a 2600 MB 5.25-inch optical drive                                                                                                                                                                                                                                                                                                        |
| 5200MB | 5200 MB            | Using a 5200 MB 5.25-inch optical drive                                                                                                                                                                                                                                                                                                        |
| 9100MB | 9100 MB            | Using a 9100 MB 5.25-inch optical drive                                                                                                                                                                                                                                                                                                        |
| 23GB   | 23 GB              | Using Sony Blue Laser read-write and WORM media                                                                                                                                                                                                                                                                                                |
| 30GB   | 30 GB              | Using Plasmon UDO1 read-write and WORM media                                                                                                                                                                                                                                                                                                   |
| 60GB   | 60 GB              | Using Plasmon or IBM UDO2 read-write and WORM media                                                                                                                                                                                                                                                                                            |

**Restriction:** If you are considering mixing different generations of UDO media and drives, be aware of the following:

- UDO1 drives can read and write UDO1 media only.
- UDO2 drives can read from, but not write to, UDO1 media.

### ESTCAPacity

Specifies the estimated capacity for the sequential access volumes categorized by this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes). For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB.

The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K). To force the IBM Tivoli Storage Manager server to determine the estimated capacity for the volumes assigned to this device class, specify ESTCAPACITY="".

For more information on estimated capacities, see Table 410.

### MOUNTRetention

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types, setting this parameter to a low value (for example, two minutes) enhances device sharing between applications.

### MOUNTWait

Specifies the maximum number of minutes the server will wait for an operator

to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number from 0 to 9999.

### **MOUNTLimit**

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

### **DRIVES**

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true **MOUNTLIMIT** value (including online status).

**Note:** For EXTERNAL library types, do not specify **DRIVES** for the **MOUNTLIMIT** value. Specify the number of drives for the library as the **MOUNTLIMIT** value.

### *number*

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

### **0 (zero)**

Specifies that no new transactions can gain access to the storage pool. Any current transaction will continue and complete but new transactions will be terminated.

## UPDATE DEVCLASS — QIC


### UPDATE DEVCLASS (Update a QIC device class)

Use the QIC device class when you are using QIC tape devices.

#### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

#### Syntax



#### Parameters

##### `device_class_name` (Required)

Specifies the name of the device class to be defined.

##### `LIBRARY`

Specifies the name of the defined library object that contains the tape drives that can be used by this device class. For information about defining a library object, see the `DEFINE LIBRARY` command.

**FORMAT**

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional.

The following tables list the recording formats, estimated capacities and recording format options for QIC devices:

*Table 411. Recording format and default estimated capacity for quarter-inch cartridge (QIC) tape*

| <b>Format</b> | <b>Estimated Capacity</b> | <b>Description</b>                                                                                                                                                                                                                                                                                                                                                               |
|---------------|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DRIVE         | –                         | The server selects the highest format that can be supported by the sequential access drive on which a volume is mounted.<br><br>Avoid specifying the DRIVE value when a mixture of devices is used within the same library.<br><br>For example, if you have drives that support recording formats superior to other drives in a library, do not specify the FORMAT=DRIVE option. |
| QIC120        | 26 MB – 172 MB            | 120 QIC format, depending on media                                                                                                                                                                                                                                                                                                                                               |
| QIC150        | 31 MB – 207 MB            | 150 QIC format, depending on media                                                                                                                                                                                                                                                                                                                                               |
| QIC525        | 65 MB – 427 MB            | 525 QIC format, depending on media                                                                                                                                                                                                                                                                                                                                               |
| QIC1000       | 169 MB – 1.1 GB           | 1000 QIC format, depending on media                                                                                                                                                                                                                                                                                                                                              |
| QIC2GB        | 2 GB                      | Uncompressed 2000 QIC format                                                                                                                                                                                                                                                                                                                                                     |
| QIC2GBC       | See note<br>4 GB          | Compressed 2000 QIC format                                                                                                                                                                                                                                                                                                                                                       |
| QIC4GB        | 4 GB                      | Uncompressed 4000 QIC format                                                                                                                                                                                                                                                                                                                                                     |
| QIC4GBC       | See note<br>8 GB          | Compressed 4000 QIC format                                                                                                                                                                                                                                                                                                                                                       |
| QIC12GB       | 12 GB                     | Uncompressed 12000 QIC format, using 343-meter tape                                                                                                                                                                                                                                                                                                                              |
| QIC12GBC      | See note<br>24 GB         | Compressed 12000 QIC format, using 343-meter tape                                                                                                                                                                                                                                                                                                                                |
| QIC5010       | 13 GB – 16 GB             | Uncompressed 5010 QIC format, depending on media                                                                                                                                                                                                                                                                                                                                 |
| QIC5010C      | See note<br>26 GB – 32 GB | Compressed 5010 QIC format, depending on media                                                                                                                                                                                                                                                                                                                                   |
| QIC20GB       | 20 GB                     | Uncompressed 20000 QIC format                                                                                                                                                                                                                                                                                                                                                    |
| QIC20GBC      | See note<br>40 GB         | Compressed 20000 QIC format                                                                                                                                                                                                                                                                                                                                                      |
| QIC25GB       | 25 GB                     | Uncompressed 25000 QIC format                                                                                                                                                                                                                                                                                                                                                    |
| QIC25GBC      | See note<br>50 GB         | Compressed 25000 QIC format                                                                                                                                                                                                                                                                                                                                                      |
| QIC30GB       | 30 GB                     | Uncompressed 30000 QIC format                                                                                                                                                                                                                                                                                                                                                    |

Table 411. Recording format and default estimated capacity for quarter-inch cartridge (QIC) tape (continued)

| Format   | Estimated Capacity | Description                  |
|----------|--------------------|------------------------------|
| QIC30GBC | See note<br>60 GB  | Compressed 30000 QIC format  |
| QIC50GB  | 50 GB              | Uncompressed 50GB QIC format |
| QIC50GBC | See note<br>100 GB | Compressed 50GB QIC format   |
| QIC70GB  | 70 GB              | Uncompressed 70GB QIC format |
| QIC70GBC | See note<br>140 GB | Compressed 70GB QIC format   |

**Note:** If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity may be greater than the listed value.

Table 412. QIC tape recording format options

| Tape Format | QIC-120    | QIC-150    | QIC-525    | QIC-1000   |
|-------------|------------|------------|------------|------------|
| 3M DC300XLP | -          | -          | -          | -          |
| 3M DC600A   | Read       | -          | -          | -          |
| 3M DC600XTD | Read/Write | Read/Write | -          | -          |
| 3M DC6150   | Read/Write | Read/Write | -          | -          |
| 3M DC6320   | Read/Write | Read/Write | Read/Write | -          |
| 3M DC6525   | Read/Write | Read/Write | Read/Write | -          |
| 3M DC9100   | -          | -          | -          | Read/Write |
| 3M DC9120XL | -          | -          | -          | Read/Write |

**Note:** The server cannot use 3M DC300XLP and 3M DC600A tapes.

**ESTCAPacity**

Specifies the estimated capacity for the sequential access volumes categorized by this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by **K** (kilobytes), **M** (megabytes), or **G** (gigabytes). For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB.

The smallest value allowed is 100 KB (that is, ESTCAPACITY=100K). To force the IBM Tivoli Storage Manager server to determine the estimated capacity for the volumes assigned to this device class, specify ESTCAPACITY="".

For more information on the default estimated capacity for QIC tapes, see Table 411 on page 1239.

**PREFIX**

Specifies the high level qualifier of the file name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, consider using a volume prefix that conforms to your naming conventions.

**MOUNTRetention**

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types (that is, a library managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

**MOUNTWait**

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number from 0 to 9999.

**MOUNTLimit**

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

**DRIVES**

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true MOUNTLIMIT value (including online status).

**Note:** For EXTERNAL library types, do not use the DRIVES default for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

*number*

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

**0 (zero)**

Specifies that no new transactions can gain access to the storage pool. Any current transaction will continue and complete but new transactions will be terminated.


### UPDATE DEVCLASS (Update a REMOVABLEFILE device class)

Use the REMOVABLEFILE device class for removable media devices that are attached as local, removable file systems.

#### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

#### Syntax



#### Parameters

##### *device\_class\_name* (Required)

Specifies the name of the device class to be updated.

##### LIBRARY

Specifies the name of the defined library object that contains the removable media drives used by this device class. This parameter is optional.

##### MAXCAPACITY

Specifies the maximum size of any volumes defined to a storage pool categorized by this device class. This parameter is optional.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes).

For example, MAXCAPACITY=5M specifies that the maximum capacity for a volume in this device class is 5 MB. The smallest value allowed is 100 KB (that is, MAXCAPACITY=100K).

##### MOUNTRETENTION

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

##### MOUNTWAIT

Specifies the maximum number of minutes the server will wait for an operator to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. You can specify a number from 0 to 9999.

**MOUNTLimit**

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. You can specify a number 0 - 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

**DRIVES**

Specifies that every time a mount point is allocated, the number of drives defined to the library is used to calculate the true MOUNTLIMIT value (including online status).

**Note:** For EXTERNAL library types, do not specify DRIVES for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

*number*

Specifies the maximum number of drives used concurrently in this device class by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class.

**0 (zero)**

Specifies that no new transactions can gain access to the storage pool. Any current transaction will continue and complete but new transactions will be terminated.


### UPDATE DEVCLASS (Update a SERVER device class)

Use the SERVER device class to use storage volumes or files archived in another Tivoli Storage Manager server.

#### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

#### Syntax



#### Parameters

##### *device\_class\_name* (Required)

Specifies the name of the device class to be updated.

##### SERVERName

Specifies the name of the server. The SERVERNAME parameter must match a defined server.

**Note:** If you change the SERVERNAME of an existing server to a new name, data on the volumes under the old SERVERNAME will no longer be accessible with this device class.

##### MAXCAPacity

Specifies the maximum size that objects can be when created on the target server. This parameter is optional.

Specify this value as an integer followed by **K** (kilobytes), **M** (megabytes), or **G** (gigabytes). The minimum value allowed is 100 KB (MAXCAPACITY=100K).

##### PREFIX

Specifies the beginning portion of the high-level archive file name on the target server. This parameter is optional. The maximum length of this prefix is 8 characters.

If you have already established a media label naming convention that supports your current management system, use a volume prefix that conforms to your naming conventions.

##### RETRYPeriod

Specifies the retry period in minutes. The retry period is the interval during which the server attempts to contact a target server if there is a suspected communications failure. This parameter is optional.

**RETRYInterval**

Specifies the retry interval in seconds. The retry interval is how often retries are done within a given time period. This parameter is optional.

**MOUNTRetention**

Specifies the number of minutes to retain an idle connection with the target server before closing the connection. This parameter is optional. You can specify a number from 0 to 9999.

**MOUNTLimit**

Specifies the maximum number of simultaneous sessions between the source server and the target server. Any attempts to access more sessions than indicated by the mount limit cause the requester to wait. This parameter is optional. You can specify a number from 1 to 4096.

The following are possible values:

*number*

Specifies the maximum number of simultaneous sessions between the source server and the target server.

**1 (zero)**

Specifies the number of simultaneous sessions between the source server and the target server.


### UPDATE DEVCLASS (Update a VOLSAFE device class)

Use the VOLSAFE device type to work with StorageTek VolSafe brand media and drives. This technology uses media that cannot be overwritten. Therefore, do not use these media for short-term backups of client files, the server database, or export tapes.

#### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

#### Syntax



#### Parameters

##### device\_class\_name (Required)

Specifies the name of the device class to be updated. The maximum length of the device class name is 30 characters.

##### LIBRARY

Specifies the name of the defined library object that contains the VolSafe drives that can be used by this device class. If any drives in a library are VolSafe-enabled, all drives in the library must be VolSafe-enabled. For more information about the VolSafe device type, see “DEFINE DEVCLASS (Define a VOLSAFE device class)” on page 208.

##### FORMAT

Specifies the recording format to be used when data is written to sequential access media. This parameter is optional; the default value is DRIVE.

**Attention:** If you specify DRIVE for a device class that has non-compatible sequential access devices, then you must mount volumes on devices that are capable of reading or writing the format established when the volume was first mounted. This can cause delays if the only sequential access device that can access the volume is already in use.

The following table lists the recording formats and estimated capacities for VolSafe devices:

Table 413. Recording formats and default estimated capacities for volsafe tapes

| Format | Estimated Capacity | Description                                                                                                                                                                                                                                                                                                                                       |
|--------|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DRIVE  | –                  | The server selects the highest format that is supported by the drive on which a volume is mounted.<br><br><b>Attention:</b> Avoid specifying DRIVE when a mixture of drives is used within the same library. For example, do not use this option for a library that contains some drives that support recording formats superior to other drives. |
| 9840   | 20 GB              | Uncompressed (standard) format, using a 20 GB cartridge with 270 meters (885 feet) of tape                                                                                                                                                                                                                                                        |
| 9840-C | 80 GB              | LZ-1 Enhanced (4:1) compressed format, using an 80 GB cartridge with 270 meters (885 feet) of tape                                                                                                                                                                                                                                                |

**ESTCAPacity**

Specifies the estimated capacity for the volumes that are assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate due to compression of data.

You must specify this value as an integer followed by K (kilobytes), M (megabytes), or G (gigabytes). For example, ESTCAPACITY=5M specifies that the estimated capacity for a volume in this device class is 5 MB.

The smallest value that is allowed is 100 KB (that is, ESTCAPACITY=100K). To force the IBM Tivoli Storage Manager server to determine the estimated capacity for the volumes that are assigned to this device class, specify ESTCAPACITY="".

For more information about the default estimated capacity for cartridge tapes, see Table 413.

**PREFIX**

Specifies the high-level qualifier of the data set name that the server writes into the sequential access media labels. For each sequential access volume assigned to this device class, the server uses this prefix to create the data set name. This parameter is optional.

If you already established a tape label naming convention that supports your current tape management system, consider using a tape volume prefix that conforms to your naming conventions.

**MOUNTRetention**

Specifies the number of minutes to retain an idle sequential access volume before dismounting it. This parameter is optional. The default value is 60. You can specify a number from 0 to 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online. However, for EXTERNAL library types (that is, a library managed by an external media management system), set this parameter to a low value (for example, two minutes) to enhance device sharing between applications.

**MOUNTWait**

Specifies the maximum number of minutes the server will wait for an operator

to respond to a request to either mount a volume in a drive in a manual library or check in a volume to be mounted in an automated library. This parameter is optional. If the mount request is not satisfied within the specified amount of time, the mount request is canceled. The default value is 60 minutes. You can specify a number from 0 to 9999.

### **MOUNTLimit**

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional. The default is DRIVES. You can specify a number from 0 to 4096.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

The following are possible values:

### **DRIVES**

Specifies that every time a mount point is allocated, the number of drives that are defined to the library is used to calculate the true value (including online status).

**Note:** For EXTERNAL library types, do not use the DRIVES default for the MOUNTLIMIT value. Specify the number of drives for the library as the MOUNTLIMIT value.

### *number*

Specifies the maximum number of drives that are used concurrently in this device class by the server. This value must never be allowed to exceed the number of drives that are defined and online in the library that services this device class.

### **0 (zero)**

Specifies that no new transactions can gain access to the storage pool.

## UPDATE DEVCLASS - z/OS media server (Update device class for z/OS media server)

Use this command to update a device class. A limited set of device class types is available for devices that are accessed through a z/OS media server.

- “UPDATE DEVCLASS (Update a 3590 device class for z/OS media server)” on page 1250
- “UPDATE DEVCLASS (Update a 3592 device class for z/OS media server)” on page 1255
- “UPDATE DEVCLASS (Update a cartridge device class for z/OS media server)” on page 1260
- “UPDATE DEVCLASS (Update an ECARTRIDGE device class for z/OS media server)” on page 1265
- “UPDATE DEVCLASS (Update a FILE device class for z/OS media server)” on page 1272

*Table 414. Commands related to UPDATE DEVCLASS*

| Command                             | Description                                                           |
|-------------------------------------|-----------------------------------------------------------------------|
| BACKUP DEVCONFIG                    | Backs up Tivoli Storage Manager device information to a file.         |
| DEFINE DEVCLASS (z/OS media server) | Defines a device class to use storage managed by a z/OS media server. |
| DEFINE LIBRARY                      | Defines an automated or manual library.                               |
| DELETE DEVCLASS                     | Deletes a device class.                                               |
| QUERY DEVCLASS                      | Displays information about device classes.                            |
| UPDATE LIBRARY                      | Changes the attributes of a library.                                  |


## UPDATE DEVCLASS (Update a 3590 device class for z/OS media server)

Use this command to update a device class that you defined to use a z/OS media server to access 3590 devices. The device class that targets storage for the z/OS media server requires a library definition of type ZOSMEDIA.

### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

### Syntax



### Notes:

- 1 You must specify at least one optional parameter on this command.
- 2 You cannot update the **PREFIX** parameter with this command. You must create a device class with the value that you require for the **PREFIX** parameter.

### Parameters

#### *device\_class\_name* (Required)

Specifies the name of the device class to be updated.

**LIBRARY**

Specifies the name of a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server. This parameter is optional.

For information about defining a library, see the **DEFINE LIBRARY** command.

**FORMAT**

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional.

The following table lists the recording format options for 3590 devices:

*Table 415. Recording formats for 3590*

| Format                                                                                                                                                       | Description                                              |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------|
| 3590B                                                                                                                                                        | Uncompressed (basic) format                              |
| 3590C                                                                                                                                                        | Compressed format                                        |
| 3590E-B                                                                                                                                                      | Uncompressed (basic) format, similar to the 3590B format |
| 3590E-C                                                                                                                                                      | Compressed format, similar to the 3590C format           |
| 3590H-B                                                                                                                                                      | Uncompressed (basic) format, similar to the 3590B format |
| 3590H-C                                                                                                                                                      | Compressed format, similar to the 3590C format           |
| <b>Note:</b> If the format uses the tape drive hardware compression feature the actual capacity can increase, depending on the effectiveness of compression. |                                                          |

**ESTCAPacity**

Specifies the estimated capacity for the sequential access volumes categorized by this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: K (KB), M (MB), or G (GB). For example, specify that the estimated capacity is 9 GB with the parameter **ESTCAPACITY=9G**. The smallest value that is accepted is 100 KB (**ESTCAPACITY=100K**).

**COMPression**

Specifies whether file compression is used for this device class. This parameter is optional.

Possible values are:

**Yes**

Specifies that the data for each tape volume is compressed.

**No**

Specifies that the data for each tape volume is not compressed.

**MOUNTRetention**

Specifies the number of minutes that an idle tape volume is retained before it is dismounted. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. Specify a number, 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

### **MOUNTWait**

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. Specify a number, 1 - 9999.

### **MOUNTLimit**

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

Possible values are:

### **DRIVES**

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

### *number*

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

### **0 (zero)**

Specifies that no new transactions can gain access to the storage pool.

### **EXPIRATION**

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

Specify the expiration date using the format, *yyyymmdd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as 2014007 (the seventh day of year 2014).

If you specify the **EXPIRATION** parameter, you cannot specify the **RETENTION** parameter.

### **RETENTION**

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

If you specify the **RETENTION** parameter, you cannot specify the **EXPIRATION** parameter.

**Tip:** You can specify a value of zero for this parameter. However, do so only if you also want to specify a value for the **EXPIRATION** parameter. You cannot specify a value for the **EXPIRATION** parameter if you specify a non-zero value for the **RETENTION** parameter.

### **PROTECTION**

Specifies whether the RACF program, if installed, protects volumes that are assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. Possible values are:

**No** Specifies that the RACF program does not protect volumes that are assigned to this device class.

#### **Yes**

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

**Tip:** If sensitive data is stored on volumes that are assigned to this device class, use **PROTECTION=YES** and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

#### **Automatic**

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

**Important:** If you specify **PROTECTION=AUTOMATIC**, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify **PROTECTION=AUTOMATIC**, the z/OS media server issues **RACROUTE** commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

Do not change the setting to **PROTECTION=AUTOMATIC** for a device class that has been set to **PROTECTION=NO**. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for **PROTECTION** is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to **AUTOMATIC**, the server attempts to delete profiles when volumes are deleted.

## UPDATE DEVCLASS — 3590, for z/OS media server

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

### **UNIT**

Specifies an esoteric unit name to specify a group of tape devices that support 3590 tape. This parameter is optional. The unit name can be up to 8 characters.



## UPDATE DEVCLASS — 3592, for z/OS media server

### Parameters

#### *device\_class\_name* (Required)

Specifies the name of the device class to be updated. The maximum length of the device class name is 30 characters.

#### LIBRARY

Specifies the name of a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server. This parameter is optional.

For information about defining a library, see the **DEFINE LIBRARY** command.

#### FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional. The default value is **DRIVE**.

See the following table for the recording formats.

Table 416. Recording formats for 3592

| Format       | Description                                                                                                                                                                                                                                                                                                                                           |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3592         | Uncompressed (basic) format                                                                                                                                                                                                                                                                                                                           |
| 3592C        | Compressed format                                                                                                                                                                                                                                                                                                                                     |
| 3592-2       | Uncompressed (basic) format, similar to the 3592 format                                                                                                                                                                                                                                                                                               |
| 3592-C       | Compressed format, similar to the 3592C format                                                                                                                                                                                                                                                                                                        |
| 3592-3       | Uncompressed (basic) format, similar to the 3592 format                                                                                                                                                                                                                                                                                               |
| 3592-3C      | Compressed format, similar to the 3592C format                                                                                                                                                                                                                                                                                                        |
| 3592-4       | Uncompressed (basic) format, similar to the 3592 format                                                                                                                                                                                                                                                                                               |
| 3592-4C      | Compressed format, similar to the 3592C format                                                                                                                                                                                                                                                                                                        |
| <b>DRIVE</b> | The server selects the highest format that is supported by the drive on which a volume is mounted.<br><br><b>Attention:</b> Avoid specifying <b>DRIVE</b> when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives. |

**Note:** If this format uses the tape drive hardware compression feature, depending on the effectiveness of compression, the actual capacity might be different from the listed value.

If the drives are in a library that includes drives of different tape technology, do not use the **DRIVE** value. Use the specific format that the drives use. For optimal results, do not mix generations of drives in the same library. If a library contains mixed generations, media problems can result. For example, generation 1 and generation 2 drives cannot read generation 3 media. If possible, upgrade all drives to 3592 generation 3. If you cannot upgrade all drives to 3592 generation 3, you must use a special configuration.

#### ESTCAPacity

Specifies the estimated capacity for the volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value

to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: K (KB), M (MB), or G (GB). For example, specify that the estimated capacity is 9 GB with the parameter **ESTCAPACITY=9G**. The smallest value that is accepted is 100 KB (**ESTCAPACITY=100K**).

### **COMP**ression

Specifies whether file compression is used for this device class. This parameter is optional. The default value is **YES**.

Possible values are:

#### **Yes**

Specifies that the data for each tape volume is compressed.

**No** Specifies that the data for each tape volume is not compressed.

### **MOUNT**Retention

Specifies the number of minutes that an idle tape volume is retained before it is dismantled. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. Specify a number, 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

### **MOUNT**Wait

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. Specify a number, 1 - 9999.

### **MOUNT**Limit

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

Possible values are:

### **DRIVES**

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

#### *number*

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

#### **0 (zero)**

Specifies that no new transactions can gain access to the storage pool.

### EXPIRATION

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

Specify the expiration date using the format, *yyyymmdd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as 2014007 (the seventh day of year 2014).

If you specify the **EXPIRATION** parameter, you cannot specify the **RETENTION** parameter.

### RETENTION

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

If you specify the **RETENTION** parameter, you cannot specify the **EXPIRATION** parameter.

**Tip:** You can specify a value of zero for this parameter. However, do so only if you also want to specify a value for the **EXPIRATION** parameter. You cannot specify a value for the **EXPIRATION** parameter if you specify a non-zero value for the **RETENTION** parameter.

### PROTECTION

Specifies whether the RACF program, if installed, protects volumes that are assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. Possible values are:

**No** Specifies that the RACF program does not protect volumes that are assigned to this device class.

#### Yes

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

**Tip:** If sensitive data is stored on volumes that are assigned to this device class, use **PROTECTION=YES** and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

#### Automatic

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

**Important:** If you specify **PROTECTION=AUTOMATIC**, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify **PROTECTION=AUTOMATIC**, the z/OS media server issues **RACROUTE** commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

Do not change the setting to **PROTECTION=AUTOMATIC** for a device class that has been set to **PROTECTION=NO**. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for **PROTECTION** is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to **AUTOMATIC**, the server attempts to delete profiles when volumes are deleted.

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

### **UNIT**

Specifies an esoteric unit name to specify a group of tape devices that support 3592 tape. This parameter is optional. This name can be as many as 8 characters.


## UPDATE DEVCLASS (Update a cartridge device class for z/OS media server)

Use this command to update a device class that you defined to use a z/OS media server to access IBM 3480 and 3490 cartridge tape devices. The device class that targets storage for the z/OS media server requires a library definition of type ZOSMEDIA.

### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

### Syntax



### Notes:

- 1 You must specify at least one optional parameter on this command.
- 2 You cannot update the **PREFIX** parameter with this command. You must create a device class with the value that you require for the **PREFIX** parameter.

### Parameters

#### *device\_class\_name* (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

#### LIBRARY

Specifies the name of a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server. This parameter is optional.

## UPDATE DEVCLASS — CARTRIDGE, for z/OS media server

For information about defining a library, see the **DEFINE LIBRARY** command.

### **ESTCAPacity**

Specifies the estimated capacity for volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: K (KB), M (MB), or G (GB). For example, specify that the estimated capacity is 9 GB with the parameter **ESTCAPACITY=9G**. The smallest value that is accepted is 100 KB (**ESTCAPACITY=100K**).

See the following table for the recording formats.

*Table 417. Recording formats and default estimated capacities for cartridge tapes*

| Format    | Media type       | Estimated capacity |
|-----------|------------------|--------------------|
| 3480      | (Not applicable) | 180 MB             |
| 3480XF    | (Not applicable) | 180 MB             |
| 3490 Base | CST              | 180 MB             |
| 3490B     | CST              | 360 MB             |
| 3490C     | CST              | 360 MB             |
| 3490B     | ECCST            | 720 MB             |
| 3490C     | ECCST            | 720 MB             |

**Note:** 3490 Base devices include models A01, A02, A10, A20, B02, B04, B20, B40, D41, D42.

### **MOUNTRetention**

Specifies the number of minutes that an idle tape volume is retained before it is dismantled. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. Specify a number, 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

### **MOUNTWait**

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. Specify a number, 1 - 9999.

### **MOUNTLimit**

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

## UPDATE DEVCLASS — CARTRIDGE, for z/OS media server

Possible values are:

### **DRIVES**

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

### *number*

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

### **0 (zero)**

Specifies that no new transactions can gain access to the storage pool.

### **COMPression**

Specifies whether file compression is used for this device class. This parameter is optional.

Possible values are:

### **Yes**

Specifies that the data for each tape volume is compressed.

**No** Specifies that the data for each tape volume is not compressed.

### **EXPIration**

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

Specify the expiration date using the format, *yyyymmdd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as 2014007 (the seventh day of year 2014).

If you specify the **EXPIRATION** parameter, you cannot specify the **RETENTION** parameter.

### **RETention**

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

If you specify the **RETENTION** parameter, you cannot specify the **EXPIRATION** parameter.

**Tip:** You can specify a value of zero for this parameter. However, do so only if you also want to specify a value for the **EXPIRATION** parameter. You cannot specify a value for the **EXPIRATION** parameter if you specify a non-zero value for the **RETENTION** parameter.

### **PROtection**

Specifies whether the RACF program, if installed, protects volumes that are

## UPDATE DEVCLASS — CARTRIDGE, for z/OS media server

assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. Possible values are:

**No** Specifies that the RACF program does not protect volumes that are assigned to this device class.

### Yes

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

**Tip:** If sensitive data is stored on volumes that are assigned to this device class, use **PROTECTION=YES** and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

### Automatic

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

**Important:** If you specify **PROTECTION=AUTOMATIC**, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify **PROTECTION=AUTOMATIC**, the z/OS media server issues **RACROUTE** commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

Do not change the setting to **PROTECTION=AUTOMATIC** for a device class that has been set to **PROTECTION=NO**. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for **PROTECTION** is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to **AUTOMATIC**, the server attempts to delete profiles when volumes are deleted.

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

## UPDATE DEVCLASS — CARTRIDGE, for z/OS media server

### UNIT

Specifies an esoteric unit name to specify a group of tape devices that support cartridge system tape. This parameter is optional. The unit name can be up to 8 characters.

### Example: Update compression and protection for a device class

Update a device class named SQUARE1 with RACF protection active for all tape volumes assigned to this device class. All data is compressed for this device class.

```
update devclass square1 compression=yes protection=yes
```

### Example: Update the mount retention period for a device class

Update the amount of time to retain idle mounted volumes for the ROUND1 device class to 75 minutes.

```
update devclass round1 mountretention=75
```


### Example: Update the mount limit and protection for a device class

Update the STOR1 device class to have RACF protection and a maximum mount limit of 1.

```
update devclass stor1 mountlimit=1 protection=yes
```



## UPDATE DEVCLASS — ECARTRIDGE, for z/OS media server



### Notes:

- 1 You must specify at least one optional parameter on this command.
- 2 You cannot update the **PREFIX** parameter with this command. You must create a device class with the value that you require for the **PREFIX** parameter.

### Parameters

#### *device\_class\_name* (Required)

Specifies the name of the device class to be updated.

#### LIBRARY

Specifies the name of a library that has been defined with the **LIBTYPE=ZOSMEDIA** parameter. The library and the tape drives that can be used by this device class are controlled by the z/OS media server. This parameter is optional.

For information about defining a library, see the **DEFINE LIBRARY** command.

#### FORMAT

Specifies the recording format to be used when writing data to sequential access media. This parameter is optional.

See the following table for the recording formats.

Table 418. Recording formats for ECARTRIDGE tapes

| Format       | Estimated Capacity | Description                                                                                                                                                                                                                                                                                                                                                                              |
|--------------|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>DRIVE</b> | –                  | The server selects the highest format that is supported by the drive on which a volume is mounted. <b>DRIVE</b> is the default value.<br><br><b>Attention:</b> Avoid specifying <b>DRIVE</b> when a mixture of drives is used within the same library. For example, do not use this option for a library containing some drives that support recording formats superior to other drives. |
| 18T          | 360 MB             | 18-track uncompressed (standard) (read-only) format                                                                                                                                                                                                                                                                                                                                      |

## UPDATE DEVCLASS — ECARTRIDGE, for z/OS media server

Table 418. Recording formats for ECARTRIDGE tapes (continued)

| Format    | Estimated Capacity | Description                                                                                                                                                                                 |
|-----------|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 18TC      | 1.44 GB            | 18-track extended (read-only), compressed format                                                                                                                                            |
| 36T       | 720 MB             | 36-track extended (read and write) format                                                                                                                                                   |
| 36TC      | 2.88 GB            | 36-track extended (read and write), compressed format                                                                                                                                       |
| SD3A      | 10 GB              | Uncompressed (standard) format, using a 10 GB 'A' cartridge with 91 meters (298 feet) of tape                                                                                               |
| SD3AC     | 40 GB              | Compressed format, using a 10 GB 'A' cartridge with 91 meters (298 feet) of tape                                                                                                            |
| SD3B      | 25 GB              | Uncompressed (standard) format, using a 25 GB 'B' cartridge with 204 meters (668 feet) of tape                                                                                              |
| SD3BC     | 100 GB             | Compressed format, using a 25 GB 'B' cartridge with 204 meters (668 feet) of tape                                                                                                           |
| SD3C      | 50 GB              | Uncompressed (standard) format, using a 50 GB 'C' cartridge with 392 meters (1286 feet) of tape                                                                                             |
| SD3CC     | 200 GB             | Compressed format, using a 50 GB 'C' cartridge with 392 meters (1286 feet) of tape                                                                                                          |
| 9840      | 20 GB              | Uncompressed 9840 format, using a StorageTek 9840 cartridge                                                                                                                                 |
| 9840-C    | 80 GB              | LZ-1 Enhanced (4:1) compressed 9840 format, using a StorageTek 9840 cartridge                                                                                                               |
| 9940      | 60 GB              | Uncompressed 9940 format, using a StorageTek 9940 cartridge                                                                                                                                 |
| 9940-C    | 120 GB             | Compressed 9940 format, using a StorageTek 9940 cartridge                                                                                                                                   |
| 9940B     | 200 GB             | Uncompressed 9940B format, using a StorageTek 9940 cartridge                                                                                                                                |
| 9940B-C   | 400 GB             | Compressed 9940B format, using a StorageTek 9940 cartridge                                                                                                                                  |
| M8100     | 10 GB              | Uncompressed (standard) format, using a 10 GB cartridge                                                                                                                                     |
| M8100C    | 10 GB              | Compressed format, using a 10 GB cartridge. Because there is no mechanism to determine the type of cartridge in an M8100 drive, the server assumes a 10 GB estimated uncompressed capacity. |
| T9840C    | 40 GB              | Uncompressed T9840C format, using a StorageTek 9840 cartridge                                                                                                                               |
| T9840C-C  | 80 GB              | Compressed T9840C format, using a StorageTek 9840 cartridge                                                                                                                                 |
| T9840D    | 75 GB              | Uncompressed T9840D format, using a StorageTek 9840 cartridge                                                                                                                               |
| T9840D-C  | 150 GB             | Compressed T9840D format, using a StorageTek 9840 cartridge                                                                                                                                 |
| T10000A   | 500 GB             | Uncompressed T10000A format, using a StorageTek T10000 cartridge                                                                                                                            |
| T10000A-C | 1 TB               | Compressed T10000A format, using a StorageTek T10000 cartridge                                                                                                                              |
| T10000B   | 1 TB               | Uncompressed T10000B format, using an Oracle StorageTek T10000 cartridge                                                                                                                    |

## UPDATE DEVCLASS — ECARTRIDGE, for z/OS media server

Table 418. Recording formats for ECARTRIDGE tapes (continued)

| Format                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Estimated Capacity | Description                                                                 |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|-----------------------------------------------------------------------------|
| T10000B-C                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 2 TB               | Compressed T10000B format, using an Oracle StorageTek T10000 cartridge      |
| T10000C                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 5 TB               | Uncompressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge |
| T10000C-C                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 10 TB              | Compressed T10000C format, using an Oracle StorageTek T10000 T2 cartridge   |
| <b>Note:</b> <ul style="list-style-type: none"><li>• Some formats use a compression feature of the tape drive hardware. Depending on the effectiveness of compression, the actual capacity might be double or more than the listed value.</li><li>• T10000A drives can read and write the T10000A format only. T10000B drives can read, but cannot write, the T10000A format. T10000C drives can read, but cannot write, the T10000A and T10000B formats.</li></ul> |                    |                                                                             |

### **ESTCAPacity**

Specifies the estimated capacity for the sequential access volumes assigned to this device class. This parameter is optional.

You can specify this parameter if the default estimated capacity for the device class is inaccurate because of compression of data. The value does not determine the amount of data stored on the volume. The server uses the value to estimate the usage before a volume is filled. After a volume is full, the actual amount of data stored on the tape is used for the usage calculation.

Specify the value as an integer with one of the following unit indicators: K (KB), M (MB), or G (GB). For example, specify that the estimated capacity is 9 GB with the parameter **ESTCAPACITY=9G**. The smallest value that is accepted is 100 KB (**ESTCAPACITY=100K**).

### **MOUNTRetention**

Specifies the number of minutes that an idle tape volume is retained before it is dismounted. The time span for mount retention begins after the idle timeout period has expired. This parameter is optional. Specify a number, 0 - 9999.

This parameter can improve response time for sequential access media mounts by leaving previously mounted volumes online.

### **MOUNTWait**

Specifies the maximum number of minutes that the z/OS media server waits for a volume mount. If the mount request is not satisfied within the specified time, the mount request fails. If a device is successfully allocated and the device-open request does not complete within the specified time, the device-open request ends and the mount request fails.

This parameter is optional. Specify a number, 1 - 9999.

### **MOUNTLimit**

Specifies the maximum number of sequential access volumes that can be simultaneously mounted for the device class. This parameter is optional.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.

## UPDATE DEVCLASS — ECARTRIDGE, for z/OS media server

Possible values are:

### DRIVES

Specifies that every time a mount point is allocated, the number of drives that are defined and online in the library is used to calculate the true value.

### *number*

Specifies the maximum number of drives in this device class that are used concurrently by the server. This value must never exceed the number of drives that are defined and online in the library that services this device class. You can specify a number, 0 - 4096.

### 0 (zero)

Specifies that no new transactions can gain access to the storage pool.

### COMPression

Specifies whether file compression is used for this device class. This parameter is optional.

Possible values are:

### Yes

Specifies that the data for each tape volume is compressed.

**No** Specifies that the data for each tape volume is not compressed.

### EXPIration

Specifies the expiration date that is placed on the tape labels for this device class. This parameter is optional.

Specify the date when the server no longer requires the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

Specify the expiration date using the format, *yyyymmdd* (four digits for the year and three digits for the day). For example, January 7, 2014 is specified as 2014007 (the seventh day of year 2014).

If you specify the **EXPIRATION** parameter, you cannot specify the **RETENTION** parameter.

### RETention

Specifies the number of days to retain the tape. This parameter is optional.

Specify the number of days (1 - 9999) that the server is expected to use the tape. The server does not use this information, but this information is passed to the z/OS media server for use by z/OS or tape management systems. See the *Administrator's Guide*.

If you specify the **RETENTION** parameter, you cannot specify the **EXPIRATION** parameter.

**Tip:** You can specify a value of zero for this parameter. However, do so only if you also want to specify a value for the **EXPIRATION** parameter. You cannot specify a value for the **EXPIRATION** parameter if you specify a non-zero value for the **RETENTION** parameter.

### PROtection

Specifies whether the RACF program, if installed, protects volumes that are

## UPDATE DEVCLASS — ECARTRIDGE, for z/OS media server

assigned to this device class. If protection is provided, RACF profiles are created when volumes are first used. This parameter is optional. Possible values are:

**No** Specifies that the RACF program does not protect volumes that are assigned to this device class.

### **Yes**

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes, but the profiles are not deleted when volumes are deleted from the server. Profiles must be manually deleted.

**Tip:** If sensitive data is stored on volumes that are assigned to this device class, use **PROTECTION=YES** and manually delete RACF profiles only after tape volumes have been erased.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

### **Automatic**

Specifies that the RACF program protects volumes that are assigned to this device class. RACF profiles are created for volumes when the server first uses the volumes. RACF profiles are deleted when volumes are deleted from the server.

The profiles that are created for volumes depend on the system RACF settings. The protection that is provided is the same as using **PROTECT=YES** in JCL. If the RACF program is active and both TAPEVOL and TAPEDSN are inactive, allocation of tapes fails.

**Important:** If you specify **PROTECTION=AUTOMATIC**, when a volume is deleted, its RACF profile is deleted. The volume therefore is no longer protected by the RACF program. The data on these volumes can be accessed by other users.

If you specify **PROTECTION=AUTOMATIC**, the z/OS media server issues **RACROUTE** commands to delete profiles when a volume is deleted from the server. The deletion commands that are issued depend on the current system settings for TAPEVOL and TAPEDSN. If the system settings are changed, the z/OS media server might not delete existing profiles.

Do not change the setting to **PROTECTION=AUTOMATIC** for a device class that has been set to **PROTECTION=NO**. Volumes without profiles might exist and error messages are generated when such volumes are deleted. If a different value for **PROTECTION** is required, define a new device class.

Profile creation and deletion occur based on the protection setting when the volume is first used and when it is deleted. The server does not attempt to create profiles for volumes that it has already used. If protection is set to **AUTOMATIC**, the server attempts to delete profiles when volumes are deleted.

See the documentation for the RACF program for details about the TAPEVOL and TAPEDSN settings and the profiles that are created when these settings are active.

### **UNIT**

Specifies an esoteric unit name to specify a group of tape devices that support

## UPDATE DEVCLASS — ECARTRIDGE, for z/OS media server

**ECARTRIDGE** tapes. Use the unit name that represents the subset of drives in the library that are attached to the z/OS system. This parameter is optional. The unit name can be up to 8 characters.

### UPDATE DEVCLASS (Update a FILE device class for z/OS media server)

Use this command to update a device class that you defined to use a z/OS media server to access files on magnetic disk storage as sequential-access volumes (like tape). The device class that targets storage for the z/OS media server requires a library definition of type ZOSMEDIA.

A volume in this device class is a Virtual Storage Access Method (VSAM) linear data set accessed by the z/OS media server. SCRATCH volumes can be used with a device class and the z/OS media server will dynamically allocate the VSAM LDS. It is not necessary to define volumes for the server to use the device class. If you define volumes, set the high level qualifier (HLQ) so that SMS recognizes the allocation request by the z/OS media server. If you are using defined volumes, the format volume function is not supported for the server when using this device class. The z/OS media server z/OS media server uses a FormatWrite feature of DFSMS Media Manager when filling FILE volumes.

You can define volumes for the FILE device class by using the **DEFINE VOLUME** command. However, the z/OS media server does not allocate space for a defined volume until the volume is opened for its first use.

For information about disk subsystem requirements for **FILE** disk storage, see the *Administrator's Guide*.

### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

### Syntax

```

▶▶—UPDate DEVclass—device_class_name—————▶
 └─MAXCAPacity—size─┘
▶└─PRIMARYalloc—size─┘ └─SECONDARYalloc—size─┘
▶└─PREFIX—file_volume_prefix─┘ └─MOUNTLimit—number─┘

```

### Parameters

#### *device\_class\_name* (Required)

Specifies the name of the device class to be defined. The maximum length of the device class name is 30 characters.

#### MAXCAPacity

Specifies the maximum size of file volumes that are defined to a storage pool in this device class. This parameter is optional.

Specify this value as an integer followed by K (KB), M (MB), or G (GB). The minimum size is 1 MB (**MAXCAPACITY=1M**). The maximum size is 16384 GB (**MAXCAPACITY=16384G**).

#### PRIMARYalloc

Specifies the initial amount of space that is dynamically allocated when a new volume is opened. Enough space must be available to satisfy the primary

allocation amount. Storage Management Subsystem (SMS) policy determines whether multiple physical volumes can be used to satisfy the primary allocation request.

This parameter is optional. Specify this value as an integer followed by K (KB), M (MB), or G (GB). The minimum size is 100 KB (**PRIMARYALLOC=1K**). The maximum size is 16384 GB (**MAXCAPACITY=16384G**). All values are rounded to the next higher multiple of 256 KB.

To avoid wasted space, the dynamic allocation operation uses the smaller of the values specified in the two parameters, **PRIMARYALLOC** and **MAXCAPACITY**.

SMS automatic class selection (ACS) routines can affect whether the **PRIMARYALLOC** and **SECONDARYALLOC** parameter values are used.

#### **SECONDARYALLOC**

Specifies the amount of space by which a file volume is extended when space that is already allocated to the file volume is used up. The data set for a file volume is extended up to the size set by the **MAXCAPACITY** parameter, then the volume is marked full.

Because secondary allocation of a linear data set cannot span a physical volume, consider the size of the physical volume when selecting a secondary allocation size. For example, physical volumes for a 3390 Model 3 are approximately 2.8 GB. To ensure that each extend request occupies nearly an entire physical volume but not more, use a secondary allocation size that is just less than 2.8 GB. A secondary allocation amount of 2600 MB allots enough space for the VSAM volume data set (VVDS), the volume label, and the volume table of contents (VTOC).

This parameter is optional. Specify this value as an integer followed by K (KB), M (MB), or G (GB). The minimum value is 0 KB (**SECONDARYALLOC=0K**). The maximum value is 16384 GB. Except for 0, all values are rounded to the next higher multiple of 256 KB.

If you specify 0 (**SECONDARYALLOC=0**), the file volume cannot be extended beyond the primary allocation amount.

SMS automatic class selection (ACS) routines can affect whether the **PRIMARYALLOC** and **SECONDARYALLOC** parameter values are used.

If you specify a value for the **SECONDARYALLOCATION** parameter that is not 0, or if you allow the value to default to 2600M, the SMS DATACLAS associated with the PREFIX identifier (for example, High Level Qualifier) must have the Extended Addressability (EA) attribute specified. Without the EA attribute, the SMS DATACLAS limits the allocation of the VSAM LDS FILE volume to the primary extent. (See the description of the **PRIMARYALLOCATION** parameter). With the dataset limited to primary allocation size, the dataset cannot be extended by the z/OS media server, and the volume is marked FULL before the maximum capacity is reached.

**Restriction:** Ensure that the values that you specify for the **PRIMARYALLOC** and **SECONDARYALLOC** parameters are within practical limits for the storage device. The server cannot check whether the values exceed practical device limits, and does not check whether the two values together exceed the current **MAXCAPACITY** setting.

## UPDATE DEVCLASS — FILE, for z/OS media server

**Tip:** To fill volumes when you specify a large value for the **MAXCAPACITY** parameter, specify large values for the **PRIMARYALLOC** and **SECONDARYALLOC** parameters. Use larger MVS volume sizes to reduce the chance of extend failure.

### **PREFIX**

Specifies the high-level qualifier of the data set name that is used to allocate scratch volume data sets. For all scratch file volumes created in this device class, the server uses this prefix to create the data set name. This parameter is optional. The maximum length of the prefix, including periods, is 32 characters.

Values specified for this parameter must meet the following conditions:

- The value is to be made up of qualifiers, which can be a maximum of eight characters including periods. For example, the following value is acceptable:

AB.CD2.E

- The qualifiers must be separated by a single period.
- The first letter of each qualifier must be alphabetic or national (@,#,\$), followed by alphabetic, national, hyphen, or numeric characters.

An example of a file volume data set name using the default prefix is `ADSM.B0000021.BFS`.

If you have a data set naming convention, use a prefix that conforms to your naming conventions. For example, the following value is acceptable: `TSM.SERVER2.VSAMFILE`.

If you are running multiple server instances for either Tivoli Storage Manager or Tivoli Storage Manager for z/OS Media you must use a unique value for the **PREFIX** parameter for each device class that you update.

### **MOUNTLimit**

Specifies the maximum number of **FILE** volumes that can be open concurrently for this device class. This parameter is optional. For 3995 devices emulating 3390 devices, the value must not be set higher than the numbers of concurrent input and output streams possible on the media storing the volumes.

The value that you specify in this parameter is important if there is a significant penalty switching from one volume to another. For example, switching can take place when using IBM 3995 devices to emulate 3390 devices. The value that you specify must be no higher than the number of physical drives available on the device.

If you plan to use the simultaneous-write function, ensure that sufficient drives are available for the write operation. If the number of drives needed for a simultaneous-write operation is greater than the value of the **MOUNTLIMIT** parameter for a device class, the transaction fails.

For details about the simultaneous-write function, see the *Administrator's Guide*.


## UPDATE DOMAIN (Update a policy domain)

Use this command to change a policy domain.

### Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the specified policy domain.

### Syntax



### Parameters

#### *domain\_name* (Required)

Specifies the name of the policy domain.

#### DESCRiption

Describes the policy domain using a text string. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove a previously defined description, specify a null string ("").

#### BACKRETention

Specifies the number of days (from the date the backup versions became inactive) to retain backup versions that are no longer on the client file system. This parameter is optional. You can specify an integer from 0 to 9999. The server uses the backup retention value to manage inactive versions of files when any of the following conditions occur:

- A file is rebound to a new management class, but neither the new management class nor the default management class contains a backup copy group.
- The management class to which a file is bound no longer exists, and the default management class does not contain a backup copy group.
- The backup copy group is deleted from the management class to which a file is bound and the default management class does not contain a backup copy group.

#### ARCHRETention

Specifies the number of days (from the date of archive) to retain archive copies. This parameter is optional. You can specify an integer from 0 to 30000. The server uses the archive retention value to manage archive copies of files when either of the following conditions occur:

- The management class to which a file is bound no longer exists, and the default management class does not contain an archive copy group.

## UPDATE DOMAIN

- The archive copy group is deleted from the management class to which a file is bound and the default management class does not contain an archive copy group.

### ACTIVEDESTINATION

Specifies the names of active-data pools that store active versions of backup data for nodes assigned to the domain. This parameter is optional. Spaces between the names of the active-data pools are not permitted. You cannot specify more than ten active-data pools for a domain.

Before the Tivoli Storage Manager server writes data to an active-data pool, it verifies that the node owning the data is assigned to a domain that has the active-data pool listed in the ACTIVEDESTINATION list. If the server verifies that the node meets this criteria, the data is stored in the active-data pool. If the node does not meet the criteria, then the data is not stored in the active-data pool. If the simultaneous-write function is used to write data to an active-data pool, the server performs the verification during backup operations by Tivoli Storage Manager backup-archive clients or by application clients using the Tivoli Storage Manager API. The verification is also performed when active-data is being copied using the **COPY ACTIVE**DATA command.

### Example: Update the backup retention period for a policy domain

For the EMPLOYEE\_RECORDS policy domain, set the backup retention to 45 days and the archive retention to 75 days.

```
update domain employee_records
backretention=45 archretention=75
```

### Related commands

Table 419. Commands related to UPDATE DOMAIN

| Command          | Description                                                                 |
|------------------|-----------------------------------------------------------------------------|
| COPY DOMAIN      | Creates a copy of a policy domain.                                          |
| DEFINE DOMAIN    | Defines a policy domain that clients can be assigned to.                    |
| DEFINE POLICYSET | Defines a policy set within the specified policy domain.                    |
| DELETE DOMAIN    | Deletes a policy domain along with any policy objects in the policy domain. |
| QUERY DOMAIN     | Displays information about policy domains.                                  |

## UPDATE DRIVE (Update a drive)

Use this command to update a drive.


### Privilege class

For detailed and current drive support information, see the Supported Devices website for your operating system:

[http://www.ibm.com/software/sysmgmt/products/support/IBM\\_TSM\\_Supported\\_Devices\\_for\\_Linux.html](http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html)

To issue this command, you must have system privilege or unrestricted storage privilege.

### Syntax



### Notes:

- 1 The `ACSDRVID` parameter is valid only for drives in ACSLS libraries.
- 2 The `CLEANFREQUENCY` parameter is valid only for drives in SCSI libraries.
- 3 The `CLEANFREQUENCY=ASNEEDED` parameter value does not work for all tape drives. See the parameter description for more information.

### Parameters

#### *library\_name* (Required)

Specifies the name of the library to which the drive is assigned.

#### *drive\_name* (Required)

Specifies the name that is assigned to the drive.

#### **SERIAL**

Specifies the serial number for the drives that are being updated. This parameter is valid only for drives in a SCSI or virtual tape library (VTL). This parameter is optional. The possible values are:

#### *serial\_number*

Specifies the serial number for the drive that is being updated.

## UPDATE DRIVE

**Note:** If a path to this drive is already defined, then the number you enter here is compared to the number detected by Tivoli Storage Manager. If the numbers do not match, the command fails.

### **AUTODETECT**

Specifies that the serial number is automatically detected and used by Tivoli Storage Manager if a path is already defined to this drive.

If a path to this drive is not defined, then the serial number is not detected.

### **ONLine**

Specifies whether the drive is available for use. This parameter specifies whether drives can be taken offline and used for another activity, such as maintenance. This parameter is optional.

You can issue this command when the drive is involved in an active process or session, but it is not recommended. If you issue a command to take the drive offline while it is in use, an error message is issued. The mounted volume completes its current process. If this volume was part of a series of volumes for a specific transaction, the drive is not available to complete mounting the series. If no other drives are available, the process fails.

**Attention:** When a drive is in use, do not specify the **ELEMENT** parameter with the **ONLINE** parameter. The drive is not updated, and the command fails.

The drive state is not changed even if the server is halted and restarted. If a drive is offline when the server is restarted, a warning message is issued stating that the drive must be manually brought online. If all of the drives in a library are updated to be offline, processes that need a library mount point fail, rather than queue up for a mount point.

### **YES**

Specifies that the drive is available for use (online).

**No** Specifies that the drive is not available for use (offline).

### **ELEMent**

Specifies the element address of the drive within a SCSI or VTL library. The server uses the element address to connect the physical location of the drive to the SCSI address of the drive. This parameter is valid only for a drive in a SCSI or VTL library when the command is issued from a Tivoli Storage Manager library manager server. The possible values are:

#### *address*

Specifies the element address for the drive that is being updated.

To find the element address for your library configuration, consult the information from the manufacturer.

**Remember:** If a path to this drive is already defined, then the number you enter here is compared to the number previously detected by Tivoli Storage Manager. If the numbers do not match, then this command fails.

### **AUTODETECT**

Specifies that the element number is automatically detected and used by Tivoli Storage Manager if a path is already defined to this drive.

If a path to this drive is not defined, then the element number is not detected.

**Restriction:** If the library in which the drive is located does not support the Read Element Status SCSI command, and ELEMENT=AUTODETECT, the command fails with a Tivoli Storage Manager error message.

#### ACSDRVID

Specifies the ID of the drive that is being accessed in an ACSLS library. The drive ID is a set of numbers that indicates the physical location of a drive within an ACSLS library. This drive ID must be specified as *a,l,p,d*, where *a* is the ACSID, *l* is the LSM (library storage module), *p* is the panel number, and *d* is the drive ID. The server needs the drive ID to connect the physical location of the drive to the drive's SCSI address. See your StorageTek documentation for details.

#### CLEANFREquency

Specifies how often the server activates drive cleaning. This parameter is optional. For the most complete automation of cleaning for an automated library, you must have a cleaner cartridge checked into the volume inventory for the library. This parameter is valid only for drives in SCSI libraries, and not valid for externally managed libraries, such as 3494 libraries or StorageTek libraries that are managed under ACSLS.

**Important:** There are special considerations if you plan to use server-activated drive cleaning with a SCSI library that provides automatic drive cleaning support in its device hardware.

#### NONE

Specifies that the server does not track cleaning for this drive. Use this parameter for libraries that have their own automatic cleaning.

#### ASNEEDED

Specifies that the server loads the drive with a checked-in cleaner cartridge only when a drive reports to the device driver that it needs cleaning.

The CLEANFREQUENCY=ASNEEDED parameter value does not work for all tape drives. Visit the Supported Devices website for your operating system to view detailed drive information. If **ASNEEDED** is not supported, you can use the *gigabytes* value for automatic cleaning.

**Restriction:** Tivoli Storage Manager does not control the drives that are connected to the NAS file server. If a drive is attached only to a NAS file server (no connection to a storage agent or server), do not specify **ASNEEDED** for the cleaning frequency.

#### *gigabytes*

Specifies, in gigabytes, how much data is processed on the drive before the server loads the drive with a cleaner cartridge. The server resets the gigabytes-processed counter each time it loads a cleaner cartridge in the drive.

**Important:** When CLEANFREQUENCY=gigabyte, drive cleaning can occur before the gigabyte setting is reached, if the drive notifies the device driver that a cleaning is necessary.

Consult the information from the drive manufacturer for cleaning recommendations. If the information gives recommendations for cleaning frequency in terms of hours of use, convert to a gigabytes value by doing the following:

1. Use the bytes-per-second rating for the drive to determine a gigabytes-per-hour value.

## UPDATE DRIVE

2. Multiply the gigabytes-per-hour value by the recommended hours of use between cleanings.
3. Use the result as the cleaning frequency value.

**Tip:** For IBM 3590, specify a value for the cleaning frequency to ensure that the drives receive adequate cleaning. Consult the information from the drive manufacturer for cleaning recommendations. Using the cleaning frequency that is recommended by IBM does not over clean the drives.

### Example: Update the element address for a drive

Update DRIVE3, in the library named AUTO, by changing the element address to 119.

```
update drive auto drive3 element=119
```

### Example: Take a drive offline

Update DRIVE3, in the library named MANLIB, to take it offline.

```
update drive manlib drive3 online=no
```

### Related commands

Table 420. Commands related to UPDATE DRIVE

| Command       | Description                                       |
|---------------|---------------------------------------------------|
| CLEAN DRIVE   | Marks a drive for cleaning.                       |
| DEFINE DRIVE  | Assigns a drive to a library.                     |
| DEFINE PATH   | Defines a path from a source to a destination.    |
| DELETE DRIVE  | Deletes a drive from a library.                   |
| QUERY DRIVE   | Displays information about drives.                |
| QUERY LIBRARY | Displays information about one or more libraries. |
| UPDATE PATH   | Changes the attributes associated with a path.    |

## UPDATE FILESPACE (Update file-space node-replication rules)


Use this command to update file-space replication rules. You can also enable or disable replication of data to which a file space rule applies.

Issue this command on the server that acts as a source for replicated data.

### Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node with the file space to be updated belongs.

### Syntax



### Notes:

- 1 You cannot specify a file space identifier (FSID) if you use wildcard characters for the client node name.
- 2 You can specify each rule only once.

## UPDATE FILESPACE

- 3 You must specify either the **REPLRULE** or the **REPLSTATE** parameter on this command.
- 4 The **ACTIVE\_DATA** and **ACTIVE\_DATA\_HIGH\_PRIORITY** rules are valid only if you specify **DATATYPE=BACKUP**.

### Parameters

#### *node\_name* (Required)

Specifies the client node to which the file space belongs. You can use wildcard characters to specify this name. However, file space identifiers can be different among client nodes for the same file space. Therefore, you cannot specify wildcard characters for the client node name and FSID as the value for the **NAMETYPE** parameter.

#### *file\_space\_name* (Required)

Specifies the name of the file space to be updated. You can use wildcard characters or a comma-delineated list to specify names.

For a server that has clients with Unicode-enabled file spaces, you might have to make the server convert the file space name that you enter. For example, you might have to make the server convert a name from the server code page to Unicode. For details, see the **NAMETYPE** parameter. If you specify only a single wildcard character for the name, you can use the **CODETYPE** parameter to limit the operation to Unicode file spaces or to non-Unicode file spaces.

File space names are case-sensitive. To determine the correct capitalization for the file space to be updated, use the **QUERY FILESPACE** command.

#### **NAMETYPE**

Specifies how you want the server to interpret the file space names that you enter. You can use this parameter for Tivoli Storage Manager clients that Unicode-enabled and that have Windows, Macintosh OS X, or NetWare operating systems.

Use this parameter only when you enter a partly-qualified or fully-qualified file space name. The default value is **SERVER**. You can specify one of the following values:

##### **SERVER**

The server uses the server code page to interpret file space names.

##### **UNICODE**

The server converts file space names from the server code page to the UTF-8 code page. The success of the conversion depends on the operating system, on the characters in the name, and the server code page.

Conversion can fail if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion fails, the name can contain question marks, blanks, or ellipses (...).

##### **FSID**

The server interprets file space names as file space identifiers.

#### **CODETYPE**

Specifies the type of file spaces to be included in node replication processing. The default value is **BOTH**, meaning that file spaces are included regardless of code page type. Use this parameter only when you enter a single wildcard character for the file space name. You can specify one of the following values:

**UNICODE**

Specifies only file spaces that are in Unicode.

**NONUNICODE**

Specifies only file spaces that are not in Unicode.

**BOTH**

Specifies all file spaces regardless of code page type.

**DATATYPE (Required)**

Specifies the data type to which a replication rule applies. To specify multiple data types, separate the names with commas and no intervening spaces. You can specify the following values:

**BACKUP**

Specifies the backup data type.

**ARCHIVE**

Specifies the archive data type.

**SPACEMANAGED**

Specifies the space-managed data type.

**REPLRULE**

Specifies the replication rule that applies to a data type. You cannot use wildcards. If you specify multiple data types, the replication rule applies to each data type. For example, if you specify `DATATYPE=BACKUP,ARCHIVE`, the replication rule applies to backup data and to archive data.

**Restriction:** The **REPLRULE** parameter is optional. However, if you do not specify it, you must specify the **REPLSTATE** parameter.

You can specify normal-priority replication or high-priority replication rules. In a replication process that includes both normal and high-priority data, high-priority data is replicated first. Before you specify a rule, consider the order in which you want the data to be replicated.

For example, suppose that a file space contains active backup data and archive data. Replication of the active backup data is a higher priority than the archive data. To prioritize the active backup data, specify `DATATYPE=BACKUP REPLRULE=ACTIVE_DATA_HIGH_PRIORITY`. To assign a normal priority to archive data, issue the **UPDATE FILESPACE** command again, and specify `DATATYPE=ARCHIVE REPLRULE=ALL_DATA`.

You can specify the following rules:

**ALL\_DATA**

Replicates backup, archive, or space-managed data. The data is replicated with a normal priority.

**ACTIVE\_DATA**

Replicates only the active backup data in a file space. The data is replicated with a normal priority.

## UPDATE FILESPACE

### Attention:

- If you specify `ACTIVE_DATA`, you cannot specify `ARCHIVE` or `SPACEMANAGED` as values for the **DATATYPE** parameter in the same command instance.
- If you specify `ACTIVE_DATA`, inactive backup data that belongs to client nodes is not replicated, and inactive backup data on the target replication server is deleted.

### **ALL\_DATA\_HIGH\_PRIORITY**

Replicates backup, archive, or space-managed data. The data is replicated with a high priority.

### **ACTIVE\_DATA\_HIGH\_PRIORITY**

Replicates active backup data in a file space. The data is replicated with a high priority.

### Attention:

- If you specify `ACTIVE_DATA_HIGH_PRIORITY`, you cannot specify `ARCHIVE` or `SPACEMANAGED` as values for the **DATATYPE** parameter in the same command instance.
- If you specify `ACTIVE_DATA_HIGH_PRIORITY`, inactive backup data that belongs to client nodes is not replicated, and inactive backup data on the target replication server is deleted.

### **DEFAULT**

Data is replicated according to the client node rule for the data type.

For example, suppose that you want to replicate the archive data in all the file spaces that belong to a client node. Replication of the archive data is a high priority. One method to accomplish this task is to specify `DATATYPE=ARCHIVE REPLRULE=DEFAULT` for each file space. Ensure that the client replication rule for archive data is set to `ALL_DATA_HIGH_PRIORITY` or to `DEFAULT`. If the client replication rule is `DEFAULT`, the server replication rule for archive data must be set to `ALL_DATA_HIGH_PRIORITY`.

### **NONE**

Data is not replicated. For example, if you do not want to replicate the space-managed data in a file space, specify `DATATYPE=SPACEMANAGED REPLRULE=NONE`.

### **REPLState**

Specifies the replication state for a data type. If you specified multiple data types, the state applies to all the data types. For example, if you specified `DATATYPE=BACKUP,ARCHIVE`, the state applies to backup data and archive data.

The **REPLSTATE** parameter is optional. However, if you do not specify it, you must specify the **REPLRULE** parameter. You can specify one of the following values for the **REPLSTATE** parameter:

#### **ENabled**

Specifies that the data type is ready for replication.

#### **DISabled**

Specifies that replication does not occur until you enable it.

#### **PURGEdata**

Specifies that data is deleted from the target replication server. The type of data deleted is the type of data specified by the **DATATYPE** parameter. For example, if you specify `DATATYPE=BACKUP,ARCHIVE` and

REPLSTATE=PURGEDATA, backup data and archive data are deleted from the file space on the target replication server.

After the data is deleted, the **REPLSTATE** parameter is set to DISABLED, preventing future replication of the data type or types. The replication rule for the data type is set to DEFAULT.

**Remember:** **PURGEDATA** processing does not delete file spaces. Only data is deleted. The file space shows as empty in the output of the **QUERY OCCUPANCY** command.

### Example: Update replication rules for two data types

NODE1 has three file spaces: /a, /b, and /c. The replication rules for all file spaces are set to ALL\_DATA. However, you want to replicate the backup and archive data in file space /a before the data in other file spaces is replicated.

```
update filesystem node1 /a datatype=backup,archive replrule=
 all_data_high_priority
```

### Example: Update replication rules for two data types

NODE2 has two file spaces: /a and /b. You want to temporarily suspend replication of all data in file space /b.

```
update filesystem node2 /b datatype=backup,archive,spacemanaged
 replstate=disabled
```

### Related commands

Table 421. Commands related to UPDATE FILESPACE

| Command              | Description                                                                             |
|----------------------|-----------------------------------------------------------------------------------------|
| QUERY FILESPACE      | Displays information about data in file spaces that belong to a client.                 |
| QUERY NODE           | Displays partial or complete information about one or more clients.                     |
| QUERY REPLICATION    | Displays information about node replication processes.                                  |
| QUERY STATUS         | Displays the settings of server parameters, such as those selected by the SET commands. |
| REPLICATE NODE       | Replicates data in file spaces that belong to a client node.                            |
| SET REPLRETENTION    | Specifies the retention period for replication history records.                         |
| UPDATE NODE          | Changes the attributes associated with a client node.                                   |
| UPDATE REPLRULE      | Enables or disables replication rules.                                                  |
| VALIDATE REPLICATION | Verifies replication for file spaces and data types.                                    |

## UPDATE LIBRARY (Update a library)

Use this command to update a library definition.

To update the device name or the external manager path name of a library, you must use the UPDATE PATH command.

Syntax and parameter descriptions are available for the following library types.

- “UPDATE LIBRARY (Update a 349X library)” on page 1287
- “UPDATE LIBRARY (Update an ACSLS library)” on page 1289
- “UPDATE LIBRARY (Update an EXTERNAL library)” on page 1291
- “UPDATE LIBRARY (Update a FILE library)” on page 1292
- “UPDATE LIBRARY (Update a manual library)” on page 1293
- “UPDATE LIBRARY (Update a SCSI library)” on page 1295
- “UPDATE LIBRARY (Update a shared library)” on page 1298
- “UPDATE LIBRARY (Update a VTL library)” on page 1299

For detailed and current library support information, see the Supported Devices website for your operating system:

[http://www.ibm.com/software/sysmgmt/products/support/IBM\\_TSM\\_Supported\\_Devices\\_for\\_Linux.html](http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html)

### Related commands

Table 422. Commands related to UPDATE LIBRARY


| Command            | Description                                                         |
|--------------------|---------------------------------------------------------------------|
| AUDIT LIBRARY      | Ensures that an automated library is in a consistent state.         |
| CHECKIN LIBVOLUME  | Checks a storage volume into an automated library.                  |
| CHECKOUT LIBVOLUME | Checks a storage volume out of an automated library.                |
| DEFINE DRIVE       | Assigns a drive to a library.                                       |
| DEFINE LIBRARY     | Defines an automated or manual library.                             |
| DEFINE PATH        | Defines a path from a source to a destination.                      |
| DELETE DRIVE       | Deletes a drive from a library.                                     |
| DELETE LIBRARY     | Deletes a library.                                                  |
| DELETE PATH        | Deletes a path from a source to a destination.                      |
| LABEL LIBVOLUME    | Labels volumes in manual or automated libraries.                    |
| QUERY DRIVE        | Displays information about drives.                                  |
| QUERY LIBRARY      | Displays information about one or more libraries.                   |
| QUERY PATH         | Displays information about the path from a source to a destination. |
| UPDATE DRIVE       | Changes the attributes of a drive.                                  |
| UPDATE LIBVOLUME   | Changes the status of a storage volume.                             |
| UPDATE PATH        | Changes the attributes associated with a path.                      |

**UPDATE LIBRARY (Update a 349X library)**

Use this syntax to update a 349X library.

**Privilege class**

To issue this command, you must have system privilege or unrestricted storage privilege.

**Syntax****Parameters****library\_name (Required)**

Specifies the name of the library to be updated.

**SHARED**

Specifies that this library is shared with other servers in a storage area network (SAN). You must issue this command from the server defined as the primary library manager for the shared library. This parameter is required for libraries defined to a library manager and for libraries used for NDMP operations. Specify SHARED=YES to update a library that is not currently shared.

**Important:** If a library has a path from a data mover (such as a NAS file server) but no connection to the Tivoli Storage Manager server, the library cannot be shared with another Tivoli Storage Manager server.

**AUTOLabel**

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the **CHECKIN LIBVOLUME** command.

**No** Specifies that the server does not attempt to label any volumes.

**Yes**

Specifies that the server only labels unlabeled volumes.

**OVERWRITE**

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

**WORMSCRatchcategory**

Specifies the category number to be used for WORM scratch volumes in the library. This parameter is required if you use WORM volumes. You can specify a number from 1 to 65279. This number must be unique. It cannot be shared

## UPDATE LIBRARY - 349X

with other applications or defined libraries, and it must be different from the other category numbers in this library. This parameter is only valid when 3592 WORM volumes are used.

**Restriction:** This parameter can only be updated if the device class **WORM** parameter is set to YES and the **WORMSCRATCHCATEGORY** currently has no defined value.

### RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established.

Support for persistent reservation has the following limitations:

- If you are using the Tivoli Storage Manager device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 at <http://www.ibm.com/support/docview.wss?uid=swg21470319> for details.
- If you are using the IBM device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972> for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

#### Yes

Specifies that drive preemption through persistent reserve is used.

**No** Specifies that drive preemption through persistent reserve is not used.

**Note:** A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

### Example: Add new devices to a shared library

Update a 3494 shared library named 3494LIB2 with new device names.


```
update library 3494lib2 device=/dev/lmcp1,/dev/lmcp2,/dev/lmcp3
```

**UPDATE LIBRARY (Update an ACSLS library)**

Use this syntax to update an ACSLS library.

**Privilege class**

To issue this command, you must have system privilege or unrestricted storage privilege.

**Syntax****Parameters****library\_name (Required)**

Specifies the name of the library to be updated.

**SHARED**

Specifies that this library is shared with other servers in a storage area network (SAN). You must issue this command from the server defined as the primary library manager for the shared library. This parameter is required for libraries defined to a library manager and for libraries used for NDMP operations. Specify SHARED=YES to update a library that is not currently shared.

**Important:** If a library has a path from a data mover (such as a NAS file server) but no connection to the Tivoli Storage Manager server, the library cannot be shared with another Tivoli Storage Manager server.

**RESETDrives**

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established.

Support for persistent reservation has the following limitations:

- If you are using the Tivoli Storage Manager device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 at <http://www.ibm.com/support/docview.wss?uid=swg21470319> for details.
- If you are using the IBM device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972> for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

**Yes**

Specifies that drive preemption through persistent reserve is used.

**No** Specifies that drive preemption through persistent reserve is not used.

## UPDATE LIBRARY - ACSLS

**Note:** A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

### **AUTOLabel**

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the **CHECKIN LIBVOLUME** command.

**No** Specifies that the server does not attempt to label any volumes.

### **Yes**

Specifies that the server only labels unlabeled volumes.

### **OVERWRITE**

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

### **ACSID**

Specifies the number of this StorageTek library assigned by the ACSSA (Automatic Cartridge System System Administrator). This can be a number from 0 to 126. Issue QUERY ACS on your system to get the number for your library ID. This parameter is required.

See your StorageTek documentation for more information.

### **Example: Update an ID number for an ACSLS library**

Update an ACSLS library named ACSLSLIB with a new ID number.

```
update library acslslib acsid=1
```



## UPDATE LIBRARY - FILE

### UPDATE LIBRARY (Update a FILE library)

Use this syntax to update a FILE library

#### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

#### Syntax

```
►►—UPDate LIBRARY—library_name—SHARed—=—Yes—►►
```

#### Parameters

##### *library\_name* (Required)

Specifies the name of the library to be updated.

##### SHARed

Specifies that this library is shared with other servers in a storage area network (SAN). You must issue this command from the server defined as the primary library manager for the shared library. This parameter is required for libraries defined to a library manager and for libraries used for NDMP operations. Specify SHARED=YES to update a library that is not currently shared.

**Important:** If a library has a path from a data mover (such as a NAS file server) but no connection to the Tivoli Storage Manager server, the library cannot be shared with another Tivoli Storage Manager server.

#### Example: Update a FILE library to be shared

Update a file library named FILE2, so that it is shared:

```
update library file2 shared=yes
```


## UPDATE LIBRARY (Update a manual library)

Use this syntax to update a manual library.

### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

### Syntax



### Parameters

#### *library\_name* (Required)

Specifies the name of the library to be updated.

#### RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established.

Support for persistent reservation has the following limitations:

- If you are using the Tivoli Storage Manager device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 at <http://www.ibm.com/support/docview.wss?uid=swg21470319> for details.
- If you are using the IBM device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972> for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

#### Yes

Specifies that drive preemption through persistent reserve is used.

**No** Specifies that drive preemption through persistent preserve is not used.

**Note:** A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

#### AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional.

To use this option, you must check in the tapes with CHECKLABEL=BARCODE on the **CHECKIN LIBVOLUME** command.

**No** Specifies that the server does not attempt to label any volumes.

## UPDATE LIBRARY - MANUAL

### **Yes**

Specifies that the server labels only unlabeled volumes.

### **OVERWRITE**

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.


## UPDATE LIBRARY (Update a SCSI library)

Use this syntax to update a SCSI library.

### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

### Syntax



### Parameters

#### *library\_name* (Required)

Specifies the name of the library to be updated.

#### **LIBType** (Required)

Specifies the library type that you want to update to. Possible values are:

##### VTL

Specifies that the library has a SCSI-controlled media changer device that is represented by a Virtual Tape Library. To mount volumes on drives in this type of library, Tivoli Storage Manager uses the media changer device. This value is effective when specified for libraries with a current library type of SCSI.

**Note:** Selecting the VTL library type assumes that the following conditions are true:

- Your environment does not include mixed-media
- Paths are defined between all drives in the library and all defined servers, including storage agents, that use the library

If both conditions are not met, performance can degrade to the same levels as the SCSI library type especially during times of high stress when most drives are in use concurrently.

##### SCSI

Specifies that the library has a SCSI-controlled media changer device. To mount volumes on drives in this type of library, Tivoli Storage Manager uses the media changer device. This value is effective when specified for libraries with a current library type of VTL.

##### SHARED

Specifies that this library is shared with other servers in a storage area network (SAN). You must issue this command from the server defined as the primary library manager for the shared library. This parameter is required for libraries

## UPDATE LIBRARY - SCSI

defined to a library manager and for libraries used for NDMP operations. Specify `SHARED=YES` to update a library that is not currently shared.

**Important:** If a library has a path from a data mover (such as a NAS file server) but no connection to the Tivoli Storage Manager server, the library cannot be shared with another Tivoli Storage Manager server.

### RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established.

Support for persistent reservation has the following limitations:

- If you are using the Tivoli Storage Manager device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 at <http://www.ibm.com/support/docview.wss?uid=swg21470319> for details.
- If you are using the IBM device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972> for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

#### Yes

Specifies that drive preemption through persistent reserve is used.

**No** Specifies that drive preemption through persistent reserve is not used.

**Note:** A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

### AUTOLabel

Specifies whether the server attempts to automatically label tape volumes.

To use this option, you must check in the tapes with `CHECKLABEL=BARCODE` on the **CHECKIN LIBVOLUME** command.

**No** Specifies that the server does not attempt to label any volumes.

#### Yes

Specifies that the server only labels unlabeled volumes.

### OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

### SERial

Specifies the serial number for the library being updated. This parameter is optional. The possible values are:

*serial\_number*

Specifies the serial number for the library being updated.

If a path to this library has already been defined, then the number you enter here is compared to the number detected by Tivoli Storage Manager. If the numbers do not match, the command fails. If a path has not been defined, this serial number is verified when a path is defined.

### **AUTODetect**

Specifies that the serial number is automatically detected and used by Tivoli Storage Manager if a path has already been defined to this library.

If a path to this library has not been defined, then the serial number is not detected.

### **RELABELSCRatch**

Specifies whether the server relabels volumes that have been deleted and returned to scratch. When this parameter is set to YES, a LABEL LIBVOLUME operation is started and the existing volume label is overwritten. This parameter is optional and intended for use with a Virtual Tape Library (VTL).

**Note:** If you have both virtual and real volumes in your VTL, both types are relabeled when this parameter is enabled. If the VTL includes real volumes, specifying this option might affect performance.

**No** Specifies that the server does not relabel volumes that are deleted and returned to scratch.

### **Yes**

Specifies that the server relabels volumes that are deleted and returned to scratch.

## UPDATE LIBRARY - SHARED

### UPDATE LIBRARY (Update a shared library)

Use this syntax to update a shared library.

#### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

#### Syntax

```
►►—UPDate LIBRARY—library_name—PRIMarylibmanager—==—server_name—————►◄
```

#### Parameters

##### *library\_name* (Required)

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

##### PRIMarylibmanager

Specifies the name of the Tivoli Storage Manager server that is responsible for controlling access to library resources. You must define this server with the **DEFINE SERVER** command before you can use it as a library manager.

#### Example: Change the library manager server for a library

For a library client server, change the name of the library manager server to CASTOR.


```
update library ltolib primarylibmanager=castor
```

**UPDATE LIBRARY (Update a VTL library)**

Use this syntax to update a library that is defined as VTL.

**Privilege class**

To issue this command, you must have system privilege or unrestricted storage privilege.

**Syntax****Parameters****library\_name (Required)**

Specifies the name of the library to be defined. The maximum length of this name is 30 characters.

**LIBType (Required)**

Specifies the type of library that is being defined. Possible values are:

**SCSI**

Specifies that the library has a SCSI-controlled media changer device. To mount volumes on drives in this type of library, Tivoli Storage Manager uses the media changer device. This value is effective when specified for libraries with a current library type of VTL.

**VTL**

Specifies that the library has a SCSI-controlled media changer device that is represented by a Virtual Tape Library. To mount volumes on drives in this type of library, Tivoli Storage Manager uses the media changer device. This value is effective when specified for libraries with a current library type of SCSI.

**Note:** Select the VTL library type only if the following conditions are true:

- Your environment does not include mixed-media
- Paths are defined between all drives in the library and all defined servers, including storage agents, that use the library

If both conditions are not met, performance can degrade to the same levels as the SCSI library type especially during times of high stress when most drives are in use concurrently.

**SHARED**

Specifies that this library is shared with other servers in a storage area network (SAN). You must issue this command from the server defined as the primary library manager for the shared library. This parameter is required for libraries

## UPDATE LIBRARY - VTL

defined to a library manager and for libraries used for NDMP operations. Specify `SHARED=YES` to update a library that is not currently shared.

**Important:** If a library has a path from a data mover (such as a NAS file server) but no connection to the Tivoli Storage Manager server, the library cannot be shared with another Tivoli Storage Manager server.

### RESETDrives

Specifies whether the server preempts a drive reservation with persistent reserve when the server is restarted or when a library client or storage agent reconnection is established.

Support for persistent reservation has the following limitations:

- If you are using the Tivoli Storage Manager device driver, persistent reserve is only supported on some tape drives. See Technote 1470319 at <http://www.ibm.com/support/docview.wss?uid=swg21470319> for details.
- If you are using the IBM device driver, persistent reserve must be enabled at the device driver level. See the *IBM Tape Device Drivers Installation and User's Guide* at <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972> for information about driver configuration.
- If you are using a virtual tape library that is emulating a supported drive, it might not support persistent reserve.

#### Yes

Specifies that drive preemption through persistent reserve is used.

**No** Specifies that drive preemption through persistent reserve is not used.

**Note:** A library manager will not be able to break a drive reservation if the system that has the drive reservation is not configured to use persistent reservation.

### AUTOLabel

Specifies whether the server attempts to automatically label tape volumes. This parameter is optional.

To use this option, you must check in the tapes with `CHECKLABEL=BARCODE` on the **CHECKIN LIBVOLUME** command.

**No** Specifies that the server does not attempt to label any volumes.

#### Yes

Specifies that the server only labels unlabeled volumes.

### OVERWRITE

Specifies that the server attempts to overwrite an existing label. The server overwrites existing labels *only* if both the existing label and the bar code label are not already defined in any server storage pool or volume history list.

### RELABELScratch

Specifies whether the server relabels volumes that have been deleted and returned to scratch. When this parameter is set to YES, a **LABEL LIBVOLUME** operation is started and the existing volume label is overwritten.

**Note:** If you have both virtual and real volumes in your VTL, both types are relabeled when this parameter is enabled. If the VTL includes real volumes, specifying this option might affect performance.

**Yes**

Specifies that the server relabels volumes that are deleted and returned to scratch.

**No** Specifies that the server does not relabel volumes that are deleted and returned to scratch.

**SERial**

Specifies the serial number for the library being updated. This parameter is optional. The possible values are:

*serial\_number*

Specifies the serial number for the library being updated.

If a path to this library has already been defined, then the number you enter here is compared to the number detected by Tivoli Storage Manager. If the numbers do not match, then the command fails. If a path has not been defined, this serial number is verified when a path is defined.

**AUTODetect**

Specifies that the serial number is automatically detected and used by Tivoli Storage Manager if a path has already been defined to this library.

If a path to this library has not been defined, then the serial number is not detected.

## UPDATE LIBVOLUME (Change the status of a storage volume)

Use this command to change the status of a sequential access storage volume in a library.

### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

### Syntax

```

▶▶—UPDate LIBVolume—library_name—volume_name—
└─STATUS—=—┬─PRIVate
 └─SCRatch
└─OWNER—=—server_name

```

### Parameters

*library\_name* **(Required)**

Specifies the name of the library.

*volume\_name* **(Required)**

Specifies the volume name of the storage volume.

**STATUS**

Specifies a change to the status of a storage volume. Possible values are:

**PRIVate**

Specifies that the server updates the storage volume to a private volume.

**SCRatch**

Specifies that the server updates the storage volume to a scratch volume.

**OWNER**

Specifies which server owns a private volume in a shared library that is shared across a SAN. You can change the owner of a private volume in a shared library (SAN) when you issue the command from the library manager server. If you do not specify this parameter, the library manager server owns the private volume.

**Important:** Do not use OWNER as a value for scratch volumes. However, you can use OWNER when changing a scratch volume to private.

### Example: Update a volume's status

Update the volume named WPDV00 located in the library named AUTO to reflect a status of PRIVATE.

```
update libvolume auto wpdv00 status=private
```

### Related commands

Table 423. Commands related to UPDATE LIBVOLUME

| Command       | Description                                                 |
|---------------|-------------------------------------------------------------|
| AUDIT LIBRARY | Ensures that an automated library is in a consistent state. |

Table 423. Commands related to UPDATE LIBVOLUME (continued)

| Command            | Description                                                              |
|--------------------|--------------------------------------------------------------------------|
| CHECKIN LIBVOLUME  | Checks a storage volume into an automated library.                       |
| CHECKOUT LIBVOLUME | Checks a storage volume out of an automated library.                     |
| DEFINE VOLUME      | Assigns a volume to be used for storage within a specified storage pool. |
| LABEL LIBVOLUME    | Labels volumes in manual or automated libraries.                         |
| QUERY LIBRARY      | Displays information about one or more libraries.                        |
| QUERY LIBVOLUME    | Displays information about a library volume.                             |


## UPDATE MACHINE (Update machine information)

Use this command to update machine information. This information will be included in the plan file to help you to recover the client machines.

### Privilege class

To issue this command, you must have system privilege.

### Syntax



### Parameters

#### *machine\_name* (Required)

Specifies the name of the machine to be updated.

#### DESCRiption

Specifies a description of the machine. This parameter is optional. The text can be up to 255 characters. Enclose the text in quotation marks if it contains any blank characters. To remove existing text, specify a null string ("").

#### BUilding

Specifies the name or number of the building that this machine is in. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters. To remove existing text, specify a null string ("").

#### FLoor

Specifies the name or number of the floor that this machine is on. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters. To remove existing text, specify a null string ("").

#### R0om

Specifies the name or number of the room that this machine is in. This parameter is optional. The text can be up to 16 characters. Enclose the text in quotation marks if it contains any blank characters. To remove existing text, specify a null string ("").

#### PRIority

Specifies the restore priority for the machine as an integer from 1 to 99. The highest priority is 1. This parameter is optional. Tivoli Storage Manager uses this value to prioritize client machine recovery.

#### ADSMServer

Specifies whether the machine contains a Tivoli Storage Manager server. This parameter is optional. Possible values are:

**No** This machine does not contain a Tivoli Storage Manager server.

**Yes**

This machine contains a Tivoli Storage Manager server. Only one machine can be defined as containing a Tivoli Storage Manager server.

**Example: Update information for a specific machine**

Update the DISTRICT5 machine information to reflect that it contains the server.

```
update machine district5 admsserver=yes
```

**Related commands**

*Table 424. Commands related to UPDATE MACHINE*

| Command        | Description                                                                                        |
|----------------|----------------------------------------------------------------------------------------------------|
| DEFINE MACHINE | Defines a machine for DRM.                                                                         |
| DELETE MACHINE | Deletes a machine.                                                                                 |
| INSERT MACHINE | Inserts machine characteristics or recovery instructions into the Tivoli Storage Manager database. |
| QUERY MACHINE  | Displays information about machines.                                                               |

## UPDATE MGMTCLASS (Update a management class)


Use this command to change a management class. To allow clients to use the updated management class, you must activate the policy set that contains the management class.

**Important:** The **UPDATE MGMTCLASS** command fails if a copy storage pool is specified as the destination for files that were migrated by a Tivoli Storage Manager for Space Management client.

### Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

### Syntax



### Parameters

#### *domain\_name* (Required)

Specifies the policy domain to which the management class belongs.

#### *policy\_set\_name* (Required)

Specifies the policy set to which the management class belongs. You cannot update a management class that belongs to the ACTIVE policy set.

#### *class\_name* (Required)

Specifies the management class to update.

#### SPACEMGTECHnique

Specifies whether a file using this management class is eligible for migration. This parameter is optional. This parameter is effective only for Tivoli Storage Manager for Space Management clients, not for backup-archive clients or application clients. Possible values are:

##### **AUTOMATIC**

Specifies that the file is eligible for both automatic migration and selective migration.

##### **SElective**

Specifies that the file is eligible for selective migration only.

##### **NONE**

Specifies that the file is not eligible for migration.

**AUTOMIGNonuse**

Specifies the number of days that must elapse since a file was last used before it is eligible for automatic migration. This parameter is optional. If **SPACEMGTECHNIQUE** is not **AUTOMATIC**, the server ignores this attribute. You can specify an integer from 0 to 9999.

This parameter is effective only for Tivoli Storage Manager for Space Management clients, not for backup-archive clients or application clients.

**MIGREQUIRESBkup**

Specifies whether a backup version of a file must exist before a file can be migrated. This parameter is optional. This parameter is effective only for Tivoli Storage Manager for Space Management clients, not for backup-archive clients or application clients. Possible values are:

**Yes**

Specifies that a backup version must exist.

**No** Specifies that a backup version is optional.

**MIGDESTination**

Specifies the primary storage pool where the server initially stores files migrated by Tivoli Storage Manager for Space Management clients. This parameter is effective only for Tivoli Storage Manager for Space Management clients, not for backup-archive clients or application clients.

The command fails if you specify a copy storage pool as the destination.

**DESCRiption**

Specifies a description of the management class. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove a previously defined description, specify a null string ("").

**Example: Update the policy domain and storage pool of a specific management class**

For the management class **ACTIVEFILES**, in policy set **VACATION** in the **EMPLOYEE\_RECORDS** policy domain, change the storage pool where migrated files are stored.

```
update mgmtclass employee_records vacation
activefiles migdestination=diskpool2
```

**Related commands**

Table 425. Commands related to **UPDATE MGMTCLASS**

| Command             | Description                                                                                |
|---------------------|--------------------------------------------------------------------------------------------|
| ASSIGN DEFMGMTCLASS | Assigns a management class as the default for a specified policy set.                      |
| COPY MGMTCLASS      | Creates a copy of a management class.                                                      |
| DEFINE COPYGROUP    | Defines a copy group for backup or archive processing within a specified management class. |
| DEFINE MGMTCLASS    | Defines a management class.                                                                |
| DEFINE POLICYSET    | Defines a policy set within the specified policy domain.                                   |
| DELETE MGMTCLASS    | Deletes a management class and its copy groups from a policy domain and policy set.        |

## UPDATE MGMTCLASS

Table 425. Commands related to UPDATE MGMTCLASS (continued)

| Command          | Description                                     |
|------------------|-------------------------------------------------|
| QUERY COPYGROUP  | Displays the attributes of a copy group.        |
| QUERY MGMTCLASS  | Displays information about management classes.  |
| QUERY POLICYSET  | Displays information about policy sets.         |
| UPDATE COPYGROUP | Changes one or more attributes of a copy group. |

## UPDATE NODE (Update node attributes)

Use this command to modify the attributes of a registered node.

You must use the **RENAME NODE** command to change the name of a registered node.

Administrators with the same name as the node are sometimes created during a **REGISTER NODE** command. If you do not specify a name, the administrator is given the same name as the node. To keep the same-named node and administrator synchronized, the authentication method and the **SSLREQUIRED** setting for the administrator ID are updated to match the node.

If you update the node authentication method or the node **SSLREQUIRED** setting and there is a same-named administrator, those administrator ID settings change.


You must have system level authority to update the node authentication method or the node **SSLREQUIRED** setting and also update a same-named administrator ID. If the same-named administrator ID has client owner authority over the node that is being updated, then system level authority is not required. You must have either unrestricted policy privilege or restricted policy privilege for the policy domain to which the client node belongs.

**Restriction:** When updating a password with the **UPDATE NODE** command, you cannot use a wildcard character with the `node_name` parameter.


### Privilege class


To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the client node belongs.

### Syntax



## UPDATE NODE



**Notes:**

- 1 You must specify at least one optional parameter on this command.
- 2 Passwords are optional for this command, except when changing the authentication method from LDAP to LOCAL.
- 3 If you specify the **REPLSTATE** parameter and you do not specify the **REPLMODE** parameter, the replication mode of the node is set to SEND.
- 4 If you specify the **REPLMODE** parameter, you must also specify the **REPLSTATE** parameter.
- 5 The **SYNCLDAPDELETE** parameter applies only if a node authenticating to an LDAP directory server reverts to local authentication. If other servers use the same namespace, specify SYNCLDAPDELETE=NO. If the LDAP entry is to be deleted to synchronize with the Tivoli Storage Manager server database, select **YES**.

**Parameters*****node\_name* (Required)**

Specifies the name of the client node to be updated. You can use wildcard characters to specify this name.

***password***

Specifies the new password for the client node. This parameter is optional for most cases. If the node authentication method is changed from LDAP to LOCAL, a password is required. The maximum length of the password is 64 characters. Passwords remain current for a period determined by the password expiration period.

## UPDATE NODE

### **FORCEPwreset**

Specifies whether to force a client to change or reset the password. This parameter is optional. You can specify one of the following values:

**No** Specifies that the password expiration period is set by the **SET PASSEXP** command. The administrator does not force a client to change or reset the password while attempting to logon to the server.

**Remember:** This value is valid only when you specify a password.

### **Yes**

Specifies that the client node or administrator password will expire at the next logon. The client must change or reset the password at that time.

### **PASSExp**

Specifies the number of days the password remains valid. You can set the password expiration period from 0 to 9999 days. A value of 0 means that the password never expires. This parameter is optional. If you do not specify this parameter, the password expiration period is unchanged.

You can change the password expiration period by using the **UPDATE NODE** or **SET PASSEXP** commands. To set a common expiration period for all administrators and client nodes, issue the **SET PASSEXP** command. You can also use the **SET PASSEXP** command to selectively set password expiration periods. If you selectively set a password expiration period by using the **REGISTER NODE** command, the **UPDATE NODE** command, or the **SET PASSEXP** command, the expiration period is excluded from common password expiration periods that were created using the **SET PASSEXP** command.

You can use the **RESET PASSEXP** command to reset the password expiration period to the common expiration period. This parameter does not apply to passwords that authenticate with an LDAP directory server.

### **CLOptset**

Specifies the name of the option set to be used by the client. This parameter is optional. To remove a client option set, specify the **CLOPTSET** parameter with a null string ("").

### **CONtact**

Specifies a text string of information identifying the client node. This parameter is optional. The maximum length of the text string is 255 characters. Enclose the contact information in quotation marks if it contains any blanks. To remove previously defined contact information, specify a null string ("").

### **DOmain**

Specifies the name of the policy domain to which you want to register the client node. This parameter is optional.

**Restriction:** For servers with data retention protection enabled, an archived registered node cannot be reassigned to a different policy domain. See the *Administrator's Guide* for more information.

### **COMPression**

Specifies whether the client node compresses its files before sending them to the server for backup and archive. This parameter is optional.

**Restriction:** This parameter cannot be specified for a NAS node.

You can specify one of the following values:

**Client**

Specifies that the client determines whether files will be compressed.

**Yes**

Specifies that the client node compresses its files before sending them to the server for backup and archive.

**No** Specifies that the client node does not compress its files before sending them to the server for backup and archive.

**ARCHDelete**

Specifies whether the client node can delete its own archived files from the server. This parameter is optional. You can specify one of the following values:

**Yes**

Specifies that the client node can delete its own archive files from the server.

**No** Specifies that the client node cannot delete its own archive files from the server.

**BACKDelete**

Specifies whether the client node can delete its own backup files from the server. This parameter is optional. You can specify one of the following values:

**No** Specifies that the client node cannot delete its own backup files from the server.

**Yes**

Specifies that the client node can delete its own backup files from the server.

**WHEREDomain**

Specifies the name of the policy domain to be used as a filter in combination with the node name to select nodes to update. This parameter is optional.

**WHEREPlatform**

Specifies the name of the client platform to be used as a filter in combination with the node name to select nodes to update. This parameter is optional.

**MAXNUMMP**

Specifies the maximum number of mount points a node can use on the server or storage agent only for operations such as backup, archive, and Tivoli Storage Manager for Space Management migration. The parameter is optional and does not apply to nodes with a type of NAS or SERVER. The default value is 1. You can specify an integer from 0 - 999. A value of 0 specifies that a node cannot acquire any mount point for a client data store operation. The **MAXNUMMP** value is not evaluated or enforced during client data read operations such as restore, retrieve, and Tivoli Storage Manager for Space Management recall. However, mount points in use for data read operations are evaluated against attempted concurrent data store operations for the same client node and might prevent the data store operations from being able to acquire mount points.

For volumes in a storage pool associated with the FILE device type, the server can have multiple sessions to read and one process to write to the same volume concurrently. To increase concurrency and provide efficient access for nodes with data in FILE storage pools, increase the value of the **MAXNUMMP** parameter.

For nodes that store data into primary storage pools with the simultaneous-write function enabled, you have to adjust the value of the **MAXNUMMP** parameter to specify the correct number of mount points for each

## UPDATE NODE

client session. A client session requires one mount point for the primary storage pool and one mount point for each copy storage pool and each active-data pool.

### URL

Specifies the client Uniform Resource Locator (URL) address that the administrator can use to administer the Tivoli Storage Manager client. If you want to remove the value for this field, specify empty single quotation marks or empty double quotation marks with no spaces (" for single quotation marks, or "" for double quotation marks).

### KEEPMP

Specifies whether the client node keeps the mount point for the entire session. The parameter is optional. You can specify one of the following values:

**No** Specifies that the client node release the mount point during the session. If policy definitions cause data to be stored to a disk storage pool after storing data to a sequential access storage pool, any mount points held by the session will be released.

### Yes

Specifies that the client node must retain the mount point during the entire session. If policy definitions cause data to be stored to a disk storage pool after storing data to a sequential access storage pool, any mount points held by the session will not be released.

### AUTOFSRename

Specifies whether the client is prompted for renaming file spaces when the client system upgrades to a client that supports Unicode. The prompting and renaming, if allowed, occur only when the client runs one of the following operations: archive, selective backup, full incremental backup, or partial incremental backup. The renaming changes the names of existing backed-up file spaces that are not in Unicode in server storage. Then the file spaces are backed up in Unicode. You can use this parameter for Unicode-enabled Tivoli Storage Manager clients using Windows, Macintosh OS X, and NetWare operating systems.

**Important:** After the client with support for Unicode is installed, any new file spaces that the client backs up are stored in server storage using the UTF-8 code page. UTF-8 is a byte-oriented encoding form specified by the Unicode Standard.

You can specify one of the following values:

### Yes

The server automatically renames existing file spaces when the client system upgrades to a client that supports Unicode, and the client runs one of the following operations: archive, selective backup, full incremental backup, or partial incremental backup. The renaming occurs whether the client uses the graphical user interface, the command line, or the client scheduler.

For example, the server renames a drive as follows:

|                |             |
|----------------|-------------|
| Original name: | D_DRIVE     |
| New name:      | D_DRIVE_OLD |

The new name indicates that the file space is stored on the server in format that is not Unicode.

**No** The server does not rename file spaces automatically when the client

system upgrades to a client that supports Unicode, and the client runs one of the following operations: archive, selective backup, full incremental backup, or partial incremental backup.

### Client

The option `AUTOFSRENAME` in the client option file determines whether file spaces are renamed.

By default, the client option is set to `PROMPT`. When the client system upgrades to a client that supports Unicode and the client runs a Tivoli Storage Manager operation with the graphical user interface or the command line, the program displays a one-time prompt to the user about whether to rename file spaces.

When the client scheduler runs an operation, the program does not prompt for a choice about renaming, and does not rename file spaces. Backups of existing file spaces are sent as before (not in Unicode).

Refer to the appropriate *Backup-Archive Clients Installation and User's Guide* book for more information about the client option.

### VALIDATEprotocol

Specify whether Tivoli Storage Manager performs a cyclic redundancy check to validate the data sent between the client and the server. The parameter is optional. You can specify one of the following values:

**No** Specifies that Tivoli Storage Manager does not perform data validation on any data sent between the client and server.

#### Dataonly

Specifies that Tivoli Storage Manager performs data validation only on file data that is sent between the client and server. The data does not include the file metadata. This mode affects performance because additional resources are required to calculate and compare cyclic redundancy check (CRC) values between the client and the server.

#### All

Specifies that Tivoli Storage Manager performs data validation on all client file data, client file metadata, and server metadata that is sent between the client and server. This mode affects performance because additional resources are required to calculate and compare CRC values between the client and the server.

### TXNGROUPmax

Specifies the number of files that are transferred as a group between a client and a server between transaction commit points. Client performance might be improved by using a larger value for this option.

Specifying 0 indicates the node uses the server global value that is set in the server options file. To use a value other than the server global value, specify a value of 4 through 65,000 for this parameter. The node value takes precedence over the server value.

**Tip:** Increasing the `TXNGROUPMAX` value increases recovery log utilization. Higher recovery log utilization might increase the risk of running out of log space. Evaluate the performance of each node before changing this parameter. For more information about managing the recovery log, see the *Administrator's Guide*.

## UPDATE NODE

### **DATAWritepath**

Specifies the transfer path used when the client sends data to the server, storage agent, or both, during storage operations such as backup or archive. The parameter is optional.

**Remember:** If a path is unavailable, the node cannot send any data. For example, if you select the LAN-free option but a LAN-free path is not defined, the operation will not work.

You can specify one of the following values:

#### **ANY**

Specifies that data is sent to the server, storage agent, or both, using any available path. A LAN-free path will be used if one is available. If a LAN-free path is unavailable, the data will be moved using the LAN.

#### **LAN**

Specifies that data is sent over the LAN.

#### **LANFree**

Specifies that data is sent over a LAN-free path.

### **DATAReadpath**

Specifies the transfer path used when the server, storage agent, or both read data for a client, during operations such as restore or retrieve. The parameter is optional.

**Remember:** If a path is unavailable, data cannot be read. For example, if you select the LAN-free option but a LAN-free path is not defined, the operation will not work.

You can specify one of the following values:

#### **ANY**

Specifies that the server, storage agent, or both use any available path to read data. A LAN-free path will be used if one is available. If a LAN-free path is unavailable, the data will be read using the LAN.

#### **LAN**

Specifies that data is read using the LAN.

#### **LANFree**

Specifies that data is read using a LAN-free path.

### **SESSIONINITiation**

Controls whether the server or the client initiates sessions. The parameter is optional.

#### **Clientorserver**

Specifies that the client might initiate sessions with the server by communicating on the TCP/IP port defined with the server option TCPPOINT. Server-prompted scheduling might also be used to prompt the client to connect to the server.

#### **SERVEROnly**

Specifies that the server does not accept client requests for sessions. All sessions must be initiated by server-prompted scheduling on the port defined for the client with the **REGISTER** or **UPDATE NODE** commands. You cannot use the client acceptor (dsmcad) to start the scheduler when **SESSIONINITIATION** is set to **SERVERONLY**.

### **HLAddress**

Specifies the client IP address that the server contacts to initiate scheduled events. This optional parameter is used only when **SESSIONINITIATION** is set to **SERVERONLY**, regardless of any addresses previously used by the client to contact the server. If **SESSIONINITIATION SERVERONLY** is not in use, this option has no effect.

The address can be specified either in numeric or host name format. If a numeric address is used, it is saved without verification by a domain name server. If the address is not correct, it can cause failures when the server attempts to contact the client. Host name format addresses are verified with a domain name server. Verified names are saved and resolved with Domain Name Services when the server contacts the client.

#### **LLAddress**

Specifies the client port number on which the client listens for sessions from the server. This optional parameter is used only when **SESSIONINITIATION** is set to **SERVERONLY**, regardless of any addresses previously used by the client to contact the server. If **SESSIONINITIATION SERVERONLY** is not in use, this option has no effect.

The value for this parameter must match the value of client option **TCPCCLIENTPORT**. The default value is 1501.

#### **EMAILAddress**

This parameter is used for additional contact information. The information specified by this parameter is not acted upon by Tivoli Storage Manager.

#### **DEDUPLICATION**

Specifies where data deduplication can occur for this node. You can specify one of the following values:

##### **SERVEROnly**

Specifies that data stored by this node can be deduplicated on the server only.

##### **Clientorserver**

Specifies that data stored by this node can be deduplicated on either the client or the server. For data deduplication to take place on the client, you must also specify a value of **YES** for the **DEDUPLICATION** client option. You can specify this option in the client option file or in the client option set on the Tivoli Storage Manager server.

#### **TARGETLevel**

Specifies the client deployment package that is targeted for this node. You can substitute an applicable release package for V.R.M.F, which stands for Version.Release.Modification.Fix Level. For example: **TARGETLevel=6.2.0.0**.

You must specify each segment with a number that is applicable to a deployment package. You cannot use an asterisk in any field as a substitution for a valid number. To remove an existing value, specify a null string (" "). The parameter is optional.

**Restriction:** The **TARGETLEVEL** parameter does not apply to nodes with a type of **NAS** or **SERVER**.

#### **BACKUPINITiation**

Specifies whether the non-root user ID on the client node can back up files to the server. The parameter is optional. The default value is **ALL**, indicating that non-root user IDs can back up data to the server. You can select one of the following values:

### **A11**

Specifies that non-root user IDs can back up files to the server. ALL is the default if BACKUPINITIATION is not specified.

### **ROOT**

Specifies that only the root user ID can back up files to the server.

**Restriction:** The attribute is ignored by the server if the backup-archive client connects from a platform other than AIX, HP-UX, Linux, Solaris, or Mac OS.

### **BKREPLRuledefault, ARREPLRuledefault, and SPREPLRuledefault**

Specifies the replication rule that applies to a data type if the file space rules for the data type are set to DEFAULT:

#### **BKREPLRuledefault**

Specifies the replication rule for backup data.

#### **ARREPLRuledefault**

Specifies the replication rule for archive data.

#### **SPREPLRuledefault**

Specifies the replication rule for space-managed data.

You can specify normal-priority replication or high-priority replication rules. In a replication process that includes both normal and high-priority data, high-priority data is replicated first. Before you specify a rule, consider the order in which you want the data to be replicated.

For example, suppose that a client node contains active backup data and archive data. Replication of the active backup data is a higher priority than the archive data. To prioritize both types of data, specify `BKREPLRULEDEFAULT=ACTIVE_DATA_HIGH_PRIORITY ARREPLRULEDEFAULT=ALL_DATA`.

You can specify the following rules:

#### **ALL\_DATA**

Replicates active and inactive backup data, archive data, or space-managed data. The data is replicated with a normal priority.

#### **ACTIVE\_DATA**

Replicates only active backup data. The data is replicated with a normal priority. This rule is valid only for **BKREPLRULEDEFAULT**.

#### **Attention:**

- If you specify `ACTIVE_DATA`, inactive backup data that belongs to client nodes is not replicated.
- If the replication rule for backup data in any file spaces that belong to the client node is `DEFAULT`, inactive backup data in those file spaces on the target replication server is deleted.

#### **ALL\_DATA\_HIGH\_PRIORITY**

Replicates active and inactive backup data, archive data, or space-managed data. Data is replicated with a high priority.

#### **ACTIVE\_DATA\_HIGH\_PRIORITY**

Replicates only active backup data. The data is replicated with a high priority. This rule is valid only for **BKREPLRULEDEFAULT**.

**Attention:**

- If you specify `ACTIVE_DATA_HIGH_PRIORITY`, inactive backup data that belongs to client nodes is not replicated.
- If the replication rule for backup data in any file spaces that belong to the client node is `DEFAULT`, inactive backup data in those file spaces on the target replication server is deleted.

**DEFAULT**

Replicates data according to the server replication rule for backup data.

For example, suppose that you want to replicate the archive data in all the file spaces that belongs to a client node. Replication of the archive data is a high priority. One method to accomplish this task is to specify `ARREPLRULEDEFAULT=DEFAULT`. Ensure that the file space rules for archive data are also set to `DEFAULT` and that the server rule for archive data is set to `ALL_DATA_HIGH_PRIORITY`.

**Restriction:** If a node is configured for replication, the file space rules are set to `DEFAULT` after the node stores data on the source replication server.

**NONE**

Data of the specified type is not replicated.

For example, if you do not want to replicate space-managed data that belongs to a client node, specify `SPREPLRULEDEFAULT=NONE`

**REPLState**

Specifies whether data that belongs to the client node is ready to be replicated. This parameter is optional. You can specify one of the following values:

**ENabled**

Specifies that the client node is ready for replication.

**DISabled**

Specifies that replication does not occur until you enable it.

The system response to these settings depends on the following factors:

**Whether the client node definition exists only on the source replication server and you are configuring the client node for replication for the first time**

If you set the replication state to `ENABLED` or `DISABLED`, the replication mode of the node on the source replication server is automatically set to `SEND` after the **UPDATE NODE** command is issued. When replication first occurs, a client node definition on the target server is automatically created. The replication state of the client node on the target server is automatically set to `ENABLED`. The replication mode is set to `RECEIVE`.

**Whether the client node definition exists on the source and the target replication servers, and the node data was previously replicated**

For replication to occur, the replication state of the client node on both the source and the target servers must be set to `ENABLED`. For example, if the replication state of a client node on the source server is `ENABLED` and the replication state on the target server is `DISABLED`, replication does not occur.

**Whether the client node definition exists on the source and the target replication servers, and the node data was previously exported from the source replication server and imported to the target replication server**

In this case, you are configuring the client nodes to synchronize the

## UPDATE NODE

data between the two servers. When replication first occurs, the replication state of the client node on the target server is automatically set to **ENABLED**. Data on the source and target servers is synchronized.

**Restriction:** To synchronize data, you must specify the **REPLMODE** parameter in addition to the **REPLSTATE** parameter.

You can specify the **REPLMODE** parameter only if the client node has never been replicated:

- If the client node definition exists only on the source replication server, the replication mode of the node on the source replication server is automatically set to **SEND** when the **UPDATE NODE** command is issued. The replication mode of the node on the target replication server is automatically set to **RECEIVE**.
- If data that belongs to the node was previously replicated, the replication mode of the node on the source replication server is **SEND**. The replication mode of the node on the target replication server is **RECEIVE**.

### **REPLMode**

Specifies whether to synchronize the data that belongs to this client node. Specify this parameter only if data that belongs to the client node was exported from the source replication server and imported to the target replication server. Synchronization occurs during replication.

To synchronize data, you must issue the **UPDATE NODE** command on both the source and target replication servers and specify the **REPLMODE** and **REPLSTATE** parameters. The value that you specify for the **REPLMODE** parameter depends on whether the server is a source of or a target for replicated data.

You can specify one of the following values:

### **SYNCSEnd**

Specifies that data that belongs to this client node is synchronized with data on a target server during replication. Specify this value only on the server that exported the data. When the synchronization is complete, the replication mode for the client node on the source server is automatically set to **SEND**. The replication mode remains **SEND** unless you remove the node by issuing the **REMOVE REPLNODE** command.

### **SYNCRECeive**

Specifies that data that belongs to this client node is synchronized with data on a source server during replication. Specify this value only on the server that imported the data. When the synchronization is complete, the replication mode for the client node on the target server is automatically set to **RECEIVE**. The replication mode remains **RECEIVE** unless you remove the node by issuing the **REMOVE REPLNODE** command.

### **Restrictions:**

- You can set the **REPLMODE** parameter only if the initial replication state is **NONE**. To synchronize data, you change the replication state to **ENABLED** or **DISABLED** and specify a value for the **REPLMODE** parameter.
- Data can be synchronized only if you specified **DATES=ABSOLUTE** on the **IMPORT NODE** command. If you specified **DATES=RELATIVE** to import data, you must rename the node or delete its data before replication. If you do not take one of these steps, you can lose data.
- If the **REPLMODE** parameter was set incorrectly, you must issue the **REMOVE REPLNODE** command before updating the client node definition. For example,

suppose that you updated the definition of a client node whose data you wanted to replicate. The data that belongs to the node was previously exported to the target replication server. You specified **ENABLED** as the setting of the **REPLSTATE** parameter. However, you did not specify **SYNCSEND** on the source replication server. As a result, the **REPLMODE** parameter was automatically set to **SEND**, and data that belongs to the node could not be synchronized or replicated.

Issuing **REMOVE REPLNODE** sets the replication state and the replication mode to **NONE**. After the **REMOVE REPLNODE** command is completed, reissue the **UPDATE NODE** command with the correct parameters and values.

### **ROLEOVERRIDE**

Specifies whether to override the reported role of the client for processor value unit (PVU) estimation reporting. The default is **USERREPORTED**.

The role reported by the client is either client-device (for example, a workstation) or server-device (for example, file/print server, application server, database). By default, the client reports its role based on the client type and the operating system. All clients initially report their role as server-device, except for Tivoli Storage Manager backup-archive clients running Microsoft Windows workstation distributions (Windows Vista) and Macintosh OS X.

Specify one of the following values:

#### **Client**

Specifies a client-device.

#### **Server**

Specifies a server-device.

#### **Other**

Specifies that this node is not to be used for PVU estimation reporting. This can be useful when there are multiple nodes deployed for a physical system (for example, virtual environments, test nodes, retired nodes, and nodes not in production or clustering).

#### **Usereported**

Use the reported role provided by the client.

### **AUTHentication**

This parameter determines the password authentication method that you use; either **LDAP** or **LOCAL**.

#### **Local**

Specifies that the node uses the local Tivoli Storage Manager server database to store passwords.

#### **Ldap**

Specifies that the node uses an LDAP directory server to authenticate passwords. Passwords are not stored in the Tivoli Storage Manager database.

### **SSLrequired**

Specifies whether the node must use Secure Sockets Layer (SSL) to communicate with the Tivoli Storage Manager server. When you authenticate passwords with an LDAP directory server, you must protect the sessions by using SSL or another network security method.

#### **Yes**

SSL is required.

**No** SSL is not required.

## UPDATE NODE

### DEFault

SSL is required for a node if its password authenticates with an LDAP directory server. SSL is not required for a node that authenticates its password with the Tivoli Storage Manager server (LOCAL).

### SERVERonly

SSL is required between the backup-archive client and the server. SSL is not required between the backup-archive client and the storage agent.

### SYNCLdapdelete

This parameter applies only if you want a node that authenticates with an LDAP directory server, to change to authenticate with the Tivoli Storage Manager. The default is NO.

### Yes

Select YES if you want to synchronize an entry on the Tivoli Storage Manager server, and the entry is not used by another server and not sharing a namespace on the LDAP directory server.

**Remember:** When you select YES you update all other servers sharing the same namespace, which keeps those nodes in sync.

**No** Select NO if an entry is being used by another server and shares a namespace on the LDAP directory server.

### Example: Update node SIMON to authenticate with an LDAP directory server and connect using SSL

```
update node simon authentication=ldap sslrequired=yes
```

Specifying **SSLREQUIRED** does not configure the server for SSL. You must follow the instructions for connecting with SSL in order for the example to work.

### Example: Update a node with software release information for a future deployment

The client deployment feature helps you update a backup-archive client to a newer release. The information that is generated from the **UPDATE NODE** command can help you when you plan a deployment. The information is stored for a future deployment and can be viewed by issuing the **QUERY NODE** command. After a deployment, you can issue the **QUERY NODE** command to see the current level and the target level. For our example, we want to update node LARRY to backup-archive client Version 6.3.0.0.

```
update node LARRY targetlevel=6.3.0.0
```

### Example: Update a node backup to compress data and keep the client from deleting archived files

Update node LARRY so that the data on node LARRY is compressed when it is backed up or archived by Tivoli Storage Manager and so that the client cannot delete archived files.

```
update node larry compression=yes archdelete=no
```

### Example: Update a node's number of files that can be transferred as a group

Update node LARRY and increase the TXNGroupmax value to 1,000.

```
update node larry txngroupmax=1000
```

**Example: Update a node and allow it to deduplicate on the client**

Update a node BOB so that it can deduplicate on the client.

```
update node bob deduplication=clientorserver
```

**Example: Update the role of node BOB to a server-device for PVU estimation reporting**

If you want to accumulate PVU values, only server device roles are recorded. You can update a node from client-device to server-device by issuing the **UPDATE NODE** command. For this example, node BOB is updated to a server-device.

```
update node bob role=server
```

**Example: Update a node definition on a source replication server**

NODE1 is defined to a source replication server. The data that belongs to NODE1 was previously exported to a target replication server. Update the replication rule for backup data that belongs to NODE1 so that active backup data is replicated with a high priority. Enable replication for the node. Set up data synchronization with the target replication server.

```
update node node1 bkreplruledefault=active_data_high_priority
replstate=enabled replmode=syncsend
```

**Related commands**

Table 426. Commands related to UPDATE NODE

| Command                    | Description                                                                                  |
|----------------------------|----------------------------------------------------------------------------------------------|
| QUERY FILESPACE            | Displays information about data in file spaces that belong to a client.                      |
| QUERY NODE                 | Displays partial or complete information about one or more clients.                          |
| QUERY PVUESTIMATE          | Displays an estimate of the client-devices and server-devices being managed.                 |
| QUERY REPLNODE             | Displays information about the replication status of a client node.                          |
| REGISTER ADMIN             | Defines a new administrator without granting administrative authority.                       |
| REGISTER NODE              | Defines a client node to the server and sets options for that user.                          |
| REMOVE NODE                | Removes a client from the list of registered nodes for a specific policy domain.             |
| REMOVE REPLNODE            | Removes a node from replication.                                                             |
| RENAME NODE                | Changes the name for a client node.                                                          |
| RESET PASSEXP              | Resets the password expiration for nodes or administrators.                                  |
| SET DEDUPVERIFICATIONLEVEL | Specifies the percentage of extents verified by the server during client-side deduplication. |
| SET PASSEXP                | Specifies the number of days after which a password is expired and must be changed.          |
| UPDATE ADMIN               | Changes the password or contact information associated with any administrator.               |
| UPDATE FILESPACE           | Changes file-space node-replication rules.                                                   |

## UPDATE NODE

## UPDATE NODEGROUP (Update a node group)

Use this command to modify the description of a node group.

### Privilege class

To issue this command, you must have system or unrestricted policy privilege.

### Syntax

```
►►—UPDate NODEGroup—group_name—DESCription—=description—◀◀
```

### Parameters

*group\_name*

Specifies the name of the node group whose description you want to update.

### DESCription

Specifies a description of the node group. This parameter is required. The maximum length of the description is 255 characters. If the description contains any blanks, enclose the entire description in quotation marks.

### Example: Update a node group's description

Update the node group, *group1*, with a new description.

```
update nodegroup group1 description="Human Resources"
```

### Related commands

Table 427. Commands related to UPDATE NODEGROUP

| Command                | Description                                             |
|------------------------|---------------------------------------------------------|
| DEFINE BACKUPSET       | Defines a previously generated backup set to a server.  |
| DEFINE NODEGROUP       | Defines a group of nodes.                               |
| DEFINE NODEGROUPMEMBER | Adds a client node to a node group.                     |
| DELETE BACKUPSET       | Deletes a backup set.                                   |
| DELETE NODEGROUP       | Deletes a node group.                                   |
| DELETE NODEGROUPMEMBER | Deletes a client node from a node group.                |
| GENERATE BACKUPSET     | Generates a backup set of a client's data.              |
| QUERY BACKUPSET        | Displays backup sets.                                   |
| QUERY NODEGROUP        | Displays information about node groups.                 |
| UPDATE BACKUPSET       | Updates a retention value associated with a backup set. |

### UPDATE PATH (Change a path)

Use this command to update a path definition.

Syntax and parameter descriptions are available for the following path types.

- “UPDATE PATH (Change a path when the destination is a drive)” on page 1327
- “UPDATE PATH (Change a path when the destination is a library)” on page 1331
- “UPDATE PATH (Update a path when the destination is a ZOSMEDIA library)” on page 1334

For detailed and current device support information, see the Supported Devices website for your operating system:

[http://www.ibm.com/software/sysmgmt/products/support/IBM\\_TSM\\_Supported\\_Devices\\_for\\_Linux.html](http://www.ibm.com/software/sysmgmt/products/support/IBM_TSM_Supported_Devices_for_Linux.html)

### Related commands

*Table 428. Commands related to UPDATE PATH*

| Command          | Description                                                         |
|------------------|---------------------------------------------------------------------|
| DEFINE DATAMOVER | Defines a data mover to the Tivoli Storage Manager server.          |
| DEFINE DRIVE     | Assigns a drive to a library.                                       |
| DEFINE LIBRARY   | Defines an automated or manual library.                             |
| DEFINE PATH      | Defines a path from a source to a destination.                      |
| DELETE PATH      | Deletes a path from a source to a destination.                      |
| QUERY PATH       | Displays information about the path from a source to a destination. |
| UPDATE DATAMOVER | Changes the definition for a data mover.                            |


## UPDATE PATH (Change a path when the destination is a drive)

Use this syntax when updating a path definition to a drive.

### Privilege class

To issue this command you must have system privilege or unrestricted storage privilege.

### Syntax



### Parameters

#### *source\_name* (Required)

Specifies the name of source for the path. This parameter is required.

#### *destination\_name* (Required)

Specifies the name of the destination. This parameter is required.

#### SRCType (Required)

Specifies the type of the source. This parameter is required. Possible values are:

##### DATAMover

Specifies that a data mover is the source.

##### SERVer

Specifies that a server or a storage agent is the source.

#### AUTODetect

Specifies whether the serial number for a drive or library will be automatically detected, reported, and updated in Tivoli Storage Manager. This parameter is optional. This parameter is only valid for paths defined from the local server to a drive or a library. Possible values are:

**No** Specifies that the serial number is not automatically updated.

##### Yes

Specifies that the serial number is automatically updated to reflect the same serial number that the drive reports to Tivoli Storage Manager.

#### Important:

1. If you have not previously entered a serial number, then AUTODETECT defaults to YES. If you have previously entered a serial number, then AUTODETECT defaults to NO.

## UPDATE PATH - Destination is a drive

2. AUTODETECT=YES in this command overrides the serial number set in the **DEFINE DRIVE** command.
3. If you set DESTTYPE=DRIVE and AUTODETECT=YES, then the drive element number in the IBM Tivoli Storage Manager database will be automatically changed to reflect the same element number that corresponds to the serial number of that drive. This is true for drives in a SCSI library. For more information about the element number, see the **DEFINE DRIVE** command.
4. Depending on the capabilities of the device, the AUTODETECT parameter may not be supported.

### **DESTType=DRive (Required)**

Specifies that a drive is the destination. When the destination is a drive, you must specify a library name. This parameter is required.

### **LIBRARY**

Specifies the name of the library to which the drive is assigned. The library and its drives must already be defined to the Tivoli Storage Manager server. If the path is from a NAS data mover to a library, the library must have LIBTYPE of SCSI, 349x, or ACSLS.

### **DEVIce**

Specifies the name of the device as known to the source, or FILE if the device is a logical drive in a FILE library.

The source uses the device name to access the drive. See Table 429 for examples.

Table 429. Examples of device names

| Source to destination                                                        | Example                                                                                                         |
|------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| Server to a drive (not a FILE drive)                                         | /dev/tsm SCSI/mt3                                                                                               |
| Storage agent to a drive (not a FILE drive)                                  | /dev/tsm SCSI/mt3                                                                                               |
| Storage agent to a drive when the drive is a logical drive in a FILE library | FILE                                                                                                            |
| NAS data mover to a drive                                                    | NetApp NAS file server: rst01<br>EMC Celerra NAS file server:<br>c436t011<br>IBM System Storage N Series: rst01 |

### **Important:**

- For more complete information about device names when the source is a server, see the *Administrator's Guide*.
- For information about the device name when the source is a storage agent, see the *Storage Agent User's Guide*.
- For 349X libraries, the alias name is a symbolic name that is specified in the /etc/ibmatl.conf file. For more information, refer to the *IBM Tape Device Drivers Installation and User's Guide*. The guide can be downloaded from the IBM Systems support site at the following URL: <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972>.
- For information about how to obtain names for devices that are connected to a NAS file server, consult the product information for the file server. For

example, for a NetApp file server, connect to the file server using Telnet and issue the **SYSCONFIG** command. Use this command to determine device names for drives:

```
sysconfig -t
```

### **ONLine**

Specifies whether the path is available for use. This parameter is optional. Possible values are:

#### **Yes**

Specifies that the path is available for use.

**No** Specifies that the path is not available for use.

The source and the destination must both be available to use the path.

For example, if the path from a data mover to a drive is online, but either the data mover or the drive is offline, you cannot use the path.

### **DIRectory**

Specifies the directory location or locations for a storage agent to access the files in a FILE library. The DIRECTORY parameter is also used for devices of type REMOVABLEFILE. For REMOVABLEFILE devices, the DIRECTORY parameter provides information for the server (not a storage agent) along with the DRIVE parameter to describe access to the device. This parameter is optional.

On storage agents, this parameter is only valid when *all* of the following conditions are true:

- The source type is SERVER (meaning a storage agent that has been defined as a server to this server).
- The source name is the name of a storage agent, *not* the server.
- The destination is a logical drive that is part of a FILE library.
- If multiple directories were specified for the device class associated with the FILE library, the same number of directories must be specified with the DIRectory parameter of the **DEFINE PATH** command, for each drive in the FILE library. Storage agent directories are not validated on the Tivoli Storage Manager server. Specifying incorrect directories can cause a run-time failure.

The directory name or names identify the locations where the storage agent reads and writes the files that represent storage volumes for the FILE device class that is associated with the FILE library. The default value for DIRECTORY is the directory of the server at the time the command is issued.

Use a naming convention that you can use to associate the directory with a particular physical drive. This can help ensure that your configuration is valid for sharing the FILE library between the server and storage agent. If the storage agent is on a Windows system, use a universal naming convention (UNC) name. When the storage agent lacks permission to access remote storage, the storage agent will experience mount failures.

#### **Important:**

- Tivoli Storage Manager does not create shares or permissions, or mount the target file system. You must perform these actions before starting the storage agent.
- You can modify a list of directories only by replacing the entire list.
- You must ensure that storage agents can access newly created FILE volumes. To access FILE volumes, storage agents replace names from the directory list

## UPDATE PATH - Destination is a drive

in the device-class definition with the names in the directory list for the associated path definition. The following illustrates the importance of matching device classes and paths to ensure that storage agents can access newly created FILE volumes.

Suppose you want to use these three directories for a FILE library:

```
/opt/tivoli1
/opt/tivoli2
/opt/tivoli3
```

1. You use the following command to set up a FILE library named CLASSA with one drive named CLASSA1 on SERVER1:

```
define devclass classa devtype=file
directory="/opt/tivoli1,/opt/tivoli2,/opt/tivoli3"
shared=yes mountlimit=1
```

2. You want the storage agent STA1 to be able to use the FILE library, so you define the following path for storage agent STA1:

```
define path server1 stal srctype=server desttype=drive device=file
directory="/opt/ibm1,/opt/ibm2,/opt/ibm3" library=classa
```

In this scenario, the storage agent, STA1, will replace the directory name /opt/tivoli1 with the directory name /opt/ibm1/ to access FILE volumes that are in the /opt/tivoli1 directory on the server.

3. If file volume /opt/tivoli1/file1.dsm is created on SERVER1, and if the following command is issued,

```
update devclass classa directory="/opt/otherdir,/opt/tivoli2,
/opt/tivoli3"
```

SERVER1 will still be able to access file volume /opt/tivoli1/file1.dsm, but the storage agent STA1 will not be able to access it because a matching directory name in the PATH directory list no longer exists. If a directory name is not available in the directory list associated with the device class, the storage agent can lose access to a FILE volume in that directory. Although the volume will still be accessible from the Tivoli Storage Manager server for reading, failure of the storage agent to access the FILE volume can cause operations to be retried on a LAN-only path or to fail.

### Example: Update a path from a data mover NAS file server to a tape drive

Update a path from a data mover that is a NAS file server to the drive TAPEDRV2 that the data mover uses for backup and restore operations. In this example, the NAS data mover is NAS1, the library is NASLIB, and the device name for the drive is rst01.


```
update path nas1 tapedrv2 srctype=datamover desttype=drive library=naslib
device=rst01
```

**UPDATE PATH (Change a path when the destination is a library)**

Use this syntax when updating a path definition to a library.

**Privilege class**

To issue this command you must have system privilege or unrestricted storage privilege.

**Syntax****Parameters****source\_name (Required)**

Specifies the name of source for the path. This parameter is required.

**destination\_name (Required)**

Specifies the name of the destination. This parameter is required.

**Important:** To define a path from a NAS data mover to a library, the library must have LIBTYPE of SCSI, 349X, or Automated Cartridge System Library Software (ACSLs).

**SRCType (Required)**

Specifies the type of the source. This parameter is required. Possible values are:

**DATAMover**

Specifies that a data mover is the source.

**SERVER**

Specifies that a server or a storage agent is the source.

**AUTODetect**

Specifies whether the serial number for a drive or library is automatically detected, reported, and updated in Tivoli Storage Manager. This parameter is optional. This parameter is only valid for paths defined from the local server to a library. Possible values are:

**No** Specifies that the serial number is not automatically updated.

**Yes**

Specifies that the serial number is automatically updated to reflect the same serial number that the drive reports to Tivoli Storage Manager.

**Important:**

1. If you have not previously entered a serial number, then AUTODETECT defaults to YES. If you have previously entered a serial number, then AUTODETECT defaults to NO.

## UPDATE PATH - Destination is a library

2. AUTODETECT=YES in this command overrides the serial number set in the **DEFINE DRIVE** command.
3. Depending on the capabilities of the device, the AUTODETECT parameter may not be supported.

### **DESTType=LIBRARY (Required)**

Specifies that a library is the destination.. This parameter is required.

### **DEVIce**

Specifies the name of the device as known to the source, or FILE if the device is a logical drive in a FILE library.

The source uses the device name to access the drive or library. See Table 430 for examples.

Table 430. Examples of device names

| Source to destination       | Example          |
|-----------------------------|------------------|
| Server to a library         | /dev/tmsmcsi/lb4 |
| NAS data mover to a library | mc0              |

### **Important:**

- For more complete information about device names when the source is a server, see the *Administrator's Guide*.
- For information about the device name when the source is a storage agent, see the *Storage Agent User's Guide*.
- For 349X libraries, the alias name is a symbolic name that is specified in the /etc/ibmatl.conf file. For more information, refer to the *IBM Tape Device Drivers Installation and User's Guide*. The guide can be downloaded from the IBM Systems support site at the following URL: <http://www.ibm.com/support/docview.wss?uid=ssg1S7002972>.
- For information about how to obtain names for devices that are connected to a NAS file server, consult the product information for the file server. For example, for a NetApp file server, connect to the file server using Telnet and issue the **SYSCONFIG** command. Use this command to determine the device name for a library:

```
sysconfig -m
```

### **EXTERNALManager**

Specifies the location of the external library manager where Tivoli Storage Manager can send media access requests. Use single quotation marks around the value of this parameter. For example, enter:

```
/opt/GESedt-acsls/bin/elmdt
```

This parameter is required when the library name is an external library.

### **ONLine**

Specifies whether the path is available for use. This parameter is optional. Possible values are:

#### **Yes**

Specifies that the path is available for use.

**No** Specifies that the path is not available for use.

The source and the destination must both be available to use the path.

## UPDATE PATH - Destination is a library

**Important:** If the path to a library is offline, the server will not be able to access the library. If the server is halted and restarted while the path to the library is offline, the library will not be initialized.

For additional information, see the *Administrator's Guide*.

## UPDATE PATH - Destination is a ZOSMEDIA library

### UPDATE PATH (Update a path when the destination is a ZOSMEDIA library)

Use this syntax when you update a path to a ZOSMEDIA library.

#### Privilege class

To issue this command, you must have system privilege or unrestricted storage privilege.

#### Syntax

```
►►—UPDate PATH—source_name—destination_name—SRCType—=—SERVer—————►
►—DESTType—=—LIBRary—ZOSMEDIASERVer—=—server_name—————►
└─ONLine—=—┬─Yes
 └─No
```

#### Parameters

##### *source\_name* (Required)

Specifies the name of source for the path.

##### *destination\_name* (Required)

Specifies the name of the destination.

##### SRCType=SERVer (Required)

Specifies that the Tivoli Storage Manager server or a storage agent is the source.

##### DESTType=LIBRary (Required)

Specifies that a library is the destination.

##### ZOSMEDIAServer (Required)

Specifies the server name that represents a Tivoli Storage Manager for z/OS Media server.

##### ONLine

Specifies whether the path is available for use. This parameter is optional. Possible values are:

##### Yes

Specifies that the path is available for use.

**No** Specifies that the path is not available for use.

The source and the destination must both be available to use the path.

**Important:** If the path to a library is offline, the server cannot access the library. If the server is halted and restarted while the path to the library is offline, the library is not initialized during server initialization. The path must be updated to ONLINE=YES to access the library.

For more information, see the *Administrator's Guide*.

## UPDATE POLICYSET (Update a policy set description)

Use this command to change the description of a policy set. You cannot change the description of the ACTIVE policy set.

### Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

### Syntax

```
►►—UPDate Policyset—domain_name—policy_set_name—————►
►—DESCRiption—=description—————►◄◄
```

### Parameters

#### *domain\_name* (Required)

Specifies the policy domain to which the policy set belongs.

#### *policy\_set\_name* (Required)

Specifies the policy set to update. You cannot change the ACTIVE policy set.

#### DESCRiption (Required)

Specifies text that describes the policy set. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove a previously defined description, specify a null string ("").

### Example: Update a policy set

Update a policy set called VACATION for the EMPLOYEE\_RECORDS policy domain with a description of "Schedule Planning Information."

```
update policyset employee_records vacation
description="schedule planning information"
```

### Related commands

Table 431. Commands related to UPDATE POLICYSET

| Command            | Description                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------|
| ACTIVATE POLICYSET | Validates and activates a policy set.                                                         |
| COPY MGMTCLASS     | Creates a copy of a management class.                                                         |
| DEFINE DOMAIN      | Defines a policy domain that clients can be assigned to.                                      |
| DEFINE MGMTCLASS   | Defines a management class.                                                                   |
| DEFINE POLICYSET   | Defines a policy set within the specified policy domain.                                      |
| DELETE POLICYSET   | Deletes a policy set, including its management classes and copy groups, from a policy domain. |
| QUERY POLICYSET    | Displays information about policy sets.                                                       |

## UPDATE POLICYSET

*Table 431. Commands related to UPDATE POLICYSET (continued)*

| <b>Command</b>     | <b>Description</b>                                                                                   |
|--------------------|------------------------------------------------------------------------------------------------------|
| VALIDATE POLICYSET | Verifies and reports on conditions the administrator must consider before activating the policy set. |

## UPDATE PROFILE (Update a profile description)

Use this command on a configuration manager to update a profile description.

### Privilege class

To issue this command, you must have system privilege.

### Syntax

```
►►—UPDate PROFILE—profile_name—DESCRiption—=description—◄◄
```

### Parameters

#### *profile\_name* (Required)

Specifies the profile to update.

#### DESCRiption (Required)

Specifies a description for the profile. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove a description, specify a null string ("").

### Example: Update a profile's description

Update the description for profile DELTA.

```
update profile delta description="PAYROLL domain"
```

### Related commands

Table 432. Commands related to UPDATE PROFILE

| Command                | Description                                                        |
|------------------------|--------------------------------------------------------------------|
| COPY PROFILE           | Creates a copy of a profile.                                       |
| DEFINE PROFASSOCIATION | Associates objects with a profile.                                 |
| DEFINE PROFILE         | Defines a profile for distributing information to managed servers. |
| DELETE PROFASSOCIATION | Deletes the association of an object with a profile.               |
| DELETE PROFILE         | Deletes a profile from a configuration manager.                    |
| LOCK PROFILE           | Prevents distribution of a configuration profile.                  |
| QUERY PROFILE          | Displays information about configuration profiles.                 |
| SET CONFIGMANAGER      | Specifies whether a server is a configuration manager.             |
| UNLOCK PROFILE         | Enables a locked profile to be distributed to managed servers.     |


## UPDATE RECOVERYMEDIA (Update recovery media)

Use this command to update information about recovery media.

### Privilege class

To issue this command, you must have system privilege.

### Syntax



### Parameters

#### *media\_name* (Required)

Specifies the name of the recovery media to be updated.

#### **VOLumentnames**

Specifies the names of volumes that contain the recoverable data (for example, operating system image copies). If you specify a `TYPE=BOOT`, you must specify the boot media volume names in the order in which they are to be loaded at recovery time. The volume names list can be up to 255 characters. Enclose the list in quotation marks if it contains any blank characters. To remove all volume names, specify a null string ("").

#### **DESCRiption**

Specifies the description of the recovery media. This parameter is optional. You can use up to 255 characters. Enclose the text in quotation marks if it contains any blank characters.

#### **L0cation**

Describes the location of the recovery media. This parameter is optional. You can use up to 255 characters. Enclose the text in quotation marks if it contains any blank characters. To remove a location description, specify a null string ("") for the value.

#### **Type**

Specifies the type of recovery media. This parameter is optional. Possible values are:

##### **BOOT**

Specifies that this is boot media. You must specify volume names if the type is `BOOT`.

##### **OTHER**

Specifies that this is not boot media. For example, a CD that contains operating system manuals.

**PROduct**

Specifies the name of the product that wrote to this media. This parameter is optional. You can use up to 16 characters. Enclose the text in quotation marks if it contains any blank characters. To remove a product name, specify a null string ("" ) for the value.

**PRODUCTInfo**

Specifies any information about the product that wrote to the media that you may need to restore the machine. This parameter is optional. You can use up to 255 characters. Enclose the text in quotation marks if it contains any blank characters. To remove previously defined product information, specify a null string ("" ) for the value.

**Example: Update a recovery media's location description**

Update the location description for recovery media DIST5RM to "Corporate Headquarters Data Vault."

```
update recoverymedia dist5rm
location="Corporate Headquarters Data Vault"
```

**Related commands**

*Table 433. Commands related to UPDATE RECOVERYMEDIA*

| Command              | Description                                      |
|----------------------|--------------------------------------------------|
| DEFINE RECOVERYMEDIA | Defines the media required to recover a machine. |
| DELETE RECOVERYMEDIA | Deletes recovery media.                          |
| QUERY RECOVERYMEDIA  | Displays media available for machine recovery.   |



Table 434. Commands related to UPDATE REPLRULE (continued)

| Command               | Description                                                        |
|-----------------------|--------------------------------------------------------------------|
| SET ARREPLRULEDEFAULT | Specifies the server node-replication rule for archive data.       |
| SET BKREPLRULEDEFAULT | Specifies the server node-replication rule for backup data.        |
| SET SPREPLRULEDEFAULT | Specifies the server node-replication rule for space-managed data. |
| UPDATE FILESPACE      | Changes file-space node-replication rules.                         |
| UPDATE NODE           | Changes the attributes associated with a client node.              |
| VALIDATE REPLICATION  | Verifies replication for file spaces and data types.               |

## UPDATE SCHEDULE (Update a schedule)

Use this command to update a client or administrative command schedule.

The UPDATE SCHEDULE command takes two forms, depending on whether the schedule applies to client operations or administrative commands. Within these two forms, you can select either classic or enhanced style schedules. The syntax and parameters for each form are defined separately.

- “UPDATE SCHEDULE (Update an administrative schedule)” on page 1356
- “UPDATE SCHEDULE (Update a client schedule)” on page 1343

*Table 435. Commands related to UPDATE SCHEDULE*

| Command               | Description                                                                                     |
|-----------------------|-------------------------------------------------------------------------------------------------|
| COPY SCHEDULE         | Creates a copy of a schedule.                                                                   |
| DEFINE SCHEDULE       | Defines a schedule for a client operation or an administrative command.                         |
| DELETE SCHEDULE       | Deletes a schedule from the database.                                                           |
| QUERY EVENT           | Displays information about scheduled and completed events for selected clients.                 |
| QUERY SCHEDULE        | Displays information about schedules.                                                           |
| SET MAXCMDRETRIES     | Specifies the maximum number of retries after a failed attempt to execute a scheduled command.  |
| SET MAXSCHEDESESSIONS | Specifies the maximum number of client/server sessions available for processing scheduled work. |
| SET RETRYPERIOD       | Specifies the time between retry attempts by the client scheduler.                              |

## UPDATE SCHEDULE (Update a client schedule)

Use the **UPDATE SCHEDULE** to update selected parameters for a client schedule.


This command does not change the client associations that have been made to this schedule. Any clients that are associated with the original schedule, process the modified schedule.

Not all clients can run all scheduled operations, even though you can define the schedule on the server and associate it with the client. For example, a Macintosh client cannot run a schedule when the action is to restore or retrieve files, or run an executable script. An executable script is also known as a command file, a batch file, or a script on different client operating systems.

### Privilege class

To update a client schedule, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the schedule belongs.

### Syntax for a classic client schedule




## UPDATE SCHEDULE




### Notes:

- 1 You must specify at least one optional parameter on this command.
- 2 You can define or update a SYSTEMSTATE or VM subaction on the Administration Center, but those actions are not saved.


### Syntax for an enhanced client schedule



## UPDATE SCHEDULE



## UPDATE SCHEDULE



### Notes:

- 1 You must specify at least one optional parameter on this command.

### Parameters

#### *domain\_name* (Required)

Specifies the name of the policy domain to which this schedule belongs.

#### *schedule\_name* (Required)

Specifies the name of the schedule to be updated.

#### Type=Client

Specifies that a client schedule is updated. This parameter is optional. The default is CLIENT.

### DESCription

Specifies a description of the schedule. This parameter is optional. You can specify up to 255 characters for the description. Enclose the description in quotation marks if it contains blank characters. To remove a previously defined description, specify a null string ("") for this value.

### ACTion

Specifies the action that occurs when this schedule is processed. Possible values are:

#### **Incremental**

Specifies that the schedule backs up all files that are new or that have changed since the last incremental backup. Incremental also backs up any file for which all existing backups might have expired.

#### **Selective**

Specifies that the schedule backs up only files that are specified with the OBJECTS parameter.

#### **Archive**

Specifies that the schedule archives files that are specified with the OBJECTS parameter.

#### **Backup**

Specifies that the schedule backs up files that are specified with the OBJECTS parameter.

#### **REStore**

Specifies that the schedule restores files that are specified with the OBJECTS parameter.

When you specify ACTION=RESTORE for a scheduled operation, and the REPLACE option is set to PROMPT, no prompting occurs. If you set the option to PROMPT, the files are skipped.

If you specify a second file specification, this second file specification acts as the restore destination. If you need to restore multiple groups of files, schedule one for each file specification that you need to restore.

#### RETrieve

Indicates that the schedule retrieves files that are specified with the OBJECTS parameter.

**Remember:** A second file that is specified acts as the retrieve destination. If you need to retrieve multiple groups of files, create a separate schedule for each group of files.

#### IMAGEBACKup

Specifies that the schedule backs up logical volumes that are specified with the OBJECTS parameter.

#### IMAGEREStore

Specifies that the schedule restores logical volumes that are specified with the OBJECTS parameter.

#### Command

Specifies that the schedule processes a client operating system command or script that is specified with the OBJECTS parameter.

#### Macro

Specifies that a client processes a macro whose file name is specified with the OBJECTS parameter.

#### SUBACTioN

Possible values are:

"" When a null string (two double quotes) is specified with ACTION=BACKUP the backup is an incremental.

#### FASTBacK

Specifies that a FastBack client operation that is identified by the ACTION parameter is to be scheduled for processing. The ACTION parameter must be either ARCHIVE or BACKUP.

#### SYSTEMStAte

Specifies that a client Systemstate backup is scheduled.

VM Specifies that a client VMware backup operation is scheduled.

#### DeploY

Specifies whether to update client workstations with deployment packages that are specified with the OBJECTS parameter. The OBJECTS parameter must contain two specifications, the package files to retrieve and the location from which to retrieve them. Ensure that the objects are in the order *files location*. For example:

```
define schedule standard deploy_1 action=DEPLOY objects=
"\\IBM_ANR_WIN\c$\tsm\maintenance\client\v6r2\Windows\X32\v620\v6200*
..\IBM_ANR_WIN"
```

Values for the following options are restricted when you specify ACTION=DEPLOY:

## UPDATE SCHEDULE

### PERUNITS

Specify PERUNITS=ONETIME. If you specify PERUNITS=PERIOD, the parameter is ignored.

### DURUNITS

Specify MINUTES, HOURS, or DAYS for the DURUNITS parameter. Do not specify INDEFINITE.

### SCHEDSTYLE

Specify the default style, CLASSIC.

The SCHEDULE command fails if the parameters do not conform to the required parameter values, such as the V.R.M.F.

### OPTions

Specifies the client options that you specify to the scheduled command at the time the schedule is processed. This parameter is optional.

Only those options that are valid on the scheduled command can be specified for this parameter. Refer to the appropriate client manual for information about options that are valid from the command line. All options described there as valid only on the initial command line result in an error or are ignored when running the schedule from the server. For example, do not include the following options because they have no impact when the client processes the scheduled command:

MAXCMDRETRIES  
OPTFILE  
QUERYSCHEDPERIOD  
RETRYPERIOD  
SCHEDLOGNAME  
SCHEDMODE  
SERVERNAME  
TCPCLIENTADDRESS  
TCPCLIENTPORT

If the option string contains multiple options or options with embedded spaces, surround the entire option string with one pair of apostrophes. Enclose individual options that contain spaces in quotation marks. A leading minus sign is required in front of the option. Errors can occur if the option string contains spaces that are not quoted correctly.

The following examples show how to specify some client options:

- To specify subdir=yes and domain all-local -systemobject, enter:  
options='-subdir=yes -domain="all-local -c: -systemobject"'
- To specify domain all-local -c: -d:, enter:  
options='-domain="all-local -c: -d:"'

### OBjects

Specifies the objects for which the specified action is performed. Use a single space between each object. This parameter is required except when ACTION=INCREMENTAL. If the action is a backup, archive, retrieve, or restore operation, the objects are file spaces, directories, or logical volumes. See the *Backup-Archive Clients Installation and User's Guide* for command syntax information. If the action is to run a command or macro, the object is the name of the command or macro to run.

When you specify ACTION=INCREMENTAL without specifying a value for this parameter, the scheduled command is invoked without specified objects and attempts to process the objects as defined in the client option file. To select all file spaces or directories for an action, explicitly list them in the object string. Entering only an asterisk in the object string causes the backup to occur only for the directory where the scheduler was started.

### Important:

- If you specify a second file specification, and it is not a valid destination, you receive this error:  
ANS1082E Invalid destination file specification <filespec> entered.
- If you specify more than two file specifications, you receive this error:  
ANS1102E Excessive number of command line arguments passed to the program!

When you specify ACTION=ARCHIVE, INCREMENTAL, or SELECTIVE for this parameter, you can list a maximum of twenty (20) file specifications.

Enclose the object string in double quotes if it contains blank characters (spaces), and then surround the double quotes with single quotes. If the object string contains multiple file names, enclose each file name with its own pair of double quotes, then surround the entire string with one pair of single quotes. Errors can occur if file names contain a space that is not quoted correctly.

The following examples show how to specify some file names:

- To specify /home/file 2, /home/gif files, and /home/my test file, enter:  
OBJECTS="'/home/file 2' '/home/gif files' '/home/my test file'"
- To specify /home/test file, enter:  
OBJECTS="'/home/test file'"

### PRIority

Specifies the priority value for a schedule. This parameter is optional. You can specify an integer from 1 to 10, with 1 being the highest priority and 10 being the lowest. The default is 5.

If two or more schedules have the same window start time, the value you specify determines when Tivoli Storage Manager processes the schedule. The schedule with the highest priority starts first. For example, a schedule with PRIORITY=3 starts before a schedule with PRIORITY=5.

### STARTDate

Specifies the date for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current date. Use this parameter with the **STARTTIME** parameter to specify when the initial startup window of the schedule starts.

You can specify the date using one of the values below:

| Value                    | Description                                                                               | Example         |
|--------------------------|-------------------------------------------------------------------------------------------|-----------------|
| MM/DD/YYYY               | A specific date                                                                           | 09/15/1998      |
| TODAY                    | The current date                                                                          | TODAY           |
| TODAY+days or +days      | The current date plus days specified. The maximum number of days you can specify is 9999. | TODAY +3 or +3. |
| EOLM (End Of Last Month) | The last day of the previous month.                                                       | EOLM            |

## UPDATE SCHEDULE

| Value                             | Description                                              | Example                                                                                          |
|-----------------------------------|----------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| <i>EOLM-days</i>                  | The last day of the previous month minus days specified. | EOLM-1<br><br>To include files that were active a day before the last day of the previous month. |
| BOTM<br>(Beginning Of This Month) | The first day of the current month.                      | BOTM                                                                                             |
| <i>BOTM+days</i>                  | The first day of the current month, plus days specified. | BOTM+9<br><br>To include files that were active on the 10th day of the current month.            |

### STARTTime

Specifies the time for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current time. This parameter is used in conjunction with the **STARTDATE** parameter to specify when the initial startup window begins.

You can specify the time using one of the values below:

| Value                                | Description                                        | Example                                                                                                                                                     |
|--------------------------------------|----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>HH:MM:SS</i>                      | A specific time                                    | 10:30:08                                                                                                                                                    |
| <b>NOW</b>                           | The current time                                   | NOW                                                                                                                                                         |
| <b>NOW+HH:MM</b><br>or <b>+HH:MM</b> | The current time plus hours and minutes specified  | NOW+02:00 or +02:00.<br><br>If you issue this command at 5:00 with STARTTIME=NOW+02:00 or STARTTIME=+02:00, the beginning of the startup window is at 7:00. |
| <b>NOW-HH:MM</b><br>or <b>-HH:MM</b> | The current time minus hours and minutes specified | NOW-02:00 or -02:00.<br><br>If you issue this command at 5:00 with STARTTIME=NOW-02:00 or STARTTIME=-02:00, the beginning of the startup window is at 3:00. |

### DURation

Specifies the number of units that define the length of the startup window of the scheduled operation. This parameter is optional. This value must be from 1 to 999. The default is 1.

Use this parameter with the **DURUNITS** parameter to specify the length of the startup window. For example, if you specify **DURATION=20** and **DURUNITS=MINUTES**, the schedule must be started within 20 minutes of the start date and start time. The default length of the startup window is 1 hour. The duration of the window must be shorter than the period between windows.

This value is ignored if you specify **DURUNITS=INDEFINITE**.

**Tip:** Define schedules with durations longer than 10 minutes. Doing this will give the Tivoli Storage Manager scheduler enough time to process the schedule and prompt the client.

**DURUnits**

Specifies the time units used to determine the duration of the window in which the schedule can start. This parameter is optional. The default is HOURS.

Use this parameter with the **DURATION** parameter to specify how long the startup window remains open to process the schedule. For example, if DURATION=20 and DURUNITS=MINUTES, the schedule must be started within 20 minutes of the start date and start time. The schedule may not necessarily complete processing within this window. If the schedule needs to be retried for any reason, the retry attempts must begin before the startup window elapses, or the operation does not restart.

The default value for the length of the startup window is 1 hour. Possible values are:

**Minutes**

Specifies that the duration of the window is defined in minutes.

**Hours**

Specifies that the duration of the window is defined in hours.

**Days**

Specifies that the duration of the window is defined in days.

**INDefinite**

Specifies that the startup window of the scheduled operation has an indefinite duration. The schedule can run any time after the scheduled start time, until the schedule expires. You cannot specify DURUNITS=INDEFINITE, unless you specify PERUNITS=ONETIME. The INDEFINITE value is not allowed with enhanced schedules.

**SCHEDStyle**

This parameter is optional. SCHEDSTYLE defines either the interval between times when a schedule can run, or the days on which it can run. The style can be either **classic** or **enhanced**. This parameter must be specified when you change a schedule from classic to enhanced or back to classic. Otherwise, the value for the existing schedule is used.

For classic schedules, these parameters are allowed: PERIOD, PERUNITS, and DAYOFWEEK. These parameters are not allowed: MONTH, DAYOFMONTH, and WEEKOFMONTH. If the previous schedule style was enhanced, the MONTH, DAYOFMONTH, WEEKOFMONTH, and DAYOFWEEK parameters are reset. DAYOFWEEK, PERIOD, and PERUNITS are set to default values unless they are specified with the update command.

For enhanced schedules, these parameters are allowed: MONTH, DAYOFMONTH, WEEKOFMONTH, and DAYOFWEEK. These parameters are not allowed: PERIOD and PERUNITS. If the previous schedule style was classic, the DAYOFWEEK, PERIOD, and PERUNITS parameters are reset. MONTH, DAYOFMONTH, WEEKOFMONTH, and DAYOFWEEK are set to default values unless they are specified with the update command.

**PERiod**

Specifies the length of time between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. You can specify an integer from 1 to 999. The default is 1.

Use this parameter with the **PERUNITS** parameter to specify the period between startup windows. For example, if you specify PERIOD=5 and PERUNITS=DAYS (assuming that DAYOFWEEK=ANY), the operation is

## UPDATE SCHEDULE

scheduled every five days after the initial start date and start time. The period between startup windows must exceed the duration of each window. The default is 1 day.

This value is ignored if you specify `PERUNITS=ONETIME`.

### **PERUnits**

Specifies the time units used to determine the period between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. The default is `DAYS`.

Use this parameter with the **PERIOD** parameter to specify the period between startup windows. For example, if you specify `PERIOD=5` and `PERUNITS=DAYS` (assuming that `DAYOFWEEK=ANY`), the operation is scheduled every 5 days after the initial start date and start time. The default is 1 day. Possible values are:

### **Hours**

Specifies that the time between startup windows is in hours.

### **Days**

Specifies that the time between startup windows is in days.

### **Weeks**

Specifies that the time between startup windows is in weeks.

### **Months**

Specifies that the time between startup windows is in months.

When you specify `PERUNITS=MONTHS`, the scheduled operation will be processed each month on the same date. For example, if the start date for the scheduled operation is 02/04/1998, the schedule will process on the 4th of every month thereafter. However, if the date is not valid for the next month, then the scheduled operation will be processed on the last valid date in the month. Thereafter, subsequent operations are based on this new date. For example, if the start date is 03/31/1998, the next month's operation will be scheduled for 04/30/1998. Thereafter, all subsequent operations will be on the 30th of the month until February. Because February has only 28 days, the operation will be scheduled for 02/28/1999. Subsequent operations will be processed on the 28th of the month.

### **Years**

Specifies that the time between startup windows for the schedule is in years.

When you specify `PERUNITS=YEARS`, the scheduled operation will be processed on the same month and date of each year. For example, if the start date for the scheduled operation is 02/29/2004, the next year's scheduled operation will be 02/28/2005 because February only has 28 days. Thereafter, subsequent operations will be scheduled for February 28th.

### **Onetime**

Specifies that the schedule processes once. This value overrides the value you specified for the **PERIOD** parameter.

### **DAYofweek**

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can specify different options for the **DAYofweek** parameter, depending on whether the schedule style has been defined as Classic or Enhanced:

**Classic Schedule**

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can either specify one day of the week, or WEEKDAY, WEEKEND, or ANY. If the start date and start time fall on a day that does not correspond to a day you specify, the start date and start time will be shifted forward in 24-hour increments until the **DAYOFWEEK** parameter is satisfied.

If you select a value for **DAYOFWEEK** other than ANY, and depending on the values for PERIOD and PERUNITS, schedules may not be processed when you would expect. The default is ANY.

**Enhanced Schedule**

Specifies the days of the week on which to run the schedule. You can either specify multiple days separated by commas and no intervening blanks, or WEEKDAY, WEEKEND, or ANY. If you specify multiple days, the schedule will run on each of the specified days. If you specify WEEKDAY or WEEKEND, you must also specify either WEEKOFMONTH=FIRST or WEEKOFMONTH=LAST, and the schedule will run just once per month.

The default value is ANY, meaning the schedule will run every day of the week or on the day or days determined by other enhanced schedule parameters. **DAYOFWEEK** must have a value of ANY (either by default or specified with the command) when used with the **DAYOFMONTH** parameter.

Possible values for the **DAYofweek** parameter are:

**ANY**

Specifies that the startup window can begin on any day of the week.

**WEEKDay**

Specifies that the startup window can begin on Monday, Tuesday, Wednesday, Thursday, or Friday.

**WEEKEnd**

Specifies that the startup window can begin on Saturday or Sunday.

**SUNday**

Specifies that the startup window begins on Sunday.

**Monday**

Specifies that the startup window begins on Monday.

**Tuesday**

Specifies that the startup window begins on Tuesday.

**Wednesday**

Specifies that the startup window begins on Wednesday.

**THursday**

Specifies that the startup window begins on Thursday.

**Friday**

Specifies that the startup window begins on Friday.

**SATurday**

Specifies that the startup window begins on Saturday.

**MONth**

Specifies the months of the year during which to run the schedule. This parameter is used only with enhanced schedules. Specify multiple values by

## UPDATE SCHEDULE

using commas and no intervening blanks. The default value is ANY, which means that the schedule runs during every month of the year.

### **DAYOFMonth**

Specifies the day of the month to run the schedule. This parameter is used only with enhanced schedules. You can either specify ANY or a number from -31 through 31, excluding zero. Negative values are a day from the end of the month, counting backwards. For example, the last day of the month is -1, the next-to-the-last day of the month is -2. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule runs on each of the specified days of the month. If multiple values resolve to the same day, the schedule runs only once that day.

The default value is ANY, which means that the schedule runs on every day of the month or on the days determined by other enhanced schedule parameters. DAYOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFWEEK or WEEKOFMONTH parameters.

If an existing schedule specifies a value other than ANY for DAYOFWEEK and WEEKOFMONTH, and DAYOFMONTH is updated, DAYOFWEEK and WEEKOFMONTH are reset to ANY.

### **WEEKofmonth**

Specifies the week of the month in which to run the schedule. This parameter is used only with enhanced schedules. A week is considered any seven-day period which does not start on a particular day of the week. You can specify FIRST, SECOND, THIRD, FOURTH, LAST, or ANY. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule runs during each of the specified weeks of the month. If multiple values resolve to the same week, the schedule runs only once during that week.

The default value is ANY. ANY means that the schedule runs during every week of the month or on the day or days determined by other enhanced schedule parameters. WEEKOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFMONTH parameter.

### **EXPIration**

Specifies the date after which this schedule is no longer used. This parameter is optional. The default is NEVER. Possible values are:

#### **Never**

Specifies that the schedule never expires.

#### *expiration\_date*

Specifies the date on which this schedule expires, in MM/DD/YYYY format. If you specify an expiration date, the schedule expires at 23:59:59 on the date you specify.

### **Example: Update the priority of a schedule**

Update the MONTHLY\_BACKUP schedule that belongs to the STANDARD policy domain by setting its priority value to 1.

```
update schedule standard monthly_backup priority=1
```

### **Example: Update the expiration date of a schedule**

Update the WEEKLY\_BACKUP schedule that belongs to the EMPLOYEE\_RECORDS policy domain to expire on March 29, 1999 (03/29/1999).

```
update schedule employee_records weekly_backup expiration=03/29/1999
```

### **Example: Update a schedule to archive on the last Friday of a month**

Update a schedule from archiving files quarterly on the last Friday of the month to archiving on the last day of the specified months.

```
update schedule employee_records quarterly_archive dayofmonth=-1
```

WEEKOFMONTH and DAYOFWEEK are reset to ANY.

## UPDATE SCHEDULE

### UPDATE SCHEDULE (Update an administrative schedule)

Use this command to update selected parameters for an administrative command schedule.

You cannot schedule **MACRO** or **QUERY ACTLOG** commands.


A managed administrative schedule that is updated by a configuration manager is set to an inactive state on the managed servers during configuration refresh processing. It remains in an inactive state until it is updated to an active state on those servers.

#### Privilege class

To update an administrative schedule, you must have system privilege.

#### Syntax

##### Classic administrative schedule






**Notes:**


- 1 You must specify at least one optional parameter on this command.

**Syntax**

**Enhanced administrative schedule**



## UPDATE SCHEDULE



### Notes:

- 1 You must specify at least one optional parameter on this command.

### Parameters

#### *schedule\_name* (Required)

Specifies the name of the schedule to be updated.

#### Type=Administrative (Required)

Specifies that an administrative command schedule is updated.

#### CMD

Specifies the administrative command to be scheduled for processing. This parameter is optional. The command you specify can contain up to 512 characters. Enclose the command in quotation marks if it contains blanks.

You cannot specify redirection characters with this parameter.

#### ACTIVE

Specifies whether the administrative command is eligible for processing. This parameter is optional. An administrative command schedule will not be processed unless it is set to the active state. Possible values are:

##### YES

Specifies that the administrative command is eligible for processing.

**NO** Specifies that the administrative command is not eligible for processing.

#### DESCRiption

Specifies a description of the schedule. This parameter is optional. You can specify up to 255 characters for the description. Enclose the description in quotation marks if it contains blanks. To remove a previously defined description, specify a null string ("") for this value.

#### PRIority

Specifies the priority value for a schedule. This parameter is optional. You can specify an integer from 1 to 10, with 1 being the highest priority and 10 being the lowest. The default is 5.

If two or more schedules have the same window start time, the value you specify determines when Tivoli Storage Manager processes the schedule. The schedule with the highest priority starts first. For example, a schedule with PRIORITY=3 starts before a schedule with PRIORITY=5.

#### STARTDate

Specifies the date for the beginning of the window in which the schedule is

first processed. This parameter is optional. The default is the current date. Use this parameter with the **STARTTIME** parameter to specify when the initial startup window of the schedule starts.

You can specify the date using one of the values below:

| <b>Value</b>                      | <b>Description</b>                                                                        | <b>Example</b>                                                                                   |
|-----------------------------------|-------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| <i>MM/DD/YYYY</i>                 | A specific date                                                                           | 09/15/1998                                                                                       |
| <b>TODAY</b>                      | The current date                                                                          | TODAY                                                                                            |
| <b>TODAY+days</b> or <i>+days</i> | The current date plus days specified. The maximum number of days you can specify is 9999. | TODAY +3 or +3.                                                                                  |
| EOLM (End Of Last Month)          | The last day of the previous month.                                                       | EOLM                                                                                             |
| <i>EOLM-days</i>                  | The last day of the previous month minus days specified.                                  | EOLM-1<br><br>To include files that were active a day before the last day of the previous month. |
| BOTM (Beginning Of This Month)    | The first day of the current month.                                                       | BOTM                                                                                             |
| <i>BOTM+days</i>                  | The first day of the current month, plus days specified.                                  | BOTM+9<br><br>To include files that were active on the 10th day of the current month.            |

**STARTTime**

Specifies the time for the beginning of the window in which the schedule is first processed. This parameter is optional. The default is the current time. This parameter is used in conjunction with the **STARTDATE** parameter to specify when the initial startup window begins.

You can specify the time using one of the values below:

| <b>Value</b>                      | <b>Description</b>                                 | <b>Example</b>                                                                                                                                              |
|-----------------------------------|----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>HH:MM:SS</i>                   | A specific time                                    | 10:30:08                                                                                                                                                    |
| <b>NOW</b>                        | The current time                                   | NOW                                                                                                                                                         |
| <b>NOW+HH:MM</b> or <i>+HH:MM</i> | The current time plus hours and minutes specified  | NOW+02:00 or +02:00.<br><br>If you issue this command at 5:00 with STARTTIME=NOW+02:00 or STARTTIME=+02:00, the beginning of the startup window is at 7:00. |
| <b>NOW-HH:MM</b> or <i>-HH:MM</i> | The current time minus hours and minutes specified | NOW-02:00 or -02:00.<br><br>If you issue this command at 5:00 with STARTTIME=NOW-02:00 or STARTTIME=-02:00, the beginning of the startup window is at 3:00. |

**DURation**

Specifies the number of units that define the length of the startup window of the scheduled operation. This parameter is optional. This value must be from 1 to 999. The default is 1.

## UPDATE SCHEDULE

Use this parameter with the **DURUNITS** parameter to specify the length of the startup window. For example, if you specify **DURATION=20** and **DURUNITS=MINUTES**, the schedule must be started within 20 minutes of the start date and start time. The default length of the startup window is 1 hour. The duration of the window must be shorter than the period between windows.

This value is ignored if you specify **DURUNITS=INDEFINITE**.

### **DURUnits**

Specifies the time units used to determine the duration of the window in which the schedule can start. This parameter is optional. The default is **HOURS**.

Use this parameter with the **DURATION** parameter to specify how long the startup window remains open to process the schedule. For example, if **DURATION=20** and **DURUNITS=MINUTES**, the schedule must be started within 20 minutes of the start date and start time. The schedule may not necessarily complete processing within this window. If the schedule needs to be retried for any reason, the retry attempts must begin before the startup window elapses, or the operation does not restart.

The default value for the length of the startup window is 1 hour. Possible values are:

### **Minutes**

Specifies that the duration of the window is defined in minutes.

### **Hours**

Specifies that the duration of the window is defined in hours.

### **Days**

Specifies that the duration of the window is defined in days.

### **INDefinite**

Specifies that the startup window of the scheduled operation has an indefinite duration. The schedule can run any time after the scheduled start time, until the schedule expires. You cannot specify **DURUNITS=INDEFINITE**, unless you specify **PERUNITS=ONETIME**. The **INDEFINITE** value is not allowed with enhanced schedules.

### **SCHEDStyle**

This parameter is optional. **SCHEDSTYLE** defines either the interval between times when a schedule should run, or the days on which it should run. The style can be either **classic** or **enhanced**. This parameter must be specified when you change a schedule from classic to enhanced or back to classic. Otherwise, the value for the existing schedule is used.

For classic schedules, these parameters are allowed: **PERIOD**, **PERUNITS**, and **DAYOFWEEK**. These parameters are not allowed: **MONTH**, **DAYOFMONTH**, and **WEEKOFMONTH**. If the previous schedule style was enhanced, the **MONTH**, **DAYOFMONTH**, **WEEKOFMONTH**, and **DAYOFWEEK** parameters will be reset. **DAYOFWEEK**, **PERIOD**, and **PERUNITS** will be set to default values unless they are specified with the update command.

For enhanced schedules, these parameters are allowed: **MONTH**, **DAYOFMONTH**, **WEEKOFMONTH**, and **DAYOFWEEK**. These parameters are not allowed: **PERIOD** and **PERUNITS**. If the previous schedule style was classic, the **DAYOFWEEK**, **PERIOD**, and **PERUNITS** parameters will be reset. **MONTH**, **DAYOFMONTH**, **WEEKOFMONTH**, and **DAYOFWEEK** will be set to default values unless they are specified with the update command.

**PERiod**

Specifies the length of time between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. You can specify an integer from 1 to 999. The default is 1.

Use this parameter with the **PERUNITS** parameter to specify the period between startup windows. For example, if you specify **PERIOD=5** and **PERUNITS=DAYS** (assuming that **DAYOFWEEK=ANY**), the operation is scheduled every five days after the initial start date and start time. The period between startup windows must exceed the duration of each window. The default is 1 day.

This value is ignored if you specify **PERUNITS=ONETIME**.

**PERUnits**

Specifies the time units used to determine the period between startup windows for this schedule. This parameter is optional. This parameter is used only with classic schedules. The default is **DAYS**.

Use this parameter with the **PERIOD** parameter to specify the period between startup windows. For example, if you specify **PERIOD=5** and **PERUNITS=DAYS** (assuming that **DAYOFWEEK=ANY**), the operation is scheduled every 5 days after the initial start date and start time. The default is 1 day. Possible values are:

**Hours**

Specifies that the time between startup windows is in hours.

**Days**

Specifies that the time between startup windows is in days.

**Weeks**

Specifies that the time between startup windows is in weeks.

**Months**

Specifies that the time between startup windows is in months.

When you specify **PERUNITS=MONTHS**, the scheduled operation will be processed each month on the same date. For example, if the start date for the scheduled operation is 02/04/1998, the schedule will process on the 4th of every month thereafter. However, if the date is not valid for the next month, then the scheduled operation will be processed on the last valid date in the month. Thereafter, subsequent operations are based on this new date. For example, if the start date is 03/31/1998, the next month's operation will be scheduled for 04/30/1998. Thereafter, all subsequent operations will be on the 30th of the month until February. Because February has only 28 days, the operation will be scheduled for 02/28/1999. Subsequent operations will be processed on the 28th of the month.

**Years**

Specifies that the time between startup windows for the schedule is in years.

When you specify **PERUNITS=YEARS**, the scheduled operation will be processed on the same month and date of each year. For example, if the start date for the scheduled operation is 02/29/2004, the next year's scheduled operation will be 02/28/2005 because February only has 28 days. Thereafter, subsequent operations will be scheduled for February 28th.

## UPDATE SCHEDULE

### **Onetime**

Specifies that the schedule processes once. This value overrides the value you specified for the **PERIOD** parameter.

### **DAYofweek**

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can specify different options for the **DAYofweek** parameter, depending on whether the schedule style has been defined as Classic or Enhanced:

#### **Classic Schedule**

Specifies the day of the week on which the startup window for the schedule begins. This parameter is optional. You can either specify one day of the week, or **WEEKDAY**, **WEEKEND**, or **ANY**. If the start date and start time fall on a day that does not correspond to a day you specify, the start date and start time will be shifted forward in 24-hour increments until the **DAYOFWEEK** parameter is satisfied.

If you select a value for **DAYOFWEEK** other than **ANY**, and depending on the values for **PERIOD** and **PERUNITS**, schedules may not be processed when you would expect. The default is **ANY**.

#### **Enhanced Schedule**

Specifies the days of the week on which to run the schedule. You can either specify multiple days separated by commas and no intervening blanks, or **WEEKDAY**, **WEEKEND**, or **ANY**. If you specify multiple days, the schedule will run on each of the specified days. If you specify **WEEKDAY** or **WEEKEND**, you must also specify either **WEEKOFMONTH=FIRST** or **WEEKOFMONTH=LAST**, and the schedule will run just once per month.

The default value is **ANY**, meaning the schedule will run every day of the week or on the day or days determined by other enhanced schedule parameters. **DAYOFWEEK** must have a value of **ANY** (either by default or specified with the command) when used with the **DAYOFMONTH** parameter.

Possible values for the **DAYofweek** parameter are:

#### **ANY**

Specifies that the startup window can begin on any day of the week.

#### **WEEKDay**

Specifies that the startup window can begin on Monday, Tuesday, Wednesday, Thursday, or Friday.

#### **WEEKEnd**

Specifies that the startup window can begin on Saturday or Sunday.

#### **SUday**

Specifies that the startup window begins on Sunday.

#### **Monday**

Specifies that the startup window begins on Monday.

#### **TUesday**

Specifies that the startup window begins on Tuesday.

#### **Wednesday**

Specifies that the startup window begins on Wednesday.

#### **THursday**

Specifies that the startup window begins on Thursday.

**Friday**

Specifies that the startup window begins on Friday.

**SAaturday**

Specifies that the startup window begins on Saturday.

**MONth**

Specifies the months of the year during which to run the schedule. This parameter is used only with enhanced schedules. Specify multiple values by using commas and no intervening blanks. The default value is ANY. This means the schedule will run during every month of the year.

**DAYOFMonth**

Specifies the day of the month to run the schedule. This parameter can only be specified with enhanced schedules. You can either specify ANY or a number from -31 through 31, excluding zero. Negative values are a day from the end of the month, counting backwards. For example, the last day of the month is -1, the next-to-the-last day of the month is -2, etc. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule will run on each of the specified days of the month. If multiple values resolve to the same day, the schedule will run only once that day.

The default value is ANY. This means the schedule will run on every day of the month or on the days determined by other enhanced schedule parameters. DAYOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFWEEK or WEEKOFMONTH parameters.

**WEEKofmonth**

Specifies the week of the month in which to run the schedule. This parameter can only be specified with enhanced schedules. A week is considered any seven-day period which does not start on a particular day of the week. You can specify FIRST, SECOND, THIRD, FOURTH, LAST, or ANY. You can specify multiple values separated by commas and no intervening blanks. If you specify multiple values, the schedule will run during each of the specified weeks of the month. If multiple values resolve to the same week, the schedule will run only once during that week.

The default value is ANY, meaning the schedule will run during every week of the month or on the day or days determined by other enhanced schedule parameters. WEEKOFMONTH must have a value of ANY (either by default or specified with the command) when used with the DAYOFMONTH parameter.

**EXPIration**

Specifies the date after which this schedule is no longer used. This parameter is optional. The default is NEVER. Possible values are:

**Never**

Specifies that the schedule never expires.

*expiration\_date*

Specifies the date on which this schedule expires, in MM/DD/YYYY format. If you specify an expiration date, the schedule expires at 23:59:59 on the date you specify.

## UPDATE SCHEDULE

### Example: Update a backup schedule to every three days

Update existing administrative schedule named BACKUP\_BACKUPPOOL so that starting today, the BACKUPPOOL primary storage pool is backed up to the COPYSTG copy storage pool every three days at 10:00 p.m.

```
update schedule backup_backuppool type=administrative cmd="backup stgpool
 backuppool copystg" active=yes starttime=22:00 period=3
```

### Example: Update a backup schedule to every first and third Friday

Update a schedule named BACKUP\_ARCHIVEPOOL that backs up the primary storage pool ARCHIVEPOOL to the copy storage pool RECOVERYPOOL. The existing schedule runs on the first and tenth day of every month. Update it to run the first and third Friday of every month.

```
update schedule backup_archivepool
 dayofweek=friday weekofmonth=first,third
```

DAYOFMONTH will be reset to ANY.

## UPDATE SCRIPT (Update a Tivoli Storage Manager script)

Use this command to change a command line or to add a new command line to an IBM Tivoli Storage Manager script.

The Administration Center only supports ASCII characters for input. If you need to enter characters that are not ASCII, issue the **DEFINE SCRIPT** and **UPDATE SCRIPT** commands from the server console.


**Restriction:** You cannot redirect the output of a command within a Tivoli Storage Manager script. Instead, run the script and then specify command redirection. For example, to direct the output of **script1** to the `c:\temp\test.out` directory, run the script and specify command redirection as in the following example:

```
run script1 > c:\temp\test.out
```

### Privilege class

To issue this command, the administrator must have previously defined the script or must have system privilege.

### Syntax



### Parameters

#### *script\_name* (Required)

Specifies the name of the script to be updated.

#### *command\_line*

Specifies a new or updated command to be processed in a script. You must update a command, a description, or both when you issue this command.

Command can contain substitution variables and may be continued across multiple lines if you specify a continuation character (-) as the last character in the command. You can specify up to 1200 characters for the command. Enclose the command in quotation marks if it contains blanks. If you specify this parameter, you can optionally specify the following parameter.

You have the options of running commands serially, in parallel, or serially and in parallel by specifying the **SERIAL** or **PARALLEL** script commands for this parameter. You can run multiple commands in parallel and wait for them to complete before proceeding to the next command. Commands will run serially until the parallel command is encountered.

Conditional logic flow statements can be used. These statements include IF, EXIT, and GOTO.

#### **Line**

Specifies the line number for the command. If you do not specify a line number, the command line is appended to the existing series of command lines. The appended command line is assigned a line number of five

## UPDATE SCRIPT

greater than the last command line number in the sequence. For example, if the last line in your script is 015, the appended command line is assigned a line number of 020.

If you specify a line number, the command will replace an existing line (if the number is the same as an existing line). Or the command will insert the specified line (if the line number does not correspond to an existing line number for the command line sequence).

### DEScription

Specifies a description for the script. You can specify up to 255 characters for the description. Enclose the description in quotation marks if it contains blank characters.

### Example: Add a command to the end of a script

Assume that you have defined the following three line script, named QSAMPLE, and that you want to add the **QUERY SESSION** command to the end of the script.

```
001 /* This is a sample script */
005 QUERY STATUS
010 QUERY PROCESS
update script qsample "query session"
```

After the command processes, the script now consists of the following lines:

```
001 /* This is a sample script */
005 QUERY STATUS
010 QUERY PROCESS
015 QUERY SESSION
```

### Example: Update a specific line a script

Using the script from the prior example, change line 010 so that it processes the **QUERY STGPOOL** command instead of the **QUERY PROCESS** command:

```
update script qsample "query stgpool" line=010
```

After the command processes, the script now consists of the following lines:

```
001 /* This is a sample script */
005 QUERY STATUS
010 QUERY STGPOOL
015 QUERY SESSION
```

### Example: Insert a command in the middle of a script

Using the script from the prior example, insert a new command line (**SET REGISTRATION OPEN**) after the **QUERY STATUS** command line in the QSAMPLE script:

```
update script qsample "set registration open"
line=007
```

After the command processes, the script now consists of the following lines:

```
001 /* This is a sample script */
005 QUERY STATUS
007 SET REGISTRATION OPEN
010 QUERY STGPOOL
015 QUERY SESSION
```

Refer to the *Administrator's Guide* for more information on the following topics:

- Running commands in parallel or serially
- Using logic flow statements in scripts

**Related commands***Table 436. Commands related to UPDATE SCRIPT*

| <b>Command</b> | <b>Description</b>                                      |
|----------------|---------------------------------------------------------|
| COPY SCRIPT    | Creates a copy of a script.                             |
| DEFINE SCRIPT  | Defines a script to the Tivoli Storage Manager server.  |
| DELETE SCRIPT  | Deletes the script or individual lines from the script. |
| QUERY SCRIPT   | Displays information about scripts.                     |
| RENAME SCRIPT  | Renames a script to a new name.                         |
| RUN            | Runs a script.                                          |

## UPDATE SERVER (Update a server defined for server-to-server communications)

Use this command to update a server definition.

**Important:** If this server is a source server for a virtual volume operation, changing any of these values can affect the ability of the source server to access and manage the data it has stored on the corresponding target server. Changing the server name by using the **SET SERVERNAME** command might have additional implications, varying by platform. The following are some examples:


- Passwords might be invalidated
- Device information might be affected
- Registry information about Windows platforms might change

### Privilege class

To issue this command, you must have system privilege.

### Syntax for:


- Enterprise configuration
- Enterprise event logging
- Command routing
- Storage agent
- Node replication source and target servers
- z/OS media server




### Notes:

- 1 The VALIDATEPROTOCOL parameter applies only to storage agent definitions.

### Syntax for virtual volumes





## Parameters

### *server\_name* (Required)

Specifies the name of the server to be updated. This parameter is required.

### PAssword

Specifies the password used to sign on to the target server for virtual volumes. This parameter is optional.

### SERVERPAssword

Specifies the server password, which is used for enterprise configuration, command routing, and server-to-server event logging functions. The password must match the server password set by the SET SERVERPASSWORD command. This parameter is optional.

### HLAddress

Specifies the IP address (in dotted decimal format) of the server. This parameter is optional.

### LLAddress

Specifies the low-level address of the server. This address is usually the same as the address in the TCPPOINT server option of the target server. When SSL=YES, the port must already be designated for SSL communications on the target server.

### COMMethod

Specifies the communication method used to connect to the server. This parameter is optional.

### URL

Specifies the URL address used to access this server from the Administration Center. The parameter is optional.

### DELgraceperiod

Specifies a number of days that an object remains on the target server after it has been marked for deletion. You can specify a value 0 - 9999. This parameter is optional.

### NODENAME

Specifies a node name to be used by the server to connect to the target server. This parameter is optional.

### DESCription

Specifies a description of the server. This parameter is optional. The description can be up to 255 characters. Enclose the description in quotation marks if it contains blank characters. To remove an existing description, specify a null string (").

## UPDATE SERVER

### **FORCESync**

Specifies whether to reset the server verification key when the source server next signs on to the target server. A valid verification key enables a source server to put objects on the target server, manage the grace deletion period value, and update the password, if the current password is known and the verification key matches. The parameter is optional. You can specify one of the following values:

#### **Yes**

Specifies that a new verification key will be sent to and accepted by the target server if a valid password is received.

**No** Specifies that a new verification key will not be sent to the target server.

### **VALIDateprotocol**

Specify a cyclic redundancy check to validate the data sent between the storage agent and Tivoli Storage Manager server. The parameter is optional. The default is NO. You can specify one of the following values:

**No** Specifies that data validation is not to be performed on any data sent between the storage agent and server.

#### **All**

Specifies that data validation is to be performed on all client file data, client file metadata, and the Tivoli Storage Manager server metadata that is sent between the storage agent and server. This mode affects performance as additional resources are required to calculate and compare CRC values between the storage agent and the server.

### **ALLOWReplac**

Specifies whether a server definition that was defined by a managed server can be replaced with a definition from the configuration manager. This parameter is optional. You can specify one of the following values:

#### **Yes**

Specifies that a server definition can be replaced by a definition from the configuration manager.

**No** Specifies that a server definition cannot be replaced by the definition from the configuration manager.

### **SSL**

Specifies the communication mode of the server. You can specify one of the following values:

**No** Specifies that unencrypted communication with the specified server occurs over a TCP/IP session without SSL.

#### **Yes**

Specifies an SSL session for communication with the specified server. The following conditions apply if you specify YES:

- Before starting the servers, self-signed certificates of the partner servers must be in the key database file (`cert.kdb`) of each of the servers.
- You can define multiple server names with different parameters for the same target server.
- SSL support is active if the server options file contains the `SSLTCPPORT` or `SSLTCPADMINPORT` option or if a server is defined with `SSL=YES` at startup.

**Example: Update a deletion grace period for a server**

Update the definition of SERVER2 to specify that objects remain on the target server for 10 days after they have been marked for deletion.

```
update server server2 delgraceperiod=10
```

**Example: Update the URL for a server**

Update the definition of NEWSERVER to specify its URL address to be `http://newserver:1580/`.

```
update server newserver url=http://newserver:1580/
```

**Related commands**

*Table 437. Commands related to UPDATE SERVER*

| Command           | Description                                                                            |
|-------------------|----------------------------------------------------------------------------------------|
| DEFINE DEVCLASS   | Defines a device class.                                                                |
| DEFINE SERVER     | Defines a server for server-to-server communications.                                  |
| DELETE DEVCLASS   | Deletes a device class.                                                                |
| DELETE FILESPACE  | Deletes data associated with clients file spaces.                                      |
| DELETE SERVER     | Deletes the definition of a server.                                                    |
| QUERY NODE        | Displays partial or complete information about one or more clients.                    |
| QUERY SERVER      | Displays information about servers.                                                    |
| RECONCILE VOLUMES | Reconciles source server virtual volume definitions and target server archive objects. |
| REGISTER NODE     | Defines a client node to the server and sets options for that user.                    |
| REMOVE NODE       | Removes a client from the list of registered nodes for a specific policy domain.       |
| UPDATE DEVCLASS   | Changes the attributes of a device class.                                              |
| UPDATE NODE       | Changes the attributes associated with a client node.                                  |

## UPDATE SERVERGROUP

### UPDATE SERVERGROUP (Update a server group description)

Use this command to update the description of a server group.

#### Privilege class

To issue this command, you must have system privilege.

#### Syntax

```
►►—UPDate SERVERGRoup—group_name—DESCRiption—=description—◄◄
```

#### Parameters

##### *group\_name* (Required)

Specifies the server group to update.

##### DESCRiption (Required)

Specifies a description of the server group. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains blank characters.

#### Example: Update the description of a server group

Update the description of the server group named WEST\_COMPLEX to "Western Region Complex".

```
update servergroup west_complex
description="western region complex"
```

#### Related commands

Table 438. Commands related to UPDATE SERVERGROUP

| Command            | Description                               |
|--------------------|-------------------------------------------|
| COPY SERVERGROUP   | Creates a copy of a server group.         |
| DEFINE SERVERGROUP | Defines a new server group.               |
| DELETE SERVERGROUP | Deletes a server group.                   |
| QUERY SERVERGROUP  | Displays information about server groups. |
| RENAME SERVERGROUP | Renames a server group.                   |



## UPDATE SPACETRIGGER

### SPACEexpansion

For space triggers for sequential-access FILE-type storage pools, this parameter is used in determining the number of additional volumes that are created in the storage pool. Volumes are created using the MAXCAPACITY value from the storage pool's device class. For space triggers for random-access DISK storage pools, the space trigger creates a single volume using the EXPANSIONPREFIX.

### EXPansionprefix

This specifies the prefix that the server uses to create new storage pool files. This parameter is optional and applies only to random-access DISK device classes. The default prefix is the server installation path.

The prefix can include one or more directory separator characters, for example:

```
/opt/tivoli/tsm/server/bin/
```

You can specify up to 250 characters. If you specify a prefix that is not valid, automatic expansion can fail.

This parameter is not valid for space triggers for sequential-access FILE storage pools. Prefixes are obtained from the directories specified with the associated device class.

### STGPOOL

Specifies the storage pool associated with this space trigger. If the STGPOOL parameter is not specified, the default storage pool space trigger is updated.

This parameter does not apply to storage pools with the parameter RECLAMATIONTYPE=SNAPLOCK.

### Example: Increase the amount of space for a storage pool

Increase the amount of space in a storage pool by 50 percent when it is filled to 80 percent utilization of existing volumes. Space will be created in the directories associated with the device class.

```
update spacetrigger stg spaceexpansion=50 stgpool=file
```

### Related commands

Table 439. Commands related to UPDATE SPACETRIGGER

| Command             | Description                                                     |
|---------------------|-----------------------------------------------------------------|
| DEFINE SPACETRIGGER | Defines a space trigger to expand the space for a storage pool. |
| DELETE SPACETRIGGER | Deletes the storage pool space trigger.                         |
| QUERY SPACETRIGGER  | Displays information about a storage pool space trigger.        |

## UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)


Use this command to update an existing status monitoring threshold.

Status monitoring thresholds compare the defined conditions to the status monitoring server queries and inserts the results in the status monitoring table.

Multiple thresholds can be defined for an activity. For example, you can create a threshold that provides a warning status if storage pool capacity utilization is greater than 80%. You can then create another threshold that provides error status if storage pool capacity utilization is greater than 90%.

**Note:** If a threshold is already defined for an EXISTS condition, you cannot define another threshold with one of the other condition types.

### Syntax



### Parameters

#### *threshold\_name* (Required)

Specifies the threshold name that you want to update. The name cannot exceed 48 characters in length.

#### *activity* (Required)

Specify this value to change the activity for an existing threshold. Specify one of the following values:

##### **PROCESSSUMMARY**

Specifies the number of processes that are currently active.

##### **SESSIONSUMMARY**

Specifies the number of sessions that are currently active.

##### **CLIENTSESSIONSUMMARY**

Specifies the number of client sessions that are currently active.

##### **SCHEDCLIENTSESSIONSUMMARY**

Specifies the number of scheduled client sessions.

##### **DBUTIL**

Specifies the database utilization percentage. The default warning threshold value is 80%, and the default error threshold value is 90%.

## UPDATE STATUSTHRESHOLD

### **DBFREESPACE**

Specifies the free space available in the database in gigabytes.

### **DBUSEDSPACE**

Specifies the amount of database space that is used, in gigabytes.

### **ARCHIVELOGFREESPACE**

Specifies the free space that is available in the archive log, in gigabytes.

### **STGPOOLUTIL**

Specifies the storage pool utilization percentage. The default warning threshold value is 80%, and the default error threshold value is 90%.

### **STGPOOLCAPACITY**

Specifies the storage pool capacity in gigabytes.

### **AVGSTGPOOLUTIL**

Specifies the average storage pool utilization percentage across all storage pools. The default warning threshold value is 80%, and the default error threshold value is 90%.

### **TOTSTGPOOLCAPACITY**

Specifies the total storage pool capacity in gigabytes for all available storage pools.

### **TOTSTGPOLS**

Specifies the number of defined storage pools.

### **TOTRWSTGPOLS**

Specifies the number of defined storage pools that are readable or writeable.

### **TOTNOTRWSTGPOLS**

Specifies the number of defined storage pools that are not readable or writeable.

### **STGPOOLINUSEANDDEFINED**

Specifies the total number of defined volumes that are in use.

### **ACTIVELOGUTIL**

Specifies the current percent utilization of the active log. The default warning threshold value is 80%, and the default error threshold value is 90%.

### **ARCHLOGUTIL**

Specifies the current utilization of the archive log. The default warning threshold value is 80%, and the default error threshold value is 90%.

### **CPYSTGPOOLUTIL**

Specifies the percent utilization for a copy storage pool. The default warning threshold value is 80%, and the default error threshold value is 90%.

### **PMRYSTGPOOLUTIL**

Specifies the percent utilization for a primary storage pool. The default warning threshold value is 80%, and the default error threshold value is 90%.

### **DEVCLASSPCTDRVOFFLINE**

Specifies the percent utilization of drives that are offline, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

**DEVCLASSPCTDRVPOLLING**

Specifies the drives polling, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

**DEVCLASSPCTLIBPATHSOFFLINE**

Specifies the library paths that are offline, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

**DEVCLASSPCTPATHSOFFLINE**

Specifies the percentage of device class paths that are offline, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

**DEVCLASSPCTDISKSNOTRW**

Specifies the percentage of disks that are not writable for the disk device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

**DEVCLASSPCTDISKSUNAVAILABLE**

Specifies the percentage of the disk volumes that are unavailable, by device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

**FILEDEVCLASSPCTSCRUNALLOCATABLE**

Specifies the percentage of scratch volumes that the server cannot allocate for a given non-shared file device class. The default warning threshold value is 25%, and the default error threshold value is 50%.

**Condition**

Specify this value to change the condition of an existing threshold. This parameter is optional. Specify one of the following values:

**EXists**

Creates a status monitoring indicator if the activity exists.

**GT** Creates a status monitoring indicator if the activity outcome is greater than the specified value.

**GE** Creates a status monitoring indicator if the activity outcome is greater than or equal to the specified value.

**LT** Creates a status monitoring indicator if the activity outcome is less than the specified value.

**LE** Creates a status monitoring indicator if the activity outcome is less than or equal to the specified value.

**EQual**

Creates a status monitoring indicator if the activity outcome is equal to the specified value.

**Value**

Specify this value to change the value to compare to the activity output with the specified condition. You can specify an integer from 0 to 9223372036854775807.

**Status**

Specify this value to change the status of the indicator that is created in status monitoring if the condition that is being evaluated passes. This parameter is optional. Specify one of the following values:

## UPDATE STATUSTHRESHOLD

### Normal

Specifies that the status indicator has a normal status value.

### Warning

Specifies that the status indicator has a warning status value.

### Error

Specifies that the status indicator has an error status value.

## Update an existing status threshold

Update a status threshold for average storage pool utility percentage by issuing the following command:

```
update statusthreshold avgstgpl "AVGSTGPOOLUTIL" value=90 condition=gt status=error
```

## Related commands

Table 440. Commands related to UPDATE STATUSTHRESHOLD

| Command                                                                                                              | Description                                                                  |
|----------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|
| "DELETE STATUSTHRESHOLD (Delete a status monitoring threshold)" on page 451                                          | Deletes a status monitoring threshold.                                       |
| "QUERY MONITORSTATUS (Query the monitoring status)" on page 799                                                      | Displays information about monitoring alerts and server status settings.     |
| "QUERY MONITORSETTINGS (Query the configuration settings for monitoring alerts and server status)" on page 796       | Displays information about monitoring alerts and server status settings.     |
| "QUERY STATUSTHRESHOLD (Query status monitoring thresholds)" on page 914                                             | Displays information about a status monitoring thresholds.                   |
| "SET STATUSMONITOR (Specifies whether to enable status monitoring)" on page 1144                                     | Specifies whether to enable status monitoring.                               |
| "SET STATUSATRISKINTERVAL (Specifies whether to enable client at-risk activity interval evaluation)" on page 1142    | Specifies whether to enable client at-risk activity interval evaluation.     |
| "SET STATUSREFRESHINTERVAL (Set refresh interval for status monitoring)" on page 1146                                | Specifies the refresh interval for status monitoring.                        |
| "SET STATUSSKIPASFAILURE (Specifies whether to use client at-risk skipped files as failure evaluation)" on page 1147 | Specifies whether to use client at-risk skipped files as failure evaluation. |
| "UPDATE STATUSTHRESHOLD (Update a status monitoring threshold)" on page 1375                                         | Changes the attributes of an existing status monitoring threshold.           |

## UPDATE STGPOOL (Update a storage pool)

Use this command to change a storage pool.

**Restriction:** If a client is using the simultaneous-write function and data deduplication, the data deduplication feature is disabled during backups to a storage pool.

The UPDATE STGPOOL command takes four forms. The syntax and parameters for each form are defined separately.

- “UPDATE STGPOOL (Update a primary random access storage pool)” on page 1381
- “UPDATE STGPOOL (Update a primary sequential access pool)” on page 1390
- “UPDATE STGPOOL (Update a copy sequential access storage pool)” on page 1405
- “UPDATE STGPOOL (Update an active-data sequential access)” on page 1412

*Table 441. Commands related to UPDATE STGPOOL*

| Command                   | Description                                                                                   |
|---------------------------|-----------------------------------------------------------------------------------------------|
| BACKUP STGPOOL            | Backs up a primary storage pool to a copy storage pool.                                       |
| COPY ACTIVEDATA           | Copies active backup data.                                                                    |
| DEFINE COLLOGROUP         | Defines a collocation group.                                                                  |
| DEFINE COLLOCMEMBER       | Adds a client node to a collocation group.                                                    |
| DEFINE STGPOOL            | Defines a storage pool as a named collection of server storage media.                         |
| DELETE COLLOGROUP         | Deletes a collocation group.                                                                  |
| DELETE COLLOCMEMBER       | Deletes a client node from a collocation group.                                               |
| DELETE STGPOOL            | Deletes a storage pool from server storage.                                                   |
| MOVE DRMEDIA              | Moves DRM media on-site and off-site.                                                         |
| MOVE MEDIA                | Moves storage pool volumes that are managed by an automated library.                          |
| QUERY COLLOGROUP          | Displays information about collocation groups.                                                |
| QUERY DRMEDIA             | Displays information about disaster recovery volumes.                                         |
| QUERY NODEDATA            | Displays information about the location and size of data for a client node.                   |
| QUERY SHREDSTATUS         | Displays information about data waiting to be shredded.                                       |
| QUERY STGPOOL             | Displays information about storage pools.                                                     |
| RESTORE STGPOOL           | Restores files to a primary storage pool from copy storage pools.                             |
| RESTORE VOLUME            | Restores files stored on specified volumes in a primary storage pool from copy storage pools. |
| SET DRMDBBACKUPEXPIREDAYS | Specifies criteria for database backup series expiration.                                     |

## UPDATE STGPOOL

*Table 441. Commands related to UPDATE STGPOOL (continued)*

| <b>Command</b>    | <b>Description</b>                                     |
|-------------------|--------------------------------------------------------|
| SHRED DATA        | Manually starts the process of shredding deleted data. |
| UPDATE COLLOGROUP | Updates the description of a collocation group.        |


## UPDATE STGPOOL (Update a primary random access storage pool)

Use this command to update a random access storage pool.

### Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to be updated.

### Syntax



## UPDATE STGPOOL

### Parameters

#### *pool\_name* (Required)

Specifies the storage pool to update. This parameter is required.

#### DEScRiption

Specifies a description of the storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove an existing description, specify a null string ("").

#### ACCess

Specifies how client nodes and server processes (such as migration and reclamation) can access files in the storage pool. This parameter is optional. Possible values are:

##### READWrite

Specifies that client nodes and server processes can read and write to files stored on volumes in the storage pool.

##### READOnly

Specifies that client nodes can only read files from the volumes in the storage pool.

Server processes can move files within the volumes in the storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool has been specified as a subordinate storage pool (with the **NEXTSTGPOOL** parameter) and is defined as *readonly*, the storage pool is skipped when server processes attempt to write files to the storage pool.

##### UNAVailable

Specifies that client nodes cannot access files stored on volumes in the storage pool.

Server processes can move files within the volumes in the storage pool and can also move or copy files from this storage pool to another storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool has been specified as a subordinate storage pool (with the **NEXTSTGPOOL** parameter) and is defined as *unavailable*, the storage pool is skipped when server processes attempt to write files to the storage pool.

#### MAXSIze

Specifies the maximum size for a physical file that the server can store in the storage pool. This parameter is optional. Possible values are:

##### NOLimit

Specifies that there is no maximum size limit for physical files stored in the storage pool.

##### *maximum\_file\_size*

Limits the maximum physical file size. Specify an integer from 1 to 999999, followed by a scale factor. For example, **MAXSIZE=5G** specifies that the maximum file size for this storage pool is 5 gigabytes. Scale factors are:

##### Scale factor Meaning

K kilobyte

M megabyte

G gigabyte

| Scale factor | Meaning  |
|--------------|----------|
| T            | terabyte |

If a file exceeds the maximum size and no pool is specified as the next storage pool in the hierarchy, the server does not store the file. If a file exceeds the maximum size and a pool is specified as the next storage pool, the server stores the file in the next storage pool that can accept the file size. If you specify the next storage pool parameter, at least one storage pool in your hierarchy should have no limit on the maximum size of a file. By having no limit on the size for at least one pool, you ensure that no matter what its size, the server can store the file.

For logical files that are part of an aggregate, the server considers the size of the aggregate to be the file size. Therefore, the server does not store logical files that are smaller than the maximum size limit if the files are part of an aggregate that is larger than the maximum size limit.

#### **CRCDATA**

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an **AUDIT VOLUME** command you can continually ensure the integrity of data stored in your storage hierarchy. Possible values are:

##### **Yes**

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because additional overhead is required to calculate and compare CRC values between the storage pool and the server.

**No** Specifies that data is stored without CRC information.

#### **NEXTstgpool**

Specifies a primary storage pool to which files are migrated. This parameter is optional.

To remove an existing storage pool from the storage hierarchy, specify a null string ("") for this value.

If this storage pool does not have a next storage pool, the server cannot migrate files from this storage pool and cannot store files that exceed the maximum size for this storage pool in another storage pool.

You cannot create a chain of storage pools that leads to an endless loop through the NEXTSTGPOOL parameter. At least one storage pool in the hierarchy must have no value specified for NEXTSTGPOOL.

If you specify a sequential access pool as the NEXTSTGPOOL, the pool can only be NATIVE or NONBLOCK dataformat.

#### **HIGHmig**

Specifies that the server starts migration for this storage pool when the amount of data in the pool reaches this percentage of the pool's estimated capacity. This parameter is optional. You can specify an integer from 0 to 100.

When the storage pool exceeds the high migration threshold, the server can start migration of files by node to the next storage pool, as defined with the NEXTSTGPOOL parameter. You can specify HIGHMIG=100 to prevent migration for this storage pool.

## UPDATE STGPOOL

### **LOWmig**

Specifies that the server stops migration for this storage pool when the amount of data in the pool reaches this percentage of the pool's estimated capacity. This parameter is optional. You can specify an integer from 0 to 99.

When the storage pool reaches the low migration threshold, the server does not start migration of another node's files. Because all file spaces that belong to a node are migrated together, the occupancy of the storage pool can fall below the value you specified for this parameter. You can set LOWMIG=0 to permit migration to empty the storage pool.

### **CACHE**

Specifies whether the migration process leaves a cached copy of a file in this storage pool after migrating the file to the next storage pool. This parameter is optional. Possible values are:

#### **Yes**

Specifies that caching is enabled.

**No** Specifies that caching is disabled.

Using cache may improve the retrievability of files, but may affect the performance of other processes. See the *Administrator's Guide* for details.

### **MIGProcess**

Specifies the number of processes that are used for migrating files from this storage pool. This parameter is optional. You can specify an integer from 1 to 999.

During migration, these processes are performed in parallel to provide the potential for improved migration rates.

#### **Tips:**

- The number of migration processes is dependent upon the setting of the MIGPROCESS parameter and the number of nodes or the number of collocation groups with data in the migrating storage pool. For example, if the MIGPROCESS parameter is equal to six, but there are only two nodes with data on the storage pool, migration processing only consists of two processes, not six.
- When specifying this parameter, consider whether the simultaneous-write function is enabled for server data migration. Each migration process requires a mount point and a drive for each copy storage pool and active-data pool defined to the target storage pool.

### **MIGDelay**

Specifies the minimum number of days a file must remain in a storage pool before it becomes eligible for migration. To calculate a value to compare to the specified MIGDELAY value, the server counts the number of days that the file has been in the storage pool and the number of days, if any, since the file was retrieved by a client. The lesser of the two values is compared to the specified MIGDELAY value. For example, if all the following conditions are true, a file is not migrated:

- A file has been in a storage pool for five days.
- The file was accessed by a client within the past three days.
- The value specified for the MIGDELAY parameter is four days.

This parameter is optional. You can specify an integer from 0 to 9999.

If you want the server to count the number of days based only on when a file was stored and not when it was retrieved, use the `NORETRIEVEDATE` server option.

#### **MIGContinue**

Specifies whether you allow the server to migrate files that do not satisfy the migration delay time. This parameter is optional.

Because you can require that files remain in the storage pool for a minimum number of days, the server may migrate all eligible files to the next storage pool yet not meet the low migration threshold. This parameter allows you to specify whether the server is allowed to continue the migration process by migrating files that do not satisfy the migration delay time.

Possible values are:

##### **Yes**

Specifies that, when necessary to meet the low migration threshold, the server continues to migrate files that do not satisfy the migration delay time.

If you allow more than one migration process for the storage pool, some files that do not satisfy the migration delay time may be migrated unnecessarily. As one process migrates files that satisfy the migration delay time, a second process could begin migrating files that do not satisfy the migration delay time to meet the low migration threshold. The first process that is still migrating files that satisfy the migration delay time might have, by itself, caused the low migration threshold to be met.

**No** Specifies that the server stops migration when no eligible files remain to be migrated, even before reaching the low migration threshold. The server does not migrate files unless the files satisfy the migration delay time.

#### **AUTOCopy**

Specifies when Tivoli Storage Manager writes data simultaneously to copy storage pools and active-data pools. This parameter affects the following operations:

- Client store sessions
- Server import processes
- Server data-migration processes

If an error occurs while data is being simultaneously written to a copy storage pool or active-data pool during a migration process, the server stops writing to the failing storage pools for the remainder of the process. However, the server continues to store files into the primary storage pool and any remaining copy storage pools or active-data pools. These pools remain active for the duration of the migration process. Copy storage pools are specified using the `COPYSTGPOOLS` parameter. Active-data pools are specified using the `ACTIVEDATAPOOLS` parameter.

Possible values are:

##### **None**

Specifies that the simultaneous-write function is disabled.

##### **Client**

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions or server import processes. During server import processes, data is written simultaneously to only copy storage pools. Data is not written to active-data pools during server import processes.

### **MIGRation**

Specifies that data is written simultaneously to copy storage pools and active-data pools only during migration to this storage pool. During server data-migration processes, data is written simultaneously to copy storage pools and active-data pools only if the data does not exist in those pools. Nodes whose data is being migrated must be in a domain associated with an active-data pool. If the nodes are not in a domain associated with an active pool, the data cannot be written to the pool.

### **A11**

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions, server import processes, or server data-migration processes. Specifying this value ensures that data is written simultaneously whenever this pool is a target for any of the eligible operations.

### **COPYSTGpools**

Specifies the names of copy storage pools where the server writes data simultaneously. You can specify a maximum of three copy pool names separated by commas. Spaces between the names of the copy pools are not permitted. To add or remove one or more copy storage pools, specify the pool name or names that you want to include in the updated list. For example, if the existing copy pool list includes COPY1 and COPY2 and you want to add COPY3, specify COPYSTGPOOLS=COPY1,COPY2,COPY3. To remove all existing copy storage pools associated with the primary storage pool, specify a null string ("" ) for the value (for example, COPYSTGPOOLS="").

When specifying a value for the COPYSTGPOOLS parameter, you can also specify a value for the COPYCONTINUE parameter. For additional information, see the COPYCONTINUE parameter.

The combined total number of storage pools specified in the COPYSGTPOOLS and ACTIVE DATAPOLS parameters cannot exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of copy storage pools and the COPYCONTINUE value from the primary storage pool. The primary storage pool is specified by the copy group of the management class that is bound to the data. For details, refer to information about the simultaneous-write function in the *Administrator's Guide*.

The server can write data simultaneously to copy storage pools for the following operations:

- Backup and archive operations by Tivoli Storage Manager backup-archive clients or application clients using the Tivoli Storage Manager API
- Migration operations by Tivoli Storage Manager for Space Management clients
- Import operations that involve copying exported file data from external media to a primary storage pool associated with a copy storage pool list

**Restrictions:** The simultaneous-write function is not supported for the following store operations:

- When the operation is using LAN-free data movement. Simultaneous-write operations take precedence over LAN-free operations, causing the operations to go over the LAN. However, the simultaneous-write configuration is honored.
- NAS backup operations. If the primary storage pool specified in the DESTINATION or TOCDESTINATION in the copy group of the

management class has copy storage pools defined, the copy storage pools are ignored and the data is stored into the primary storage pool only.

**Attention:** The function provided by the COPYSTGPOOLS parameter is not intended to replace the **BACKUP STGPOOL** command. If you use the COPYSTGPOOLS parameter, continue to use the **BACKUP STGPOOL** command to ensure that the copy storage pools are complete copies of the primary storage pool. There are cases when a copy might not be created. For more information, see the COPYCONTINUE parameter description.

#### **COPYContinue**

Specifies how the server should react to a copy storage pool write failure for any of the copy storage pools listed in the COPYSTGPOOLS parameter. This parameter is optional. When specifying the COPYCONTINUE parameter, either a COPYSTGPOOLS list must already exist or the COPYSTGPOOLS parameter must also be specified.

Possible values are:

#### **Yes**

If the **COPYCONTINUE** parameter is set to YES, the server will stop writing to the failing copy pools for the remainder of the session, but continue storing files into the primary pool and any remaining copy pools. The copy storage pool list is active only for the life of the client session and applies to all the primary storage pools in a particular storage pool hierarchy.

For additional information about the **COPYCONTINUE** parameter, refer to the information about the simultaneous-write function in the *Administrator's Guide*.

**No** If the **COPYCONTINUE** parameter is set to NO, the server will fail the current transaction and discontinue the store operation.

#### **Restrictions:**

- The setting of the **COPYCONTINUE** parameter does not affect active-data pools. If a write failure occurs for any of the active-data pools, the server stops writing to the failing active-data pool for the remainder of the session, but continues storing files into the primary pool and any remaining active-data pools and copy storage pools. The active-data pool list is active only for the life of the session and applies to all the primary storage pools in a particular storage pool hierarchy.
- The setting of the **COPYCONTINUE** parameter does not affect the simultaneous-write function during server import. If data is being written simultaneously and a write failure occurs to the primary storage pool or any copy storage pool, the server import process fails.
- The setting of the **COPYCONTINUE** parameter does not affect the simultaneous-write function during server data migration. If data is being written simultaneously and a write failure occurs to any copy storage pool or active-data pool, the failing storage pool is removed and the data migration process continues. Write failures to the primary storage pool cause the migration process to fail.

#### **ACTIVEDATApools**

Specifies the names of active-data pools where the server writes data simultaneously during a client backup operation. The ACTIVEDATAPOOLS parameter is optional. Spaces between the names of the active-data pools are not permitted.

The combined total number of storage pools specified in the `COPYSGTPOOLS` and `ACTIVEDATAPOOLS` parameters cannot exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of active-data pools from the destination storage pool specified in the copy group. The primary storage pool is specified by the copy group of the management class that is bound to the data. For details, refer to information about the simultaneous-write function in the *Administrator's Guide*.

The server can write data simultaneously to active-data pools only during backup operations by Tivoli Storage Manager backup-archive clients or application clients using the Tivoli Storage Manager API.

### Restrictions:

1. This parameter is available only to primary storage pools that use `NATIVE` or `NONBLOCK` data format. This parameter is not available for storage pools that use the following data formats:
  - `NETAPPDUMP`
  - `CELERRADUMP`
  - `NDMPDUMP`
2. Writing data simultaneously to active-data pools is not supported when the operation is using LAN-free data movement. Simultaneous-write operations take precedence over LAN-free operations, causing the operations to go over the LAN. However, the simultaneous-write configuration is honored.
3. Simultaneous-write operations are not supported when a NAS backup operation is writing a TOC file. If the primary storage pool specified in the `TOCDESTINATION` in the copy group of the management class has active-data pools defined, the active-data pools are ignored and the data is stored into the primary storage pool only.
4. Data being imported will not be stored in active-data pools. After an import operation, use the `COPY ACTIVEDATA` command to store the imported data in an active-data pool.

**Attention:** The function provided by the `ACTIVEDATAPOOLS` parameter is not intended to replace the `COPY ACTIVEDATA` command. If you use the `ACTIVEDATAPOOLS` parameter, use the `COPY ACTIVEDATA` command to ensure that the active-data pools contain all active data of the primary storage pool.

### SHRED

Specifies whether data will be physically overwritten when it is deleted. This parameter is optional. You can specify an integer from 0 to 10.

If you specify a value of 0, the Tivoli Storage Manager server will delete the data from the database. However, the storage used to contain the data will not be overwritten, and the data will still exist in storage until that storage is reused for other data. It might be possible to discover and reconstruct the data after it has been deleted. Changing the value (for example, resetting it to 0) will not affect data that was deleted and is waiting to be overwritten.

If you specify a value greater than 0, the Tivoli Storage Manager server deletes the data both logically and physically. The server overwrites the storage used to contain the data the specified number of times. This prevents any attempts to discover and reconstruct the data after it has been deleted.

To ensure that all copies of the data are shredded, specify a `SHRED` value greater than 0 for the storage pool specified in the `NEXTSTGPOOL` parameter,

and do not specify either the COPYSTGPOOLS or ACTIVEDATAPOOLS. Specifying relatively high values for the overwrite count will generally improve the level of security, but could affect performance adversely.

Overwriting of deleted data is performed asynchronously after the delete operation is complete. Therefore, the space occupied by the deleted data will remain occupied for some period of time and will not be available as free space for new data.

A SHRED value greater than zero cannot be used if the value of the CACHE parameter is YES. If you want to enable shredding for an existing storage pool for which caching is already enabled, you must change the value of the CACHE parameter to NO. Note that existing cached files will remain in storage so that subsequent retrieval requests can be satisfied quickly. If space is needed to store new data, the existing cached files are erased so that the space they occupied can be used for the new data. The existing cached files will not be shredded when they are erased.

**Important:** After an export operation finishes identifying files for export, any changes to the storage pool SHRED value is ignored. An export operation that is suspended retains the original SHRED value throughout the operation. You might want to consider cancelling your export operation if changes to the storage pool SHRED value jeopardize the operation. You can reissue the export command after any needed cleanup.

### **Example: Update a random access storage pool to permit caching**

Update the random access storage pool named BACKUPPOOL to permit caching when the server migrates files to the next storage pool.

```
update stgpool backuppool cache=yes
```

## UPDATE STGPOOL

### UPDATE STGPOOL (Update a primary sequential access pool)

Use this command to update a primary sequential access storage pool.


#### Restrictions:


1. You cannot use this command to change the data format for the storage pool.
2. If the value for DATAFORMAT is NETAPPDUMP, CELERRADUMP, or NDMPDUMP, you can modify only the following attributes:
  - DESCRIPTION
  - ACCESS
  - COLLOCATE
  - MAXSCRATCH
  - REUSEDELAY

#### Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to be updated.

#### Syntax





## Notes:

- 1 This parameter is not available for storage pools that use the data formats NETAPPDUMP, CELERRADUMP, or NDMPDUMP.
- 2 This parameter is valid only for storage pools that are defined with a FILE-type device class.
- 3 This parameter is only available if the value of the DEDUPLICATE parameter is YES.

## Parameters

### *pool\_name* (Required)

Specifies the name of the storage pool to be updated.

### DESCRIPTiON

Specifies a description of the storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove an existing description, specify a null string ("").

### ACCess

Specifies how client nodes and server processes (such as migration and reclamation) can access files in the storage pool. This parameter is optional. Possible values are:

#### READWrite

Specifies that client nodes and server processes can read and write to files stored on volumes in the storage pool.

#### READOnly

Specifies that client nodes can only read files from the volumes in the storage pool.

## UPDATE STGPOOL

Server processes can move files within the volumes in the storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool has been specified as a subordinate storage pool (with the **NEXTSTGPOOL** parameter) and is defined as *readonly*, the storage pool is skipped when server processes attempt to write files to the storage pool.

### **UNAVailable**

Specifies that client nodes cannot access files stored on volumes in the storage pool.

Server processes can move files within the volumes in the storage pool and can also move or copy files from this storage pool to another storage pool. However, no new writes are permitted to volumes in the storage pool from volumes outside the storage pool.

If this storage pool has been specified as a subordinate storage pool (with the **NEXTSTGPOOL** parameter) and is defined as *unavailable*, the storage pool is skipped when server processes attempt to write files to the storage pool.

### **MAXSize**

Specifies the maximum size for a physical file that the server can store in the storage pool. This parameter is optional. Possible values are:

#### **NOLimit**

Specifies that there is no maximum size limit for physical files stored in the storage pool.

#### *maximum\_file\_size*

Limits the maximum physical file size. Specify an integer from 1 to 999999, followed by a scale factor. For example, **MAXSIZE=5G** specifies that the maximum file size for this storage pool is 5 gigabytes. Scale factors are:

#### **Scale factor Meaning**

|   |          |
|---|----------|
| K | kilobyte |
| M | megabyte |
| G | gigabyte |
| T | terabyte |

If a file exceeds the maximum size and no pool is specified as the next storage pool in the hierarchy, the server does not store the file. If a file exceeds the maximum size and a pool is specified as the next storage pool, the server stores the file in the next storage pool that can accept the file size. If you specify the next storage pool parameter, at least one storage pool in your hierarchy should have no limit on the maximum size of a file. By having no limit on the size for at least one pool, you ensure that no matter what its size, the server can store the file.

For logical files that are part of an aggregate, the server considers the size of the aggregate to be the file size. Therefore, the server does not store logical files that are smaller than the maximum size limit if the files are part of an aggregate that is larger than the maximum size limit.

**Restriction:** This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

**CRCData**

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is only valid for NATIVE data format storage pools. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an **AUDIT VOLUME** command you can continually ensure the integrity of data stored in your storage hierarchy. Possible values are:

**Yes**

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because additional overhead is required to calculate and compare CRC values between the storage pool and the server.

**No** Specifies that data is stored without CRC information.

**Tip:** For storage pools that are associated with the 3592, LTO, or ECARTRIDGE device type, logical block protection provides better protection against data corruption than CRC validation for a storage pool. If you specify CRC validation for a storage pool, data is validated only during volume auditing operations. Errors are identified after data is written to tape.

To enable logical block protection, specify a value of READWRITE for the **LBPROTECT** parameter on the **DEFINE DEVCLASS** and **UPDATE DEVCLASS** commands for the 3592, LTO, or ECARTRIDGE device types. Logical block protection is supported only on the following types of drives and media:

- IBM LTO5 and later.
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.
- Oracle StorageTek T10000C drives.

**NEXTstgpool**

Specifies a primary storage pool to which files are migrated. You cannot migrate data from a sequential access storage pool to a random access storage pool. This parameter is optional. The next storage pool must be a primary storage pool.

To remove an existing value, specify a null string ("").

If this storage pool does not have a next storage pool, the server cannot migrate files from this storage pool and cannot store files that exceed the maximum size for this storage pool in another storage pool.

When there is insufficient space available in the current storage pool, the NEXTSTGPOOL parameter for sequential access storage pools does not allow data to be stored into the next pool. In this case the server issues a message and the transaction fails.

For next storage pools with a device type of FILE, the server performs a preliminary check to determine whether sufficient space is available. If space is not available, the server skips to the next storage pool in the hierarchy. If space is available, the server attempts to store data in that pool. However, it is possible that the storage operation could fail because, at the time the actual storage operation is attempted, the space is no longer available.

You cannot create a chain of storage pools that leads to an endless loop through the NEXTSTGPOOL parameter. At least one storage pool in the hierarchy must have no value specified for NEXTSTGPOOL.

## UPDATE STGPOOL

**Restriction:** This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

If you specify a sequential access pool as the NEXTSTGPOOL, the pool can only be NATIVE or NONBLOCK dataformat.

### **HIghmig**

Specifies that the server starts migration when storage pool utilization reaches this percentage. For sequential-access disk (FILE) storage pools, utilization is the ratio of data in a storage pool to the pool's total estimated data capacity, including the capacity of all scratch volumes specified for the pool. For storage pools that use tape or optical media, utilization is the ratio of volumes that contain data to the total number of volumes in the storage pool. The total number of volumes includes the maximum number of scratch volumes. This parameter is optional. You can specify an integer from 0 to 100.

When the storage pool exceeds the high migration threshold, the server can start migration of files by volume to the next storage pool defined for the storage pool. You can set the high migration threshold to 100 to prevent migration for the storage pool.

**Restriction:** This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

### **LOwmig**

Specifies that the server stops migration when storage pool utilization is at or below this percentage. For sequential-access disk (FILE) storage pools, utilization is the ratio of data in a storage pool to the pool's total estimated data capacity, including the capacity of all scratch volumes specified for the pool. For storage pools that use tape or optical media, utilization is the ratio of volumes that contain data to the total number of volumes in the storage pool. The total number of volumes includes the maximum number of scratch volumes. This parameter is optional. You can specify an integer from 0 to 99.

When the storage pool reaches the low migration threshold, the server does not start migration of files from another volume. You can set the low migration threshold to 0 to permit migration to empty the storage pool.

**Restriction:** This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

### **REClaim**

Specifies when the server reclaims a volume, based on the percentage of reclaimable space on a volume. Reclaimable space is the amount of space occupied by files that are expired or deleted from the Tivoli Storage Manager database.

Reclamation makes the fragmented space on volumes usable again by moving any remaining unexpired files from one volume to another volume, thus making the original volume available for reuse. This parameter is optional. You can specify an integer from 1 to 100.

When determining which volumes in a storage pool to reclaim, the Tivoli Storage Manager server first determines the reclamation threshold indicated by the **RECLAIM**. The server then examines the percentage of reclaimable space for each volume in the storage pool. If the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool, the volume is a candidate for reclamation.

For example, suppose storage pool FILEPOOL has a reclamation threshold of 70 percent. This value indicates that the server can reclaim any volume in the storage pool that has a percentage of reclaimable space that is greater than 70 percent. The storage pool has three volumes:

- FILEVOL1 with 65 percent reclaimable space
- FILEVOL2 with 80 percent reclaimable space
- FILEVOL3 with 95 percent reclaimable space

When reclamation begins, the server compares the percent of reclaimable space for each volume with the reclamation threshold of 70 percent. In this example, FILEVOL2 and FILEVOL3 are candidates for reclamation because their percentages of reclaimable space are greater than 70. To determine the percentage of reclaimable space for a volume, issue the **QUERY VOLUME** command and specify **FORMAT=DETAILED**. The value in the field Pct. Reclaimable Space is the percentage of reclaimable space for the volume.

Specify a value of 50 percent or greater for this parameter so that files stored on two volumes can be combined onto a single output volume.

**Restriction:** This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

#### **RECLAIMProcess**

Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool. This parameter is optional. Enter a value from 1 to 999.

When calculating the value for this parameter, consider the number of sequential storage pools that will be involved with the reclamation and the number of logical and physical drives that can be dedicated to the operation. To access a sequential access volume, IBM Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the reclamation.

For example, suppose that you want to reclaim the volumes from two sequential storage pools simultaneously and that you want to specify four processes for each of the storage pools. The storage pools have the same device class. Assuming that the RECLAIMSTGPOOL parameter is not specified or that the reclaim storage pool has the same device class as the storage pool being reclaimed, each process requires two mount points and, if the device type is not FILE, two drives. (One of the drives is for the input volume, and the other drive is for the output volume.) To run eight reclamation processes

simultaneously, you need a total of at least 16 mount points and 16 drives. The device class for the two storage pools must have a mount limit of at least 16.

If the number of reclamation processes you specify is more than the number of available mount points or drives, the processes that do not obtain mount points or drives will wait for mount points or drives to become available. If mount points or drives do not become available within the MOUNTWAIT time, the reclamation processes will end. For information about specifying the MOUNTWAIT time, see “DEFINE DEVCLASS (Define a device class)” on page 148.

The Tivoli Storage Manager server will start the specified number of reclamation processes regardless of the number of volumes that are eligible for reclamation. For example, if you specify ten reclamation processes and only six volumes are eligible for reclamation, the server will start ten processes and four of them will complete without processing a volume.

**Restriction:** This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

### RECLAIMSTGpool

Specifies another primary storage pool as a target for reclaimed data from this storage pool. This parameter is optional. When the server reclaims volumes for the storage pool, unexpired data is moved from the volumes being reclaimed to the storage pool named with this parameter.

To remove an existing value, specify a null string (“”).

A reclaim storage pool is most useful for a storage pool that has only one drive in its library. When you specify this parameter, the server moves all data from reclaimed volumes to the reclaim storage pool regardless of the number of drives in the library.

To move data from the reclaim storage pool back to the original storage pool, use the storage pool hierarchy. Specify the original storage pool as the next storage pool for the reclaim storage pool.

**Restriction:** This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

### COLlocate

Specifies whether the server attempts to keep data belonging to a single client node, group of client nodes, or client file space assigned to as few volumes as possible. This parameter is optional.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time needed to collocate files for storing and the number of volumes required. For details, see the *Administrator's Guide*.

Possible values are:

**No** Specifies that collocation is disabled.

**GRoup**

Specifies that collocation is enabled at the group level for client nodes. The server attempts to put data for nodes that belong to the same collocation group on as few volumes as possible. If the nodes in the collocation group have multiple file spaces, the server does not attempt to collocate those file spaces.

If you specify COLLOCATE=GROUP but do not define any collocation groups or if you specify COLLOCATE=GROUP but do not add nodes to a collocation group, data is collocated by node. Be sure to consider tape usage when organizing client nodes into collocation groups. For example, if a tape-based storage pool consists of data from grouped and ungrouped nodes and you specify COLLOCATE=GROUP, the server performs the following actions:

- Collocates by group the data for grouped nodes only. Whenever possible, the server collocates data that belongs to a group of nodes on a single tape or on as few tapes as possible. Data for a single node can also be spread across several tapes associated with a group.
- Collocates by node the data for ungrouped nodes. Whenever possible, the server stores the data for a single node on a single tape. All available tapes that already have data for the node are used before available space on any other tape is used.

**NODe**

Specifies that collocation is enabled at the client node level. The server attempts to put data for one node on as few volumes as possible. If the node has multiple file spaces, the server does not attempt to collocate those file spaces. For backward compatibility, COLLOCATE=YES is still accepted by the server to specify collocation at the client node level.

If a storage pool contains data for a node that is a member of a collocation group and you specify COLLOCATE=NODE, the data will be collocated by node not by group.

**FIlespace**

Specifies that collocation is enabled at the file space level for client nodes. The server attempts to put data for one node and file space on as few volumes as possible. If a node has multiple file spaces, the server attempts to put data for different file spaces on different volumes.

**MAXSCRatch**

Specifies the maximum number of scratch volumes that the server can request. This parameter is optional. You can specify an integer from 0 to 100000000. By allowing the server to request scratch volumes, you avoid having to define each volume to be used.

The value specified for this parameter is used to estimate the total number of volumes available in the storage pool and the corresponding estimated capacity for the storage pool.

Scratch volumes are automatically deleted from the storage pool when they become empty. When scratch volumes with the device type of FILE are deleted, the space that the volumes occupied is freed by the server and returned to the file system.

**Tip:** For server-to-server operations that utilize virtual volumes and that store a small amount of data, consider specifying a value for the MAXSCRATCH parameter that is higher than the value you typically specify for write operations to other types of volumes. After a write operation to a virtual

## UPDATE STGPOOL

volume, Tivoli Storage Manager marks the volume as FULL, even if the value of the MAXCAPACITY parameter on the device-class definition has not been reached. The Tivoli Storage Manager server does not keep virtual volumes in FILLING status and does not append to them. If the value of the MAXSCRATCH parameter is too low, server-to-server operations can fail.

### **REUsedelay**

Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer from 0 to 9999. The value 0 means that a volume can be rewritten or returned to the scratch pool as soon as all files are deleted from the volume.

By specifying this parameter, you can ensure that the database could be restored to an earlier level and database references to files in the storage pool would still be valid.

### **OVFLocation**

Specifies the overflow location for the storage pool. The server assigns this location name to a volume that is ejected from the library by the **MOVE MEDIA** command. This parameter is optional. The location name can be a maximum length of 255 characters. Enclose the location name in quotation marks if the location name contains any blank characters.

To remove an existing value, specify a null string ("").

**Restriction:** This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

### **MIGDelay**

Specifies the minimum number of days a file must remain in a storage pool before it becomes eligible for migration. All files on a volume must be eligible for migration before the server selects the volume for migration. To calculate a value to compare to the specified MIGDELAY, the server counts the number of days that the file has been in the storage pool.

This parameter is optional. You can specify an integer from 0 to 9999.

If you want the server to count the number of days based only on when a file was stored and not when it was retrieved, use the NORETRIEVEDATE server option.

**Restriction:** This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

### **MIGContinue**

Specifies whether you allow the server to migrate files that do not satisfy the migration delay time. This parameter is optional.

Because you can require that files remain in the storage pool for a minimum number of days, the server may migrate all eligible files to the next storage pool yet not meet the low migration threshold. This parameter allows you to

specify whether the server is allowed to continue migration by migrating files that do not satisfy the migration delay time.

Possible values are:

**Yes**

Specifies that, when necessary to meet the low migration threshold, the server continues to migrate files that have not been stored in the storage pool for the number of days specified by the migration delay period.

**No** Specifies that the server stops migration when no eligible files remain to be migrated, even before reaching the low migration threshold. The server does not migrate files unless the files have been stored in the storage pool for the number of days specified by the migration delay period.

**Restriction:** This parameter is not available for storage pools that use the following data formats:

- NETAPPDUMP
- CELERRADUMP
- NDMPDUMP

**MIGProcess**

Specifies the number of parallel processes to use for migrating the files from the volumes in this storage pool. This parameter is optional. Enter a value from 1 to 999.

When calculating the value for this parameter, consider the number of sequential storage pools that will be involved with the migration, and the number of logical and physical drives that can be dedicated to the operation. To access a sequential-access volume, Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the migration.

For example, suppose you want to simultaneously migrate the files from volumes in two primary sequential storage pools and that you want to specify three processes for each of the storage pools. The storage pools have the same device class. Assuming that the storage pool to which files are being migrated has the same device class as the storage pool from which files are being migrated, each process requires two mount points and, if the device type is not FILE, two drives. (One drive is for the input volume, and the other drive is for the output volume.) To run six migration processes simultaneously, you need a total of at least 12 mount points and 12 drives. The device class for the storage pools must have a mount limit of at least 12.

If the number of migration processes you specify is more than the number of available mount points or drives, the processes that do not obtain mount points or drives will wait for mount points or drives to become available. If mount points or drives do not become available within the MOUNTWAIT time, the migration processes will end. For information about specifying the MOUNTWAIT time, see “DEFINE DEVCLASS (Define a device class)” on page 148.

The Tivoli Storage Manager server will start the specified number of migration processes regardless of the number of volumes that are eligible for migration. For example, if you specify ten migration processes and only six volumes are eligible for migration, the server will start ten processes and four of them will complete without processing a volume.

## UPDATE STGPOOL

**Note:** When specifying this parameter, consider whether the simultaneous-write function is enabled for server data migration. Each migration process requires a mount point and a drive for each copy storage pool and active-data pool defined to the target storage pool.

### **AUTOCopy**

Specifies when Tivoli Storage Manager performs simultaneous-write operations. This parameter affects the following operations:

- Client store sessions
- Server import processes
- Server data-migration processes

If an error occurs while data is being simultaneously written to a copy storage pool or active-data pool during a migration process, the server stops writing to the failing storage pools for the remainder of the process. However, the server continues to store files into the primary storage pool and any remaining copy storage pools or active-data pools. These pools remain active for the duration of the migration process. Copy storage pools are specified using the **COPYSTGPOOLS** parameter. Active-data pools are specified using the **ACTIVEDATAPOOLS** parameter.

Possible values are:

#### **None**

Specifies that the simultaneous-write function is disabled.

#### **CLient**

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions or server import processes. During server import processes, data is written simultaneously to only copy storage pools. Data is not written to active-data pools during server import processes.

#### **MIGRation**

Specifies that data is written simultaneously to copy storage pools and active-data pools only during migration to this storage pool. During server data-migration processes, data is written simultaneously to copy storage pools and active-data pools only if the data does not exist in those pools. Nodes whose data is being migrated must be in a domain associated with an active-data pool. If the nodes are not in a domain associated with an active pool, the data cannot be written to the pool.

#### **A11**

Specifies that data is written simultaneously to copy storage pools and active-data pools during client store sessions, server import processes, or server data-migration processes. Specifying this value ensures that data is written simultaneously whenever this pool is a target for any of the eligible operations.

### **COPYSTGpools**

Specifies the names of copy storage pools where the server simultaneously writes data. You can specify a maximum of three copy pool names separated by commas. Spaces between the names of the copy pools are not permitted. To add or remove one or more copy storage pools, specify the pool name or names that you want to include in the updated list. For example, if the existing copy pool list includes COPY1 and COPY2 and you want to add COPY3, specify **COPYSTGPOOLS=COPY1,COPY2,COPY3**. To remove all existing copy storage pools associated with the primary storage pool, specify a null string (""), for the value (for example, **COPYSTGPOOLS=""**).

When specifying a value for the COPYSTGPOOLS parameter, you can also specify a value for the COPYCONTINUE parameter. For additional information, see the COPYCONTINUE parameter.

The combined total number of storage pools specified in the COPYSSTGPOOLS and ACTIVEDATAPOOLS parameters cannot exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of copy storage pools and the COPYCONTINUE value from the primary storage pool. The primary storage pool is specified by the copy group of the management class that is bound to the data. For details, refer to information about the simultaneous-write function in the *Administrator's Guide*.

The server can write data simultaneously to copy storage pools during the following operations:

- Backup and archive operations by Tivoli Storage Manager backup-archive clients or application clients using the Tivoli Storage Manager API
- Migration operations by Tivoli Storage Manager for Space Management clients
- Import operations that involve copying exported file data from external media to a primary storage pool associated with a copy storage pool list

#### Restrictions:

1. This parameter is available only to primary storage pools that use NATIVE or NONBLOCK data format. This parameter is not available for storage pools that use the following data formats:
  - NETAPPDUMP
  - CELERRADUMP
  - NDMPDUMP
2. Simultaneous-write operations takes precedence over LAN-free data movement, causing the operations go over the LAN. However, the simultaneous-write configuration is honored.
3. The simultaneous-write function is not supported for NAS backup operations. If the primary storage pool specified in the DESTINATION or TOCDESTINATION in the copy group of the management class has copy storage pools defined, the copy storage pools are ignored and the data is stored into the primary storage pool only.

**Attention:** The function provided by the COPYSTGPOOLS parameter is not intended to replace the **BACKUP STGPOOL** command. If you use the COPYSTGPOOLS parameter, continue to use the **BACKUP STGPOOL** command to ensure that the copy storage pools are complete copies of the primary storage pool. There are cases when a copy may not be created. For more information, see the COPYCONTINUE parameter description.

#### COPYContinue

Specifies how the server should react to a copy storage pool write failure for any of the copy storage pools listed in the COPYSTGPOOLS parameter. This parameter is optional. The default is YES. When specifying the COPYCONTINUE parameter, either a COPYSTGPOOLS list must already exist or the COPYSTGPOOLS parameter must also be specified.

The COPYCONTINUE parameter has no effect on the simultaneous-write function during migration.

Possible values are:

### Yes

If the **COPYCONTINUE** parameter is set to YES, the server will stop writing to the failing copy pools for the remainder of the session, but continue storing files into the primary pool and any remaining copy pools. The copy storage pool list is active only for the life of the client session and applies to all the primary storage pools in a particular storage pool hierarchy.

For additional information about the **COPYCONTINUE** parameter, refer to the information about the simultaneous-write function in the *Administrator's Guide*.

**No** If the **COPYCONTINUE** parameter is set to NO, the server will fail the current transaction and discontinue the store operation.

### Restrictions:

- The setting of the **COPYCONTINUE** parameter does not affect active-data pools. If a write failure occurs for any of the active-data pools, the server stops writing to the failing active-data pool for the remainder of the session, but continues storing files into the primary pool and any remaining active-data pools and copy storage pools. The active-data pool list is active only for the life of the session and applies to all the primary storage pools in a particular storage pool hierarchy.
- The setting of the **COPYCONTINUE** parameter does not affect the simultaneous-write function during server import. If data is being written simultaneously and a write failure occurs to the primary storage pool or any copy storage pool, the server import process fails.
- The setting of the **COPYCONTINUE** parameter does not affect the simultaneous-write function during server data migration. If data is being written simultaneously and a write failure occurs to any copy storage pool or active-data pool, the failing storage pool is removed and the data migration process continues. Write failures to the primary storage pool cause the migration process to fail.

### ACTIVEDATApools

Specifies the names of active-data pools where the server simultaneously writes data during a client backup operation. The **ACTIVEDATAPOOLS** parameter is optional. Spaces between the names of the active-data pools are not permitted.

The combined total number of storage pools specified in the **COPYSGTPOOLS** and **ACTIVEDATAPOOLS** parameters can not exceed three.

When a data storage operation switches from a primary storage pool to a next storage pool, the next storage pool inherits the list of active-data pools from the destination storage pool specified in the copy group. The primary storage pool is specified by the copy group of the management class that is bound to the data. For details, refer to information about the simultaneous-write function in the *Administrator's Guide*.

The server can write data simultaneously to active-data pools only during backup operations by Tivoli Storage Manager backup-archive clients or application clients using the Tivoli Storage Manager API.

### Restrictions:

1. This parameter is available only to primary storage pools that use **NATIVE** or **NONBLOCK** data format. This parameter is not available for storage pools that use the following data formats:
  - **NETAPPDUMP**

- CELERRADUMP
  - NDMPDUMP
2. Writing data simultaneously to active-data pools is not supported when the operation is using LAN-free data movement. Simultaneous-write operations take precedence over LAN-free data movement, causing the operations go over the LAN. However, the simultaneous-write configuration is honored.
  3. The simultaneous-write function is not supported when a NAS backup operation is writing a TOC file. If the primary storage pool specified in the TOCDESTINATION in the copy group of the management class has active-data pools defined, the active-data pools are ignored and the data is stored into the primary storage pool only.
  4. Data being imported will not be stored in active-data pools. After an import operation, use the **COPY ACTIVEDATA** command to store the imported data in an active-data pool.

**Attention:** The function provided by the ACTIVEDATAPOOLS parameter is not intended to replace the **COPY ACTIVEDATA** command. If you use the ACTIVEDATAPOOLS parameter, use the **COPY ACTIVEDATA** command to ensure that the active-data pools contain all active data of the primary storage pool.

#### DEDUPLICATE

Specifies whether the data that is stored in this storage pool will be deduplicated. This parameter is optional and is valid only for storage pools that are defined with a FILE device class.

#### IDENTIFYPROCESS

Specifies the number of parallel processes to use for server-side duplicate identification. This parameter is optional and is valid only for storage pools that are defined with a device class associated with the FILE device type. Enter a value from 1 to 20.

When calculating the value for this parameter, consider the workload on the server and the amount of data requiring data deduplication. Server-side duplicate identification requires disk I/O and processor resources, so the more processes you allocate to data deduplication, the heavier the workload that you place on your system. In addition, consider the number of volumes that require processing. Server-side duplicate-identification processes work on volumes containing data that requires deduplication. If you update a storage pool, specifying that the data in the storage pool is to be deduplicated, all the volumes in the pool require processing. For this reason, you might have to define a high number of duplicate-identification processes initially. Over time, however, as existing volumes are processed, only the volumes containing new data have to be processed. When that happens, you can reduce the number of duplicate-identification processes.

**Remember:** Duplicate-identification processes can be either active or idle. Processes that are working on files are active. Processes that are waiting for files to work on are idle. Processes remain idle until volumes with data to be deduplicated become available. The output of the **QUERY PROCESS** command for a duplicate-identification process includes the total number of bytes and files that have been processed since the process first started. For example, if a duplicate-identification process processes four files, becomes idle, and then processes five more files, then the total number of files processed is nine. Processes end only when canceled or when the number of duplicate-identification processes for the storage pool is changed to a value less than the number currently specified.

## UPDATE STGPOOL

### **Example: Update the primary sequential storage pool's mountable scratch volumes**

Update the primary sequential storage pool named TAPEPOOL1 to permit as many as 10 scratch volumes to be mounted.

```
update stgpool tapepool1 maxscratch=10
```


## UPDATE STGPOOL (Update a copy sequential access storage pool)

Use this command to update a copy sequential access storage pool.

### Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to be updated.

### Syntax



### Notes:

- 1 This parameter is valid only for storage pools that are defined with a FILE-type device class.
- 2 This parameter is only available if the value of the DEDUPLICATE parameter is YES.

### Parameters

#### *pool\_name* (Required)

Specifies the name of the copy storage pool to be updated.

#### DESCRiption

Specifies a description of the copy storage pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove an existing description, specify a null string ("").

## UPDATE STGPOOL

### ACCess

Specifies how client nodes and server processes (such as reclamation) can access files in the copy storage pool. This parameter is optional. Possible values are:

#### READWrite

Specifies that files can be read from and written to the volumes in the copy storage pool.

#### READOnly

Specifies that client nodes can only read files stored on the volumes in the copy storage pool.

Server processes can move files within the volumes in the storage pool. The server can use files in the copy storage pool to restore files to primary storage pools. However, no new writes are permitted to volumes in the copy storage pool from volumes outside the storage pool. A storage pool cannot be backed up to the copy storage pool.

#### UNAVailable

Specifies that client nodes cannot access files stored on volumes in the copy storage pool.

Server processes can move files within the volumes in the storage pool. The server can use files in the copy storage pool to restore files to primary storage pools. However, no new writes are permitted to volumes in the copy storage pool from volumes outside the storage pool. A storage pool cannot be backed up to the copy storage pool.

### COLlocate

Specifies whether the server attempts to keep data belonging to a single client node, group of client nodes, or client file space assigned to as few volumes as possible. This parameter is optional.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time needed to collocate files for storing and the number of volumes required. For details, see the *Administrator's Guide*.

Possible values are:

**No** Specifies that collocation is disabled.

#### GRoup

Specifies that collocation is enabled at the group level for client nodes. The server attempts to put data for nodes that belong to the same collocation group on as few volumes as possible. If the nodes in the collocation group have multiple file spaces, the server does not attempt to collocate those file spaces.

If you specify COLLOCATE=GROUP but do not define any collocation groups or if you specify COLLOCATE=GROUP but do not add nodes to a collocation group, data is collocated by node. Be sure to consider tape usage when organizing client nodes into collocation groups. For example, if a tape-based storage pool consists of data from grouped and ungrouped nodes and you specify COLLOCATE=GROUP, the server performs the following actions:

- Collocates by group the data for grouped nodes only. Whenever possible, the server collocates data that belongs to a group of nodes on a single tape or on as few tapes as possible. Data for a single node can also be spread across several tapes associated with a group.

- Collocates by node the data for ungrouped nodes. Whenever possible, the server stores the data for a single node on a single tape. All available tapes that already have data for the node are used before available space on any other tape is used.

**NODe**

Specifies that collocation is enabled at the client node level. The server attempts to put data for one node on as few volumes as possible. If the node has multiple file spaces, the server does not attempt to collocate those file spaces. For backward compatibility, COLLOCATE=YES is still accepted by the server to specify collocation at the client node level.

If a storage pool contains data for a node that is a member of a collocation group and you specify COLLOCATE=NODE, the data will be collocated by node not by group.

**FILEspace**

Specifies that collocation is enabled at the file space level for client nodes. The server attempts to put data for one node and file space on as few volumes as possible. If a node has multiple file spaces, the server attempts to put data for different file spaces on different volumes.

**RECLaim**

Specifies when the server reclaims a volume, based on the percentage of reclaimable space on a volume. Reclaimable space is the amount of space occupied by files that are expired or deleted from the Tivoli Storage Manager database.

Reclamation makes the fragmented space on volumes usable again by moving any remaining active files from one volume to another volume, thus making the original volume available for reuse. This parameter is optional. You can specify an integer from 1 to 100. The value 100 means that reclamation is not performed.

When determining which volumes in a storage pool to reclaim, the Tivoli Storage Manager server first determines the reclamation threshold indicated by the **RECLAIM**. The server then examines the percentage of reclaimable space for each volume in the storage pool. If the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool, the volume is a candidate for reclamation.

For example, suppose storage pool FILEPOOL has a reclamation threshold of 70 percent. This value indicates that the server can reclaim any volume in the storage pool that has a percentage of reclaimable space that is greater than 70 percent. The storage pool has three volumes:

- FILEVOL1 with 65 percent reclaimable space
- FILEVOL2 with 80 percent reclaimable space
- FILEVOL3 with 95 percent reclaimable space

When reclamation begins, the server compares the percent of reclaimable space for each volume with the reclamation threshold of 70 percent. In this example, FILEVOL2 and FILEVOL3 are candidates for reclamation because their percentages of reclaimable space are greater than 70. To determine the percentage of reclaimable space for a volume, issue the **QUERY VOLUME** command and specify **FORMAT=DETAILED**. The value in the field Pct. Reclaimable Space is the percentage of reclaimable space for the volume.

If you change the value from the default of 100, specify a value of 50 percent or greater so that files stored on two volumes can be combined onto a single output volume.

When a copy pool volume that is offsite becomes eligible for reclamation, the reclamation process attempts to obtain the active files on the reclaimable volume from a primary or copy storage pool that is onsite. The process then writes these files to an available volume in the original copy storage pool. Effectively, these files are moved back to the onsite location. However, the files could be obtained from the offsite volume after a disaster if a database backup is used that references the files on the offsite volume. Because of the way reclamation works with offsite volumes, use it carefully with copy storage pools.

### **RECLAIMProcess**

Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool. This parameter is optional. Enter a value from 1 to 999.

When calculating the value for this parameter, consider the number of sequential storage pools that will be involved with the reclamation and the number of logical and physical drives that can be dedicated to the operation. To access a sequential access volume, IBM Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the reclamation.

For example, suppose that you want to reclaim the volumes from two sequential storage pools simultaneously and that you want to specify four processes for each of the storage pools. The storage pools have the same device class. Each process requires two mount points and, if the device type is not FILE, two drives. (One of the drives is for the input volume, and the other drive is for the output volume.) To run eight reclamation processes simultaneously, you need a total of at least 16 mount points and 16 drives. The device class for each storage pool must have a mount limit of at least eight.

If the number of reclamation processes you specify is more than the number of available mount points or drives, the processes that do not obtain mount points or drives will wait for mount points or drives to become available. If mount points or drives do not become available within the MOUNTWAIT time, the reclamation processes will end. For information about specifying the MOUNTWAIT time, see “DEFINE DEVCLASS (Define a device class)” on page 148.

The Tivoli Storage Manager server will start the specified number of reclamation processes regardless of the number of volumes that are eligible for reclamation. For example, if you specify ten reclamation processes and only six volumes are eligible for reclamation, the server will start ten processes and four of them will complete without processing a volume.

### **OFFSITERECLAIMLimit**

Specifies the number of offsite volumes to have their space reclaimed during reclamation for this storage pool. This parameter is optional. Possible values are:

#### **NOLimit**

Specifies that you want to have the space reclaimed in all of your offsite volumes.

#### *number*

Specifies the number of offsite volumes to have their space reclaimed. You can specify an integer from 0 to 99999. A value of zero means that none of the offsite volumes will be reclaimed.

**Important:** When determining the value for the `OFFSITERECLAIMLIMIT`, consider using the statistical information in the message issued at the end of the offsite volume reclamation operation. Alternatively, you can use the following Tivoli Storage Manager SQL select command to obtain the statistical information from the `SUMMARY` table for the offsite volume reclamation operation:

```
select * from summary where activity='OFFSITE RECLAMATION'
```

The statistical information includes the following items:

- The number of offsite volumes that were processed
- The number of parallel processes that were used
- The total amount of time required for the processing

The order in which offsite volumes are reclaimed is based on the amount of unused space in a volume. (Unused space includes both space that has never been used on the volume and space that has become empty because of file deletion.) Volumes with the largest amount of unused space are reclaimed first.

For example, suppose a copy storage pool contains three volumes: `VOL1`, `VOL2`, and `VOL3`. `VOL1` has the largest amount of unused space, and `VOL3` has the least amount of unused space. Suppose further that the percentage of unused space in each of the three volumes is greater than the value of the `RECLAIM` parameter. If you do not specify a value for the `OFFSITERECLAIMLIMIT` parameter, all three volumes will be reclaimed when the reclamation runs. If you specify a value of 2, only `VOL1` and `VOL2` will be reclaimed when the reclamation runs. If you specify a value of 1, only `VOL1` will be reclaimed.

#### **MAXSCRatch**

Specifies the maximum number of scratch volumes that the server can request for this storage pool. This parameter is optional. You can specify an integer from 0 to 100000000. By allowing the server to request scratch volumes as needed, you avoid having to define each volume to be used.

The value specified for this parameter is used to estimate the total number of volumes available in the copy storage pool and the corresponding estimated capacity for the copy storage pool.

Scratch volumes are automatically deleted from the storage pool when they become empty. However, if the access mode for a scratch volume is `OFFSITE`, the volume is not deleted from the copy storage pool until the access mode is changed. This allows an administrator to query the server for empty, offsite scratch volumes and return these to the onsite location.

When scratch volumes with the device type of `FILE` become empty and are deleted, the space that the volumes occupied is freed by the server and returned to the file system.

**Tip:** For server-to-server operations that utilize virtual volumes and that store a small amount of data, consider specifying a value for the `MAXSCRATCH` parameter that is higher than the value you typically specify for write operations to other types of volumes. After a write operation to a virtual volume, Tivoli Storage Manager marks the volume as `FULL`, even if the value of the `MAXCAPACITY` parameter on the device-class definition has not been reached. The Tivoli Storage Manager server does not keep virtual volumes in `FILLING` status and does not append to them. If the value of the `MAXSCRATCH` parameter is too low, server-to-server operations can fail.

## UPDATE STGPOOL

### REUsedelay

Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer from 0 to 9999. A value of 0 means that a volume can be rewritten or returned to the scratch pool as soon as all files are deleted from the volume.

**Important:** Use this parameter to help ensure that when you restore the database to an earlier level, database references to files in the copy storage pool are still valid. You must set this parameter to a value greater than the number of days you plan to retain the oldest database backup. The number of days specified for this parameter should be the same as the number specified for the SET DRMDBBACKUPEXPIREDDAYS command. For more information, see the *Administrator's Guide*.

### OVFLocation

Specifies the overflow location for the storage pool. The server assigns this location name to a volume that is ejected from the library by the MOVE MEDIA command. This parameter is optional. The location name can be a maximum length of 255 characters. Enclose the location name in quotation marks if the location name contains any blank characters.

To remove an existing value, specify a null string ("").

### CRCDATA

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is only valid for NATIVE data format storage pools. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an AUDIT VOLUME command you can continually ensure the integrity of data stored in your storage hierarchy. Possible values are:

#### Yes

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because additional overhead is required to calculate and compare CRC values between the storage pool and the server.

**No** Specifies that data is stored without CRC information.

**Tip:** For storage pools that are associated with the 3592, LTO, or ECARTRIDGE device type, logical block protection provides better protection against data corruption than CRC validation for a storage pool. If you specify CRC validation for a storage pool, data is validated only during volume auditing operations. Errors are identified after data is written to tape.

To enable logical block protection, specify a value of READWRITE for the **LBPROTECT** parameter on the **DEFINE DEVCLASS** and **UPDATE DEVCLASS** commands for the 3592, LTO, or ECARTRIDGE device types. Logical block protection is supported only on the following types of drives and media:

- IBM LTO5 and later.
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.
- Oracle StorageTek T10000C drives.

### DEDuplicate

Specifies whether the data that is stored in this storage pool will be

deduplicated. This parameter is optional and is valid only for storage pools that are defined with a FILE-type device class.

**IDENTIFYProcess**

Specifies the number of parallel processes to use for server-side duplicate identification. This parameter is optional and is valid only for storage pools that are defined with a FILE device class. Enter a value from 1 to 20.

When calculating the value for this parameter, consider the workload on the server and the amount of data requiring data deduplication. Server-side duplicate identification requires disk I/O and processor resources, so the more processes you allocate to data deduplication, the heavier the workload that you place on your system. In addition, consider the number of volumes that require processing. Server-side duplicate-identification processes work on volumes containing data that requires deduplication. If you update a storage pool, specifying that the data in the storage pool is to be deduplicated, all the volumes in the pool require processing. For this reason, you might have to define a high number of duplicate-identification processes initially. Over time, however, as existing volumes are processed, only the volumes containing new data have to be processed. When that happens, you can reduce the number of duplicate-identification processes.

**Remember:** Duplicate-identification processes can be either active or idle. Processes that are working on files are active. Processes that are waiting for files to work on are idle. Processes remain idle until volumes with data to be deduplicated become available. The output of the **QUERY PROCESS** command for a duplicate-identification process includes the total number of bytes and files that have been processed since the process first started. For example, if a duplicate-identification process processes four files, becomes idle, and then processes five more files, then the total number of files processed is nine. Processes end only when canceled or when the number of duplicate-identification processes for the storage pool is changed to a value less than the number currently specified.

**Example: Update a copy storage pool to a 30 day volume reuse and to collocate files by client node**

Update the copy storage pool named TAPEPOOL2 to change the delay for volume reuse to 30 days and to collocate files by client node.

```
update stgpool tapepool2 reusedelay=30 collocate=node
```

## UPDATE STGPOOL


### UPDATE STGPOOL (Update an active-data sequential access)

Use this command to update an active-data pool.

#### Privilege class

To issue this command, you must have system privilege, unrestricted storage privilege, or restricted storage privilege for the storage pool to be updated.

#### Syntax



#### Notes:

- 1 This parameter is valid only for storage pools that are defined with a FILE-type device class.
- 2 This parameter is only available if the value of the DEDUPLICATE parameter is YES.

#### Parameters

##### *pool\_name* (Required)

Specifies the name of the active-data pool to be updated.

##### DESCRIPTION

Specifies a description of the active-data pool. This parameter is optional. The maximum length of the description is 255 characters. Enclose the description in quotation marks if it contains any blank characters. To remove an existing description, specify a null string ("").

**ACcEss**

Specifies how client nodes and server processes (such as reclamation) can access files in the active-data pool. This parameter is optional. Possible values are:

**READWrite**

Specifies that files can be read from and written to the volumes in the active-data pool.

**READOnly**

Specifies that client nodes can only read files stored on the volumes in the active-data pool.

Server processes can move files within the volumes in the storage pool. The server can use files in the active-data pool to restore active versions of backup files to primary storage pools. However, no new writes are permitted to volumes in the active-data pool from volumes outside the storage pool. A storage pool cannot be copied to the active-data pool.

**UNAVailable**

Specifies that client nodes cannot access files stored on volumes in the active-data pool.

Server processes can move files within the volumes in the storage pool. The server can use files in the active-data pool to restore active versions of backup files to primary storage pools. However, no new writes are permitted to volumes in the active-data pool from volumes outside the storage pool. A storage pool cannot be copied to the active-data pool.

**COLlocate**

Specifies whether the server attempts to keep data belonging to a single client node, group of client nodes, or client file space assigned to as few volumes as possible. This parameter is optional.

Collocation reduces the number of sequential access media mounts for restore, retrieve, and recall operations. However, collocation increases both the amount of server time needed to collocate files for storing and the number of volumes required. For details, see the *Administrator's Guide*.

Possible values are:

**No** Specifies that collocation is disabled.

**GROup**

Specifies that collocation is enabled at the group level for client nodes. The server attempts to put data for nodes that belong to the same collocation group on as few volumes as possible. If the nodes in the collocation group have multiple file spaces, the server does not attempt to collocate those file spaces.

If you specify COLLOCATE=GROUP but do not define any collocation groups or if you specify COLLOCATE=GROUP but do not add nodes to a collocation group, data is collocated by node. Be sure to consider tape usage when organizing client nodes into collocation groups. For example, if a tape-based storage pool consists of data from grouped and ungrouped nodes and you specify COLLOCATE=GROUP, the server performs the following actions:

- Collocates by group the data for grouped nodes only. Whenever possible, the server collocates data that belongs to a group of nodes on a single tape or on as few tapes as possible. Data for a single node can also be spread across several tapes associated with a group.

## UPDATE STGPOOL

- Collocates by node the data for ungrouped nodes. Whenever possible, the server stores the data for a single node on a single tape. All available tapes that already have data for the node are used before available space on any other tape is used.

### **NODe**

Specifies that collocation is enabled at the client node level. The server attempts to put data for one node on as few volumes as possible. If the node has multiple file spaces, the server does not attempt to collocate those file spaces. For backward compatibility, COLLOCATE=YES is still accepted by the server to specify collocation at the client node level.

If a storage pool contains data for a node that is a member of a collocation group and you specify COLLOCATE=NODE, the data will be collocated by node not by group.

### **FIlespace**

Specifies that collocation is enabled at the file space level for client nodes. The server attempts to put data for one node and file space on as few volumes as possible. If a node has multiple file spaces, the server attempts to put data for different file spaces on different volumes.

### **REClaim**

Specifies when the server reclaims a volume, based on the percentage of reclaimable space on a volume. Reclaimable space is the amount of space occupied by files that are expired or deleted from the Tivoli Storage Manager database.

Reclamation makes the fragmented space on volumes usable again by moving any remaining active files from one volume to another volume, thus making the original volume available for reuse. This parameter is optional. You can specify an integer from 1 to 100. The value 100 means that reclamation is not performed.

When determining which volumes in a storage pool to reclaim, the Tivoli Storage Manager server first determines the reclamation threshold indicated by the **RECLAIM**. The server then examines the percentage of reclaimable space for each volume in the storage pool. If the percentage of reclaimable space on a volume is greater than the reclamation threshold of the storage pool, the volume is a candidate for reclamation.

For example, suppose storage pool FILEPOOL has a reclamation threshold of 70 percent. This value indicates that the server can reclaim any volume in the storage pool that has a percentage of reclaimable space that is greater than 70 percent. The storage pool has three volumes:

- FILEVOL1 with 65 percent reclaimable space
- FILEVOL2 with 80 percent reclaimable space
- FILEVOL3 with 95 percent reclaimable space

When reclamation begins, the server compares the percent of reclaimable space for each volume with the reclamation threshold of 70 percent. In this example, FILEVOL2 and FILEVOL3 are candidates for reclamation because their percentages of reclaimable space are greater than 70. To determine the percentage of reclaimable space for a volume, issue the **QUERY VOLUME** command and specify **FORMAT=DETAILED**. The value in the field Pct. Reclaimable Space is the percentage of reclaimable space for the volume.

If you change the value from the default of 60, specify a value of 50 percent or greater so that files stored on two volumes can be combined onto a single output volume.

When an active-data pool volume that is offsite becomes eligible for reclamation, the reclamation process attempts to obtain the active files on the reclaimable volume from a primary or active-data pool that is onsite. The process then writes these files to an available volume in the original active-data pool. Effectively, these files are moved back to the onsite location. However, the files could be obtained from the offsite volume after a disaster if a database backup is used that references the files on the offsite volume. Because of the way reclamation works with offsite volumes, use it carefully with active-data pools.

#### **RECLAIMProcess**

Specifies the number of parallel processes to use for reclaiming the volumes in this storage pool. This parameter is optional. Enter a value from 1 to 999.

When calculating the value for this parameter, consider the number of sequential storage pools that will be involved with the reclamation and the number of logical and physical drives that can be dedicated to the operation. To access a sequential access volume, IBM Tivoli Storage Manager uses a mount point and, if the device type is not FILE, a physical drive. The number of available mount points and drives depends on other Tivoli Storage Manager and system activity and on the mount limits of the device classes for the sequential access storage pools that are involved in the reclamation.

For example, suppose that you want to reclaim the volumes from two sequential storage pools simultaneously and that you want to specify four processes for each of the storage pools. The storage pools have the same device class. Each process requires two mount points and, if the device type is not FILE, two drives. (One of the drives is for the input volume, and the other drive is for the output volume.) To run eight reclamation processes simultaneously, you need a total of at least 16 mount points and 16 drives. The device class for each storage pool must have a mount limit of at least eight.

If the number of reclamation processes you specify is more than the number of available mount points or drives, the processes that do not obtain mount points or drives will wait for mount points or drives to become available. If mount points or drives do not become available within the MOUNTWAIT time, the reclamation processes will end. For information about specifying the MOUNTWAIT time, see “DEFINE DEVCLASS (Define a device class)” on page 148.

The Tivoli Storage Manager server will start the specified number of reclamation processes regardless of the number of volumes that are eligible for reclamation. For example, if you specify ten reclamation processes and only six volumes are eligible for reclamation, the server will start ten processes and four of them will complete without processing a volume.

#### **OFFSITERECLAIMLimit**

Specifies the number of offsite volumes to have their space reclaimed during reclamation for this storage pool. This parameter is optional. Possible values are:

##### **NOLimit**

Specifies that you want to have the space reclaimed in all of your offsite volumes.

##### *number*

Specifies the number of offsite volumes to have their space reclaimed. You can specify an integer from 0 to 99999. A value of zero means that none of the offsite volumes will be reclaimed.

**Important:** When determining the value for the OFFSITERECLAIMLIMIT, consider using the statistical information in the message issued at the end of the offsite volume reclamation operation. Alternatively, you can use the following Tivoli Storage Manager SQL select command to obtain the statistical information from the SUMMARY table for the offsite volume reclamation operation:

```
select * from summary where activity='OFFSITE RECLAMATION'
```

The statistical information includes the following items:

- The number of offsite volumes that were processed
- The number of parallel processes that were used
- The total amount of time required for the processing

The order in which offsite volumes are reclaimed is based on the amount of unused space in a volume. (Unused space includes both space that has never been used on the volume and space that has become empty because of file deletion.) Volumes with the largest amount of unused space are reclaimed first.

For example, suppose an active-data pool contains three volumes: VOL1, VOL2, and VOL3. VOL1 has the largest amount of unused space, and VOL3 has the least amount of unused space. Suppose further that the percentage of unused space in each of the three volumes is greater than the value of the RECLAIM parameter. If you do not specify a value for the OFFSITERECLAIMLIMIT parameter, all three volumes will be reclaimed when the reclamation runs. If you specify a value of 2, only VOL1 and VOL2 will be reclaimed when the reclamation runs. If you specify a value of 1, only VOL1 will be reclaimed.

### MAXSCRATCH

Specifies the maximum number of scratch volumes that the server can request for this storage pool. This parameter is optional. You can specify an integer from 0 to 100000000. By allowing the server to request scratch volumes as needed, you avoid having to define each volume to be used.

The value specified for this parameter is used to estimate the total number of volumes available in the active-data pool and the corresponding estimated capacity for the active-data pool.

Scratch volumes are automatically deleted from the storage pool when they become empty. However, if the access mode for a scratch volume is OFFSITE, the volume is not deleted from the active-data pool until the access mode is changed. This allows an administrator to query the server for empty, offsite scratch volumes and return these to the onsite location.

When scratch volumes with the device type of FILE become empty and are deleted, the space that the volumes occupied is freed by the server and returned to the file system.

**Tip:** For server-to-server operations that utilize virtual volumes and that store a small amount of data, consider specifying a value for the MAXSCRATCH parameter that is higher than the value you typically specify for write operations to other types of volumes. After a write operation to a virtual volume, Tivoli Storage Manager marks the volume as FULL, even if the value of the MAXCAPACITY parameter on the device-class definition has not been reached. The Tivoli Storage Manager server does not keep virtual volumes in FILLING status and does not append to them. If the value of the MAXSCRATCH parameter is too low, server-to-server operations can fail.

**REUsedelay**

Specifies the number of days that must elapse after all files are deleted from a volume before the volume can be rewritten or returned to the scratch pool. This parameter is optional. You can specify an integer from 0 to 9999. A value of 0 means that a volume can be rewritten or returned to the scratch pool as soon as all files are deleted from the volume.

**Important:** Use this parameter to help ensure that when you restore the database to an earlier level, database references to files in the active-data pool are still valid. You must set this parameter to a value greater than the number of days you plan to retain the oldest database backup. The number of days specified for this parameter should be the same as the number specified for the SET DRMDBBACKUPEXPIREDDAYS command. For more information, see the *Administrator's Guide*.

**OVFL0cation**

Specifies the overflow location for the storage pool. The server assigns this location name to a volume that is ejected from the library by the MOVE MEDIA command. This parameter is optional. The location name can be a maximum length of 255 characters. Enclose the location name in quotation marks if the location name contains any blank characters.

To remove an existing value, specify a null string ("").

**CRCData**

Specifies whether a cyclic redundancy check (CRC) validates storage pool data when audit volume processing occurs on the server. This parameter is only valid for NATIVE data format storage pools. This parameter is optional. The default value is NO. By setting CRCDATA to YES and scheduling an AUDIT VOLUME command you can continually ensure the integrity of data stored in your storage hierarchy. Possible values are:

**Yes**

Specifies that data is stored containing CRC information, allowing for audit volume processing to validate storage pool data. This mode impacts performance because additional overhead is required to calculate and compare CRC values between the storage pool and the server.

**No** Specifies that data is stored without CRC information.

**Tip:** For storage pools that are associated with the 3592, LTO, or ECARTRIDGE device type, logical block protection provides better protection against data corruption than CRC validation for a storage pool. If you specify CRC validation for a storage pool, data is validated only during volume auditing operations. Errors are identified after data is written to tape.

To enable logical block protection, specify a value of READWRITE for the **LBPROTECT** parameter on the **DEFINE DEVCLASS** and **UPDATE DEVCLASS** commands for the 3592, LTO, or ECARTRIDGE device types. Logical block protection is supported only on the following types of drives and media:

- IBM LTO5 and later.
- IBM 3592 Generation 3 drives and later with 3592 Generation 2 media and later.
- Oracle StorageTek T10000C drives.

**DEDUPLICATE**

Specifies whether the data that is stored in this storage pool will be

## UPDATE STGPOOL

deduplicated. This parameter is optional and is valid only for storage pools that are defined with a FILE-type device class.

### **IDENTIFYPROcess**

Specifies the number of parallel processes to use for server-side duplicate identification. This parameter is optional and is valid only for storage pools that are defined with a FILE device class. Enter a value from 1 to 20.

When calculating the value for this parameter, consider the workload on the server and the amount of data requiring data deduplication. Server-side duplicate identification requires disk I/O and processor resources, so the more processes you allocate to data deduplication, the heavier the workload that you place on your system. In addition, consider the number of volumes that require processing. Server-side duplicate-identification processes work on volumes containing data that requires deduplication. If you update a storage pool, specifying that the data in the storage pool is to be deduplicated, all the volumes in the pool require processing. For this reason, you might have to define a high number of duplicate-identification processes initially. Over time, however, as existing volumes are processed, only the volumes containing new data have to be processed. When that happens, you can reduce the number of duplicate-identification processes.

**Remember:** Duplicate-identification processes can be either active or idle. Processes that are working on files are active. Processes that are waiting for files to work on are idle. Processes remain idle until volumes with data to be deduplicated become available. The output of the **QUERY PROCESS** command for a duplicate-identification process includes the total number of bytes and files that have been processed since the process first started. For example, if a duplicate-identification process processes four files, becomes idle, and then processes five more files, then the total number of files processed is nine. Processes end only when canceled or when the number of duplicate-identification processes for the storage pool is changed to a value less than the number currently specified.

### **Example: Update an active data pool**

Update the active-data pool named TAPEPOOL2 to change the delay for volume reuse to 30 days and to collocate files by client node.

```
update stgpool tapepool3 reusedelay=30 collocate=node
```

## UPDATE VIRTUALFSMAPPING (Update a virtual file space mapping)

Use this command to update an existing virtual file space mapping definition. Refer to the documentation about your NAS device for more information about this command.


The NAS node must have an associated data mover definition because when the Tivoli Storage Manager server updates a virtual file space mapping, the server attempts to contact the NAS device to validate the virtual file system and file system name.

### Privilege class

To issue this command, you must have one of the following privilege classes:

- System privilege
- Unrestricted policy privilege
- Restricted policy privilege for the domain to which the NAS node is assigned

### Syntax



### Parameters

#### *node\_name* (Required)

Specifies the NAS node on which the file system and path reside. You cannot use wildcard characters or specify a list of names.

#### *virtual\_filespace\_name* (Required)

Specifies the virtual file space mapping to update. You cannot use wildcard characters or specify a list of names.

#### **FILESystem**

Specifies the new name of the file system in which the path is located. The file system name must exist on the specified NAS node. The file system name cannot contain wildcard characters. The file system name should only be modified when the file system name is modified on the NAS device or, for example, the directory is moved to a different file system. This parameter is optional.

#### **PATH**

Specifies the new path from the root of the file system to the directory. The path can only reference a directory. This should only be modified when the path on the NAS device has changed; for example, the directory is moved to a

## UPDATE VIRTUALFSMAPPING

different path. The maximum length of the path is 1024 characters. The path name is case sensitive. This parameter is optional.

### **NAMEType**

Specifies how the server should interpret the path name specified. Specify this parameter only if you specify a path. This parameter is useful when a path contains characters that are not part of the code page on which the server is running. The default value is SERVER.

Possible values are:

### **SERVER**

The code page in which the server is running is used to interpret the path.

### **HEXadecimal**

The server interprets the path that you enter as the hexadecimal representation of the path. This option should be used when a path contains characters that cannot be entered. For example, this could occur if the NAS file system is set to a language different from the one in which the server is running.

## **Example: Modify the path of a virtual file space mapping**

Update the virtual file space mapping named /mikeshomedir for the NAS node NAS1 by modifying the path.

```
update virtualfsmapping nas1 /mikeshomedir path=/new/home/mike
```

## **Related commands**

*Table 442. Commands related to UPDATE VIRTUALFSMAPPING*

| <b>Command</b>          | <b>Description</b>                   |
|-------------------------|--------------------------------------|
| DEFINE VIRTUALFSMAPPING | Define a virtual file space mapping. |
| DELETE VIRTUALFSMAPPING | Delete a virtual file space mapping. |
| QUERY VIRTUALFSMAPPING  | Query a virtual file space mapping.  |



## UPDATE VOLHISTORY

### **Vault**

The volume is off-site, contains valid data, and is not accessible for on-site processing.

### **COURIERRetrieve**

The volume is being moved on-site.

### **Example: Update the location of a volume used for database backup**

Update the location of a volume used for database backup, BACKUP1, to show that it has been moved to an off-site location.

```
update volhistory backup1 devclass=tapebkup
location="700 w. magee rd."
```

### **Related commands**

*Table 443. Commands related to UPDATE VOLHISTORY*

| <b>Command</b>    | <b>Description</b>                                                                    |
|-------------------|---------------------------------------------------------------------------------------|
| BACKUP VOLHISTORY | Records volume history information in external files.                                 |
| DELETE VOLHISTORY | Removes sequential volume history information from the volume history file.           |
| MOVE DRMEDIA      | Moves DRM media on-site and off-site.                                                 |
| PREPARE           | Creates a recovery plan file.                                                         |
| QUERY DRMEDIA     | Displays information about disaster recovery volumes.                                 |
| QUERY VOLHISTORY  | Displays sequential volume history information that has been collected by the server. |

## UPDATE VOLUME (Change a storage pool volume)


Use this command to change the access mode for one or more volumes in storage pools.

You can correct an error condition associated with a volume by updating the volume to an access mode of READWRITE. You can also use this command to change the location information for one or more volumes in sequential access storage pools.

### Privilege class

To issue this command, you must have system privilege or operator privilege.

### Syntax



### Notes:

- 1 You must update at least one attribute (ACCESS or LOCATION).
- 2 This value is valid only for volumes in primary storage pools.
- 3 This value is valid only for volumes in copy storage pools.
- 4 This parameter is valid only for volumes in sequential access storage pools.

### Parameters

#### *volume\_name* (Required)

Specifies the storage pool volume to update. You can use wildcard characters to specify names.

#### ACCess

Specifies how client nodes and server processes (such as migration) can access files in the storage pool volume. This parameter is optional. Possible values are:

##### READWrite

Specifies that client nodes and server processes can read from and write to files stored on the volume.

If the volume being updated is an empty scratch volume that had an access mode of offsite, the server deletes the volume from the database.

##### READOnly

Specifies that client nodes and server processes can only read files stored on the volume.

If the volume being updated is an empty scratch volume that had an access mode of offsite, the server deletes the volume from the database.

##### UNAVailable

Specifies that neither client nodes nor server processes can access files stored on the volume.

Before making a random access volume unavailable, you must vary the volume offline. After you make a random access volume unavailable, you cannot vary the volume online.

If you make a sequential access volume unavailable, the server does not attempt to mount the volume.

If the volume being updated is an empty scratch volume that had an access mode of offsite, the server deletes the volume from the database.

##### DESTroyed

Specifies that a primary storage pool volume has been permanently damaged. Neither client nodes nor server processes can access files stored on the volume. Use this access mode to indicate an entire volume that needs to be restored by using the **RESTORE STGPOOL** command. After all files on a destroyed volume have been restored to other volumes, the server automatically deletes the destroyed volume from the database.

Only volumes in primary storage pools can be updated to DESTROYED.

Before updating a random access volume to DESTROYED access, you must vary the volume offline. After you update a random access volume to DESTROYED, you cannot vary the volume online.

If you update a sequential access volume to DESTROYED, the server does not attempt to mount the volume.

If a volume contains no files and you change the access mode to DESTROYED, the server deletes the volume from the database.

##### OFFsite

Specifies that a copy storage pool volume is at an offsite location from which it cannot be mounted. Only volumes in copy storage pools can have the access mode of OFFSITE.

Use this mode to help you track volumes that you move to offsite locations. See the *Administrator's Guide* for details.

If you specify values for both the ACCESS and LOCATION parameters but the access mode cannot be updated for a particular volume, the location attribute is also not updated for that volume. For example, if you specify ACCESS=OFFSITE and a LOCATION value for a primary storage pool volume, neither the access nor location values are updated because a primary storage pool volume cannot be given an access mode of OFFSITE.

#### **Location**

Specifies the location of the volume. This parameter is optional. It can be specified only for volumes in sequential access storage pools. The maximum length of the location is 255 characters. Enclose the location in quotation marks if it contains any blank characters. To remove a previously defined location, specify the null string ("").

#### **WHEREStoragePool**

Specifies the name of the storage pool for volumes to be updated. Use this parameter to restrict the update by storage pool. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a storage pool name, volumes belonging to any storage pool are updated.

#### **WHEREDeviceClass**

Specifies the name of the device class for volumes to be updated. Use this parameter to restrict the update by device class. This parameter is optional. You can use wildcard characters to specify names. If you do not specify a device class name, volumes with any device class are updated.

#### **WHEREAccess**

Specifies the current access mode of volumes to be updated. Use this parameter to restrict the update to volumes that currently have the specified access mode. This parameter is optional. You can specify multiple access modes by separating the modes with commas and no intervening spaces. If you do not specify a value for this parameter, the update is not restricted by the current access mode of a volume. Possible values are:

##### **READWrite**

Update volumes with an access mode of READWRITE.

##### **READOnly**

Update volumes with an access mode of READONLY.

##### **UNAVailable**

Update volumes with an access mode of UNAVAILABLE.

##### **Offsite**

Update volumes with an access mode of OFFSITE.

##### **DESTroyed**

Update volumes with an access mode of DESTROYED.

#### **WHEREStatus**

Specifies the status of volumes to be updated. Use this parameter to restrict the update to volumes that have a specified status. This parameter is optional. You can specify multiple status values by separating the values with commas and no intervening spaces. If you do not specify a value for this parameter, the update is not restricted by volume status. Possible values are:

##### **ONline**

Update volumes with a status of ONLINE.

## UPDATE VOLUME

### Offline

Update volumes with a status of OFFLINE.

### EMPTy

Update volumes with a status of EMPTY.

### PENding

Update volumes with a status of PENDING. These are volumes from which all files have been deleted, but the time specified by the REUSEDELAY parameter has not elapsed.

### FILLing

Update volumes with a status of FILLING.

### FULL

Update volumes with a status of FULL.

### Preview

Specifies whether you want to preview the update operation without actually updating volumes. This parameter is optional. The default value is NO.

Possible values are:

**No** Specifies that volumes are actually updated.

### Yes

Specifies that you want only to preview the update operation. This option displays the volumes that will be updated if you actually perform the update operation.

### Example: Make a tape volume unavailable

Update a tape volume named DSMT20 to make it unavailable to client nodes and server processes.

```
update volume dsmt20 access=unavailable
```

### Example: Update the access mode of all offsite volumes in a specific storage pool

Update all empty, offsite volumes in the TAPEPOOL2 storage pool. Set the access mode to READWRITE and delete the location information for the updated volumes.

```
update volume * access=readwrite location="" wherestgpool=tapepool2
 whereaccess=offsite wherestatus=empty
```

### Related commands

Table 444. Commands related to UPDATE VOLUME

| Command       | Description                                                              |
|---------------|--------------------------------------------------------------------------|
| DEFINE VOLUME | Assigns a volume to be used for storage within a specified storage pool. |
| DELETE VOLUME | Deletes a volume from a storage pool.                                    |
| QUERY VOLUME  | Displays information about storage pool volumes.                         |
| VARY          | Specifies whether a disk volume is available to the server for use.      |

---

## VALIDATE commands

Use the **VALIDATE** command to verify that an object is complete or valid for Tivoli Storage Manager.

- “VALIDATE LANFREE (Validate LAN-Free paths)” on page 1428
- “VALIDATE POLICYSET (Verify a policy set)” on page 1430
- “VALIDATE REPLICATION (Validate replication for a client node)” on page 1432

## VALIDATE LANFREE (Validate LAN-Free paths)

Use this command to determine which destinations for a given node using a specific storage agent are capable of LAN-Free data movement.

### Privilege class

To issue this command, you must have system privilege.

### Syntax

▶▶—VALidate LANfree—*node\_name*—*stgagent\_name*—▶▶

### Parameters

***node\_name* (Required)**

The name of the node to evaluate.

***stgagent\_name* (Required)**

The name of the storage agent to evaluate.

### Example: Validate a current LAN-Free configuration

Validate the current server definitions and configuration for node TIGER to use storage agent AIX\_STA1 for LAN-free data operations.

```
validate lanfree tiger aix_sta1
```

| Node Name | Storage Agent | Operation | Mgmt Class Name | Destination Name | LAN-Free capable? | Explanation                                                                                                  |
|-----------|---------------|-----------|-----------------|------------------|-------------------|--------------------------------------------------------------------------------------------------------------|
| TIGER     | AIX_STA1      | BACKUP    | STANDARD        | OUTPOOL          | NO                | No available online paths.                                                                                   |
| TIGER     | AIX_STA1      | BACKUP    | STANDARD        | PRIMARY          | NO                | Destination storage pool is configured for simultaneous write.                                               |
| TIGER     | AIX_STA1      | BACKUP    | STANDARD        | SHRPOOL          | YES               |                                                                                                              |
| TIGER     | AIX_STA1      | BACKUP    | NOARCH          | LFFILE           | NO                | Storage pool contains data deduplicated by clients, and is not accessible by storage agents V6.1 or earlier. |
| TIGER     | AIX_STA1      | ARCHIVE   | STANDARD        | OUTPOOL          | NO                | No available online paths.                                                                                   |
| TIGER     | AIX_STA1      | ARCHIVE   | STANDARD        | PRIMARY          | NO                | Destination storage pool is configured for simultaneous write.                                               |
| TIGER     | AIX_STA1      | ARCHIVE   | STANDARD        | SHRPOOL          | YES               |                                                                                                              |

### Related commands

Table 445. Commands related to VALIDATE LANFREE

| Command         | Description                                |
|-----------------|--------------------------------------------|
| QUERY COPYGROUP | Displays the attributes of a copy group.   |
| QUERY DEVCLASS  | Displays information about device classes. |
| QUERY DOMAIN    | Displays information about policy domains. |

Table 445. Commands related to VALIDATE LANFREE (continued)

| Command         | Description                                                                             |
|-----------------|-----------------------------------------------------------------------------------------|
| QUERY DRIVE     | Displays information about drives.                                                      |
| QUERY LIBRARY   | Displays information about one or more libraries.                                       |
| QUERY MGMTCLASS | Displays information about management classes.                                          |
| QUERY NODE      | Displays partial or complete information about one or more clients.                     |
| QUERY PATH      | Displays information about the path from a source to a destination.                     |
| QUERY POLICYSET | Displays information about policy sets.                                                 |
| QUERY SERVER    | Displays information about servers.                                                     |
| QUERY STATUS    | Displays the settings of server parameters, such as those selected by the SET commands. |
| QUERY STGPOOL   | Displays information about storage pools.                                               |

### VALIDATE POLICYSET (Verify a policy set)

Use this command to verify that a policy set is complete and valid before you activate it. The command examines the management class and copy group definitions in the policy set and reports on conditions that you need to consider before activating the policy set.

The **VALIDATE POLICYSET** command fails if any of the following conditions exist:

- The policy set has no default management class.
- A copy group within the policy set specifies a copy storage pool as a destination.
- A management class specifies a copy storage pool as the destination for files that were migrated by a Tivoli Storage Manager for Space Management client.
- A TOCDESTINATION parameter is specified, and the storage pool is either a copy pool or has a data format other than NATIVE or NONBLOCK.

The server issues warning messages for the following conditions:

- A copy group specifies a storage pool that does not exist as a destination for backed-up or archived files.

If you activate a policy set with copy groups that specify nonexistent storage pools, the client backup or archive operations fail.

- A management class specifies a storage pool that does not exist as a destination for files migrated by Tivoli Storage Manager for Space Management clients.
- The policy set does not have one or more management classes that exist in the current ACTIVE policy set.

If you activate the policy set, backup files bound to the deleted management classes are rebound to the default management class in the new active policy set.

- The policy set does not have one or more copy groups that exist in the current ACTIVE policy set.

If you activate the policy set, files bound to the management classes with deleted copy groups are no longer archived or backed up.

- The default management class for the policy set does not contain a backup or archive copy group.

If you activate the policy set with this default management class, clients using the default cannot back up or archive files.

- A management class specifies that a backup version must exist before a file can be migrated from a client node (MIGREQUIRESBKUP=YES), but the management class does not contain a backup copy group.

If the server has data retention protection enabled, the following conditions must exist:

- All management classes in the policy set to be validated must contain an archive copy group.
- If a management class exists in the active policy set, a management class with the same name must exist in the policy set to be validated.
- If an archive copy group exists in the active policy set, the corresponding copy group in the policy set to be validated must have a RETVER value at least as large as the corresponding values in the active copy group.

## Privilege class

To issue this command, you must have system privilege, unrestricted policy privilege, or restricted policy privilege for the policy domain to which the policy set belongs.

## Syntax

►►—VALidate POLicyset—*domain\_name*—*policy\_set\_name*—►►

## Parameters

### *domain\_name* (Required)

Specifies the name of the policy domain to which the policy set is assigned.

### *policy\_set\_name* (Required)

Specifies the name of the policy set to be validated.

## Example: Validate a specific policy set

Validate the policy set VACATION located in the EMPLOYEE\_RECORDS policy domain.

```
validate policyset employee_records vacation
```

## Related commands

Table 446. Commands related to VALIDATE POLICYSET

| Command            | Description                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------|
| ACTIVATE POLICYSET | Validates and activates a policy set.                                                         |
| COPY POLICYSET     | Creates a copy of a policy set.                                                               |
| DEFINE COPYGROUP   | Defines a copy group for backup or archive processing within a specified management class.    |
| DEFINE MGMTCLASS   | Defines a management class.                                                                   |
| DELETE POLICYSET   | Deletes a policy set, including its management classes and copy groups, from a policy domain. |
| QUERY POLICYSET    | Displays information about policy sets.                                                       |
| UPDATE COPYGROUP   | Changes one or more attributes of a copy group.                                               |
| UPDATE POLICYSET   | Changes the description of a policy set.                                                      |

## VALIDATE REPLICATION

### VALIDATE REPLICATION (Validate replication for a client node)

Use this command to identify the replication rules that apply to file spaces in client nodes that are configured for replication. You can also use this command to verify that the source replication server can communicate with the target replication server.


Before you begin replication processing, use the **VALIDATE REPLICATION** command to determine whether your replication configuration is correct.

Issue this command on the server that acts as a source for replicated data.

#### Privilege class

To issue this command, you must have system privilege.

#### Syntax



#### Parameters

##### *node\_name* (Required)

Specifies the name of the client node whose file spaces you want to display. To specify multiple client node names, separate the names with commas and no intervening spaces. You can use wildcard characters to specify names.

Information is displayed only for client nodes that are either enabled or disabled for replication. The replication mode must be SEND. To determine whether a client node is enabled or disabled for replication and its mode, issue the **QUERY NODE** command. Look for values in the Replication State and Replication Mode fields.

##### **VERIFYconnection**

Specifies whether to check the connection to a target replication server. The version of the target replication server is also checked to verify that it is version 6.3. This parameter is optional. The default is NO. You can specify one of the following values:

**No** The connection and version of the target replication server are not checked.

**Yes**

The connection and version of the target replication server are checked.

#### Example: Validate replication for a client node

The name of the client node is NODE1. Verify the connection status between the source and the target replication servers.

```
validate replication node1 verifyconnection=yes
```

```

Node Name: NODE1
Filespace Name: \\node1\c$
FSID: 1
Type: Bkup
Controlling Replication Rule: ACTIVE_DATA
Replication Rule Level: System Level
Server Name: DRSRV
Connection Status: Valid Connection

Node Name: NODE1
Filespace Name: \\node1\c$
FSID: 1
Type: Arch
Controlling Replication Rule: ALL_DATA_HIGH_PRIORITY
Replication Rule Level: Node Level
Server Name: DRSRV
Connection Status: Valid Connection

Node Name: NODE1
Filespace Name: \\node1\c$
FSID: 1
Type: SpMg
Controlling Replication Rule: ALL_DATA
Replication Rule Level: System Level
Server Name: DRSRV
Connection Status: Valid Connection

```

Output is displayed for all data types regardless of whether a file space contains the data types. For example, if a file space contains only backup and archive data, the output of the **VALIDATE REPLICATION** command also contains information that would be relevant to space-managed data.

## Field descriptions

### Node Name

The node that owns the replicated data.

### Filespace Name

The name of the file space that belongs to the node.

File space names can be in a different code page or locale than the server. If they are, the names in the Administration Center and the administrative command-line interface might not be displayed correctly. Data is backed up and can be restored normally, but the file space name or file name might be displayed with a combination of invalid characters or blank spaces.

If the file space name is Unicode-enabled, Tivoli Storage Manager converts the name to the server code page for display. The success of the conversion depends on the operating system, the characters in the name, and the server code page. Conversion can be incomplete if the string includes characters that are not available in the server code page or if the server cannot access system conversion routines. If the conversion is incomplete, the name might contain question marks, blanks, unprintable characters, or ellipses (...).

**FSID** The file space identifier for the file space. The server assigns a unique FSID when a file space is first stored on the server.

**Type** The type of data. The following values are possible:

**Arch** Archive data

**Bkup** Backup data

## VALIDATE REPLICATION

**Spmg** Data that was migrated by a Tivoli Storage Manager for Space Management client.

### Controlling Replication Rule

The name of the replication rule that controls replication for a data type in a file space. To determine whether the controlling rule is a file space rule, a client rule, or a server rule, check the Replication Rule Level field.

### Replication Rule Level

The level of the controlling rule in the replication-rule hierarchy. The following values are possible:

#### Filespace

The controlling rule is assigned to a data type in the file space.

**Node** The controlling rule is assigned to a data type for a client node.

**Server** The controlling rule is assigned to a data type for all file spaces in all client nodes that are configured for replication.

### Server Name

The name of the target replication server to be queried.

### Connection Status

The connection status between the source and the target replication server. The following values are possible:

#### Valid Connection

Communication with the target replication server was successful, and the target replication server is a V6.3 server.

#### Target Server Not Set

The target replication server is not set. To set the target replication server, issue the **SET REPLSERVER** command.

#### Communication Failure

The source replication server was unable to contact the target replication server. Examine the activity log for error messages about failed communications. Consider the following possible causes:

- The replication configuration on the source replication server is not valid. One or more of the following problems might exist:
  - The server definition for the target replication server is incorrect.
  - If the target replication-server definition was deleted and redefined, issue the **PING SERVER** command to test the connection between the source and the target replication server. If the **PING SERVER** command is successful, issue the **UPDATE SERVER** command and specify **FORCESYNC=YES** to reset the server verification keys.
  - The server name, server low-level address, server high-level address, and server password do not match the values that are specified in the server definition on the target replication server.
- The replication configuration on the target replication server is not valid. One or more of the following problems might exist:
  - The version of the target replication server is earlier than V6.3.

- The server definition for the source replication server is incorrect.
- The server name, server low-level address, server high-level address, and server password do not match the values that are specified in the server definition on the source replication server.
- Network communications are unavailable. To test the connection between the source and target server, issue the **PING SERVER** command .
- The target replication server is unavailable.
- Sessions between the source and the target replication servers are disabled. To verify the status of sessions, issue the **QUERY STATUS** command.

### Related commands

*Table 447. Commands related to VALIDATE REPLICATION*

| Command               | Description                                                                             |
|-----------------------|-----------------------------------------------------------------------------------------|
| ENABLE SESSIONS       | Resumes server activity following the DISABLE command or the ACCEPT DATE command.       |
| QUERY FILESPACE       | Displays information about data in file spaces that belong to a client.                 |
| QUERY NODE            | Displays partial or complete information about one or more clients.                     |
| QUERY REPLRULE        | Displays information about node replication rules.                                      |
| QUERY SERVER          | Displays information about servers.                                                     |
| QUERY STATUS          | Displays the settings of server parameters, such as those selected by the SET commands. |
| REPLICATE NODE        | Replicates data in file spaces that belong to a client node.                            |
| SET ARREPLRULEDEFAULT | Specifies the server node-replication rule for archive data.                            |
| SET BKREPLRULEDEFAULT | Specifies the server node-replication rule for backup data.                             |
| SET REPLSERVER        | Specifies a target replication server.                                                  |
| SET SPREPLRULEDEFAULT | Specifies the server node-replication rule for space-managed data.                      |
| UPDATE FILESPACE      | Changes file-space node-replication rules.                                              |
| UPDATE NODE           | Changes the attributes associated with a client node.                                   |
| UPDATE REPLRULE       | Enables or disables replication rules.                                                  |
| UPDATE SERVER         | Updates information about a server.                                                     |

## VARY (Bring a random access volume online or offline)


Use this command to make a random access storage pool volume online or offline to the server.

### Privilege class

This command is valid only for volumes on random access devices. For example, use this command during maintenance or corrective action of a random access volume. You cannot vary a random access volume online that is defined as unavailable.

To issue this command, you must have system privilege or operator privilege.

### Syntax



### Parameters

#### ONline

Specifies that the server can use the random access volume.

#### OFFline

Specifies that the server cannot use the volume.

#### volume\_name (Required)

Specifies the volume identifier. Volume names cannot contain embedded blanks or equal signs.

#### Wait

Specifies whether to wait for the server to complete processing this command in the foreground. This parameter is optional. The default is NO. Possible values are:

##### No

Specifies that the server processes this command in the background, while other tasks run. The server displays messages created from the background process either in the activity log or the server console, depending on where messages are logged.

##### Yes

Specifies that the server processes this command in the foreground. Wait for the command to complete before you continue with other tasks. The server displays the output messages to the administrative client when the command completes.

You cannot specify WAIT=YES from the server console.

### Example: Bring volume online

Make volume /adsm/stgv01/1 available to the server for use as a storage pool volume.

```
vary online /adsm/stgv01/1
```

**Related commands***Table 448. Commands related to VARY*

| <b>Command</b> | <b>Description</b>                                                       |
|----------------|--------------------------------------------------------------------------|
| CANCEL PROCESS | Cancels a background server process.                                     |
| DEFINE VOLUME  | Assigns a volume to be used for storage within a specified storage pool. |
| DELETE VOLUME  | Deletes a volume from a storage pool.                                    |
| QUERY PROCESS  | Displays information about background processes.                         |
| QUERY VOLUME   | Displays information about storage pool volumes.                         |

**VARY**

---

## Chapter 3. Server options

At installation, IBM Tivoli Storage Manager provides a server options file that contains a set of default options to start the server.

The file is:

`dsmserv.opt`

See the *Installation Guide* for your platform to determine where your server options file resides.

Server options let you customize the following:

- Communication
- Server storage
- Client-server
- Date, number, time, and language
- Database and recovery log
- Data transfer
- Message
- Event logging
- Security and licensing

Several other options are available for miscellaneous purposes.

To display the current option settings, enter:

`query option`

---

### Modifying server options

The server reads the server options file at server initialization. When you update a server option by editing the file, you must stop and start the server to activate the updated server options file.

You can change some options dynamically without stopping and starting the server, by using the SETOPT command. See “SETOPT (Set a server option for dynamic update)” on page 1153 for details.

The `dsmserv.opt.smp` file (also provided at installation) contains the format of the options file and all the default settings. You can change any options in the `dsmserv.opt.smp` file. To have the server use the changed options, you must rename the file to `dsmserv.opt`. To activate an option within the server options file, remove the `*>>>` that precedes the option. The server ignores any options preceded by `*>>>`.

---

## Types of server options

Server options let you customize how some functions and processes work.

### Server communication options

You can use server options to specify server communication methods and their characteristics.

Table 449. Communication options

| Option                | Description                                                                                                                                                                                      |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ADMINCOMMTIMEOUT      | The amount of time that the server waits for an administrative client message during an operation that causes a database update                                                                  |
| ADMINIDLETIMEOUT      | The amount of time an administrative client session can be idle                                                                                                                                  |
| ADMINONCLIENTPORT     | The port that determines whether administrative sessions can use the port specified in the TCPPORT option                                                                                        |
| COMMMETHOD            | The server communication method                                                                                                                                                                  |
| DBMTCPPORT            | The port number on which the TCP/IP communication driver for the database manager waits for client session requests                                                                              |
| DNSLOOKUP             | Control of use of Domain Name Services to lookup names of systems contacting the server                                                                                                          |
| LDAPCACHEDURATION     | Determines the amount of time that authentication sessions, to the same node or administrator, are skipped. You might see a slight performance boost when skipping sessions.                     |
| LDAPURL               | Specifies the LDAP directory server. Each setting must have the LDAP directory server name, a port number, and the base distinguished name of the namespace or suffix that the server maintains. |
| NDMPCONTROLPORT       | The internal communications port used for certain Network Data Management Protocol (NDMP) operations                                                                                             |
| NDMPENABLEKEEPALIVE   | The TCP keepalive mechanism                                                                                                                                                                      |
| NDMPKEEPIDLEMINUTES   | The amount of idle time before the first TCP keepalive packet is sent                                                                                                                            |
| SHMPORT               | The TCP/IP port address of a server when using shared memory                                                                                                                                     |
| SNMPHEARTBEATINTERVAL | The interval in minutes between queries of the Tivoli Storage Manager server                                                                                                                     |
| SNMPMESSAGECATEGORY   | The trap types used when messages are forwarded from the server                                                                                                                                  |

Table 449. Communication options (continued)

| Option           | Description                                                                                                                                                                                                                                                                                                                                                        |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SNMPSUBAGENT     | The parameters needed for the Tivoli Storage Manager subagent to communicate with the SNMP daemon                                                                                                                                                                                                                                                                  |
| SNMPSUBAGENTHOST | The location of the Tivoli Storage Manager SNMP subagent                                                                                                                                                                                                                                                                                                           |
| SNMPSUBAGENTPORT | The port address of the Tivoli Storage Manager SNMP subagent                                                                                                                                                                                                                                                                                                       |
| SSLFIPSMODE      | Specifies whether the Federal Information Processing Standards (FIPS) mode is in effect for Secure Sockets Layer (SSL)                                                                                                                                                                                                                                             |
| SSLTCPADMINPORT  | The port address on which the server's TCP/IP communication driver waits for requests for SSL-enabled sessions for the command-line administrative client                                                                                                                                                                                                          |
| SSLTCPPOINT      | <p>The SSL-only port number on which the server's TCP/IP communication driver waits for requests for SSL-enabled sessions from the following sources:</p> <ul style="list-style-type: none"> <li>• Command line backup-archive client</li> <li>• Backup-archive GUI</li> <li>• Administrative client</li> <li>• Application programming interface (API)</li> </ul> |
| SSLTLS12         | If using SSL with self-signed certificates, specify and distribute a new self-signed certificate to all backup-archive clients                                                                                                                                                                                                                                     |
| TCPADMINPORT     | The TCP/IP port number for administrative sessions                                                                                                                                                                                                                                                                                                                 |
| TCPBUFSIZE       | The size of the buffer used for TCP/IP send requests                                                                                                                                                                                                                                                                                                               |
| TCPPOINT         | The TCP/IP port number for client sessions                                                                                                                                                                                                                                                                                                                         |
| TCPWINDOWSIZE    | The client node TCP/IP sliding window                                                                                                                                                                                                                                                                                                                              |

## Server storage options

You can use server options to configure server storage operations and device operations.

Table 450. Server storage options

| Option              | Description                                                                                                                                                                                                     |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 3494SHARED          | Enables sharing of a 3494 library with applications other than Tivoli Storage Manager                                                                                                                           |
| ACSACCESSID         | The ID for the ACS access control                                                                                                                                                                               |
| ACSLOCKDRIVE        | Allows the drives within the ACSLS libraries to be locked                                                                                                                                                       |
| ACSQUICKINIT        | Allows a quick or full initialization of the ACSLS library                                                                                                                                                      |
| ACSTIMEOUTX         | The multiple for the built-in timeout value for the ACSLS API                                                                                                                                                   |
| ASSISTVCRRECOVERY   | Specifies whether the server assists an IBM 3570 or 3590 drive in recovering from a lost or corrupted Vital Cartridge Records (VCR) condition                                                                   |
| CHECKTAPEPOS        | Specifies whether the server validates data position on tape                                                                                                                                                    |
| CLIENTDEDUPTXNLIMIT | Specifies the maximum size of a transaction when client-side deduplicated data is backed up or archived                                                                                                         |
| DEDUPREQUIRESBACKUP | Specifies whether volumes in primary sequential-access storage pools that are set up for data deduplication can be reclaimed and whether duplicate data can be discarded before the storage pools are backed up |
| DEDUPTIER2FILESIZE  | File size at which Tier 2 chunking is used for data deduplication.                                                                                                                                              |
| DEDUPTIER3FILESIZE  | File size at which Tier 3 chunking is used for data deduplication.                                                                                                                                              |
| DEVCONFIG           | The name of the file that store backup copies of device configuration information                                                                                                                               |
| DRIVEACQUIRERETRY   | The number of times that the server retries the acquisition of a drive in an IBM 349x library that is shared among multiple applications                                                                        |
| ENABLENASDEDUP      | Specifies whether the server deduplicates data that is stored by a NetApp network-attached storage (NAS) file server.                                                                                           |
| NUMOPENVOLSALLOWED  | The number of input FILE volumes in a deduplicated storage pool that can be open at one time                                                                                                                    |
| RECLAIMDELAY        | The number of days that the reclamation of a SnapLock volume is delayed                                                                                                                                         |
| RECLAIMPERIOD       | The number of days for the reclamation period of a SnapLock volume                                                                                                                                              |
| RESOURCETIMEOUT     | The length of time that the server waits for a resource before canceling the pending acquisition of a resource                                                                                                  |
| RETENTIONEXTENSION  | The number of days to extend the retention date of a SnapLock volume                                                                                                                                            |
| SANDISCOVERY        | Whether the Tivoli Storage Manager SAN discovery function is enabled                                                                                                                                            |
| SANDISCOVERYTIMEOUT | Amount of time before the SAN discovery process times out                                                                                                                                                       |
| SANREFRESHTIME      | Amount of time before cached SAN discovery information is refreshed                                                                                                                                             |
| SEARCHMPQUEUE       | The order in which the server satisfies requests in the mount queue                                                                                                                                             |
| SERVERDEDUPTXNLIMIT | Specifies the maximum size of objects that can be deduplicated on the server.                                                                                                                                   |

## Client-server options

You can use server options to control client-server processing.

Table 451. Client-Server options

| Option                  | Description                                                                                                              |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------|
| COMMTIMEOUT             | The number of seconds the server waits for a response from a client before timing out the client session                 |
| DISABLESCHEDS           | Whether administrative and client schedules are disabled during the Tivoli Storage Manager server recovery scenario      |
| IDLETIMEOUT             | The number of minutes the server allows a client session to remain idle before timing out the client session             |
| MAXSESSIONS             | The maximum number of simultaneous client sessions with the server                                                       |
| THROUGHPUTDATATHRESHOLD | The throughput threshold that a client session must reach to prevent being cancelled after the time threshold is reached |
| THROUGHPUTTIMETHRESHOLD | The time threshold for a session after which it may be cancelled for low throughput                                      |
| VERBCHECK               | Whether additional error checking is done for commands sent by the client                                                |

## Date, number, time, and language options

You can use server options to specify display formats for the dates, times, numbers, and national language.

Table 452. Date, number, time, and language options

| Option   | Description                                              |
|----------|----------------------------------------------------------|
| LANGUAGE | The national language is used to present client messages |

## Database options

You can use server options to control some aspects of database processing.

Table 453. Database options

| Option                   | Description                                                                                                                                |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| ACTIVELOGDIRECTORY       | The new directory for the location where the active log is stored. Use this option to change the location of the active log.               |
| ACTIVELOGSIZE            | The maximum size of the active log.                                                                                                        |
| ALLOWREORGINDEX          | Server-initiated index reorganization.                                                                                                     |
| ALLOWREORGTABLE          | Server-initiated table reorganization.                                                                                                     |
| ARCHLOGDIRECTORY         | The directory that the database manager can archive a log file into after all the transactions represented in that log file are completed. |
| ARCHFAILOVERLOGDIRECTORY | The directory in which the server tries to store archive log files that cannot be stored in the archive log directory.                     |
| DBDIAGLOGSIZE            | The maximum size of the database manager diagnostic log files                                                                              |
| DBDIAGPATHFSTHRESHOLD    | The threshold for free space on the file system or disk that contains the database manager diagnostic log files                            |
| FSUSEDTHRESHOLD          | The percentage of the file system that can be used by the database before an alert message is issued                                       |
| DBMEMPERCENT             | The percentage of system memory that is dedicated to the database.                                                                         |
| MIRRORLOGDIRECTORY       | The directory for mirroring the active log path.                                                                                           |

Table 453. Database options (continued)

| Option         | Description                                                                                         |
|----------------|-----------------------------------------------------------------------------------------------------|
| REORGBEGINTIME | The earliest time that the Tivoli Storage Manager server can start a table or index reorganization. |
| REORGDURATION  | The interval during which server-initiated table or index reorganization can start.                 |

## Data transfer options

You can use server options to control how Tivoli Storage Manager groups and transfers data.

Table 454. Group options

| Option                | Description                                                                                                                       |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| MOVEBATCHSIZE         | The number of files that are to be moved and grouped in a batch, within a transaction                                             |
| MOVESIZETHRESH        | The threshold for the amount of data moved as a batch, within the same server transaction                                         |
| NDMPPORTRANGE         | The IP address associated with the interface in which the server receives all Network Data Management Protocol (NDMP) backup data |
| NDMPPREFDATAINTERFACE | The IP address associated with the interface in which the server receives all Network Data Management Protocol (NDMP) backup data |
| REPLBATCHSIZE         | The number of files that are to be replicated in a batch, within the same server transaction                                      |
| REPLSIZETHRESH        | The threshold for the amount of data replicated as a batch, within the same server transaction                                    |
| TXNGROUPMAX           | The number of files that are transferred as a group between a client and the server between transaction commit points             |

## Message options

You can use server options to give you more flexibility in the way Tivoli Storage Manager issues messages.

Table 455. Message options

| Option        | Description                                                                                             |
|---------------|---------------------------------------------------------------------------------------------------------|
| EXPQUIET      | Whether Tivoli Storage Manager sends detailed informational messages during expiration processing       |
| MESSAGEFORMAT | Whether a message number is displayed in all lines of a multi-line message                              |
| MSGINTERVAL   | The time, in minutes, between messages prompting an operator to mount a tape for Tivoli Storage Manager |

## Event logging options

Options can help you manage event logging receivers.

Table 456. Event logging options

| Option               | Description                                                                                        |
|----------------------|----------------------------------------------------------------------------------------------------|
| EVENTSERVER          | Whether the server should try to contact the event server when the server starts up                |
| FILEEXIT             | A file to which enabled events are routed (binary format)                                          |
| FILETEXTEXIT         | A file to which enabled events are routed (readable format)                                        |
| REPORTRETRIEVE       | Record client restore and retrieve operations                                                      |
| TECBEGINEVENTLOGGING | Whether event logging for the Tivoli receiver should begin when the server starts up               |
| TECHOST              | The host name or IP address for the Tivoli event server                                            |
| TECPORT              | The TCP/IP port address on which the Tivoli event server is listening                              |
| TECUTF8EVENT         | A Tivoli Enterprise Console (TEC) event sent from the Tivoli Storage Manager server in UTF8 format |
| UNIQUETDPTECEVENTS   | Tivoli Data Protection (TDP) events sent to the Tivoli Enterprise Console as unique                |
| UNIQUETECEVENTS      | Events sent to the Tivoli Enterprise Console as unique                                             |
| USEREXIT             | A user-defined exit that will be given control to manage an event                                  |

## Security options and licensing options

You can use server options to customize server security and license audits.

Table 457. Security and licensing options

| Option            | Description                                                                                                                                                                                      |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AUDITSTORAGE      | Specifies that during a license audit operation, the server calculates, by node, the amount of backup, archive, and space management storage in use                                              |
| LDAPURL           | Specifies the LDAP directory server. Each setting must have the LDAP directory server name, a port number, and the base distinguished name of the namespace or suffix that the server maintains. |
| QUERYAUTH         | The administrative authority level required to issue QUERY or SQL SELECT commands                                                                                                                |
| REQSYSAUTHOUTFILE | Specifies if system authority is required for administrative commands that cause Tivoli Storage Manager to write to an external file                                                             |
| SHREDDING         | Specifies whether shredding of deleted sensitive data is done automatically or manually                                                                                                          |

## Miscellaneous options

You can use a variety of miscellaneous server options to customize Tivoli Storage Manager.

*Table 458. Miscellaneous options*

| <b>Option</b>   | <b>Description</b>                                                                                                                                                    |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ALIASHALT       | Allows administrators to give the Tivoli Storage Manager HALT command a different name                                                                                |
| DISPLAYLFINFO   | Specifies whether accounting records and summary table entries report the storage agent name                                                                          |
| EXPINTERVAL     | The interval between automatic inventory expiration processes                                                                                                         |
| FFDCLOGNAME     | The name for the first failure data capture (FFDC) log                                                                                                                |
| FFDCMAXLOGSIZE  | The maximum size of the first failure data capture (FFDC) log                                                                                                         |
| NOPREEMPT       | Specifies that no operation can preempt another for access to a volume and that only a database backup operation can preempt another operation for access to a device |
| NORETRIEVEDATE  | Specifies that the server does not update the retrieve date of a file in a disk storage pool when a client restores or retrieves the file                             |
| RESTOREINTERVAL | The length of time that a restartable restore session can be saved in the server database                                                                             |
| VOLUMEHISTORY   | The name of the file to be automatically updated whenever server sequential volume history information is changed                                                     |

---

## 3494SHARED

The 3494SHARED option specifies whether an IBM 3494 library can share applications other than Tivoli Storage Manager.

The default is NO, meaning that no application other than Tivoli Storage Manager can share the 3494. When you set this option to YES, for every mount request, Tivoli Storage Manager determines if each drive is in use. After the query completes, Tivoli Storage Manager selects an available drive that is not in use by another application. Enable sharing only if you have more than two drives in your library. If you are currently sharing an IBM 3494 library with other applications, you must specify this option.

### Syntax

▶▶—3494SHARED—Yes  
No

### Parameters

#### Yes

Specifies that other applications can share the 3494 library.

**No** Specifies that no other applications can share the 3494 library.

### Examples

Enable sharing of a 3494 library:

```
3494shared yes
```

## ACSACCESSID

---

### ACSACCESSID

The ACSACCESSID option specifies the ID for the ACS access control for an ACSLS library.

#### Syntax

▶▶—ACSACCESSID—*name*—————▶▶

#### Parameters

*name*

Specifies a 1 to 64 character ID. The default ID is your local host name.

#### Examples

acsaccessid region

---

## ACSLOCKDRIVE

The ACSLOCKDRIVE option specifies if the drives within the ACSLS libraries are locked. Drive locking ensures the exclusive use of the drive in the ACSLS library in a shared environment. However, there is some performance gain if libraries are not locked. When other applications do not share the Tivoli Storage Manager drives, drive locking is not required.

### Syntax

►► ACSLOCKDRIVE  Yes  No ◄◄

### Parameters

#### Yes

Specifies that drives are locked.

**No** Specifies that drives are not locked.

### Examples

```
acslockdrive yes
```

---

## ACSQUICKINIT

The ACSQUICKINIT option specifies whether, at server startup, the initialization of the ACSLS library is a quick or full initialization. The default is Yes. A quick initialization avoids the overhead associated with synchronizing the Tivoli Storage Manager server inventory with the ACSLS library inventory (through an audit of the library).

### Syntax

►► ACSQUICKINIT  Yes  No ◀◀

### Parameters

#### Yes

Specifies that a quick initialization of the ACSLS library is performed. When the option is set to Yes, Tivoli Storage Manager bypasses library inventory verification, initializing the library quickly, and making it available to Tivoli Storage Manager sooner than if a full initialization is done.

This option should be set to Yes when it is known that the physical library inventory and the Tivoli Storage Manager library inventory have not changed and an audit is not needed.

#### No

Specifies that a full initialization of the ACSLS library and library inventory is performed. When the option is set to No, Tivoli Storage Manager synchronizes its library volume inventory with what is reported by the ACSLS library manager.

### Examples

```
acsquickinit yes
```

---

## ACSTIMEOUTX

The ACSTIMEOUTX option specifies the multiple for the built-in timeout value for ACSLS APIs. The built-in timeout value for the ENTER, EJECT, and AUDIT ACS API is 1800 seconds; for all other ACSLS APIs it is 600 seconds. For example, if the multiple value specified is 5, the timeout value for audit API becomes 9000 seconds, and all other APIs become 3000 seconds.

### Syntax

▶▶—ACSTIMEOUTX—*value*—————▶▶

### Parameters

*value*

Specifies the multiple for the built-in timeout value for ACSLS API. The range is from 1 to 100. The default is 1.

### Examples

```
acstimeoutx 1
```

### ACTIVELOGDIRECTORY

The ACTIVELOGDIRECTORY option specifies the name of the directory where all active logs are stored.

This option is appended to the options file when the DSMSERV FORMAT command is run. Under normal operating conditions, the option does not need to be changed. See “DSMSERV FORMAT (Format the database and log)” on page 1566 for guidance on this option.

#### Syntax

▶▶—ACTIVELOGDirectory—*dir\_name*—————▶▶

#### Parameters

*dir\_name*

Specifies a fully qualified directory name. The directory must exist, it must be empty, and it must be accessible by the user ID of the database manager. If you change the active log directory, Tivoli Storage Manager moves the existing active logs to the location that is specified by this directory. The maximum number of characters is 175.

#### Examples

```
activelogdirectory /tsm/activelogdir
```

## ACTIVELOGSIZE

The ACTIVELOGSIZE option sets the total log size.

This option is appended to the options file when the DSMSERV FORMAT command is run. Under normal operating conditions the option does not need to be changed. See “DSMSERV FORMAT (Format the database and log)” on page 1566 for guidance on this option.

### Syntax

▶▶ ACTIVELOGSize 16GB  
megabytes ▶▶

### Parameters

*megabytes*

Specifies the size of the active log in megabytes. The minimum value is 2048 MB (2 GB); the maximum is 131,072 MB (128 GB). If you specify an odd number, the value is rounded up to the next even number. The default is 16,384 MB (16 GB).

### Examples

activelogsize 8192

---

### ADMINCOMMTIMEOUT

The ADMINCOMMTIMEOUT option specifies how long the server waits for an expected administrative client message during an operation that causes a database update.

If the length of time exceeds this time-out period, the server ends the session with the administrative client. You may want to increase the time-out value to prevent administrative client sessions from timing out.

You can update this server option without stopping and restarting the server by using the SETOPT command. See “SETOPT (Set a server option for dynamic update)” on page 1153.

#### Syntax

►► ADMINCOMMTIMEOUT 60 seconds ◀◀

#### Parameters

*seconds*

Specifies the maximum number of seconds that a server waits for an administrative client response. The default value is 60. The minimum value is 1.

#### Examples

```
admincommtimeout 60
```

---

## ADMINIDLETIMEOUT

The ADMINIDLETIMEOUT option specifies the amount of time, in minutes, that an administrative client session can be idle before the server cancels the session.

If there is a heavy network load in your environment, you might want to increase the time-out value to prevent administrative clients from timing out. However, a large number of idle sessions could prevent other users from connecting to the server.

You can update this server option without stopping and restarting the server by using the SETOPT command. See “SETOPT (Set a server option for dynamic update)” on page 1153.

### Syntax

►► ADMINIDLETIMEOUT 15  
minutes ◀◀

### Parameters

*minutes*

Specifies the maximum number of minutes that a server waits for an idle administrative client. The default value is 15 minutes. The minimum value is 1 minute.

### Examples

```
adminidletimeout 20
```

## ADMINONCLIENTPORT

---

### ADMINONCLIENTPORT

The ADMINONCLIENTPORT option defines whether or not the TCPPOINT can be used by administrative sessions. The default is YES.

#### Syntax



#### Parameters

##### YES

If the option is set to YES, or if the TCPPOINT and TCPADMINPORT are the same value (the default), then administrative sessions can use the TCPPOINT.

**NO** If the option is set to NO, and if the TCPADMINPORT value is different than the TCPPOINT value, then administrative sessions cannot use the TCPPOINT.

#### Examples

Specify that the TCPPOINT can be used by administrative sessions.

```
adminonclientport yes
```

---

## ALIASHALT

The ALIASHALT option allows administrators to give the Tivoli Storage Manager **HALT** command a different name.

The administrative client recognizes an alias for the HALT command when the client is started with the CHECKALIASHALT option specified. See “Administrative client options” on page 5 for details.

### Syntax

▶▶—ALIASHALT—*newname*—————▶▶

### Parameters

*newname*

Specifies the alias of the HALT command for shutting down the Tivoli Storage Manager server. Minimum length of *newname* is 1; maximum length is 16.

### Examples

```
aliashalt tsmhalt
```

---

### ALLOWREORGINDEX

The ALLOWREORGINDEX option specifies whether server-initiated index reorganization is enabled or disabled.

The default is NO.

#### Syntax

▶▶—ALLOWREORGINDEX—No  
Yes—▶▶

#### Parameters

##### Yes

Specifies that server-initiated index reorganization is enabled.

**No** Specifies that server-initiated index reorganization is disabled.

#### Examples

Specify that server-initiated index reorganization is enabled.

```
allowreorgindex yes
```

---

## ALLOWREORGTABLE

The ALLOWREORGTABLE option specifies whether server-initiated table reorganization is enabled or disabled.

The default is YES.

### Syntax

▶▶—ALLOWREORGTABLE  Yes  No ▶▶

### Parameters

#### Yes

Specifies that server-initiated table reorganization is enabled.

**No** Specifies that server-initiated table reorganization is disabled.

### Examples

Specify that server-initiated table reorganization is disabled.

```
allowreorgtable no
```

### ARCHFAILOVERLOGDIRECTORY

The ARCHFAILOVERLOGDIRECTORY option specifies the directory which the server uses to store archive log files that cannot be stored in the archive log directory.

This option is appended to the options file when the DSMSERV FORMAT command is run. Typically the directory does not need to be changed.

#### Syntax

►►—ARCHFailoverlogdirectory—*dir\_name*—◄◄

#### Parameters

*dir\_name*

Specifies a fully qualified directory name. The maximum number of characters is 175.

#### Examples

archfailoverlogdirectory /tsm/archfailoverlog

---

## ARCHLOGDIRECTORY

The ARCHLOGDIRECTORY option specifies a directory that the database manager can archive a log file into after all the transactions represented in that log file are completed.

This option is appended to the options file when the DSMSERV FORMAT command is run.

### Syntax

►►—ARCHLOGDirectory—*dir\_name*—————►►

### Parameters

*dir\_name*

Specifies a fully qualified directory name. The maximum number of characters is 175.

### Examples

archlogdirectory /tsm/archlog

---

### ASSISTVCRRECOVERY

The ASSISTVCRRECOVERY option specifies whether Tivoli Storage Manager assists an IBM 3570 or 3590 drive in recovering from a lost or corrupted Vital Cartridge Records (VCR) condition. If you specify YES (the default) and if Tivoli Storage Manager detects an error during the mount processing, it locates to the end-of-data during the dismount processing to allow the drives to restore the VCR. During the tape operation, there may be some small effect on performance because the drive cannot perform a fast locate with a lost or corrupted VCR. However, there is no loss of data.

#### Syntax

▶▶ ASSISTVCRREcovery  Yes  No ▶▶

#### Parameters

##### Yes

Specifies server assistance in recovery.

##### No

Specifies no server assistance in recovery.

#### Examples

Turn off recovery assistance:

```
assistvcrrecovery no
```

---

## AUDITSTORAGE

As part of a license audit operation, the server calculates, by node, the amount of server storage used for backup, archive, and space-managed files. For servers managing large amounts of data, this calculation can take a great deal of CPU time and can stall other server activity. You can use the AUDITSTORAGE option to specify that storage is not to be calculated as part of a license audit.

**Note:** This option was previously called NOAUDITSTORAGE.

### Syntax

►►—AUDITStorage—┐ Yes  
                  └─ No

### Parameters

#### Yes

Specifies that storage is to be calculated as part of a license audit. The default is Yes.

**No** Specifies that storage is not to be calculated as part of a license audit.

### Examples

```
auditstorage yes
```

---

## CHECKTAPEPOS

The CHECKTAPEPOS option specifies whether the Tivoli Storage Manager server validates the position of data blocks on tape.

The CHECKTAPEPOS option applies only to operations using tape drives. It does not apply to non-tape, sequential-access device classes such as FILE or OPTICAL. If the server information about position does not match the position detected by the drive, an error message is displayed, the transaction is rolled back, and the data is not committed to the database.

Using the CHECKTAPEPOS option, you can enable append-only mode for IBM LTO Generation 5 and later drives, and for any drives that support this feature. When it is enabled, the drive issues an error after receiving instructions to overwrite any data on the currently mounted volume. The Tivoli Storage Manager server repositions the tape to the correct block and continues writing data. Append-only mode provides added protection by preventing most data overwrite situations. If you are using a drive that supports this feature, you can validate data position on tape using both Tivoli Storage Manager and the drive or you can enable one or the other.

**Note:** The Tivoli Storage Manager CHECKTAPEPOS server option does not require an append-only capable drive.

Changes to the CHECKTAPEPOS option only affect mounts after the update to the drive is complete.

The default is YES.

### Syntax



### Parameters

#### Yes

Specifies that the Tivoli Storage Manager server validates data position on tape. For drives that support append-only mode, this parameter specifies that Tivoli Storage Manager enables the drive to also validate the data position during each WRITE operation to prevent data overwrite. This is the default.

**No** Specifies that all data position validation is turned off.

#### TSMonly

Specifies that the Tivoli Storage Manager server validates data position on tape. The server does not use append-only mode even if the drive supports the feature

#### DRIVEonly

Specifies that the Tivoli Storage Manager server enables append-only mode for drives that support this feature. The server does not validate the data position on tape.

**Example**

Validate data position on tape and enable append-only mode for a supported drive:

```
checktapepos yes
```

## CLIENTDEDUPTXNLIMIT

The CLIENTDEDUPTXNLIMIT option specifies the maximum size of a transaction when client-side deduplicated data is backed up or archived.

When you use client-side deduplication for large objects, intensive database activity can result from long-running transactions that are required to update the database. High levels of database activity can produce the following symptoms:

- Reduced throughput for client backup and archive operations
- Resource contention resulting from concurrent server operations
- Excessive recovery log activity

The extent to which these symptoms occur depends on the number and size of objects being stored using client-side data deduplication, the intensity and type of concurrent operations taking place on the Tivoli Storage Manager server, and the Tivoli Storage Manager server configuration.

With the CLIENTDEDUPTXNLIMIT server option, you can specify a maximum size, in gigabytes, for transactions when client-side deduplicated data is backed up or archived. If an object or set of objects in a single transaction exceeds the limit specified by CLIENTDEDUPTXNLIMIT, the objects are not deduplicated by the client, and the transaction can fail. You can specify a value 32 - 1024 GB. The default value is 300 GB.

If an object or set of objects in a single transaction exceeds the limit specified by CLIENTDEDUPTXNLIMIT, the objects or set of objects is not deduplicated by the client. However, the objects are sent to the server. These objects can be deduplicated on the server, depending on whether the destination storage pool is configured for data deduplication and on the value of the SERVERDEDUPTXNLIMIT option. Objects in a deduplication-enabled storage pool that are less than the value of the SERVERDEDUPTXNLIMIT are deduplicated by a server duplicate-identification process.

The appropriate value for this option depends on the Tivoli Storage Manager server configuration and concurrent server activity. You can specify a high value for this option if you minimize resource contention. To minimize resource contention, perform operations, such as backup, archive, duplicate identification (the **IDENTIFY DUPLICATES** command), and reclamation, at different times.

To update this server option without stopping and restarting the server, use the **SETOPT** command.

### Syntax

```

▶▶ CLIENTDEDUPTXN limit 300 gigabytes

```

### Parameters

*gigabytes*

Specifies the maximum size, in gigabytes, of objects that can be backed up or archived using client-side data deduplication. You can specify a value 32 - 1024. The default value is 300.

**Examples**

Disable client-side data deduplication for all objects over 80 GB:

```
clientdeduptxnlimit 80
```

## COMMMETHOD

The COMMMETHOD option specifies a communication method to be used by the server.

You can configure the server to use multiple communication methods. The more commonly used are the TCPIP, V6TCPIP and SHAREDMEM communication methods. To specify multiple communication methods, enable each method by adding a COMMMETHOD stanza to the dsmserv.opt options file.

**Important:** When enabling a communication method, you must also add the options that are specific to the communication method to the options file.

### Syntax



### Parameters

You can choose one of the following communication methods:

#### NONE

Specifies that no communication method is used. This option does not allow users to connect to the server and is useful for experimenting with policy commands.

#### SHAREDMEM

Specifies the shared memory communication method option. This method uses the same area of memory to send data between several applications at the same time. Both the server and the backup-archive client need to be configured to support the shared memory communication method, and they must be installed on the same machine.

#### SNMP

Specifies the SNMP communication method option.

#### TCPIP

Specifies the TCP/IP communication method option. This is the default. When TCPIP is specified, TCP/IP Version 4 is used exclusively.

#### V6TCPIP

Specifies the TCP/IP communication method option. If TCP/IP Version 4 and Version 6 are both configured, Tivoli Storage Manager uses both protocols simultaneously. If both COMMMETHOD TCPIP and COMMMETHOD V6TCPIP are specified, V6TCPIP overrides the specification of TCPIP. A valid domain name server (DNS) environment must be present to use either TCP/IP Version 4 or TCP/IP Version 6 if this option is specified.

### Examples

Example of specifying multiple communication methods to be used by the server (tcpip and v6tcpip):

```
commethod tcpip
commethod v6tcpip
```

---

### COMMTIMEOUT

The **COMMTIMEOUT** option specifies how long the server waits for an expected client message during an operation that causes a database update. If the length of time exceeds this time-out, the server ends the session with the client. You may want to increase the time-out value to prevent clients from timing out. Clients may time out if there is a heavy network load in your environment or they are backing up large files.

The **COMMTIMEOUT** server option is used for non-administrative sessions. See the **ADMINCOMMTIMEOUT** option for administrative client sessions.

You can update this server option without stopping and restarting the server by using the **SETOPT** command.

#### Syntax

►►—COMMTimeout—60  
seconds—►►

#### Parameters

*seconds*

Specifies the maximum number of seconds that a server waits for a client response. The default value is 60. The minimum value is 1.

#### Examples

```
commtimeout 60
```

## DBDIAGLOGSIZE

This option helps to control the maximum size of the diagnostic log files.

The database manager uses the diagnostic log files to log messages intended for troubleshooting purposes. If the log files are not controlled, they can grow to fill the file system. Use the DBDIAGLOGSIZE option to set the maximum size of the diagnostic log files.

There can be up to 10 rotating diagnostic log files. Each file name indicates the order in which the file was created. After a file is full, it is closed, and the next file is created. When the 10th file is closed, the oldest file is deleted, and a new file is created. Here is an example of how the rotating log files might look:

```
db2diag.14.log, db2diag.15.log, ... , db2diag.22.log, db2diag.23.log
```

When db2diag.23.log is full, db2diag.14.log is deleted, and db2diag.24.log is created.

The server checks the file space containing the diagnostic log files every hour. Messages are displayed every 12 hours if either of the following conditions occur:

- The available space in the file system where the diagnostic log files are located is less than 20% of the total file system space.
- The available space in the file system where the server instance directory is located is less than 1 GB.

For more information about the db2diag log files, see <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp>.

### Syntax

►►—DBDIAGLOGSize—<sup>1024</sup>—*megabytes*—►►

### Parameters

*megabytes*

Specifies the size of the diagnostic log files in megabytes. The default value is 1024. The minimum value is 2; the maximum value is 9999.

### Examples

Specify the size of the diagnostic log files as 5120 megabytes (5 GB):

```
dbdiaglogsize 5120
```

---

### DBDIAGPATHFSTHRESHOLD (Specify a threshold for free space)

The DBDIAGPATHFSTHRESHOLD option specifies the threshold for free space on the file system or disk that contains the db2diag.log file.

When the amount of free space is equal to or less than the specified threshold, the ANR1545W error message is shown. By default, the message is shown when the file system or disk has 20% or less of free disk space.

You can update this server option without stopping and restarting the server by using the **SETOPT** command. See “SETOPT (Set a server option for dynamic update)” on page 1153.

#### Syntax

►►—DBDIAGPATHFSTHreshold—*percent*—————►

#### Parameter

*percent*

Specifies the percentage of available space in the file system. Valid values are in the range 0 - 100. The default is 20.

**Tip:** For best results, do not set a low or high value for the **percent** parameter. A low value might cause the file system to become full before you can correct the issue. A full file system might corrupt the server database. A high value might result in many ANR1545W messages in the server activity log.

#### Example

Set the threshold value to 10%.

```
setopt DBDIAGPATHFSTH 10
```

---

## DBMEMPERCENT

Use this option to specify the percentage of the virtual address space that is dedicated to the database manager processes.

If applications other than Tivoli Storage Manager server are running on the system, ensure that the value allows adequate memory for the other applications.

### Syntax

►►—DBMEMPERCENT—*percent*—AUTO—◄◄

### Parameters

*percent*

Set a value from 10 to 99.

**AUTO**

The database manager sets the percentage automatically to a value that is between 75 percent and 95 percent of system RAM. The default value is AUTO.

### Examples

dbmempercent 50

---

## DBMTCPPORT

The DBMTCPPORT option specifies the port number on which the TCP/IP communication driver for the database manager waits for requests for client sessions.

The specified port number must be reserved for use by the database manager.

By default, the Tivoli Storage Manager server uses interprocess communications (IPC) to establish connections for the first two connection pools, with a maximum of 480 connections for each pool. After the first 960 connections are established, the Tivoli Storage Manager server uses TCP/IP for any additional connections.

### Syntax

►—DBMTCPPort—*port\_number*—◄

### Parameters

*port\_number*

Specifies the number of the TCP/IP port on which the database manager waits for communications from the server. Valid values are integers from 1024 to 65535.

The default port number is the value of the server TCPPORT option plus 50,000. For example, if the server TCPPORT option is 1500, the default DBMTCPPORT port number would be 51500.

If the TCPPORT server option is greater than 9999, add the last four digits of its value to 50000. For example, if the TCPPORT option is 11500, 1550 is added to 50000, resulting in a DBMTCPPORT port number of 51500.

### Example

```
dbmtcport 51500
```

# DEDUPREQUIRESBACKUP

The DEDUPREQUIRESBACKUP option specifies whether volumes in primary sequential-access storage pools that are set up for data deduplication can be reclaimed and whether duplicate data can be discarded before the storage pools are backed up.

If the value of this option is YES (the default), you must back up data to copy storage pools that are not set up for data deduplication. Use the BACKUP STGPOOL command to back up data to copy storage pools.

Be aware that reclamation of a volume in a storage pool that is set up for data deduplication might not occur when the volume first becomes eligible. The server makes additional checks to ensure that data from a storage pool that is set up for data deduplication has been backed up to a copy storage pool. These checks require more than one BACKUP STGPOOL instance before the server reclaims a volume. After the server verifies that the data was backed up, the volume is reclaimed.

You can change this option dynamically using the SETOPT command.

**Attention:** To minimize the possibility of data loss, do not change the default setting for this server option. Specify a value of NO only if you do not have any copy storage pools and are not performing storage pool backups.

### Syntax



### Parameters

- Yes** Specifies that the storage pool must be backed up before volumes can be reclaimed and before duplicate data can be discarded. This is the default.
- No** Specifies that volumes in primary sequential-access storage pools that are set up for data deduplication can be reclaimed and duplicate data can be discarded if the storage pools are not backed up.

### Examples

Specify that primary sequential-access storage pools that are set up for data deduplication do not have to be backed up.

```
deduprequiresbackup no
```

---

### DEDUPTIER2FILESIZE

The DEDUPTIER2FILESIZE option specifies at what file size Tivoli Storage Manager begins to use Tier 2 data deduplication.

#### Syntax

►►—DEDUPTIER2FILESIZE—*nnn*—◄◄

#### Parameters

*nnn*

Specifies the file size, in gigabytes, at which point the Tivoli Storage Manager server begins using Tier 2 chunking for data deduplication. You can specify a value from 20 to 9999. The default is 100.

**Note:** If the value specified or defaulted to for this option is greater than the value for the SERVERDEDUPTXNLIMIT option, then this option is ignored for server data deduplication. If the value specified or defaulted to for this option is greater than the value for CLIENTDEDUPTXNLIMIT, then this option is ignored for client data deduplication.

#### Examples

```
deduptier2filesize 550
```

---

## DEDUPTIER3FILESIZE

The DEDUPTIER3FILESIZE option specifies at what file size Tivoli Storage Manager begins to use Tier 3 data deduplication.

### Syntax

►►—DEDUPTIER3FILESIZE—*nnn*—◄◄

### Parameters

*nnn*

Specifies the file size, in gigabytes, at which point the Tivoli Storage Manager server begins using Tier 3 chunking for data deduplication. You can specify a value from 90 to 9999. The default is 400.

- If the value specified or defaulted to for this option is greater than the value for the SERVERDEDUPTXNLIMIT option, then this option is ignored for server data deduplication.
- If the value specified or defaulted to for this option is greater than the value for CLIENTDEDUPTXNLIMIT, then this option is ignored for client data deduplication.
- If the value specified or defaulted to for this option is less than the value specified or defaulted to for DEDUPTIER2FILESIZE then the value of DEDUPTIER2FILESIZE is used for this option.

### Examples

```
deduptier3filesize 1150
```

---

## DEVCONFIG

The DEVCNFIG option specifies the name of a file in which you want Tivoli Storage Manager to store a backup copy of device configuration information.

Tivoli Storage Manager stores the following information in the device configuration file:

- Device class definitions created by using the **DEFINE DEVCLASS** command
- Drive definitions created by using the **DEFINE DRIVE** command
- Library definitions created by using the **DEFINE LIBRARY** command
- Library inventory information for the LIBTYPE=SCSI automated libraries
- Path definitions created by using the **DEFINE PATH** command
- Server definitions created with the **DEFINE SERVER** command
- Server name created with the **SET SERVERNAME** command
- Server password created with the **SET SERVERPASSWORD** command

### Note:

- Only path definitions with SRCTYPE=SERVER are backed up to the device configuration file. Paths of SRCTYPE=DATAMOVER are *not* written to the file.
- Library volume location information is stored as comments (*/\*...\*/*) in the device configuration file whenever **CHECKIN LIBVOLUME**, **CHECKOUT LIBVOLUME**, and **AUDIT LIBRARY** commands are issued for SCSI libraries.

**Attention:** To restore the database after a disaster, you must have a copy of the current device configuration file. The device configuration file cannot be recreated.

You can include one or more DEVCNFIG options in the server options file. When you use multiple DEVCNFIG options, Tivoli Storage Manager automatically updates and stores a backup copy of device configuration information in each file you specify.

## Syntax

►►—DEVCONFig—*file\_name*—◄◄

## Parameters

*file\_name*

Specifies the name of a file in which to store a backup copy of device configuration information.

## Examples

devconfig devices.sav

---

## DISABLESCHEDS

The DISABLESCHEDS option specifies whether administrative and client schedules are disabled during Tivoli Storage Manager server recovery.

### Syntax

►►—DISABLEScheds—

|     |
|-----|
| Yes |
| No  |

—————►

### Parameters

#### Yes

Specifies that administrative and client schedules are disabled.

#### No

Specifies that administrative and client schedules are enabled.

### Examples


```
disablescheds no
```

## DISPLAYLFINFO

The DISPLAYLFINFO option specifies how the accounting records and summary table entries report the node name.

When this option is enabled, the accounting records and summary table entries report *node\_name(storage\_agent\_name)* for the node name. If the option is not enabled, the accounting records and summary table entries simply report *node\_name* for the node name. The default is No.

### Syntax



### Parameters

#### Yes

Specifies that the accounting records and summary table entries will report the storage agent name.

#### No

Specifies that the accounting records and summary table entries will not report the storage agent name. This is the default.

### Examples

```
displaylfinfo yes
```

The result shows the following accounting record with the storage agent name displayed (STA53):

```
5,0,ADSM,07/13/2004,15:35:14,COLIND-TUC(STA53),,WinNT,1,Tcp/Ip,1,0,0,0,
0,223,4063,0,0,222,7,8,3,1,4,0,0,0,0,3,0
```

The corresponding summary table also displays the storage agent name:

```
START_TIME: 2004-07-13 15:35:07.000000
END_TIME: 2004-07-13 15:35:14.000000
ACTIVITY: BACKUP
NUMBER: 8
ENTITY: COLIND-TUC(STA53)
COMMMETH: Tcp/Ip
ADDRESS: colind-tuc:2229
SCHEDULE_NAME:
EXAMINED: 0
AFFECTED: 223
FAILED: 0
BYTES: 4160875
IDLE: 8
MEDIAM: 1
PROCESSES: 1
SUCCESSFUL: YES
VOLUME_NAME:
DRIVE_NAME:
LIBRARY_NAME:
LAST_USE:
COMM_WAIT: 3
NUM_OFFSITE_VOLS:
```

---

## DNSLOOKUP

The DNSLOOKUP option specifies whether the server uses system API calls to determine the domain name server (DNS) names of systems that contact the server.

### Syntax



### Parameters

#### Yes

Specifies that the server obtains the DNS names of contacting systems. Yes is the default.

**No** Specifies that the server does not obtain the DNS names of contacting systems.

### Examples

```
dnslookup yes
```



---

## ENABLENASDEDUP

The ENABLENASDEDUP server option specifies whether the server deduplicates data that is stored by a network-attached storage (NAS) file server. This option applies only to NetApp file servers.

If the value of this option is NO, the data stored by the file server is skipped during duplicate-identification processing. If the value of this option is YES, the value of the **DEDUPLICATE** parameter in the storage pool definition must be YES.

### Syntax

►►—ENABLENASDEDUP—No—Yes—►►

### Parameters

#### Yes

Specifies that Tivoli Storage Manager server deduplicates data stored by a NetApp file server.

**No** Specifies that the server does not deduplicate data stored by a NetApp file server.

### Example

Specify that the server deduplicates data stored by a NetApp file server.

```
enablenasdedup yes
```

---

## EVENTSERVER

The EVENTSERVER option specifies whether at startup the server should try to contact the event server.

### Syntax



### Parameters

#### Yes

Specifies that, at startup, the server tries to contact the event server. Contact occurs only if a DEFINE EVENTSERVER command has already been issued. This is the default.

**No** Specifies that, at startup, the server does not try to contact the event server.

### Examples

```
eventserver yes
```

---

## EXPINTERVAL

The EXPINTERVAL option specifies the interval, in hours, between automatic inventory expiration processes by Tivoli Storage Manager. Inventory expiration removes client backup and archive file copies from the server as specified by the management classes to which the client files are bound. If expiration is not run periodically, storage pool space is not reclaimed from expired client files, and the server requires more storage space than required by policy.

You can also use the EXPIRE INVENTORY command to start inventory expiration. Expiration can make space available in your storage pools for additional client backup or archive files.

You can update this server option without stopping and restarting the server by using the SETOPT command. See “SETOPT (Set a server option for dynamic update)” on page 1153.

### Syntax

►—EXPINterval 24 hours —►

### Parameters

*hours*

Specifies the time, in hours, between automatic inventory expiration processes. You can specify from 0 to 336 (14 days). A value of 0 means that expiration must be started with the EXPIRE INVENTORY command. The default is 24.

### Examples

```
expinterval 5
```

---

## EXPQUIET

The EXPQUIET option specifies whether Tivoli Storage Manager sends detailed messages during expiration processing.

You can update this server option without stopping and restarting the server by using the SETOPT command. See “SETOPT (Set a server option for dynamic update)” on page 1153.

### Syntax

►►—EXPQUIET— [ —No—  
                  [ —Yes— ]

### Parameters

**No** Specifies that the server sends detailed messages. This is the default.

**Yes**

Specifies that the server sends only minimal messages. These messages are sent only for files that have expired based on the copy group in the default management class or retention grace period for the domain.

### Examples

```
expquiet no
```

---

## FFDCLOGNAME

The FFDCLOGNAME option specifies a name for the first failure data capture (FFDC) log.

The FFDC log is used to gather diagnostic information about the server. When an error occurs, data about the error is written to the FFDC log. This information can be provided to IBM support to help diagnose problems. The FFDC log is in the server instance directory.

### Syntax

►►—FFDCLOGNAME—dsmserv.err  
file\_name—►►

### Parameters

*file\_name*

Specifies a file name for the FFDC log. The fully qualified file name is relative to the server instance directory. The default value is dsmserv.err.

### Examples

```
ffdclogname /tsminst1/tsmffdc.log
ffdclogname tsmffdc.log
ffdclogname c:\tsmserv1\tsmffdc.log
```

### FFDCMAXLOGSIZE

The FFDCMAXLOGSIZE option specifies the size for the first failure data capture (FFDC) log.

The FFDC log is used to gather diagnostic information about the server. When an error occurs, data about the error is written to the FFDC log. This information can be provided to IBM support to help diagnose problems.

#### Syntax

►►—FFDCMAXLOGSIZE—500  
kilobytes—►►

#### Parameters

*kilobytes*

Specifies the size to which the FFDCLOG can grow before wrapping. The default value is 500.

To allow the size of the log to grow indefinitely, specify a value of -1. To disable the log, specify 0.

#### Examples

```
ffdcmaxlogsize 2000
```

## FILEEXIT

The FILEEXIT option specifies a file to which enabled events are routed. Each logged event is a record in the file.

### Syntax



### Parameters

#### Yes

Specifies that event logging to the file exit receiver begins automatically at server startup.

**No** Specifies that event logging to the file exit receiver does not begin automatically at server startup. When this parameter has been specified, you must begin event logging manually by issuing the `BEGIN EVENTLOGGING` command.

#### *file\_name*

Specifies the name of the file in which the events are stored.

#### **REPLACE**

Specifies that if the file already exists, it will be overwritten.

#### **APPEND**

Specifies that if the file already exists, data is appended to it.

#### **PRESERVE**

Specifies that if the file already exists, it will not be overwritten.

### Examples

```
fileexit yes /tsm/server/data replace
```

---

## FILETEXTXIT

The FILETEXTXIT option specifies a file to which enabled events are routed. Each logged event is a fixed-size, readable line.

### Syntax



### Parameters

#### Yes

Specifies that event logging to the file exit receiver begins automatically at server startup.

**No** Specifies that event logging to the file exit receiver does not begin automatically at server startup. When this parameter has been specified, you must begin event logging manually by issuing the BEGIN EVENTLOGGING command.

#### *file\_name*

Specifies the name of the file in which the events are stored.

#### REPLACE

Specifies that if the file already exists, it will be overwritten.

#### APPEND

Specifies that if the file already exists, data will be appended to it.

#### PRESERVE

Specifies that if the file already exists, it will not be overwritten.

### Examples

```
filetextexit yes /tsm/server/data replace
```

---

## FSUSEDTHRESHOLD

The FSUSEDTHRESHOLD option specifies what percentage of the file system can be filled up by the database before an alert message is issued.

You can update this server option without stopping and restarting the server by using the **SETOPT** command.

If this value is set to a low number, the activity log might be flooded with messages about the database space being filled, even if there is still space available. If the value is set too high, the database space might be filled before you can add more space to the file system.

### Syntax

►►—FSUSEDThresho1d—*percent*—————►►

### Parameters

*percent*

Specifies the value of used space in the database. You can specify a value from 0 to 100. The default is 90.

### Examples

```
fsusedthreshold 70
```

---

## IDLETIMEOUT

The **IDLETIMEOUT** option specifies the amount of time, in minutes, that a client session can be idle before the server cancels the session. You may want to increase the time-out value to prevent clients from timing out if there is a heavy network load in your environment. Note, however, that a large number of idle sessions could prevent other users from connecting to the server.

The **IDLETIMEOUT** server option is used for non-administrative sessions. See the **ADMINIDLETIMEOUT** option for administrative client sessions.

You can update this server option without stopping and restarting the server by using the **SETOPT** command.

### Syntax

►► IDLETimeout minutes 15 ◀◀

### Parameters

*minutes*

Specifies the maximum number of minutes that a server waits for an idle client. The default value is 15 minutes. The minimum value is 1 minute.

### Examples

```
idletimeout 15
```

## LANGUAGE

The LANGUAGE option controls the initialization of locales. A locale includes the language and the date, time, and number formats to be used for the console and server.

If your client and server are running different languages, the messages generated may not be understandable when messages are issued from the client to the server or if the server sends output to the client.

If initialization of the locale fails, the server defaults to American English.

### Syntax



### Parameters

#### en\_US

Specifies that American English will be used as the default language for the server.

#### locale

Specifies the name of the locale supported by the server. See the following tables for information on supported locales by operating system.

**Note:** Tivoli Storage Manager runs in any locale, but defaults to American English. For the locales listed, language support is available. Refer to the *Installation Guide* for information on installing language support.

Table 459. Server languages for Linux

| LANGUAGE               | LANGUAGE option value |
|------------------------|-----------------------|
| Chinese, Simplified    | zh_CN                 |
|                        | zh_CN.gb18030         |
|                        | zh_CN.utf8            |
| Chinese, Traditional   | Big5 / Zh_TW          |
|                        | zh_TW                 |
|                        | zh_TW.utf8            |
| English, United States | en_US                 |
|                        | en_US.utf8            |
| French                 | fr_FR                 |
|                        | fr_FR.utf8            |
| German                 | de_DE                 |
|                        | de_DE.utf8            |
| Italian                | it_IT                 |
|                        | it_IT.utf8            |
| Japanese               | ja_JP                 |
|                        | ja_JP.utf8            |

## LANGUAGE

Table 459. Server languages for Linux (continued)

| LANGUAGE              | LANGUAGE option value |
|-----------------------|-----------------------|
| Korean                | ko_KR                 |
|                       | ko_KR.utf8            |
| Portuguese, Brazilian | pt_BR                 |
|                       | pt_BR.utf8            |
| Russian               | ru_RU                 |
|                       | ru_RU.utf8            |
| Spanish               | es_ES                 |
|                       | es_ES.utf8            |

### Examples

```
lang ja_JP
```

## LDAPCACHEDURATION

The **LDAPCACHEDURATION** option determines the amount of time that the IBM Tivoli Storage Manager server caches LDAP password authentication information.

After a successful LDAP bind, the value that you enter determines the amount of time that information about the LDAP directory server is kept available. The higher the number, the better the performance of the LDAP directory server. During the cache period, though, changes on the LDAP directory server do not take immediate effect on the node. For example, old passwords might be available for some time, even after they were changed or locked on the LDAP server.

Include the **LDAPCACHEDURATION** option in a **SETOPT** command to have the option take effect immediately.

**Restriction:** The **LDAPCACHEDURATION** option does not apply to storage agents.

### Syntax

▶—LDAPCACHEDURATION—*minutes*—▶

### Parameters

*minutes*

Specifies the maximum amount of time after a successful LDAP bind, that subsequent sessions to the same node or administrator skip secondary LDAP bind operations. Values range from zero to 360 minutes.

### Example: Set the LDAPCACHEDURATION value to 6 hours (maximum)

In the `dsmserv.opt` file, specify the following value:

```
ldapcacheduration 360
```

After a node or administrator authenticates with an external directory server, the LDAP bind is skipped for 360 minutes on all sessions.

## LDAPURL

The **LDAPURL** option specifies the location of your LDAP directory server. Set the **LDAPURL** option after you configure the LDAP directory server.

The **LDAPURL** option cannot be used in combination with the **SETOPT** command.

### Syntax

▶—LDAPURL—*ldap\_url\_value*—▶

### Parameters

#### *ldap\_url\_value*

Specifies the URL of one LDAP directory server, or the URLs of multiple LDAP directory servers. You can enter multiple values, with each URL value up to 1024 characters. The port number is optional and defaults to 389. Each URL value must contain the following parameters:

- An LDAP directory server name
- The base distinguished name (Base DN) of the namespace or suffix that the LDAP directory server maintains. The DN format must comply with the directory server that you choose. For example:

|               |                               |
|---------------|-------------------------------|
| server name:  | mainframe.storage.memphis.com |
| ldap port:    | 389                           |
| ldap base dn: | cn=memphis_data               |

The LDAPURL must conform to the following specifications:

- When you specify multiple URLs, each URL must be on a separate line.
- If you want to specify multiple URLs, each URL must point to a different external directory and all external directories must contain the same data.
- Each URL must begin with ldap://.

**Restriction:** The URL that you designate cannot begin with ldaps://.

Tivoli Storage Manager supports LDAP connections that are secured with the standard LDAPv3 “StartTLS” operation that establishes a secure Transport Layer Security (TLS) exchange on an existing LDAP connection. The “LDAP Simple Bind” operation that Tivoli Storage Manager uses does not protect the password when it is sent. A secure TLS connection is needed to protect the password.

**Restriction:** The **LDAPURL** option does not apply to storage agents.

### Example: Set the LDAP value for one LDAP directory server

In the `dsmserv.opt` file, specify the following value:

```
ldapurl ldap://mainframe.storage.memphis.com:389/cn=memphis_data
```

### Example: Set the LDAP value for three LDAP directory servers

In the `dsmserv.opt` file, specify the following values:

```
ldapur1 ldap://mongo.host.hawaii/cn=southernhemisphere,dc=com
ldapur1 ldap://server.windows.y2k:484/cn=company,dc=com
ldapur1 ldap://server.windows.vista:232/cn=manufactured,dc=com
```

**Remember:** You do not have to enter the default port if you want to use it. The first entry in our example is using the default port (389).

---

## MAXSESSIONS

The MAXSESSIONS option specifies the maximum number of simultaneous client sessions that can connect with the server.

You can update this server option without stopping and restarting the server by using the SETOPT command. See “SETOPT (Set a server option for dynamic update)” on page 1153.

### Syntax

▶▶ MAXSessions number\_of\_sessions ◀◀

### Parameters

*number\_of\_sessions*

Specifies the maximum number of simultaneous client sessions. The default value is 25 client sessions. The minimum value is 2 client sessions. The maximum value is limited only by available virtual storage size or communication resources.

### Examples

maxsessions 25

---

## MESSAGEFORMAT

The MESSAGEFORMAT option specifies whether a message number is displayed in all lines of a multi-line message.

### Syntax

►►—MESSageformat—*number*—◄◄

### Parameters

*number*

Select a number to specify if a message number is to be displayed only on the first line of a multi-line message or is to be displayed on all lines.

- 1 The message number for a message is displayed only in the first line of the message. This is the default.
- 2 The message number for a message is displayed in all lines of a message.

### Examples

messageformat 2

---

### MIRRORLOGDIRECTORY

The MIRRORLOGDIRECTORY option specifies the directory for mirroring the active log path.

All changes made to the active log directory are also written to this mirror directory. This option is appended to the options file when the DSMSEVER FORMAT command is run. Typically, the directory does not need to be changed.

#### Syntax

▶—MIRRORlogdirectory—*dir\_name*—▶

#### Parameters

*dir\_name*

Specifies a fully qualified directory name for the active log mirror. The maximum number of characters is 175.

#### Examples

mirrorlogdirectory /tsm/mirrorlog

---

## MOVEBATCHSIZE

The MOVEBATCHSIZE option specifies the number of client files that are to be moved and grouped together in a batch, within the same server transaction. This data movement results from storage pool backups and restores, migration, reclamation, and MOVE DATA operations. This option works with the MOVESIZETHRESH option.

### Syntax

►► MOVEBatchsize number\_of\_files 1000 ◄◄

### Parameters

*number\_of\_files*

Specifies a number of files between 1 and 1000. The default is 1000.

### Examples

movebatchsize 100

---

### MOVESIZETHRESH

The MOVESIZETHRESH option specifies, in megabytes, a threshold for the amount of data moved as a batch, within the same server transaction. When this threshold is reached, no more files are added to the current batch, and a new transaction is started after the current batch is moved.

#### Syntax

►►—MOVESizethresh—<sup>4096</sup>—*megabytes*—►►

#### Parameters

*megabytes*

Specifies the number of megabytes as an integer from 1 to 32768. The default value is 4096. This option is used with the MOVEBATCHSIZE option.

#### Examples

movesizethresh 500

---

## MSGINTERVAL

The MSGINTERVAL option specifies the time, in minutes, between messages prompting an operator to mount a tape for the server.

### Syntax

►►—MSGINTERval—<sup>1</sup>minutes—►►

### Parameters

*minutes*

Specifies the time interval at which the operator is prompted by the server to mount a tape. The default value is 1 minute. The minimum value is 1 minute.

### Examples

msginterval 2

## NDMPCONTROLPORT

---

### NDMPCONTROLPORT

The NDMPCONTROLPORT option specifies the port number to be used for internal communications for certain Network Data Management Protocol (NDMP) operations. The Tivoli Storage Manager server does not function as a general purpose NDMP tape server.

#### Syntax

►►—NDMPControlport—port\_number—◄◄

#### Parameters

*port\_number*

The port number to be used for internal communications for certain NDMP operations. The port number must be from 1024 to 32767. The default is 10000.

#### Examples

```
ndmpcontrolport 9999
```

## NDMPENABLEKEEPALIVE

The NDMPENABLEKEEPALIVE server option specifies whether the Tivoli Storage Manager server enables Transmission Control Protocol (TCP) keepalive on network data-management protocol (NDMP) control connections to network-attached storage (NAS) devices. The default is NO.

TCP keepalive is implemented within the network support of an operating system. TCP keepalive prevents a long-running, inactive connection from being closed by firewall software that detects and closes inactive connections.

**Restriction:** To prevent errors, do not enable TCP keepalive in certain types of environments. One example is environments that do not have firewalls between the Tivoli Storage Manager server and a NAS device. Another example is environments with firewalls that tolerate long-running, inactive connections. Enabling TCP keepalive in this type of environment can cause an idle connection to be inadvertently closed if the connection partner temporarily fails to respond to TCP keepalive packets.

### Syntax

►►—NDMPENABLEKEEPALIVES—NO  
YES—►►

### Parameters

**NO** Disable TCP keepalive on all NDMP control connections. NO is the default.

#### YES

Enable TCP keepalive on all NDMP control connections. The default idle time before the first TCP keepalive packet is sent is 120 minutes.

To change the idle time, use the NDMPKEEPIDLEMINUTES server option.

### Example

Enable TCP keepalive on all NDMP control connections so that inactive NDMP connections are not closed:

```
ndmpenablekeepalive yes
```

---

### NDMPKEEPIDLEMINUTES

The NDMPKEEPIDLEMINUTES server option specifies the amount of time, in minutes, before the operating system transmits the first Transmission Control Protocol (TCP) keepalive packet on a network data-management protocol (NDMP) control connection. The default is 120 minutes.

**Prerequisite:** Use this option only after you set the value of the NDMPENABLEKEEPALIVES server option to YES.

#### Syntax

►►—NDMPKEEPIDLEMINUTES—<sup>120</sup>minutes—►►

#### Parameters

*minutes*

The number of minutes of inactivity on NDMP control connections before TCP keepalive packets are transmitted. The default value is 120. The minimum is 1 minute. The maximum is 600 minutes.

#### Example

Specify an idle time of 15 minutes before the first TCP keepalive packet is sent:

```
ndmpkeepidleminutes 15
```

## NDMPPORTRANGE

The NDMPPORTRANGE option specifies the range of port numbers through which Tivoli Storage Manager cycles to obtain a port number for accepting a session from a network-attached storage (NAS) device for data transfer. The default is 0,0 which means that Tivoli Storage Manager lets the operating system provide a port (ephemeral port).

If all ports specified are in use when a NAS device attempts to connect to the server, the operation fails. If a single port number is chosen (no comma and no port number for the high value), the default for the high port number is the low port number plus 100.

When Network Data Management Protocol (NDMP) data is directed to a Tivoli Storage Manager native pool, communication can be initiated from either the NDMP systems or the Tivoli Storage Manager server. If a firewall separates the server and NAS devices, it may be necessary to specify port numbers in firewall rules to allow traffic to pass to and from the NAS devices. NAS devices communicate to the Tivoli Storage Manager server the port numbers that they will use when contacting the server. The port numbers of the server are controlled with the NDMPPortrange options. Port number control for NAS devices is specific to vendors. Consult your vendor documentation.

### Syntax

```
▶▶—NDMPPortrange—port_number_low—┌───┐───▶▶
 └─,port_number_high─┘
```

### Parameters

*port\_number\_low*

The low port number from which Tivoli Storage Manager starts to cycle when needing a port number for accepting session from a NAS device for data transfer. The minimum port number value is 1024.

*port\_number\_high*

The high port number to which Tivoli Storage Manager can cycle when needing a port number for accepting session from a NAS device for data transfer. The maximum port number value is 32767. The high port number must be the same or larger than the low port number.

### Examples

Specify that Tivoli Storage Manager can cycle from port numbers 1024 - 2024.

```
ndmportrange 1024,2024
```

### NDMPPREFDATAINTERFACE

This option specifies the IP address associated with the interface in which you want the server to receive all Network Data Management Protocol (NDMP) backup data.

This option affects all subsequent NDMP filer-to-server operations, but does not affect NDMP control connections, which use the system's default network interface. The value for this option is a host name or IPV4 address that is associated with one of the active network interfaces of the system on which the Tivoli Storage Manager server is running. This interface must be IPV4 enabled.

You can update this server option without stopping and restarting the server by using the **SETOPT** command. For more information, see the Administrator's Guide.

#### Syntax

►—NDMPPREFDATAINTERFACE—*ip\_address*—◄

#### Parameters

*ip\_address*

Specify an address in either dotted decimal or host name format. If you specify a dotted decimal address, it is not verified with a domain name server. If the address is not correct, it can cause failures when the server attempts to open a socket at the start of an NDMP filer-to-server backup.

Host name format addresses are verified with a domain name server. There is no default value. If a value is not set, all NDMP operations will use the Tivoli Storage Manager server's network interface for receiving backup data during NDMP filer-to-server backup operations. To clear the option value, specify the SETOPT command with a null value, "".

#### Examples:

```
ndmpprefdatainterface net1.tucson.ibm.com
ndmpprefdatainterface 9.11.152.89
```

---

## NOPREEMPT

The server allows certain operations to preempt other operations for access to volumes and devices. You can specify the NOPREEMPT option to disable preemption. When preemption is disabled, no operation can preempt another for access to a volume, and only a database backup operation can preempt another operation for access to a device.

For example, a client data restore operation preempts a client data backup for use of a specific device or access to a specific volume.

See the *Administrator's Guide* for details.

### Syntax

▶▶—NOPREEMPT—◀◀

### Parameters

None

### Examples

Disable preemption among server operations:

```
nopreempt
```

---

### NORETRIEVEDATE

The `NORETRIEVEDATE` option specifies that the server does not update the retrieve date of a file in a disk storage pool when a client restores or retrieves the file. This option and the `MIGDELAY` storage pool parameter control when the server migrates files.

If you do not specify `NORETRIEVEDATE`, the server migrates files after they have been in the storage pool for the number of days specified by the `MIGDELAY` parameter. The number of days is counted from the day that the file was stored in the storage pool or retrieved by a client, whichever is more recent. If you specify `NORETRIEVEDATE`, the server does not update the retrieve date of a file, and the number of days is counted from the day the file entered the disk storage pool.

If you specify this option and caching is enabled for a disk storage pool, reclamation of cached space is affected. When space is needed in a disk storage pool that contains cached files, the server gets the space by selectively erasing cached copies. Files that have the oldest retrieve dates and occupy the largest amount of space are selected for removal. When you specify `NORETRIEVEDATE`, the server does not update the retrieve date when a file is retrieved. This may cause cached copies to be removed even though they have recently been retrieved by a client.

#### Syntax

▶▶—`NORETRIEVEDATE`—▶▶

#### Parameters

None.

#### Examples

Specify that the retrieve dates of files in disk storage pools are not updated when clients restore and retrieve the files:

```
noretrievevdate
```

---

## NUMOPENVOLSALLOWED

The NUMOPENVOLSALLOWED option specifies the number of input FILE volumes in a deduplicated storage pool that can be open at one time.

Input volumes contain data to be read during client-restore operations and server processes, such as reclamation and migration. Use this option to improve performance by reducing the frequency with which volumes are opened and closed.

Each session within a client operation or server process can have as many open FILE volumes as specified by this option. A session is initiated by a client operation or by a server process. Multiple sessions can be started within each.

During a client restore operation, volumes can remain open for the duration of a client restore operation and as long a client session is active. During a no-query restore operation, the volumes remain open until the no-query restore completes. At that time, all volumes are closed and released. However, for a classic restore operation started in interactive mode, the volumes might remain open at the end of the restore operation. The volumes are closed and released when the next classic restore operation is requested.

Set this value in the server options file or use the SETOPT command.

**Tip:** This option can significantly increase the number of volumes and mount points in use at any one time. To optimize performance, follow these steps:

- To set NUMOPENVOLSALLOWED, select a beginning value (the default is recommended). Monitor client sessions and server processes. Note the highest number of volumes open for a single session or process. Increase the setting of NUMOPENVOLSALLOWED if the highest number of open volumes is equal to the value specified by NUMOPENVOLSALLOWED.
- To prevent sessions or processes from having to wait for a mount point, increase the value of the MOUNTLIMIT parameter in the device-class definition. Set the value of the MOUNTLIMIT parameter high enough to allow all client sessions and server processes using deduplicated storage pools to open the number of volume specified by the NUMOPENVOLSALLOWED option. For client sessions, check the destination in the copy group definition to determine how many nodes are storing data in the deduplicated storage pool. For server processes, check the number of processes allowed for each process for the storage pool.
- A situation might occur in which a node backs up and restores or archives and retrieves concurrently to and from a deduplicated storage pool. All the mount points required for these operations increase the total number of mount points required by the node.

As a result, the node might not be able to start additional backup sessions if it already has more mount points open than what the MAXNUMMP parameter in the client-node definition allows. This can occur even though the MOUNTLIMIT for the device class was not exceeded.

To prevent backup and retrieve operations from failing, set the value of the MAXNUMMP parameter in the client-node definition to a value at least as high as the NUMOPENVOLSALLOWED option. Increase this value if you notice that the node is failing backup or retrieve operations because the MAXNUMMP value is being exceeded.

## NUMOPENVOLSALLOWED

### Syntax

▶▶—NUMOPENVOLSA11owed—*number\_of\_open\_volumes*—▶▶

### Parameters

*number\_of\_open\_volumes*

Specifies the number of input FILE volumes in a deduplicated storage pool that can be open at one time. The default is 10. The minimum value is 3. The maximum value is 999.

### Examples

Specify that up to 5 volumes in a deduplicated storage pool can be open at one time.

```
numopenvolsallowed 5
```

## QUERYAUTH

The QUERYAUTH option specifies the administrative authority level required to issue QUERY or SQL SELECT commands. By default any administrator can issue QUERY and SELECT commands. You can use this option to restrict the use of these commands.

### Syntax



### Parameters

#### None

Any administrator can issue QUERY or SELECT commands without requiring any administrative authority.

#### System

Administrators must have SYSTEM authority to issue QUERY or SELECT commands.

#### Policy

Administrators must have POLICY authority over one or more policy domains or SYSTEM authority to issue QUERY or SELECT commands.

#### Storage

Administrators must have STORAGE authority over one or more storage pools or SYSTEM authority to issue QUERY or SELECT commands.

#### Operator

Administrators must have OPERATOR or SYSTEM authority to issue QUERY or SELECT commands.

### Examples

To restrict the use of QUERY and SELECT commands to administrators with system or storage authority, enter:

```
queryauth storage
```

## RECLAIMDELAY

---

### RECLAIMDELAY

This option delays the reclamation of a SnapLock volume, allowing remaining data to expire so that there is no need to reclaim the volume.

#### Syntax

►► RECLAIMDELAY number\_of\_days ◀◀

#### Parameters

*number\_of\_days*

Specifies the number of days to delay the reclamation of a SnapLock volume.

Before a SnapLock volume is reclaimed, the Tivoli Storage Manager server allows the specified number of days to pass, so that any files remaining on the volume have a chance to expire. The default reclaim delay period is 4 days and can be set anywhere from 1 to 120 days.

#### Examples

Specify that the number of days to delay reclamation is 30 days:

```
reclaimdelay 30
```

---

## RECLAIMPERIOD

This option allows you to set the number of days for the reclamation period of a SnapLock volume.

### Syntax

```
▶▶ RECLAIMPERIOD number_of_days ◀◀
```

30

### Parameters

*number\_of\_days*

Specifies the number of days that are allowed for the reclamation period of a SnapLock volume.

After the retention of a SnapLock volume has expired, the Tivoli Storage Manager server will reclaim the volume within the specified number of days if there is still data remaining on the volume. The default reclaim period is 30 days and can be set anywhere from 7 to 365 days.

The reclamation period does not begin until the RECLAIMDELAY period has expired.

### Examples

Specify that the reclaim period is 45 days:

```
reclaimperiod 45
```

---

## REORGBEGINTIME

The REORGBEGINTIME option specifies the earliest time that the Tivoli Storage Manager server can start a table or index reorganization.

Schedule server-initiated reorganizations to start during periods when server activity is low. Use this option together with the REORGDURATION option. The REORGDURATION specifies an interval during which reorganization can start.

### Syntax

►►—REORGBEGINTime—*hh:mm*—►►

### Parameters

*hh:mm*

Specifies the time that the server can start a reorganization: The default start time 6:00 a.m. Use a 24-hour format to specify the time.

| Time      | Description            | Values                    |
|-----------|------------------------|---------------------------|
| <i>hh</i> | The hour of the day    | Specify a number 00 - 23. |
| <i>mm</i> | The minute of the hour | Specify a number 00 - 59. |

### Examples

Specify 6:00 a.m. as the earliest time that a reorganization can start.  
reorgbegintime 06:00

Specify 8:30 p.m. as the earliest time that a reorganization can start.  
reorgbegintime 20:30

Specify noon as the earliest time that a reorganization can start.  
reorgbegintime 12:00

Specify 3:30 p.m. as the earliest time that a reorganization can start.  
reorgbegintime 15:30

Specify midnight as the earliest time that a reorganization can start.  
reorgbegintime 00:00

---

## REORGURATION

The REORGURATION option specifies an interval during which server-initiated table or index reorganization can start.

Schedule server-initiated reorganizations to start during periods when server activity is low. Use this option together with the REORGBEGINTIME option. The REORGBEGINTIME option specifies the earliest time that the server can start a reorganization.

### Syntax

►►—REORGURATION—*nn*—►►

### Parameters

*nn* Specifies the number of hours during which a reorganization can start. The minimum value is 1, the maximum value is 24. The default value is 24.

### Example

Specify an interval of four hours during which a reorganization can start.

```
reorgduration 4
```

---

## REPORTRETRIEVE

The REPORTRETRIEVE option reports on restore or retrieve operations that are performed by client nodes or administrators. The default is NO.

### Syntax

►►—REPORTRETRIEVE—YES  
                  └─NO┘

### Parameters

#### YES

Specifies that messages will be issued to the server console and stored in the activity log whenever files are restored or retrieved from the Tivoli Storage Manager server. The messages will specify the name of the objects being restored or retrieved and identify the client node or administrator performing the operation.

**NO** Specifies that messages will not be issued.

### Examples

Specify that messages will be issued and stored in the activity log whenever files are restored or retrieved from the IBM Tivoli Storage Manager server:  
reportretrieve yes

The following message is issued for an administrator client session:

```
ANR0411I Session 8 for administrator COLIND-TUC logged in as node
COLIND-TUC restored or retrieved Backup object: node COLIND-TUC,
filesystem \\colind-tuc\c$, object\CODE\TESTDATA\ XXX.OUT
```

---

## REPLBATCHSIZE

The REPLBATCHSIZE option specifies the number of client files that are to be replicated in a batch, within the same server transaction. This option affects only the node replication processes and works with the REPLSIZETHRESH option to improve node replication processing.

The REPLBATCHSIZE option limits the number of files in a transaction and the REPLSIZETHRESH option limits the number of bytes in a transaction. The transaction ends when either the REPLBATCHSIZE threshold or the REPLSIZETHRESH threshold is reached.

### Syntax

►►—REPLBatchsize—<sup>4096</sup>*number\_of\_files*—►►

### Parameters

*number\_of\_files*

Specifies a number of files between 1 - 32768. The default is 4096.

### Examples

replbatchsize 25000

---

# REPLSIZETHRESH

The REPLSIZETHRESH option specifies, in megabytes, a threshold for the amount of data replicated, within the same server transaction.

The amount of data is based on the non-deduplicated size of the file, which is the original size of the file. The amount of data that is replicated is controlled by the threshold. When the amount of data exceeds the threshold, the server ends the transaction and no more files are added to the current batch. A new transaction is started after the current batch is replicated. This option is used with the REPLBATCHSIZE option.

For example, suppose that a file is 10 MB and is stored in a data-deduplication-enabled storage pool and only 2 MB of the file is transferred during replication. The amount of data replicated includes the 10 MB size of the file, and excludes the 2 MB transferred. When the amount of data replicated exceeds the value specified for the REPLSIZETHRESH threshold, the transaction ends.

### Syntax

►►—REPLSizethresh—4096—*megabytes*—►►

### Parameters

*megabytes*

Specifies the number of megabytes as an integer from 1 - 32768. The default value is 4096.

### Examples

```
replsizethresh 2000
```

## REQSYSAUTHOUTFILE

The REQSYSAUTHOUTFILE option specifies if system authority is required for administrative commands that cause Tivoli Storage Manager to write to an external file.

This option applies to the following commands:

- BACKUP DEVCONFIG with the FILENAMES parameter
- BACKUP VOLHISTORY with the FILENAMES parameter
- DEFINE BACKUPSET
- DELETE BACKUPSET
- MOVE DRMEDIA with the CMD parameter
- MOVE MEDIA with the CMD parameter
- QUERY DRMEDIA with the CMD parameter
- QUERY MEDIA with the CMD parameter
- QUERY SCRIPT with the OUTPUTFILE parameter

### Syntax

```

>>—REQSYSauthoutfile—Yes—————>>
 |
 |—No—|

```

### Parameters

#### Yes

System authority is required for administrative commands that cause Tivoli Storage Manager to write to an external file.

**No** System authority is not required for administrative commands that cause Tivoli Storage Manager to write to an external file. That is, there is no change to the authority level that is required to issue the command.

### Examples

```
reqsysauthoutfile no
```

---

### RESOURCETIMEOUT

The RESOURCETIMEOUT option specifies how long the server waits for a resource before canceling the pending acquisition of a resource. When a timeout occurs the request for the resource will be canceled.

**Note:** When managing a set of shared library resources, such as servers designated as library managers and clients, consider setting this option at the same time limit for all participants in the shared configuration. In any case of error recovery, Tivoli Storage Manager will always defer to the longest time limit.

#### Syntax

►►—RESOURCetimeout—60  
*minutes*—►►

#### Parameters

*minutes*

Specifies the maximum number of minutes that the server waits for a resource. The default value is 60 minutes. The minimum value is 1 minute.

#### Examples

Specify that the server will wait 15 minutes for a server resource:

```
resourcetimeout 15
```

---

## RESTOREINTERVAL

The RESTOREINTERVAL option specifies how long a restartable restore session can be saved in the server database. As long as the restore session is saved in the database, it can be restarted from the point at which it stopped.

You can update this server option without stopping and restarting the server by using the SETOPT command. See “SETOPT (Set a server option for dynamic update)” on page 1153.

### Syntax

►►—RESTOREINTERVAL minutes 1440—►►

### Parameters

*minutes*

Specifies how long, in minutes, that a restartable restore session can be in the database before it can be expired. The minimum value is 0. The maximum is 10080 (one week). The default is 1440 minutes (24 hours). If the value is set to 0 and the restore is interrupted or fails, the restore is still put in the restartable state. However, it is immediately eligible to be expired.

### Examples

```
restoreinterval 1440
```

---

## RETENTIONEXTENSION

The RETENTIONEXTENSION option specifies the number of days to extend the retention date of a SnapLock volume. This option allows the server to extend the retention date of a SnapLock volume in order to avoid excessive reclamation.

### Syntax

►—RETENTIONEXTENSION—*number\_of\_days*—◄

### Parameters

#### *number\_of\_days*

Specifies the number of days to extend the retention date of a SnapLock volume. The minimum value is 30 days; the maximum value is 9999 days; the default is 365.

If you specify a value of 0 (zero) for the **RETVER** parameter of an archive copy group, the actual value that is used for **RETVER** is the value of the option RETENTIONEXTENSION, if one of the following conditions is also true:

- The destination storage pool for the archive copy group is a SnapLock storage pool.
- The storage pool that is the target for a storage pool migration or of a **MOVE DATA** or **MOVE NODEDATA** command is a SnapLock storage pool.

If a SnapLock volume is the target volume for data from another SnapLock volume and if the remaining retention of the data on the volume is less than the value specified, then the retention date is set using the value specified. Otherwise, the remaining retention of the data is used to set the retention of the volume.

If a SnapLock volume has entered the reclamation period but the percentage of reclaimable space of the volume has not exceeded the reclamation threshold of the storage pool or the value specified on the **THRESHOLD** parameter of a **RECLAIM STGPOOL** command, then the retention date of the SnapLock volume is extended by the amount specified in the RETENTIONEXTENSION option.

### Examples

Specify that the retention date is extended by 60 days:

```
retentionextension 60
```

## SANDISCOVERY

The SANDISCOVERY option specifies whether the Tivoli Storage Manager SAN discovery function is enabled.


To use SAN discovery, all devices on the SAN must have a unique device serial number. When set to ON, the server performs SAN discovery in the following instances:

- When the device path has been changed
- When the **QUERY SAN** command is issued

Using SAN discovery, the server can automatically correct the special file name for a device if it has been changed for a specified tape device.

The Tivoli Storage Manager server does not require persistent binding with the SAN discovery function enabled. To display a list of devices seen by the server, you can issue the **QUERY SAN** command.

### Syntax



### Parameters

**ON** Specifies that the server performs SAN discovery when the device path has been changed, or when the **QUERY SAN** command is issued.

#### OFF

Specifies that the server does not perform SAN discovery when the device path has been changed, or when the **QUERY SAN** command is issued. If the Tivoli Storage Manager server is not able to open a device, a message is issued but the path associated with the device is not taken offline. This value is the default.

#### UNSCANNEDPATHOFF

Specifies that the server does not perform SAN discovery when the device path has been changed, or when the **QUERY SAN** command is issued. If the Tivoli Storage Manager server is not able to open a device, a message is issued and the path to the device is taken offline.

### Examples

```
sandiscovery on
```

---

### SANDISCOVERYTIMEOUT

The SANDISCOVERYTIMEOUT option specifies the amount of time allowed for host bus adapters to respond when they are queried by the SAN discovery process. Once the time specified for the SANDISCOVERYTIMEOUT is reached, the process times out.

#### Syntax

▶▶—SANDISCOVERYTIMEOUT—*value*—▶▶

#### Parameters

*value*

Specifies the amount of time to elapse before the SAN discovery process times out. The range is from 15 to 1800 seconds. The default is 15 seconds.

#### Examples

```
sandiscoverytimeout 45
```

---

## SANREFRESHTIME

The SANREFRESHTIME option specifies the amount of time that elapses before the cached SAN discovery information is refreshed. The SANREFRESHTIME option has a default value of 0, which means that there is no SAN discovery cache. The information is obtained directly from the host bus adapter (HBA) every time the server performs a SAN discovery operation.

**Note:** The QUERY SAN server command always receives SAN information at the time that the command is issued and ignores any value specified for SANREFRESHTIME.

### Syntax

►►—SANREFRESHTIME—0  
*time*—►►

### Parameters

*time*

The length of time, in seconds, before the cached SAN discovery information is refreshed. The default value is 0 and specifies that SAN discovery information is not cached. If a value other than 0 is specified, for example, 100 seconds, then the SAN discovery information is refreshed 100 seconds after the prior SAN discovery operation.

### Examples

Refresh SAN discovery information after 100 seconds.

```
sanrefreshtime 100
```

Turn off the caching of SAN discovery information.

```
sanrefreshtime 0
```

### SEARCHMPQUEUE

The SEARCHMPQUEUE option specifies the order in which the server satisfies requests in the mount queue. If the option is specified, the server first tries to satisfy requests for volumes that are already mounted. These requests may be satisfied before other requests, even if the others have been waiting longer for the mount point. If this option is not specified, the server satisfies requests in the order in which they are received.

#### Syntax

▶▶—SEARCHMPQUEUE—▶▶

#### Parameters

None

#### Examples

Specify that the server tries to first satisfy a request for a volume that is already mounted:

```
searchmpqueue
```

## SERVERDEDUPTXNLIMIT

The SERVERDEDUPTXNLIMIT option specifies the maximum size of objects that can be deduplicated on the server.

When you use duplicate-identification processes (the **IDENTIFY DUPLICATES** command) for large objects, intensive database activity can result from long-running transactions that are required to update the database. High levels of database activity can produce following symptoms:

- Reduced throughput for client backup and archive operations
- Resource contention resulting from concurrent server operations
- Excessive recovery log activity

The extent to which these symptoms occur depends on the number and size of objects being processed, the intensity and type of concurrent operations taking place on the Tivoli Storage Manager server, and the Tivoli Storage Manager server configuration.

With the SERVERDEDUPTXNLIMIT server option, you can specify a maximum size, in gigabytes, for objects that can be deduplicated on the server. If an object or set of objects in a single transaction exceeds the limit specified by SERVERDEDUPTXNLIMIT, the objects are not deduplicated by the server. You can specify a value 32 - 2048 GB. The default value is 300 GB.

Increasing the value of this option causes the Tivoli Storage Manager server to search for objects previously deferred whose size falls below the new transaction limit.

**Remember:** The search for objects previously deferred can take time. Use care when increasing the value of SERVERDEDUPTXNLIMIT. Reducing the value of this option does not cause Tivoli Storage Manager to search for deferred objects.

The appropriate value for this option depends on the Tivoli Storage Manager server configuration and concurrent server activity. You can specify a high value for this option if you minimize resource contention. To minimize resource contention, perform operations, such as backup, archive, duplicate identification, and reclamation, at different times.

To update this server option without stopping and restarting the server, use the **SETOPT** command.

### Syntax

►►—SERVERDEDUPTXNLIMIT—<sup>300</sup>—gigabytes—►►

### Parameters

*gigabytes*

Specifies the maximum size, in gigabytes, of objects that can be duplicated on the server. You can specify a value 32 - 2048. The default value is 300.

## **SERVERDEDUPTXNLIMIT**

### **Examples**

Disable server-side deduplication for all objects over 120 GB:

```
serverdeduptxnlimit 120
```

---

## SHMPORT

The SHMPORT option specifies the TCP/IP port address of a server when using shared memory. All shared memory communications start with a TCP/IP connection.

### Syntax

►►—SHMPort—*port\_number*—◄◄

### Parameters

*port\_number*

Specifies the port number. You can specify a value from 1024 to 32767. The default value is 1510.

### Examples

```
shmport 1580
```



---

## SNMPHEARTBEATINTERVAL

The SNMPHEARTBEATINTERVAL option specifies the interval in minutes between queries of the Tivoli Storage Manager server.

### Syntax

►►—SNMPHEARTBEATINTERVAL—<sup>5</sup>minutes—◄◄

### Parameters

*minutes*

Specifies the heartbeat interval in minutes. Valid values are from 0 to 1440 (one day). The default is 5 minutes.

### Examples


```
snmpheartbeatinterval 20
```

---

### SNMPMESSAGECATEGORY

The SNMPMESSAGECATEGORY option specifies the trap types used when messages are forwarded from the server, through the Simple Network Management Protocol (SNMP) subagent, to the SNMP manager.

#### Syntax



#### Parameters

##### SEVERITY

Specifies that there are four trap types based on message severity level:

- 1 Severe
- 2 Error
- 3 Warning
- 4 Information

This is the default.

##### INDIVIDUAL

Specifies that a separate trap type is used for each message. The numeric part of the message identifier indicates the trap type.


#### Examples

```
snmpmessagecategory individual
```

## SNMPSUBAGENT

The SNMPSUBAGENT option specifies the parameters needed for the Tivoli Storage Manager subagent to communicate with the Simple Network Management Protocol (SNMP) daemon. This option is only to configure the SNMP subagent for communicating with the SNMP agent; it is ignored by the server.

### Syntax



### Parameters

#### HOSTname *host\_name*

Specifies the TCP/IP name or number of the host running the SNMP agent that the Tivoli Storage Manager SNMP subagent connects to. This parameter is optional. The default name is *localhost*.

#### COMMunityname *community\_name*

Specifies the configured community name on the system running the SNMP agent. This parameter is optional. The default name is *public*.

#### TIMEOUT *seconds*

Specifies the time, in seconds, in which a request must be received. This parameter is optional. The default value is 600.

### Examples

```
snmpsubagent hostname jimbo communityname public timeout 2600
```

### SNMPSUBAGENTHOST

The SNMPSUBAGENTHOST option specifies the location of the Tivoli Storage Manager Simple Network Management Protocol (SNMP) subagent. The default for this option is 127.0.0.1.

#### Syntax

▶▶—SNMPSUBAGENTHOST—*host\_name*————▶▶

#### Parameters

*host\_name*

Specifies the TCP/IP host name or number on which the Tivoli Storage Manager SNMP subagent is located. The subagent and server must be on the same node.

#### Examples

```
snmpsubagenthost 9.116.23.450
```

---

## SNMPSUBAGENTPORT

The SNMPSUBAGENTPORT option specifies the port number of the Tivoli Storage Manager Simple Network Management Protocol (SNMP) subagent.

### Syntax

►►—SNMPSUBAGENTPORT—*port\_number*—►►

### Parameters

*port\_number*

Specifies the port number of the Tivoli Storage Manager SNMP subagent. Valid values are 1000 - 32767. The default is 1521.

### Examples

snmpsubagentport 1525

---

## SSLFIPSMODE

The SSLFIPSMODE option specifies whether the Federal Information Processing Standards (FIPS) mode is in effect for Secure Sockets Layer (SSL). The default is NO.

Because SSLv3 is not supported by FIPS mode, when using SSL with version 6.1 or version 5.5 clients, you must turn off FIPS mode.

### Syntax

►►—SSLFIPSMODE—=*No*—►►  
 ►►—SSLFIPSMODE—=*No*—►►  
                                   *Yes*—►►

### Parameters

**No** Specifies that SSL FIPS mode is not active on the server. This setting is required when Backup-Archive Client versions previous to IBM Tivoli Storage Manager 6.3 are to connect to the server with SSL.

#### Yes

A value of YES indicates that SSL FIPS mode is active on the server. This setting restricts SSL session negotiation to use FIPS-approved cipher suites. Specifying YES is suggested when SSL communication is activated and all Backup-Archive Clients are at version 6.3.

**Remember:** Set server option SSLTLS12 to NO if your environment includes Backup-Archive Clients that use server self-signed certificates that were generated by a server whose version was before 6.3.

To disable SSL FIPS mode on the server:

```
SSLFIPSMODE no
```

### SSLTCPADMINPORT

The SSLTCPADMINPORT option specifies the port address on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions. The sessions are for the command-line administrative client.

First-time use of the SSLTCPADMINPORT or SSLTCPPORT option triggers the creation of a key database file (cert.kdb) at server startup. The key database file is created in the server instance directory and is initialized with a self-signed certificate.

The following types of sessions do not use SSL (Secure Sockets Layer):

- NDMP (Network Data Management Protocol)
- ACSLS (Automated Cartridge System Library Software)
- SNMP (Simple Network Management Protocol) subagent
- Database restore operations

If the ADMINONCLIENTPORT option is set to NO, SSL-enabled sessions for the administrative client require SSLTCPADMINPORT with a port number other than one specified by the SSLTCPPORT option. The SSLTCPADMINPORT option does not affect the TCPPORT or TCPADMINPORT options and their interaction with the ADMINONCLIENTPORT option.

The TCP/IP communications driver must be enabled with COMMMETHOD TCPIP or COMMMETHOD V6TCPIP.

#### Syntax

▶▶—SSLTCPADMINPort—*port\_number*—▶▶

#### Parameters

*port\_number*

Specifies the port number of the server. Valid values are 1024 - 32767. There is no default.

#### Examples

```
ssltcpadminport 1543
```

---

## SSLTCPPORT

The SSLTCPPORT option specifies the SSL (Secure Sockets Layer) port number. The server TCP/IP communication driver waits for requests on this port for SSL-enabled sessions from the client.

First-time use of the SSLTCPADMINPORT or SSLTCPPORT options triggers the creation of a key database file (`cert.kdb`) at server startup. The key database file is created in the server instance directory and is initialized with a self-signed certificate.

The following types of sessions do not use SSL:

- NDMP (Network Data Management Protocol)
- ACSLS (Automated Cartridge System Library Software)
- SNMP (Simple Network Management Protocol) subagent
- Database restore operations

If the ADMINONCLIENTPORT option is set to NO, SSL-enabled sessions for the administrative client require SSLTCPADMINPORT with a port number different from one specified by the SSLTCPPORT option. The SSLTCPPORT option does not affect the TCPSPORT or TCPADMINPORT options and their interaction with the ADMINONCLIENTPORT option.

The TCP/IP communications driver must be enabled with COMMMETHOD TCP/IP or COMMMETHOD V6TCP/IP.

### Syntax

►►—SSLTCPPort—*port\_number*—◄◄

### Parameters

*port\_number*

Specifies the port number of the server. Valid values are 1024 - 32767. There is no default.

### Examples

```
ssltcpport 1542
```



---

## TCPADMINPORT

The TCPADMINPORT option specifies the port number on which the server TCP/IP communication driver is to wait for requests for sessions other than client sessions. This includes administrative sessions, server-to-server sessions, Simple Network Management Protocol (SNMP) subagent sessions, storage agent sessions, library client sessions, managed server sessions, and event server sessions.

Using different port numbers for the options TCPSPORT and TCPADMINPORT enables you to create one set of firewall rules for client sessions and another set for the other session types listed above. By using the **SESSIONINITIATION** parameter of REGISTER and UPDATE NODE, you can close the port specified by TCPSPORT at the firewall, and specify nodes whose scheduled sessions will be started from the server. If the two port numbers are different, separate threads will be used to service client sessions and the session types. If you allow the two options to use the same port number (by default or by explicitly setting them to the same port number), a single server thread will be used to service all session requests.

Client sessions attempting to use the port specified by TCPADMINPORT will be terminated (if TCPSPORT and TCPADMINPORT specify different ports). Administrative sessions are allowed on either port, (unless the ADMINONCLIENTPORT option is set to NO) but by default will use the port specified by TCPADMINPORT.

### Syntax

▶▶—TCPADMINPort—*port\_number*—▶▶

### Parameters

*port\_number*

Specifies the port number of the server. Valid values are 1024 - 32767. The default is the value of TCPSPORT.

### Examples

```
tcpadminport 1502
```

---

### TCPBUFSIZE

The TCPBUFSIZE option specifies the size of the buffer used for TCP/IP send requests. During a restore, client data moves from the Tivoli Storage Manager session component to a TCP communication driver. The TCPBUFSIZE option determines if the server sends the data directly from the session buffer or copies the data to the TCP buffer. A 32 KB buffer size forces the server to copy data to its communication buffer and flush the buffer when it fills.

**Note:** This option is not related to the TCPWINDOWSIZE option.

#### Syntax

►—TCPBufsize—*kilobytes*—►

#### Parameters

*kilobytes*

Specifies the size, in kilobytes, of the buffer used for TCP/IP send requests.

The value range is from 1 to 64. The default is 16.

#### Examples

```
tcpbufsize 5
```

---

## TCPNODELAY

The TCPNODELAY option specifies whether the server disables the delay of sending successive small packets on the network.

Change the value from the default of YES only under one of these conditions:

- You are directed to change the option by your service representative.
- You fully understand the effects of the TCP Nagle algorithm on network transmissions. Setting the option to NO enables the Nagle algorithm, which delays sending small successive packets.

### Syntax

►—TCPNodeLAY—Yes—  
No—►

### Parameters

#### Yes

Specifies that the server allows successive small packets to be sent immediately over the network. Setting this option to YES might improve performance in some high-speed networks. The default is YES.

**No** Specifies that the server does not allow successive small packets to be sent immediately over the network.

### Examples

```
tcpnodeLAY no
```

---

## TCPPOPT

The TCPPOPT option specifies the port number on which the server TCP/IP communication driver is to wait for requests for client sessions.

Using different port numbers for the options TCPPOPT and TCPADMINPORT enables you to create one set of firewall rules for client sessions and another set for other session types (administrative sessions, server-to-server sessions, Simple Network Management Protocol (SNMP) subagent sessions, storage agent sessions, library client sessions, managed server sessions, and event server sessions). If the two port numbers are different, separate threads will be used to service client sessions and the other session types. If you allow the two options to use the same port number (by default or by explicitly setting them to the same port number), a single server thread will be used to service all session requests.

### Syntax

▶▶—TCPPOPT—*port\_number*—————▶▶

### Parameters

*port\_number*

Specifies the port number of the server. Valid values are 1024 - 32767. The default value is 1500.

tcpport 1500

---

## TCPWINDOWSIZE

The TCPWINDOWSIZE option specifies, in kilobytes, the amount of receive data that can be buffered at one time on a TCP/IP connection. The sending host cannot send more data until it receives an acknowledgment and a TCP receive window update. Each TCP packet contains the advertised TCP receive window on the connection. A larger window lets the sender continue sending data, and may improve communication performance, especially on fast networks with high latency.

### Note:

- To improve backup performance, increase the TCPWINDOWSIZE on the server. To improve restore performance, increase the TCPWINDOWSIZE on the client.
- The TCP window acts as a buffer on the network.
- A window size larger than the buffer space on the network adapter might degrade throughput due to resending packets that were lost on the adapter.
- The TCPWINDOWSIZE option is not related to the TCPBUFFSIZE option nor to the send and receive buffers allocated in client or server memory.

### Syntax

►►—TCPWindowsize—*kilobytes*—◄◄

### Parameters

*kilobytes*

Specifies the size you want to use, in kilobytes, for the TCP/IP sliding window for your client node. You can specify a value from 0 to 2048. The default is 63. If you specify 0, the server uses the default window size set by the operating system. Values from 1 to 2048 indicate that the window size is in the range of 1 KB to 2 MB.

### Examples


```
tcpwindowsize 63
```

---

### TECBEGINEVENTLOGGING

The `TECBEGINEVENTLOGGING` option specifies whether event logging for the Tivoli receiver should begin when the server starts up. If the `TECHOST` option is specified, `TECBEGINEVENTLOGGING` defaults to `YES`.

#### Syntax



#### Parameters

##### Yes

Specifies that event logging begins when the server starts up and if a `TECHOST` option is specified.

**No** Specifies that event logging should not begin when the server starts up. To later begin event logging to the Tivoli receiver (if the `TECHOST` option has been specified), you must issue the `BEGIN EVENTLOGGING` command.

#### Examples

```
tecbegineventlogging yes
```

---

## TECHOST

The TECHOST option specifies the host name or IP address for the Tivoli event server.

### Syntax

▶▶—TECHost—*host\_name*—————▶▶

### Parameters

*host\_name*

Specifies the host name or IP address for the Tivoli event server.

### Examples

techost 9.114.22.345

---

## TECPORT

The TECPORT option specifies the TCP/IP port address on which the Tivoli event server is listening. This option is only required if the Tivoli event server is on a system that does not have a Port Mapper service running.

### Syntax

►►—TECPort—*port\_number*—————►►

### Parameters

*port\_number*

Specifies the Tivoli event server port address. The value must be between 0 and 32767. This option is not required.

### Examples

tecport 1555

---

## TECUTF8EVENT

The TECUTF8EVENT option allows the Tivoli Storage Manager administrator to send information to the Tivoli Enterprise Console (TEC) server in UTF-8 data format. The default is No. You can display whether or not this option is enabled by issuing the QUERY OPTION command.

### Syntax

►►—TECUTF8event—Yes  
No—————►

### Parameters

#### Yes

Specifies that the Tivoli Storage Manager server will encode the TEC event into UTF-8 before issuing the event to the TEC server.

**No** Specifies that Tivoli Storage Manager server will not encode the TEC event into UTF-8 and it will be issued to the TEC server in ASCII format.

### Examples

```
tecutf8event yes
```

---

## THROUGHPUTDATATHRESHOLD

The THROUGHPUTDATATHRESHOLD option specifies a throughput threshold that a client session must reach to prevent being cancelled after the time threshold is reached.

This option is used in conjunction with the THROUGHPUTTIMETHRESHOLD server option, which sets the value for the time threshold plus the media wait time. The time threshold starts when the client begins sending data to the server for storage (as opposed to setup or session housekeeping data).

You can update this server option without stopping and restarting the server by using the SETOPT command. See “SETOPT (Set a server option for dynamic update)” on page 1153.

### Syntax

►►—THROUGHPUTDatathreshold— *kilobytes\_per\_second*—►►

### Parameters

*kilobytes\_per\_second*

Specifies the throughput that client sessions must achieve to prevent cancellation after THROUGHPUTTIMETHRESHOLD minutes have elapsed. This threshold does not include time spent waiting for media mounts. A value of 0 prevents examining client sessions for insufficient throughput. Throughput is computed by adding send and receive byte counts and dividing by the length of the session. The length does not include time spent waiting for media mounts and starts at the time a client sends data to the server for storage. The default is 0. The minimum value is 0; the maximum is 99999999.

### Examples

Specify that the server is to wait until 90 minutes plus the media wait time after a session has started sending data before storage examines it as a candidate for cancellation due to low throughput. If a session is not achieving 50 KB per second in transfer rates, it will be cancelled.

```
throughputtimethreshold 90
Throughputdatathreshold 50
```

---

## THROUGHPUTTIMETHRESHOLD

The THROUGHPUTTIMETHRESHOLD option specifies the time threshold for a session after which it may be cancelled for low throughput.

You can update this server option without stopping and restarting the server by using the SETOPT command. See “SETOPT (Set a server option for dynamic update)” on page 1153.

### Syntax

►►—THROUGHPUTtimethreshold—*minutes*—◄◄

### Parameters

*minutes*

Specifies the threshold for examining client sessions and cancelling them if the data throughput threshold is not met (see the THROUGHPUTDATATHRESHOLD server option). This threshold does not include time spent waiting for media mounts. The time threshold starts when a client begins sending data to the server for storage (as opposed to setup or session housekeeping data). A value of 0 prevents examining client sessions for low throughput. The default is 0. The minimum value is 0; the maximum is 99999999.

### Examples

Specify that the server is to wait until 90 minutes plus the media wait time after a session has started sending data before examining it as a candidate for cancellation. If a session is not achieving 50 thousand bytes per second in transfer rates, it will be cancelled.

```
throughputtimethreshold 90
Throughputdatathreshold 50
```

---

## TXNGROUPMAX

The TXNGROUPMAX option specifies the number of objects that are transferred as a group between a client and the server between transaction commit points. The minimum value is 4 objects and the maximum value is 65000 objects. The default value is 4096 objects. The objects transferred are actual files, directories, or both. The server counts each file or directory as one object.

It is possible to affect the performance of client backup, archive, restore, and retrieve operations by using a larger value for this option:

1. If you increase the value of the TXNGROUPMAX option by a large amount, watch for possible effects on the recovery log. A larger value for the TXNGROUPMAX option can result in increased utilization of the recovery log, as well as an increased length of time for a transaction to commit. If the effects are severe enough, they can lead to problems with operation of the server.
2. Increasing the value of the TXNGROUPMAX option can improve throughput for operations storing data directly to tape, especially when storing a large number of objects. However, a larger value of the TXNGROUPMAX option can also increase the number of objects that must be resent in the case where the transaction is stopped because an input file changed during backup, or because a new storage volume was required. The larger the value of the TXNGROUPMAX option, the more data must be resent.
3. Increasing the TXNGROUPMAX value will affect the responsiveness of stopping the operation and the client may have to wait longer for the transaction to complete.

You can override the value of this option for individual client nodes. See the TXNGROUPMAX parameter in “REGISTER NODE (Register a node)” on page 971 and “UPDATE NODE (Update node attributes)” on page 1309.

This option is related to the TXNBYTELIMIT option in the client options file. TXNBYTELIMIT controls the number of bytes, as opposed to the number of objects, that are transferred between transaction commit points. At the completion of transferring an object, the client commits the transaction if the number of bytes transferred during the transaction reaches or exceeds the value of TXNBYTELIMIT, regardless of the number of objects transferred.

### Syntax

►►—TXNGroupmax—*number\_of\_objects*—◄◄

### Parameters

*number\_of\_objects*

Specifies a number from 4 to 65000 for the maximum number of objects per transaction. The default is 4096.

### Examples

```
txngroupmax 4096
```

---

## UNIQUEPTEEVENTS

The UNIQUEPTEEVENTS option generates a unique Tivoli Enterprise Console (TEC) event class for each individual Tivoli Storage Manager message, including client, server, and Tivoli Data Protection (TDP) agent messages. The default is No.

### Syntax

►►—UNIQUEPteevents—

|     |
|-----|
| Yes |
| No  |

—————►

### Parameters

#### Yes

Specifies that unique TDP messages are sent to the TEC event server.  
Dynamically sets UNIQUEPTEEvents to YES.

**No** Specifies that general messages are sent to the TEC event server.

### Examples

```
uniqueppteevents yes
```

---

### UNIQUETECEVENTS

The UNIQUETECEVENTS option generates a unique Tivoli Enterprise Console (TEC) event class for each individual Tivoli Storage Manager message. The default is No.

#### Syntax



#### Parameters

##### Yes

Specifies that unique messages are sent to the TEC event server.

##### No

Specifies that general messages are sent to the TEC event server.

#### Examples

```
uniquetecevents yes
```

---

## USEREXIT

The USEREXIT option specifies a user-defined exit that will be given control to manage an event.

### Syntax

►►—USEREXIT—┌ Yes  
└ No—*module\_name*—►►

### Parameters

#### Yes

Specifies that event logging to the user exit receiver begins automatically at server startup.

**No** Specifies that event logging to the user exit receiver does not begin automatically at server startup. When this parameter has been specified, you must begin event logging manually by issuing the BEGIN EVENTLOGGING command.

#### *module\_name*

Specifies the module name of the user exit.

This is the name of a shared library containing the exit. The module name can be either a fully qualified path name or just the module name itself. If it is just the module name, it is loaded from the current directory.

### Examples

```
userexit yes fevent.exit
```

### VERBCHECK

The VERBCHECK option specifies that the server will do additional error checking on the structure of commands sent by the client. This option should only be enabled when the client sends incorrectly formed requests to the server, causing the server to crash. When this option is enabled, you will get a protocol error instead of a server crash.

#### Syntax

▶—VERBCHECK—▶

#### Parameters

None

#### Examples

Enable additional error checking for commands sent by the client:

```
verbcheck
```

---

## VOLUMEHISTORY

The VOLUMEHISTORY option specifies the name of files to be automatically updated whenever server sequential volume history information is changed. There is no default for this option.

You can include one or more VOLUMEHISTORY options in the server options file. When you use multiple VOLUMEHISTORY options, the server automatically updates and stores a backup copy of the volume history information in each file you specify.

### Syntax

►►—VOLUMEHistory—*file\_name*—————►

### Parameters

*file\_name*

Specifies the name of the file where you want the server to store a backup copy of the volume history information that it collects.

### Examples

volumehistory volhist.out

## VOLUMEHISTORY

---

## Chapter 4. Server utilities

Use server utilities to perform special tasks on the server while the server is not running.

*Table 460. Server utilities*

| <b>Utility</b>                                                                            | <b>Description</b>                                                 |
|-------------------------------------------------------------------------------------------|--------------------------------------------------------------------|
| "DSMSERV (Start the server)" on page 1560                                                 | Starts the server.                                                 |
| "The server startup script: rc.dsmserv" on page 1561                                      | Automatically starts a server instance.                            |
| "DSMSERV DISPLAY DBSPACE (Display information about database storage space)" on page 1563 | Displays information about storage space defined for the database. |
| "DSMSERV DISPLAY LOG (Display recovery log information)" on page 1564                     | Displays information about recovery log storage space.             |
| "DSMSERV FORMAT (Format the database and log)" on page 1566                               | Initializes the database and recovery log.                         |
| "DSMSERV INSERTDB (Move a server database into an empty database)" on page 1568           | Inserts a server database into a new Version 6 database.           |
| "DSMSERV LOADFORMAT (Format a database)" on page 1571                                     | Formats an empty database.                                         |
| "DSMSERV REMOVEDB (Remove a database)" on page 1573                                       | Removes a Tivoli Storage Manager database.                         |
| "DSMSERV RESTORE DB (Restore the database)" on page 1574                                  | Restores a Tivoli Storage Manager database.                        |



---

## The server startup script: rc.dsmserv

The startup script can be used in your system startup to automatically start a server instance under a specific user ID.

### Syntax

```
▶▶rc.dsmserv [-u user_name] [-U user_name] [-i instance_dir] ▶▶
```

### Parameters

**-u** *user\_name*

Specifies the Tivoli Storage Manager instance user ID for which the environment is set up. The server will run under this user ID.

**-U** *user\_name*

Specifies the Tivoli Storage Manager instance user ID for which the environment is set up. The server will run under the user ID of the invoker of the command.

**-i** *instance\_dir*

Specifies an instance directory to use. This becomes the current working directory of the server.

## The server startup script: dsmserv.rc

The dsmserv.rc script can stop a server instance or manually or automatically start a server. Red Hat Enterprise Linux 6 (RHEL6) users must use the dsmserv.rc script, although other versions of Linux can also use it.

### Prerequisites

Before issuing the **DSMSERV.RC** command, meet these parameters:

1. Ensure that the Tivoli Storage Manager server instance runs under a non-root user ID with the same name as the instance owner.
2. Copy the dsmserv.rc script to the /etc/rc.d/init.d directory. The dsmserv.rc script is located in the /opt/tivoli/tsm/server/bin directory.
3. Rename the dsmserv.rc script to the same name as the server instance user. For example, if the server instance user name is tsminst1, the script must be renamed to tsminst1.
4. Configure the run level at which the server is going to be started. You can use a tool such as the chkconfig utility to configure the run level. Specify a value that corresponds to a multiuser mode, with networking enabled. Run levels are typically 2, 3, or 5, depending on the operating system configuration. Consult the documentation for your operating system for details on run levels.

### Syntax



### Parameters

#### start

Starts the server.

#### stop

Stops the server.

#### status

Shows the status of the server. If the status is “started,” the process ID of the server process is also shown.

#### restart

Combines the stop and start actions to restart the server.

## DSMSERV DISPLAY DBSPACE (Display information about database storage space)

Use this utility to display information about storage space that is defined for the database. The output of this utility is the same as the output of `QUERY DBSPACE`, but you can use this utility when the server is not running.

### Syntax

```

DSMSERV [-u user_name] [-i instance_dir] [-o options_file]
 [-noexpire] [-quiet] DISPlay DBSPace

```

### Parameters

- u** *user\_name*  
Specifies a user name to switch to before initializing the server.
- i** *instance\_dir*  
Specifies an instance directory to use. This becomes the current working directory of the server.
- o** *options\_file*  
Specifies an options file to use.
- noexpire**  
Specifies that expiration processing is suppressed when starting.
- quiet**  
Specifies that messages to the console are suppressed.

### Example: Display database space information

Display information about database storage space. See “Field descriptions” for details about the information shown in the output. Issue the command.

```
dsmserv display dbspace
```

| Location  | Total Space (MB) | Used Space (MB) | Free Space (MB) |
|-----------|------------------|-----------------|-----------------|
| /tsmdb001 | 46,080.00        | 20,993.12       | 25,086.88       |
| /tsmdb002 | 46,080.00        | 20,992.15       | 25,087.85       |

### Field descriptions

#### Location

The directory or path that is used for storing the database

#### Total Space (MB)

The total number of megabytes in the location

#### Used Space (MB)

The number of megabytes in use in the location

#### Free Space (MB)

The space remaining in the file system where the path is located

## DSMSERV DISPLAY LOG (Display recovery log information)

Use this utility to display information about recovery logs including the active log, the mirror for the active log, the failover directory for the archive log, and the overflow location for logs. Use this utility when the server is not running.

### Syntax

```

DSMSERV -u user_name -i instance_dir -o options_file
DISPLAY LOG -noexpire -quiet

```

### Parameters

- u** *user\_name*  
Specifies a user name to switch to before initializing the server.
- i** *instance\_dir*  
Specifies an instance directory to use. This becomes the current working directory of the server.
- o** *options\_file*  
Specifies an options file to use.
- noexpire**  
Specifies that expiration processing is suppressed when starting.
- quiet**  
Specifies that messages to the console are suppressed.

### Examples: Display recovery log information

Display information about the recovery logs. See “Field descriptions” for details about the information shown in the output.

```
dsmserv display log
```

```

Total Space(MB): 38,912
Used Space(MB): 401.34
Free Space(MB): 38,358.65
Active Log Directory: /activelog
Archive Log Directory: /archivelog
Mirror Log Directory: /mirrorlog
Archive Failover Log Directory: /archfailoverlog

```

### Field descriptions

#### Total Space

Specifies the maximum size of the active log.

#### Used Space

Specifies the total amount of active log space currently used in the database, in megabytes.

#### Free Space

Specifies the amount of active log space in the database that is not being used by uncommitted transactions, in megabytes.

**Active Log Directory**

Specifies the location where active log files are stored. When you change the active log directory, the server moves all archived logs to the archive log directory and all active logs to a new active log directory.

**Mirror Log Directory**

Specifies the location where the mirror for the active log is maintained.

**Archive Failover Log Directory**

Specifies the location in which the server saves archive logs if the logs cannot be archived to the archive log destination.

## DSMSERV FORMAT (Format the database and log)

Use the **DSMSERV FORMAT** utility to initialize the server database and recovery log. No other server activity is allowed while initializing the database and recovery log.

The directories that are specified in this utility should be on fast, reliable storage. Do not place the directories on file systems that might run out of space. If certain directories (for example, the active log directory) become unavailable or full, the server stops.

When a server is initially created, with the **DSMSERV FORMAT** utility or with the configuration wizard, a server database and recovery log are created. In addition, files are created to hold database information that is used by the database manager.

### Setting the server code page


Set the DB2CODEPAGE system environment variable to 819 for each server instance. Before you issue the **DSMSERV LOADFORMAT** command, log on to the system as the server instance owner and issue this command:

```
db2set -i instance_name DB2CODEPAGE=819
```

For example:

```
db2set -i tsminst1 DB2CODEPAGE=819
```

### Syntax



### Parameters

**-u user\_name**  
Specifies a user name to switch to before initializing the server. This parameter is optional.

**-i** *instance\_dir*

Specifies an instance directory to use. This directory becomes the current working directory of the server. This parameter is optional.

**-o** *options\_file*

Specifies an options file to use. This parameter is optional.

**-noexpire**

Specifies that expiration processing is suppressed when starting. This parameter is optional.

**-quiet**

Specifies that messages to the console are suppressed. This parameter is optional.

**DBDir**

Specifies the relative path names of one or more directories that are used to store database objects. Directory names must be separated by commas but without spaces. You can specify up to 128 directory names. You must specify either the **DBDIR** or the **DBFILE** parameter.

**DBFile**

Specifies the name of a file that contains the relative path names of one or more directories that are used to store database objects. Each directory name must be on a separate line in the file. You can specify up to 128 directory names. You must specify either the **DBDIR** or the **DBFILE** parameter.

**ACTIVELOGSize**

Specifies the size of the active log in megabytes. This parameter is optional. The minimum value is 2048 MB (2 GB); the maximum is 131,072 MB (128 GB). If you specify an odd number, the value is rounded up to the next even number. The default is 16384 MB.

**ACTIVELOGDirectory (Required)**

Specifies the directory in which the Tivoli Storage Manager server writes and stores active log files. There is only one active log location. The name must be a fully qualified directory name. The directory must exist, it must be empty, and it must be accessible by the user ID of the database manager. The maximum number of characters is 175.

**ARCHLogdirectory (Required)**

Specifies the directory for the archive log files. The name must be a fully qualified directory name. The maximum number of characters is 175.

**ARCHFailoverlogdirectory**

Specifies the directory to be used as an alternative storage location if the **ARCHLOGDIRECTORY** directory is full. This parameter is optional. The maximum number of characters is 175.

**MIRRORlogdirectory**

Specifies the directory in which the server mirrors the active log (those files in the **ACTIVELOGDIRECTORY** directory). This parameter is optional. The directory must be a fully qualified directory name. The maximum number of characters is 175.

### Example: Format a database

```
dsmserv format dbdir=/tsmdb001 activesize=8192
activedirectory=/active log archlogdirectory=/archlog
archfailoverlogdirectory=/archfaillog mirrorlogdirectory=/mirrorlog
```

## DSMSERV INSERTDB (Move a server database into an empty database)

Use the **DSMSERV INSERTDB** utility to move a server database into a new V6.3 database. The database can be extracted from the original server and inserted into a new database on the new server by using a network connection between the two servers. The database can also be inserted from media that contains the extracted database.


Before you use the **DSMSERV INSERTDB** utility, complete the planning and preparation tasks, such as backing up the database and saving configuration information. Ensure that you meet all requirements before you move the server database. For planning and preparation information, see the upgrade or migration procedure for your operating system.

### Requirements for insertion by using media


Before you run the utility to insert the server database into an empty database, ensure that your system meets the following requirements.

- The manifest file from the **DSMUPGRD EXTRACTDB** operation must be available.
- If the manifest file does not contain device configuration information, or if you are specifying the **CONFIGINFO=DEVCONFIG** parameter, both of the following statements must be true:
  - The server options file must contain an entry for the device configuration file.
  - The device configuration file must have information about the device class that is specified in the manifest file.
- The media that contains the extracted database must be available to the V6.3 server. Also, the permissions must be set to grant access to the media for the user ID that owns the V6.3 server instance.

### Syntax



#### A: Insert from media:



▶ MANifest *file\_name*

### B: Insert over a network:

SESSWait *60*  
SESSWait *minutes*

## Parameters

### -u *user\_name*

Specifies a user name to switch to before initializing the server. This parameter is optional.

### -i *instance\_dir*

Specifies an instance directory to use. This directory becomes the current working directory of the server. This parameter is optional.

### -o *options\_file*

Specifies an options file to use. This parameter is optional.

### -noexpire

Specifies that expiration processing is suppressed when starting. This parameter is optional.

### -quiet

Specifies that messages to the console are suppressed. This parameter is optional.

### DEVclass

Specifies a sequential-access device class. You can specify any device class except for the DISK device class. The definition for the device class must exist in either the manifest file or the device configuration file.

This parameter is optional and is used only when the database that you want to insert into the empty V6.3 database was extracted to media. If the database is on media and you do not specify a device class, the device class that is identified in the manifest file is used.

**Restriction:** You cannot use a device class with a device type of NAS or CENTERA.

### MANifest

Specifies the location of the manifest file. Use a fully qualified file name, or place in a local directory. For example: `./manifest.txt`

This parameter is required when the database that you want to insert into the empty V6.3 database was extracted to media.

### CONFiginfo

Specifies the source of the device configuration information that is used by the **DSMSERV INSERTDB** operation. The default value for this parameter is **MANIFEST**. Possible values are as follows:

#### MANifest

Specifies that device configuration information is read from the manifest file. If the manifest file does not have device configuration information, the device configuration file is used instead.

## DSMSERV INSERTDB

### **DEVConfig**

Specifies that device configuration information is read from the device configuration file.

### **SESSWait**

Specifies the number of minutes that the V6.3 server waits to be contacted by the original server. The default value is 60 minutes.

Use this parameter only if the data that is inserted into the empty V6.3 database is transmitted from the source server with a network connection.

### **PREview**

Specifies whether to preview the insertion operation. This parameter is optional. The default value is **NO**.

Use the **PREVIEW=YES** parameter to test a database. When you use this parameter, the operation includes all steps of the process, except for the actual insertion of data into the new database. When you preview the insertion operation, you can quickly verify that the source database is readable. You can also identify any data constraint violations that might prevent an upgraded database from being put into production.

## DSMSERV LOADFORMAT (Format a database)

Use the **DSMSERV LOADFORMAT** utility when upgrading from Version 5. The utility formats an empty database in preparation for inserting an extracted database into the empty database.

### Setting the server code page


Set the DB2CODEPAGE system environment variable to 819 for each server instance. Before you issue the **DSMSERV LOADFORMAT** command, log on to the system as the server instance owner and issue this command:

```
db2set -i instance_name DB2CODEPAGE=819
```

For example:

```
db2set -i tsminst1 DB2CODEPAGE=819
```

### Syntax



### Parameters

**-u** *user\_name*

Specifies a user name to switch to before initializing the server. This parameter is optional.

**-i** *instance\_dir*

Specifies an instance directory to use. This directory becomes the current working directory of the server. This parameter is optional.

**-o** *options\_file*

Specifies an options file to use. This parameter is optional.

## DSMSERV LOADFORMAT

### **-noexpire**

Specifies that expiration processing is suppressed when starting. This parameter is optional.

### **-quiet**

Specifies that messages to the console are suppressed. This parameter is optional.

### **DBDir**

Specifies the relative path names of one or more directories that are used to store database objects. Directory names must be separated by commas but without spaces. You can specify up to 128 directory names. You must specify either the **DBDIR** or the **DBFILE** parameter.

### **DBFile**

Specifies the name of a file that contains the relative path names of one or more directories that are used to store database objects. Each directory name must be on a separate line in the file. You can specify up to 128 directory names. You must specify either the **DBDIR** or the **DBFILE** parameter.

### **ACTIVELOGSize**

Specifies the size of the active log in megabytes. This parameter is optional. The minimum value is 2048 MB (2 GB); the maximum is 131,072 MB (128 GB). If you specify an odd number, the value is rounded up to the next even number. The default is 16384 MB.

### **ACTIVELOGDirectory (Required)**

Specifies the directory in which the Tivoli Storage Manager server writes and stores active log files. There is only one active log location. The name must be a fully qualified directory name. The directory must exist, it must be empty, and it must be accessible by the user ID of the database manager. The maximum number of characters is 175.

### **ARCHLogdirectory (Required)**

Specifies the directory for the archive log files. The name must be a fully qualified directory name. The maximum number of characters is 175.

### **ARCHFailoverlogdirectory**

Specifies the directory to be used as an alternative storage location if the **ARCHLOGDIRECTORY** directory is full. This parameter is optional. The maximum number of characters is 175.

### **MIRRORlogdirectory**

Specifies the directory in which the server mirrors the active log (those files in the **ACTIVELOGDIRECTORY** directory). This parameter is optional. The directory must be a fully qualified directory name. The maximum number of characters is 175.

## **Example: Format a database**

```
dsmserv loadformat dbdir=/tsmdb001 activesize=8192
activedirectory=/active logdirectory=/archlog
archfailoverlogdirectory=/archfaillog mirrorlogdirectory=/mirrorlog
```

## DSMSERV REMOVEDB (Remove a database)

Use the DSMSERV REMOVEDB utility to remove a Tivoli Storage Manager server database.

Use this command to remove a Tivoli Storage Manager server database. This command deletes all user data and log files, as well as any backup and restore history for the database. If the log files are needed for a roll-forward recovery after a restore operation, these files should be saved before you issue this command.

**Note:** The Tivoli Storage Manager server must be halted before you issue this command.

### Syntax

```

DSMSERV [-u user_name] [-i instance_dir] [-o options_file]
 [-noexpire] [-quiet] REMOVEDB database_name

```

### Parameters

**-u** *user\_name*

Specifies a user name to switch to before initializing the server.

**-i** *instance\_dir*

Specifies an instance directory to use. This becomes the current working directory of the server.

**-o** *options\_file*

Specifies an options file to use.

**-noexpire**

Specifies that expiration processing is suppressed when starting.

**-quiet**

Specifies that messages to the console are suppressed.

*database\_name*

The database name that was entered during installation. If the database was formatted manually, then this is the database name parameter in the DSMSERV FORMAT or DSMSERV LOADFORMAT utility. This database name can also be found in `dsmserv.opt` file. This parameter is required.

### Example: Remove a database

Remove the Tivoli Storage Manager server database TSMDB1 and all of its references.

```
dsmserv removedb TSMDB1
```

---

### DSMSERV RESTORE DB (Restore the database)

Use this utility to restore a database by using a database backup.

**Restriction:** You cannot restore a server database from a different Tivoli Storage Manager server if both of the following conditions are present:

- The database backup is stored on virtual volumes
- The connection to the Tivoli Storage Manager server is protected by SSL (Secure Sockets Layer)

Use this utility for the following tasks:

- “DSMSERV RESTORE DB (Restore a database to its most current state)” on page 1575
- “DSMSERV RESTORE DB (Restore a database to a point-in-time)” on page 1578

Before you restore a database to a point-in-time or to its most current state, you should be familiar with the procedures as they are described in the *Administrator's Guide*.

The restore operation uses database backups created with the BACKUP DB command.

**Important:** After a point-in-time restore operation, issue the AUDIT VOLUME command to audit all DISK volumes and resolve any inconsistencies between database information and storage pool volumes. Before restoring the database, examine the volume history file to find out about any sequential access storage pool volumes that were deleted or reused since the point in time to which the database was restored.

## DSMSERV RESTORE DB (Restore a database to its most current state)

Use the DSMSERV RESTORE DB utility to restore a database to its most current state under certain conditions.

The following conditions must be met:

- An intact volume history file is available.
- The recovery logs are available.
- A device configuration file with the applicable device information is available.


**Restriction:** You cannot restore a server database from a different Tivoli Storage Manager server if both of the following conditions are present:

- The database backup is stored on virtual volumes
- The connection to the Tivoli Storage Manager server is protected by SSL (Secure Sockets Layer)

Tivoli Storage Manager requests volume mounts to load the most recent backup series and then uses the recovery logs to update the database to its most current state.

Snapshot database backups cannot be used to restore a database to its most current state.

### Syntax



### Parameters

- u *user\_name*  
Specifies a user name to switch to before initializing the server.
- i *instance\_dir*  
Specifies an instance directory to use. This instance directory becomes the current working directory of the server.
- o *options\_file*  
Specifies an options file to use.
- noexpire  
Specifies that expiration processing is suppressed when starting.

## DSMSERV RESTORE DB

### **-quiet**

Specifies that messages to the console are suppressed.

### **RECOVdir**

Specifies a directory in which to store recovery log information from the database backup media. This directory must have enough space to hold this transaction recovery information and must be an empty directory. If this parameter is not specified, the default is to the directory specified by one of the following parameters in the DSMSERV FORMAT or DSMSERV LOADFORMAT utility:

- ARCHFAILOVERLOGDIRECTORY, if specified
- ARCHLOGDIRECTORY, if ARCHFAILOVERLOGDIRECTORY is not specified

### **ACTIVELOGDirectory**

Specifies a directory in which to store the log files used to track the active database operations. This directory must be specified only if the intent is to switch to an active log directory different from the one that had already been configured.

**On** Specifies a file listing the directories to which the database will be restored. Specify each directory on a separate line in the file. For example, the ON parameter specifies the restorelist.txt file, which contains the following list:

```
/tsmdb001
/tsmdb002
/tsmdb003
```

If this parameter is not specified, the original directories that were recorded in the database backup are used.

### **Preview**

Specifies that the volume history files be examined and that the database backup volumes from the volume history file be evaluated.

1. Which set of database backup volumes best meets the most current criteria specified for restore processing? The volume history information provides details about the backup series ID, the operation ID (full, incremental 1, incremental 2, and so on), the date of the database backup, and the device class. This information and the parameters specified in the **DSMSERV RESTORE DB** command determine what to use to perform the restore. The volume history file is examined to find the most recent database backup and then to perform the restore by using that backup.
2. Is self-describing data available for the selected set of database backup volumes? Cross-check the volume history information for this backup series. The reconciliation reports what the self-describing data contains compared to what was learned from the volume history entries. The cross-check involves mounting one or more of the volumes that are indicated by the volume history. Then, using the self-describing data that was included in the database backup volumes, that information is reconciled against what is in the volume history for the database backup. If the information from the volume history file is inconsistent with the self-describing data, then messages are issued to identify the problem. For example, not all values are specified and available, and no self-describing data is found.

If the volume history information is consistent with self-describing data from the database backup, a message is issued indicating that the database backup can be used for restore processing.

If the volume history information is inconsistent with the self-describing data from the database backup or if the self-describing data for the backup cannot be found, error messages are issued indicating what was checked and what was missing.

If the `PREVIEW` parameter is not specified or if it is set to `NO`, and if the volume history and self-describing data from the database backup are consistent, then the restore proceeds.

If the `PREVIEW` parameter is not specified or if it is set to `NO`, and the reconciliation and validation fail, the database restore is not performed. Make additional volumes available and referred to from the volume history file, or remove the incomplete backup series or operation so that the Tivoli Storage Manager server selects a different preferred series or operation and continues processing.

If the `PREVIEW` parameter is set to `YES`, the process performs only the evaluation of the volume history file and the reconciliation and validation against the selected database backup.

### **Example: Restore the database to its most current state**

Restore the database to its most current state by using the already configured active log directory.

```
dsmserv restore db
```

## DSMSERV RESTORE DB (Restore a database to a point-in-time)


To restore a database to a point in time a volume history file and a device configuration file must be available.

**Restriction:** You cannot restore a server database from a different IBM Tivoli Storage Manager server if both of the following conditions are present:

- The database backup is stored on virtual volumes
- The connection to the IBM Tivoli Storage Manager server is protected by SSL (Secure Sockets Layer)

Full and incremental database backups or snapshot database backups can be used to restore a database to a point in time.

### Syntax



### Parameters

- u user\_name**  
Specifies a user name to switch to before initializing the server.
- i instance\_dir**  
Specifies an instance directory to use. This becomes the current working directory of the server.
- o options\_file**  
Specifies an options file to use.
- noexpire**  
Specifies that expiration processing is suppressed when starting.
- quiet**  
Specifies that messages to the console are suppressed.

#### TODate (Required)

Specifies the date to which to restore the database. Possible values are:

*MM/DD/YYYY*

Specifies that you want to restore a database by using the last backup series that was created before this specified date.

**TODAY**

Specifies that you want to restore a database by using the most recent backup series that was created before today.

**TODAY-*numdays* or -*numdays***

Specifies that you want to restore a database by using the most recent backup series that was created the specified number of days before the current date.

**TOTime**

Specifies the time of day to which to restore the database. This parameter is optional. The default is the end of the day (23:59:59). Possible values are:

*HH:MM:SS*

Specifies that you want to restore the database by using the last backup series created on or before the specified time on the date that is specified on the TODATE parameter.

**NOW**

Specifies that you want to restore the database by using a backup series created on or before the current time on the date that is specified on the TODATE parameter.

For example, if you issue the DSMSERV RESTORE DB utility at 9:00 with TOTIME=NOW, the database is restored by using the last backup series created on or before 9:00 on the date that is specified on the TODATE parameter.

**NOW-*numhours:numminutes* or -*numhours:numminutes***

Specifies that you want to restore the database by using a backup series created on or before the current time minus a specified number of hours and, optionally, minutes on the date that is specified on the TODATE parameter.

For example, if you issue the DSMSERV RESTORE DB utility at 9:00 with TOTIME=NOW-3:30 or TOTIME+-3:30, the database is restored by using the last backup series created on or before 5:30 on the date that is specified on the TODATE parameter.

**Source**

Specifies whether the database is restored by using either database full and incremental backup volumes or snapshot database volumes. This parameter is optional. The default value is DBBackup. Possible values are:

**DBBackup**

Specifies that the database is restored as follows:

1. Reads the volume history file to locate the database full and incremental backup volumes that are needed.
2. Requests mounts and loads the data from the database full and incremental backup volumes as required to restore the database volume to the specified time.

**DBSnapshot**

Specifies that the database is restored as follows:

1. Reads the volume history file to locate the snapshot database volumes that are needed,

## DSMSERV RESTORE DB

2. Requests mounts and loads data from snapshot database volumes as required to restore the volume to the specified time.

### RECOV`erydir`

Specifies a directory in which to store recovery log information from the database backup media. This log information is used to establish transaction consistency of the server database as part of the recovery processing. This directory must have enough space to hold this transaction recovery information and must be an empty directory. If this parameter is not specified, the default is the directory that is specified by one of the following parameters in the DSMSERV FORMAT or DSMSERV LOADFORMAT utility:

- ARCHFAILOVERLOGDIRECTORY, if specified
- ARCHLOGDIRECTORY, if ARCHFAILOVERLOGDIRECTORY is not specified

### ACTIVELOGDirectory

Specifies a directory in which to store the log files that are used to track the active database operations. Specify this directory only if the intent is to switch to an active log directory that is different from the one that had already been configured.

- On** Specifies a file listing the directories to which the database will be restored. Specify each directory on a separate line in the file. For example, the ON parameter specifies the restorelist.txt file, which contains the following list:

```
/tsmdb001
/tsmdb002
/tsmdb003
```

If this parameter is not specified, the original directories that were recorded in the database backup are used.

### PREview

Specifies that the volume history files be examined and that the database backup volumes from the volume history file be evaluated.

1. Which set of database backup volumes best meets the point-in-time criteria specified for restore processing? The volume history information provides details about the backup series ID, the operation ID (full, incremental 1, incremental 2, and so on), the date of the database backup, and the device class. This information and the parameters specified in the **DSMSERV RESTORE DB** command determine what to use to perform the restore. The volume history file is examined to find the best database backup that meets the specified point-in-time criteria and then perform the restore by using that backup.
2. Is self-describing data available for the selected set of database backup volumes? Cross-check the volume history information for this backup series. The reconciliation reports what the self-describing data contains compared to what was learned from the volume history entries. The cross-check involves mounting one or more of the volumes that are indicated by the volume history. Then, using the self-describing data that was included in the database backup volumes, that information is reconciled against what is in the volume history for the database backup. If the information from the volume history file is inconsistent with the self-describing data, then messages are issued to identify the problem. For example, not all values are specified and available, and no self-describing data is found.

If the volume history information is consistent with self-describing data from the database backup, a message is issued indicating that the database backup can be used for restore processing.

If the volume history information is inconsistent with the self-describing data from the database backup or if the self-describing data for the backup cannot be found, error messages are issued indicating what was checked and what was missing.

If the `PREVIEW` parameter is not specified or if it is set to `NO`, and if the volume history and self-describing data from the database backup are consistent, then the restore proceeds.

If the `PREVIEW` parameter is not specified or if it is set to `NO`, and the reconciliation and validation fail, the database restore is not performed. Make additional volumes available and referred to from the volume history file, or remove the incomplete backup series or operation so that the Tivoli Storage Manager server selects a different preferred series or operation and continues processing.

If the `PREVIEW` parameter is set to `YES`, the process performs only the evaluation of the volume history file and the reconciliation and validation against the selected database backup.

### **Example: Restore the database to a specific point in time**

Restore the database to its state on May 12, 2011 at 2:25 PM.

```
dsmserv restore db todate=05/12/2011 totime=14:45
```

## DSMSERV RESTORE DB

---

## Appendix A. Return codes for use in IBM Tivoli Storage Manager scripts

You can write Tivoli Storage Manager scripts that use return codes to determine how script processing proceeds. The return codes can be one of three severities: OK, WARNING, ERROR.

Tivoli Storage Manager scripts use the symbolic return code for processing, not the numeric value. The administrative client displays the numeric values when a command is run. The return codes are shown in the following table.

Table 461. Return codes

| Return code     | Severity | Numeric value | Description                                                                                                   |
|-----------------|----------|---------------|---------------------------------------------------------------------------------------------------------------|
| RC_OK           | OK       | 0             | The command completed successfully                                                                            |
| RC_UNKNOWN      | ERROR    | 2             | The command is not found; not a known command.                                                                |
| RC_SYNTAX       | ERROR    | 3             | The command is valid, but one or more parameters were not specified correctly.                                |
| RC_ERROR        | ERROR    | 4             | An internal server error prevented the command from successfully completing.                                  |
| RC_NOMEMORY     | ERROR    | 5             | The command could not be completed because of insufficient memory on the server.                              |
| RC_NOLOG        | ERROR    | 6             | The command could not be completed because of insufficient recovery log space on the server.                  |
| RC_NODB         | ERROR    | 7             | The command could not be completed because of insufficient database space on the server.                      |
| RC_NOSTORAGE    | ERROR    | 8             | The command could not be completed because of insufficient storage space on the server.                       |
| RC_NOAUTH       | ERROR    | 9             | The command failed because the administrator is not authorized to issue the command.                          |
| RC_EXISTS       | ERROR    | 10            | The command failed because the specified object already exists on the server.                                 |
| RC_NOTFOUND     | WARNING  | 11            | Returned by a QUERY or SQL SELECT command when no objects are found that match specifications.                |
| RC_INUSE        | ERROR    | 12            | The command failed because the object to be operated upon was in use.                                         |
| RC_ISREFERENCED | ERROR    | 13            | The command failed because the object to be operated upon is still referenced by some other server construct. |
| RC_NOTAVAILABLE | ERROR    | 14            | The command failed because the object to be operated upon is not available.                                   |

Table 461. Return codes (continued)

| Return code     | Severity | Numeric value | Description                                                                                     |
|-----------------|----------|---------------|-------------------------------------------------------------------------------------------------|
| RC_IOERROR      | ERROR    | 15            | The command failed because an input/output (I/O) error was encountered on the server.           |
| RC_NOTXN        | ERROR    | 16            | The command failed because a database transaction failed on the server.                         |
| RC_NOLOCK       | ERROR    | 17            | The command failed because a lock conflict was encountered in the server database.              |
| RC_NOTHREAD     | ERROR    | 19            | The command could not be completed because of insufficient memory on the server.                |
| RC_LICENSE      | ERROR    | 20            | The command failed because the server is not in compliance with licensing.                      |
| RC_INVDEST      | ERROR    | 21            | The command failed because a destination value was invalid.                                     |
| RC_IFILEOPEN    | ERROR    | 22            | The command failed because an input file that was needed could not be opened.                   |
| RC_OFILEOPEN    | ERROR    | 23            | The command failed because it could not open a required output file.                            |
| RC_OFILEWRITE   | ERROR    | 24            | The command failed because it could not successfully write to a required output file.           |
| RC_INVADMIN     | ERROR    | 25            | The command failed because the administrator was not defined.                                   |
| RC_SQLERROR     | ERROR    | 26            | An SQL error was encountered during a SELECT statement query.                                   |
| RC_INVALIDUSE   | ERROR    | 27            | The command failed because the command is used in an invalid manner.                            |
| RC_NOTABLE      | ERROR    | 28            | The command failed because of an unknown SQL table name.                                        |
| RC_FS_NOTCAP    | ERROR    | 29            | The command failed because of Non-Compatible Filespace name types.                              |
| RC_INVALIDADDR  | ERROR    | 30            | The command failed because of an incorrect high-level address or low-level address.             |
| RC_INVALIDCG    | ERROR    | 31            | The command failed because the management class does not have an archive copy group.            |
| RC_OVERSIZE_VOL | ERROR    | 32            | The command failed because the volume size exceeds the maximum allowed.                         |
| RC_DEFVOL_FAIL  | ERROR    | 33            | The command failed because volumes cannot be defined in RECLAMATIONTYPE=SNAPLOCK storage pools. |
| RC_DELVOL_FAIL  | ERROR    | 34            | The command failed because volumes cannot be deleted in RECLAMATIONTYPE=SNAPLOCK storage pools. |
| RC_CANCELED     | WARNING  | 35            | The command is canceled.                                                                        |

Table 461. Return codes (continued)

| Return code    | Severity | Numeric value | Description                                                                     |
|----------------|----------|---------------|---------------------------------------------------------------------------------|
| RC_INVPOLICY   | ERROR    | 36            | The command failed because there is an invalid definition in the policy domain. |
| RC_INVALIDPW   | ERROR    | 37*           | The command failed because of an invalid password.                              |
| RC_UNSUPP_PARM | WARNING  | 38*           | The command failed because the command or the parameter is not supported.       |

**Note:** \*These return codes are supported in Tivoli Storage Manager Version 6 but are not supported in version 5.



---

## Appendix B. Support information

You can find support information for IBM products from various sources.

Start at the IBM Support Portal: <http://www.ibm.com/support/entry/portal/>. You can select the products that you are interested in and search for a wide variety of relevant information.

---

### Getting technical training

Information about Tivoli technical training courses is available online.

Go to the following websites to sign up for training, ask questions, and interact with others who use IBM storage products.

#### **Tivoli software training and certification**

Choose from instructor led, online classroom training, self-paced Web classes, Tivoli certification preparation, and other training options at <http://www.ibm.com/software/tivoli/education/>

#### **Tivoli Support Technical Exchange**

Technical experts share their knowledge and answer your questions in webcasts at [http://www.ibm.com/software/sysmgmt/products/support/supp\\_tech\\_exch.html](http://www.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html).

#### **Storage Management community**

Interact with others who use IBM storage management products at <http://www.ibm.com/developerworks/servicemanagement/sm/index.html>

#### **Global Tivoli User Community**

Share information and learn from other Tivoli users throughout the world at <http://www.tivoli-ug.org/>.

#### **IBM Education Assistant**

View short "how to" recordings designed to help you use IBM software products more effectively at <http://publib.boulder.ibm.com/infocenter/ieduasst/tivv1r0/index.jsp>

---

### Searching knowledge bases

If you have a problem with your Tivoli Storage Manager family product, there are several knowledge bases that you can search.

Begin by searching the Tivoli Storage Manager Information Center at <http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3>. From this website, you can search the current Tivoli Storage Manager documentation.

## Searching the Internet

If you cannot find an answer to your question in the IBM Tivoli Storage Manager information center, search the Internet for the information that might help you resolve the problem.

To search multiple Internet resources, go to the IBM support website at <http://www.ibm.com/support/entry/portal/>.

You can search for information without signing in. Sign in using your IBM ID and password if you want to customize the site based on your product usage and information needs. If you do not already have an IBM ID and password, click **Sign in** at the top of the page and follow the instructions to register.

From the support website, you can search various resources including:

- IBM technotes.
- IBM downloads.
- IBM Redbooks® publications.
- IBM Authorized Program Analysis Reports (APARs). Select the product and click **Downloads** to search the APAR list.

If you still cannot find a solution to the problem, you can search forums and newsgroups on the Internet for the latest information that might help you find problem resolution.

An independent user discussion list, ADSM-L, is hosted by Marist College. You can subscribe by sending an email to [listserv@vm.marist.edu](mailto:listserv@vm.marist.edu). The body of the message must contain the following text: `SUBSCRIBE ADSM-L your_first_name your_family_name`.

To share your experiences and learn from others in the Tivoli Storage Manager and Tivoli Storage FlashCopy Manager user communities, go to Service Management Connect (<http://www.ibm.com/developerworks/servicemanagement/sm/index.html>). From there you can find links to product wikis and user communities.

## Using IBM Support Assistant

IBM Support Assistant is a complimentary software product that can help you with problem determination. It is available for some Tivoli Storage Manager and Tivoli Storage FlashCopy Manager products.

To learn about which products are supported, go to the IBM Support Assistant download web page at <http://www.ibm.com/software/support/isa/download.html>.

IBM Support Assistant helps you gather support information when you must open a problem management record (PMR), which you can then use to track the problem. The product-specific plug-in modules provide you with the following resources:

- Support links
- Education links
- Ability to submit problem management reports

You can find more information at the IBM Support Assistant website:

<http://www.ibm.com/software/support/isa/>

You can also install the stand-alone IBM Support Assistant application on any workstation. You can then enhance the application by installing product-specific plug-in modules for the IBM products that you use. Find add-ons for specific products at <http://www.ibm.com/support/docview.wss?uid=swg27012689>.

## Finding product fixes

A product fix to resolve your problem might be available from the IBM software support website.

You can determine what fixes are available by checking the IBM software support website at <http://www.ibm.com/support/entry/portal/>.

- If you previously customized the site based on your product usage:
  1. Click the link for your product, or a component for which you want to find a fix.
  2. Click **Downloads**, and then click **Fixes by version**.
- If you have not customized the site based on your product usage, click **Downloads** and search for your product.

## Receiving notification of product fixes

You can receive notifications about fixes, flashes, upgrades, and other news about IBM products.

To sign up to receive notifications about IBM products, follow these steps:

1. From the support page at <http://www.ibm.com/support/entry/portal/>, click **Sign in to create, manage, or view your subscriptions** in the **Notifications** pane.
2. Sign in using your IBM ID and password. If you do not have an ID and password, click **register now** and complete the registration process.
3. Click **Manage all my subscriptions** in the **Notifications** pane.
4. Click the **Subscribe** tab and then click **Tivoli**.
5. Select the products for which you want to receive notifications and click **Continue**.
6. Specify your notification preferences and click **Submit**.

---

## Contacting IBM Software Support

You can contact IBM Software Support if you have an active IBM subscription and support contract and if you are authorized to submit problems to IBM.

To obtain help from IBM Software Support, complete the following steps:

1. Ensure that you have completed the following prerequisites:
  - a. Set up a subscription and support contract.
  - b. Determine the business impact of your problem.
  - c. Describe your problem and gather background information.
2. Follow the instructions in “Submitting the problem to IBM Software Support” on page 1591.

## Setting up a subscription and support contract

Set up a subscription and support contract. The type of contract that you need depends on the type of product you have.

For IBM distributed software products (including, but not limited to, IBM Tivoli, Lotus, and Rational® products, as well as IBM DB2 and IBM WebSphere® products that run on Microsoft Windows or on operating systems such as AIX or Linux), enroll in IBM Passport Advantage® in one of the following ways:

- **Online:** Go to the Passport Advantage website at <http://www.ibm.com/software/lotus/passportadvantage/>, click **How to enroll**, and follow the instructions.
- **By telephone:** You can call 1-800-IBMSERV (1-800-426-7378) in the United States. For the telephone number to call in your country, go to the IBM Software Support Handbook web page at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.

## Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you must understand and assess the business impact of the problem you are reporting.

|                   |                                                                                                                                                                  |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Severity 1</b> | <b>Critical</b> business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution. |
| <b>Severity 2</b> | <b>Significant</b> business impact: The program is usable but is severely limited.                                                                               |
| <b>Severity 3</b> | <b>Some</b> business impact: The program is usable with less significant features (not critical to operations) unavailable.                                      |
| <b>Severity 4</b> | <b>Minimal</b> business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem has been implemented.               |

## Describing the problem and gathering background information

When explaining a problem to IBM, it is helpful to be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently.

To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be re-created? If so, what steps led to the failure?
- Have any changes been made to the system? For example, hardware, operating system, networking software, and so on.
- Are you using a workaround for this problem? If so, be prepared to explain it when you report the problem.

## Submitting the problem to IBM Software Support

You can submit the problem to IBM Software Support online or by telephone.

### Online

Go to the IBM Software Support website at [http://www.ibm.com/support/entry/portal/Open\\_service\\_request/Software/Software\\_support\\_\(general\)](http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_(general)). Sign in to access IBM Service Requests and enter your information into the problem submission tool.

### By telephone

For the telephone number to call in your country, go to the IBM Software Support Handbook at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.



---

## Appendix C. Accessibility features for the Tivoli Storage Manager product family

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

### Accessibility features

The following list includes the major accessibility features in the Tivoli Storage Manager family of products:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

If you install the IBM Tivoli Storage Manager Operations Center in console mode, the installation is fully accessible.

The accessibility features of the Operations Center are fully supported only in the Mozilla Firefox browser that is running on a Windows system.

The Tivoli Storage Manager Information Center, and its related publications, are accessibility-enabled. For information about the accessibility features of the information center, see the following topic: [http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3/topic/com.ibm.help.ic.doc/iehs36\\_accessibility.html](http://pic.dhe.ibm.com/infocenter/tsminfo/v6r3/topic/com.ibm.help.ic.doc/iehs36_accessibility.html).

### Keyboard navigation

On Windows, the Tivoli Storage Manager product family follows Microsoft conventions for all keyboard navigation and access. Drag and Drop support is managed by using the Microsoft Windows Accessibility option known as *MouseKeys*. For more information about MouseKeys and other Windows accessibility options, see the Windows online help, citing the keyword "MouseKeys".

On other operating systems, these products follow the operating-system conventions for keyboard navigation and access.

### Vendor software

The Tivoli Storage Manager product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

### IBM and accessibility

See the IBM Human Ability and Accessibility Center (<http://www.ibm.com/able>) for information about the commitment that IBM has to accessibility.



---

## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd  
1623-14, Shimotsuruma, Yamato-shi  
Kanagawa 242-8502 Japan*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:**

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who want to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation  
2Z4A/101  
11400 Burnet Road  
Austin, TX 78758  
U.S.A.*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample

programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. \_enter the year or years\_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)<sup>®</sup> are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

LTO and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.



---

## Glossary

This glossary includes terms and definitions for IBM Tivoli Storage Manager and IBM Tivoli Storage FlashCopy Manager products.

To view glossaries for other IBM products, go to <http://www.ibm.com/software/globalization/terminology/>.

The following cross-references are used in this glossary:

- *See* refers the reader from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- *See also* refers the reader to a related or contrasting term.

### A

#### **absolute mode**

In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup even if the file has not changed since the last backup. See also *modified mode*.

#### **access control list (ACL)**

In computer security, a list associated with an object that identifies all the subjects that can access the object and their access rights. For example, an access control list is associated with a file that identifies the users who can access that file and their access rights.

#### **access mode**

An attribute of a storage pool or a storage volume that specifies whether the server can write to or read from the storage pool or storage volume. The access mode can be read/write, read-only, or unavailable. Volumes in primary storage pools can also have an access mode of destroyed. Volumes in copy storage pools can also have an access mode of offsite.

#### **acknowledgment**

The transmission of acknowledgment characters as a positive response to a data transmission.

**ACL** See *access control list*.

#### **activate**

To validate the contents of a policy set and then make it the active policy set.

#### **active-data pool**

A named set of storage pool volumes that contain only active versions of client backup data.

#### **active file system**

A file system to which space management has been added. With space management, tasks for an active file system include automatic migration, reconciliation, selective migration, and recall. Contrast with *inactive file system*.

#### **active policy set**

The activated policy set that contains the policy rules in use by all client nodes that are assigned to the policy domain. See also *policy domain* and *policy set*.

#### **active version**

The most recent backup copy of a file stored. The active version of a file cannot be deleted until a backup process detects that the user has either replaced the file with a newer version or has deleted the file from the file server or workstation. Contrast with *inactive version*.

#### **activity log**

A log that records normal activity messages that are generated by the server. These messages include information about server and client operations, such as the start time of sessions or device I/O errors.

#### **adaptive subfile backup**

A type of backup that sends only changed portions of a file to the server, instead of sending the entire file. Adaptive subfile backup reduces network traffic and increases the speed of the backup.

#### **administrative client**

A program that runs on a file server, workstation, or mainframe that administrators use to control and monitor the Tivoli Storage Manager server. Contrast with *backup-archive client*.

#### **administrative command schedule**

A database record that describes the

planned processing of an administrative command during a specific time period. See also *client schedule*.

**administrative privilege class**

See *privilege class*.

**administrative session**

A period of time during which an administrator user ID communicates with a server to perform administrative tasks. Contrast with *client node session*.

**administrator**

A user who is registered to the server as an administrator, and who is authorized to perform tasks and issue commands through the assignment of an administrative privilege class.

**Advanced Program-to-Program Communication (APPC)**

An implementation of the SNA LU 6.2 protocol that allows interconnected systems to communicate and share the processing of programs.

**agent node**

A client node that has been granted proxy authority to perform operations on behalf of another client node, which is the target node.

**aggregate**

An object, stored in one or more storage pools, consisting of a group of logical files that are packaged together. See also *logical file* and *physical file*.

**aggregate data transfer rate**

A performance statistic that indicates the average number of bytes that were transferred per second while processing a given operation.

**APPC** See *Advanced Program-to-Program Communication*.

**application client**

A program that is installed on a system to protect an application. The Tivoli Storage Manager server provides backup services to an application client.

**archive**

To copy programs, data, or files to other storage media, usually for long-term storage or security. Contrast with *retrieve*.

**archive copy**

A file or group of files that was archived to server storage.

**archive copy group**

A policy object containing attributes that control the generation, destination, and expiration of archived files.

**archive-retention grace period**

The number of days that the storage manager retains an archived file when the server is unable to rebind the file to an appropriate management class. See also *bind*.

**association**

(1) The defined relationship between a client node and a client schedule. An association identifies the name of a schedule, the name of the policy domain to which the schedule belongs, and the name of a client node that performs scheduled operations.

(2) On a configuration manager, the defined relationship between a profile and an object such as a policy domain. Profile associations define the configuration information that is distributed to a managed server when it subscribes to the profile.

**audit**

To check for logical inconsistencies between information that the server has and the actual condition of the system. The storage manager can audit information about items such as volumes, libraries, and licenses. For example, when a storage manager audits a volume, the server checks for inconsistencies between information about backed-up or archived files that are stored in the database and the actual data that are associated with each backup version or archive copy in server storage.

**authentication**

The process of checking a user's password before permitting user access to the Tivoli Storage Manager server. Authentication can be turned on or off by an administrator with system privilege.

**authentication rule**

A specification that another user can use to either restore or retrieve files from storage.

**authority**

The right to access objects, resources, or functions. See also *privilege class*.

**authorization rule**

A specification that permits another user to either restore or retrieve a user's files from storage.

**authorized user**

A user who has administrative authority for the Tivoli Storage Manager client on a workstation. This user changes passwords, performs open registrations, and deletes file spaces.

**AutoFS**

See *automounted file system*.

**automatic detection**

A feature that detects, reports, and updates the serial number of a drive or library in the database when the path from the local server is defined.

**automatic migration**

The process that is used to automatically move files from a local file system to storage, based on options and settings that are chosen by a root user on a workstation. See also *threshold migration* and *demand migration*.

**automatic reconciliation**

The process that is used to reconcile file systems at regular intervals. The intervals are set by a user with root user authority. See also *reconciliation*.

**automounted file system (AutoFS)**

A file system that is managed by an automounter daemon. The automounter daemon monitors a specified directory path, and automatically mounts the file system to access data.

**B****backup-archive client**

A program that runs on a workstation or file server and provides a means for users to back up, archive, restore, and retrieve files. Contrast with *administrative client*.

**backup copy group**

A policy object containing attributes that control the generation, destination, and expiration of backup versions of files. A backup copy group belongs to a management class.

**backup-retention grace period**

The number of days the storage manager retains a backup version after the server is unable to rebind the file to an appropriate management class.

**backup set**

A portable, consolidated group of active versions of backup files that are generated for a backup-archive client.

**backup set collection**

A group of backup sets that are created at the same time and which have the same backup set name, volume names, description, and device classes. The server identifies each backup set in the collection by its node name, backup set name, and file type.

**backup version**

A file or directory that a client node backed up to server storage. More than one backup version can exist in server storage, but only one backup version is the active version. See also *active version* and *inactive version*.

**bind** To associate all versions of a file with a management class name. See *rebind*.

**bindery**

A database that consists of three system files for a NetWare server. The files contain user IDs and user restrictions.

**C**

**cache** To place a duplicate copy of a file on random access media when the server migrates a file to another storage pool in the hierarchy.

**cache file**

A snapshot of a logical volume created by Logical Volume Snapshot Agent. Blocks are saved immediately before they are modified during the image backup and their logical extents are saved in the cache files.

**CAD** See *client acceptor*.

**central scheduler**

A function that permits an administrator to schedule client operations and administrative commands. The operations can be scheduled to occur periodically or on a specific date. See *client schedule* and *administrative command schedule*.

**client** A software program or computer that requests services from a server.

**client acceptor**

An HTTP service that serves the applet for the web client to web browsers. On Windows systems, the client acceptor is installed and run as a service. On AIX, UNIX, and Linux systems, the client acceptor is run as a daemon, and is also called the *client acceptor daemon (CAD)*.

**client acceptor daemon (CAD)**

See *client acceptor*.

**client domain**

The set of drives, file systems, or volumes that the user selects to back up or archive data, using the backup-archive client.

**client node**

A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

**client node session**

A session in which a client node communicates with a server to perform backup, restore, archive, retrieve, migrate, or recall requests. Contrast with *administrative session*.

**client option set**

A group of options that are defined on the server and used on client nodes in conjunction with client options files.

**client options file**

An editable file that identifies the server and communication method, and provides the configuration for backup, archive, hierarchical storage management, and scheduling.

**client-polling scheduling mode**

A method of operation in which the client queries the server for work. Contrast with *server-prompted scheduling mode*.

**client schedule**

A database record that describes the planned processing of a client operation during a specific time period. The client operation can be a backup, archive, restore, or retrieve operation, a client operating system command, or a macro. See also *administrative command schedule*.

**client/server**

Pertaining to the model of interaction in

distributed data processing in which a program on one computer sends a request to a program on another computer and awaits a response. The requesting program is called a client; the answering program is called a server.

**client system-options file**

A file, used on AIX, UNIX, or Linux system clients, containing a set of processing options that identify the servers to be contacted for services. This file also specifies communication methods and options for backup, archive, hierarchical storage management, and scheduling. This file is also called the *dsm.sys* file. See also *client user-options file*.

**client user-options file**

A file that contains the set of processing options that the clients on the system use. The set can include options that determine the server that the client contacts, and options that affect backup operations, archive operations, hierarchical storage management operations, and scheduled operations. This file is also called the *dsm.opt* file. For AIX, UNIX, or Linux systems, see also *client system-options file*.

**closed registration**

A registration process in which only an administrator can register workstations as client nodes with the server. Contrast with *open registration*.

**collocation**

The process of keeping all data belonging to a single-client file space, a single client node, or a group of client nodes on a minimal number of sequential-access volumes within a storage pool. Collocation can reduce the number of volumes that must be accessed when a large amount of data must be restored.

**collocation group**

A user-defined group of client nodes whose data is stored on a minimal number of volumes through the process of collocation.

**commit point**

A point in time when data is considered consistent.

**Common Programming Interface for Communications (CPI-C)**

A call-level interface that provides a consistent application programming interface (API) for applications that use program-to-program communications. CPI-C uses LU 6.2 architecture to create a set of interprogram services that can establish and end a conversation, send and receive data, exchange control information, and notify a partner program of errors.

**communication method**

The method by which a client and server exchange information. See also *Transmission Control Protocol/Internet Protocol*.

**communication protocol**

A set of defined interfaces that permit computers to communicate with each other.

**compression**

A function that removes repetitive characters, spaces, or strings of characters from the data being processed and replaces the repetitive characters with control characters. Compression reduces the amount of storage space that is required for the data.

**configuration manager**

A server that distributes configuration information, such as policies and schedules, to managed servers according to their profiles. Configuration information can include policy and schedules. See also *managed server* and *profile*.

**conversation**

A connection between two programs over a session that allows them to communicate with each other while processing a transaction.

**copy backup**

A full backup in which the transaction log files are not deleted so that backup procedures that use incremental or differential backups are not disrupted.

**copy group**

A policy object containing attributes that control how backup versions or archive copies are generated, where backup versions or archive copies are initially

located, and when backup versions or archive copies expire. A copy group belongs to a management class. See also *archive copy group*, *backup copy group*, *backup version*, and *management class*.

**copy storage pool**

A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class (for space-managed files). See also *primary storage pool* and *destination*.

**CPI-C** See *Common Programming Interface for Communications*.

**D****daemon**

A program that runs unattended to perform continuous or periodic functions, such as network control.

**damaged file**

A physical file in which Tivoli Storage Manager has detected read errors.

**data access control mode**

A mode that controls whether a command can access a migrated file, see a migrated file as zero-length, or receive an input/output error if it attempts to access a migrated file. See also *execution mode*.

**database backup series**

One full backup of the database, plus up to 32 incremental backups made since that full backup. Each full backup that is run starts a new database backup series. A number identifies each backup series.

**database snapshot**

A complete backup of the entire database to media that can be taken off-site. When a database snapshot is created, the current database backup series is not interrupted. A database snapshot cannot have incremental database backups associated with it. See also *database backup series*. Contrast with *full backup*.

**data deduplication**

A method of reducing storage needs by eliminating redundant data. Only one instance of the data is retained on storage

media. Other instances of the same data are replaced with a pointer to the retained instance.

**data manager server**

A server that collects metadata information for client inventory and manages transactions for the storage agent over the local area network. The data manager server informs the storage agent with applicable library attributes and the target volume identifier.

**data mover**

A device that moves data on behalf of the server. A network-attached storage (NAS) file server is a data mover.

**data storage-management application-programming interface (DSMAPI)**

A set of functions and semantics that can monitor events on files, and manage and maintain the data in a file. In an HSM environment, a DSMAPI uses events to notify data management applications about operations on files, stores arbitrary attribute information with a file, supports managed regions in a file, and uses DSMAPI access rights to control access to a file object.

**deduplication**

See *data deduplication*.

**default management class**

A management class that is assigned to a policy set. This class is used to govern backed up or archived files when a file is not explicitly associated with a specific management class through the include-exclude list.

**demand migration**

The process that is used to respond to an out-of-space condition on a file system for which hierarchical storage management (HSM) is active. Files are migrated to server storage until space usage drops to the low threshold that was set for the file system. If the high threshold and low threshold are the same, one file is migrated.

**desktop client**

The group of backup-archive clients that includes clients on Microsoft Windows, Apple, and Novell NetWare operating systems.

**destination**

A copy group or management class attribute that specifies the primary storage pool to which a client file will be backed up, archived, or migrated.

**device class**

A named set of characteristics that are applied to a group of storage devices. Each device class has a unique name and represents a device type of disk, file, optical disk, or tape.

**device configuration file**

(1) For a server, a file that contains information about defined device classes, and, on some servers, defined libraries and drives. The information is a copy of the device configuration information in the database.

(2) For a storage agent, a file that contains the name and password of the storage agent, and information about the server that is managing the SAN-attached libraries and drives that the storage agent uses.

**device driver**

A program that provides an interface between a specific device and the application program that uses the device.

**disaster recovery manager (DRM)**

A function that assists in preparing and using a disaster recovery plan file for the server.

**disaster recovery plan**

A file that is created by the disaster recovery manager (DRM) that contains information about how to recover computer systems if a disaster occurs and scripts that can be run to perform some recovery tasks. The file includes information about the software and hardware that is used by the server, and the location of recovery media.

**domain**

A grouping of client nodes with one or more policy sets, which manage data or storage resources for the client nodes. See *policy domain* or *client domain*.

**DRM** See *disaster recovery manager*.

**DSMAPI**

See *data storage-management application-programming interface*.

**dynamic serialization**

A type of copy serialization in which a file or folder is backed up or archived on the first attempt regardless of whether it changes during a backup or archive.

**E**

**EA** See *extended attribute*.

**EB** See *exabyte*.

**EFS** See *Encrypted File System*.

**Encrypted File System (EFS)**

A file system that uses file system-level encryption.

**enterprise configuration**

A method of setting up servers so that the administrator can distribute the configuration of one of the servers to the other servers, using server-to-server communication. See also *configuration manager*, *managed server*, *profile*, and *subscription*.

**enterprise logging**

The process of sending events from a Tivoli Storage Manager server to a designated event server. The event server routes the events to designated receivers, such as to a user exit. See also *event*.

**error log**

A data set or file that is used to record error information about a product or system.

**estimated capacity**

The available space, in megabytes, of a storage pool.

**event** (1) An administrative command or a client operation that is scheduled to be run using Tivoli Storage Manager scheduling.

(2) A message that an Tivoli Storage Manager server or client issues. Messages can be logged using Tivoli Storage Manager event logging.

**event record**

A database record that describes actual status and results for events.

**event server**

A server to which other servers can send events for logging. The event server routes the events to any receivers that are enabled for the sending server's events.

**exabyte (EB)**

For processor storage, real and virtual storage, and channel volume, 1 152 921 504 606 846 976 bytes. For disk storage capacity and communications volume, 1 000 000 000 000 000 000 bytes.

**exclude**

The process of identifying files in an include-exclude list. This process prevents the files from being backed up or migrated whenever a user or schedule enters an incremental or selective backup operation. A file can be excluded from backup and space management, backup only, or space management only.

**exclude-include list**

See *include-exclude list*.

**execution mode**

A mode that controls the space-management related behavior of commands that run under the **dsmmode** command.

**expiration**

The process by which files, data sets, or objects are identified for deletion because their expiration date or retention period has passed.

**expiring file**

A migrated or premigrated file that has been marked for expiration and removal from storage. If a stub file or an original copy of a premigrated file is deleted from a local file system, or if the original copy of a premigrated file is updated, the corresponding migrated or premigrated file is marked for expiration the next time reconciliation is run.

**extend**

To increase the portion of available space that can be used to store database or recovery log information.

**extended attribute (EA)**

Names or value pairs that are associated with files or directories. There are three classes of extended attributes: user attributes, system attributes, and trusted attributes.

**extent** The part of a file that is created during the data-deduplication process. Extents are compared with other file extents to identify duplicates.

**external library**

A type of library that is provided by Tivoli Storage Manager that permits LAN-free data movement for StorageTek libraries that are managed by Automated Cartridge System Library Software (ACSL). To activate this function, the Tivoli Storage Manager library type must be EXTERNAL.

**F****file access time**

On AIX, UNIX, or Linux systems, the time when the file was last accessed.

**file age**

For migration prioritization purposes, the number of days since a file was last accessed.

**file device type**

A device type that specifies the use of sequential access files on disk storage as volumes.

**file server**

A dedicated computer and its peripheral storage devices that are connected to a local area network that stores programs and files that are shared by users on the network.

**file space**

A logical space in server storage that contains a group of files that have been backed up or archived by a client node, from a single logical partition, file system, or virtual mount point. Client nodes can restore, retrieve, or delete their file spaces from server storage. In server storage, files belonging to a single file space are not necessarily stored together.

**file space ID (FSID)**

A unique numeric identifier that the server assigns to a file space when it is stored in server storage.

**file state**

The space management mode of a file that resides in a file system to which space management has been added. A file can be in one of three states: resident, premigrated, or migrated. See also *resident file*, *premigrated file*, and *migrated file*.

**file system migrator (FSM)**

A kernel extension that intercepts all file system operations and provides any space

management support that is required. If no space management support is required, the operation is passed to the operating system, which performs its normal functions. The file system migrator is mounted over a file system when space management is added to the file system.

**file system state**

The storage management mode of a file system that resides on a workstation on which the hierarchical storage management (HSM) client is installed. A file system can be in one of these states: native, active, inactive, or global inactive.

**frequency**

A copy group attribute that specifies the minimum interval, in days, between incremental backups.

**FSID** See *file space ID*.

**FSM** See *file system migrator*.

**full backup**

The process of backing up the entire server database. A full backup begins a new database backup series. See also *database backup series* and *incremental backup*. Contrast with *database snapshot*.

**fuzzy backup**

A backup version of a file that might not accurately reflect what is currently in the file because the file was backed up at the same time as it was being modified.

**fuzzy copy**

A backup version or archive copy of a file that might not accurately reflect the original contents of the file because it was backed up or archived the file while the file was being modified. See also *backup version* and *archive copy*.

**G****General Parallel File System**

A high-performance shared-disk file system that can provide data access from nodes in a cluster environment.

**gigabyte (GB)**

In decimal notation, 1 073 741 824 when referring to memory capacity; in all other cases, it is defined as 1 000 000 000.

**global inactive state**

The state of all file systems to which

space management has been added when space management is globally deactivated for a client node. When space management is globally deactivated, hierarchical storage management (HSM) cannot perform migration, recall, or reconciliation. However, a root user can update space management settings and add space management to additional file systems. Users can access resident and premigrated files.

**Globally Unique Identifier (GUID)**

An algorithmically determined number that uniquely identifies an entity within a system.

**GPFS™**

See *General Parallel File System*.

**GPFS node set**

A mounted, defined group of GPFS file systems.

**group backup**

The backup of a group containing a list of files from one or more file space origins.

**GUID** See *Globally Unique Identifier*.

**H**

**hierarchical storage management (HSM)**

A function that automatically distributes and manages data on disk, tape, or both by regarding devices of these types and potentially others as levels in a storage hierarchy that range from fast, expensive devices to slower, cheaper, and possibly removable devices. The objectives are to minimize access time to data and maximize available media capacity.

**hierarchical storage management (HSM) client**

A client program that works with the Tivoli Storage Manager server to provide hierarchical storage management (HSM) for a system. See also *hierarchical storage management* and *space manager client*.

**HSM** See *hierarchical storage management*.

**HSM client**

See *hierarchical storage management client*.

**I**

**ILM** See *information lifecycle management*.

**image** A file system or raw logical volume that is backed up as a single object.

**image backup**

A backup of a full file system or raw logical volume as a single object.

**inactive file system**

A file system for which space management has been deactivated. Contrast with *active file system*.

**inactive version**

A backup version of a file that is either not the most recent backup version, or that is a backup version of a file that no longer exists on the client system. Inactive backup versions are eligible for expiration processing according to the management class assigned to the file. Contrast with *active version*.

**include-exclude file**

A file containing statements to determine the files to back up and the associated management classes to use for backup or archive. See also *include-exclude list*.

**include-exclude list**

A list of options that include or exclude selected files for backup. An exclude option identifies files that should not be backed up. An include option identifies files that are exempt from the exclusion rules or assigns a management class to a file or a group of files for backup or archive services.

**incremental backup**

(1) A copy of all database data that has changed since the most recent successful full backup operation. An incremental backup is also known as a *cumulative backup image* because each incremental backup includes the contents of the previous incremental backup.

(2) The process of backing up information in the database that is new or changed since the last full backup. Contrast with *full backup*. See also *database backup series*.

(3) For Data Protection for Microsoft Exchange Server, a backup in which the transaction logs are backed up and then cleared.

**individual mailbox restore**

See *mailbox restore*.

**information lifecycle management (ILM)**

A policy-based file management system

for storage pools and file sets. See also *General Parallel File System*.

**inode** The internal structure that describes the individual files on AIX, UNIX, or Linux systems. An inode contains the node, type, owner, and location of a file.

**inode number**  
A number specifying a particular inode file in the file system.

**IP address**  
A unique address for a device or logical unit on a network that uses the IP standard.

## J

**job file**  
A generated file that contains configuration information for a migration job. The file is XML format and can be created and edited in the hierarchical storage management (HSM) client for Windows client graphical user interface.

**journal-based backup**  
A method for backing up Windows clients and AIX clients that exploits the change notification mechanism in a file to improve incremental backup performance by reducing the need to fully scan the file system.

**journal daemon**  
On AIX, UNIX, or Linux systems, a program that tracks change activity for files residing in file systems.

**journal service**  
In Microsoft Windows, a program that tracks change activity for files residing in file systems.

## K

**kilobyte (KB)**  
For processor storage, real and virtual storage, and channel volume, 210 or 1 024 bytes. For disk storage capacity and communications volume, 1 000 bytes.

## L

**LAN** See *local area network*.

**LAN-free data movement**  
The movement of client data between a client system and a storage device on a storage area network (SAN), bypassing

the local area network. This process is also referred to as *LAN-free data transfer*.

**LAN-free data transfer**  
See *LAN-free data movement*.

**leader data**  
Bytes of data, from the beginning of a migrated file, that are stored in the file's corresponding stub file on the local file system. The amount of leader data that is stored in a stub file depends on the stub size that is specified.

**library**  
(1) A repository for demountable recorded media, such as magnetic disks and magnetic tapes.  
(2) A collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

**library client**  
A server that uses server-to-server communication to access a library that is managed by another storage management server. See also *library manager*.

**library manager**  
A server that controls device operations when multiple storage management servers share a storage device. See also *library client*.

**local** (1) Pertaining to a device, file, or system that is accessed directly from a user system, without the use of a communication line.  
(2) For HSM products, pertaining to the destination of migrated files that are being moved.

**local area network (LAN)**  
A network that connects several devices in a limited area (such as a single building or campus) and that can be connected to a larger network.

**local shadow volumes**  
Data that is stored on shadow volumes localized to a disk storage subsystem.

**LOFS** See *loopback virtual file system*.

**logical file**  
A file that is stored in one or more server storage pools, either by itself or as part of an aggregate. See also *aggregate* and *physical file*.

**logical occupancy**

The space that is used by logical files in a storage pool. This space does not include the unused space created when logical files are deleted from aggregate files, so it might be less than the physical occupancy.

**logical unit (LU)**

An access point through which a user or application program accesses the Systems Network Architecture (SNA) network to communicate with another user or application program.

**logical unit number (LUN)**

In the Small Computer System Interface (SCSI) standard, a unique identifier that is used to differentiate devices, each of which is a logical unit (LU).

**logical volume**

A portion of a physical volume that contains a file system.

**logical volume backup**

A backup of a file system or logical volume as a single object.

**Logical Volume Snapshot Agent (LVSA)**

Software that can act as the snapshot provider for creating a snapshot of a logical volume during an online image backup.

**loopback virtual file system (LOFS)**

A file system that is created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

**LU** See *logical unit*.

**LUN** See *logical unit number*.

**LVSA** See *Logical Volume Snapshot Agent*.

**M****macro file**

A file that contains one or more storage manager administrative commands, which can be run only from an administrative client using the MACRO command. Contrast with *Tivoli Storage Manager command script*.

**mailbox restore**

A function that restores Microsoft Exchange Server data (from IBM Data

Protection for Microsoft Exchange backups) at the mailbox level or mailbox-item level.

**managed object**

In Tivoli Storage Manager, a definition in the database of a managed server that was distributed to the managed server by a configuration manager. When a managed server subscribes to a profile, all objects that are associated with that profile become managed objects in the database of the managed server. In general, a managed object cannot be modified locally on the managed server. Objects can include policy, schedules, client option sets, server scripts, administrator registrations, server definitions, and server group definitions.

**managed server**

A Tivoli Storage Manager server that receives configuration information from a configuration manager using a subscription to one or more profiles. Configuration information can include definitions of objects such as policy and schedules. See also *configuration manager*, *subscription*, and *profile*.

**management class**

A policy object that users can bind to each file to specify how the server manages the file. The management class can contain a backup copy group, an archive copy group, and space management attributes. See also *copy group*, *space manager client*, *bind*, and *rebind*.

**maximum transmission unit**

The largest possible unit of data that can be sent on a given physical medium in a single frame. For example, the maximum transmission unit for Ethernet is 1500 bytes.

**MB** See *megabyte*.

**media server**

In a z/OS environment, a program that provides access to z/OS disk and tape storage for Tivoli Storage Manager servers that run on operating systems other than z/OS.

**megabyte (MB)**

(1) 1 048 576 bytes (2 to the 20th power) when used in this publication.

(2) For processor storage, real and virtual storage, and channel volume, 2 to the power of 20 or 1 048 576 bits. For disk storage capacity and communications volume, 1 000 000 bits.

**metadata**

Data that describes the characteristics of data; descriptive data.

**migrate**

To move data from one storage location to another. In Tivoli Storage Manager products, migrating can mean moving data from a client node to server storage, or moving data from one storage pool to the next storage pool defined in the server storage hierarchy. In both cases the movement is controlled by policy, such as thresholds that are set. See also *migration threshold*.

**migrated file**

A file that has been copied from a local file system to Tivoli Storage Manager storage. For HSM clients on UNIX or Linux systems, the file is replaced with a stub file on the local file system. On Windows systems, creation of the stub file is optional. See also *stub file* and *resident file*. For HSM clients on UNIX or Linux systems, contrast with *premigrated file*.

**migrate-on-close recall mode**

A mode that causes a migrated file to be recalled back to its originating file system temporarily. Contrast with *normal recall mode* and *read-without-recall recall mode*.

**migration job**

A specification of files to migrate, and actions to perform on the original files after migration. See also *job file*.

**migration threshold**

High and low capacities for storage pools or file systems, expressed as percentages, at which migration is set to start and stop.

**mirroring**

The process of writing the same data to multiple locations at the same time. Mirroring data protects against data loss within the recovery log.

**mode** A copy group attribute that specifies whether to back up a file that has not

been modified since the last time the file was backed up. See *modified mode* and *absolute mode*.

**modified mode**

In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup only if it has changed since the last backup. A file is considered a changed file if the date, size, owner, or permissions of the file have changed. See also *absolute mode*.

**mount limit**

The maximum number of volumes that can be simultaneously accessed from the same device class. The mount limit determines the maximum number of mount points. See also *mount point*.

**mount point**

On the Tivoli Storage Manager server, a logical drive through which volumes in a sequential access device class are accessed. For removable-media device types, such as tape, a mount point is a logical drive that is associated with a physical drive. For the file device type, a mount point is a logical drive that is associated with an I/O stream. The number of mount points for a device class is defined by the value of the mount limit attribute for that device class. See also *mount limit*.

**mount retention period**

The maximum number of minutes that the server retains a mounted sequential-access media volume that is not being used before it dismounts the sequential-access media volume.

**mount wait period**

The maximum number of minutes that the server waits for a sequential-access volume mount request to be satisfied before canceling the request.

**MTU** See *maximum transmission unit*.

**N****Nagle algorithm**

An algorithm that reduces congestion of TCP/IP networks by combining smaller packets and sending them together.

**named pipe**

A type of interprocess communication

that permits message data streams to pass between peer processes, such as between a client and a server.

**NAS** See *network-attached storage*.

**NAS node**

A client node that is a network-attached storage (NAS) file server. Data for the NAS node is transferred by a NAS file server that is controlled by the network data management protocol (NDMP). A NAS node is also called a NAS file server node.

**native file system**

A file system that is locally added to the file server and is not added for space management. The hierarchical storage manager (HSM) client does not provide space management services to the file system.

**native format**

A format of data that is written to a storage pool directly by the Tivoli Storage Manager server. Contrast with *non-native data format*.

**NDMP**

See *Network Data Management Protocol*.

**NetBIOS**

See *Network Basic Input/Output System*.

**network-attached storage (NAS) file server**

A dedicated storage device with an operating system that is optimized for file-serving functions. A NAS file server can have the characteristics of both a node and a data mover.

**Network Basic Input/Output System (NetBIOS)**

A standard interface to networks and personal computers that is used on local area networks to provide message, print-server, and file-server functions. Application programs that use NetBIOS do not have to handle the details of LAN data link control (DLC) protocols.

**Network Data Management Protocol (NDMP)**

A protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server, without installing vendor-acquired software on that file server.

**network data-transfer rate**

A rate that is calculated by dividing the

total number of bytes that are transferred by the data transfer time. For example, this rate can be the time that is spent transferring data over a network.

**node** A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

**node name**

A unique name that is used to identify a workstation, file server, or PC to the server.

**node privilege class**

A privilege class that gives an administrator the authority to remotely access backup-archive clients for a specific client node or for all clients in a policy domain. See also *privilege class*.

**non-native data format**

A format of data that is written to a storage pool that differs from the format that the server uses for operations.

**normal recall mode**

A mode that causes a migrated file to be copied back to its originating file system when it is accessed.

**O**

**offline volume backup**

A backup in which the volume is locked so that no other system applications can access it during the backup operation.

**online volume backup**

A backup in which the volume is available to other system applications during the backup operation.

**open registration**

A registration process in which users can register their workstations as client nodes with the server. Contrast with *closed registration*.

**operator privilege class**

A privilege class that gives an administrator the authority to disable or halt the server, enable the server, cancel server processes, and manage removable media. See also *privilege class*.

**options file**

A file that contains processing options. On Windows and NetWare systems, the file is

called `dsm.opt`. On AIX, UNIX, Linux, and Mac OS X systems, the file is called `dsm.sys`.

#### **originating file system**

The file system from which a file was migrated. When a file is recalled using normal or migrate-on-close recall mode, it is always returned to its originating file system.

#### **orphaned stub file**

A file for which no migrated file can be found on the Tivoli Storage Manager server that the client node is contacting for space management services. For example, a stub file can be orphaned when the client system-options file is modified to contact a server that is different than the one to which the file was migrated.

#### **out-of-space protection mode**

A mode that controls whether the program intercepts out-of-space conditions. See also *execution mode*.

### **P**

#### **pacing**

In SNA, a technique by which the receiving system controls the rate of transmission of the sending system to prevent overrun.

**packet** In data communication, a sequence of binary digits, including data and control signals, that is transmitted and switched as a composite whole.

**page** A defined unit of space on a storage medium or within a database volume.

#### **partial-file recall mode**

A recall mode that causes the hierarchical storage management (HSM) function to read just a portion of a migrated file from storage, as requested by the application accessing the file.

#### **password generation**

A process that creates and stores a new password in an encrypted password file when the old password expires. Automatic generation of a password prevents password prompting. Password generation can be set in the options file (`passwordaccess` option). See also *options file*.

**path** An object that defines a one-to-one

relationship between a source and a destination. Using the path, the source accesses the destination. Data can flow from the source to the destination, and back. An example of a source is a data mover (such as a network-attached storage [NAS] file server), and an example of a destination is a tape drive.

#### **pattern-matching character**

See *wildcard character*.

#### **physical file**

A file that is stored in one or more storage pools, consisting of either a single logical file, or a group of logical files that are packaged together as an aggregate. See also *aggregate* and *logical file*.

#### **physical occupancy**

The amount of space that is used by physical files in a storage pool. This space includes the unused space that is created when logical files are deleted from aggregates. See also *physical file*, *logical file*, and *logical occupancy*.

#### **plug-in**

A self-contained software component that modifies (adds, or changes) the function in a particular system. When a plug-in is added to a system, the foundation of the original system remains intact.

#### **policy domain**

A grouping of policy users with one or more policy sets, which manage data or storage resources for the users. The users are client nodes that are associated with the policy domain.

#### **policy privilege class**

A privilege class that gives an administrator the authority to manage policy objects, register client nodes, and schedule client operations for client nodes. Authority can be restricted to certain policy domains. See also *privilege class*.

#### **policy set**

A group of rules in a policy domain. The rules specify how data or storage resources are automatically managed for client nodes in the policy domain. Rules can be contained in management classes. See also *active policy set* and *management class*.

**premigrated file**

A file that has been copied to Tivoli Storage Manager storage, but has not been replaced with a stub file on the local file system. An identical copy of the file resides both on the local file system and in Tivoli Storage Manager storage. Premigrated files occur on UNIX and Linux file systems to which space management has been added. Contrast with *migrated file* and *resident file*.

**premigrated files database**

A database that contains information about each file that has been premigrated to Tivoli Storage Manager storage. The database is stored in a hidden directory named `.SpaceMan` in each file system to which space management has been added.

**premigration**

The process of copying files that are eligible for migration to Tivoli Storage Manager storage, but leaving the original file intact on the local file system.

**premigration percentage**

A space management setting that controls whether the next eligible candidates in a file system are premigrated following threshold or demand migration.

**primary storage pool**

A named set of volumes that the server uses to store backup versions of files, archive copies of files, and files migrated from client nodes. See also *destination* and *copy storage pool*.

**privilege class**

A level of authority that is granted to an administrator. The privilege class determines which administrative tasks the administrator can perform. See also *node privilege class*, *operator privilege class*, *policy privilege class*, *storage privilege class*, and *system privilege class*.

**profile**

A named group of configuration information that can be distributed from a configuration manager when a managed server subscribes. Configuration information can include registered administrator IDs, policies, client schedules, client option sets, administrative schedules, storage manager command scripts, server definitions, and

server group definitions. See also *configuration manager* and *managed server*.

**Q**

**quota** (1) For HSM on AIX, UNIX, or Linux systems, the limit (in megabytes) on the amount of data that can be migrated and premigrated from a file system to server storage.

(2) For HSM on Windows systems, a user-defined limit to the space that is occupied by recalled files.

**R****randomization**

The process of distributing schedule start times for different clients within a specified percentage of the schedule's startup window.

**raw logical volume**

A portion of a physical volume that is comprised of unallocated blocks and has no journaled file system (JFS) definition. A logical volume is read/write accessible only through low-level I/O functions.

**read-without-recall recall mode**

A mode that causes hierarchical storage management (HSM) to read a migrated file from storage without storing it back on the local file system. The last piece of information read from the file is stored in a buffer in memory on the local file system. Contrast with *normal recall mode* and *migrate-on-close recall mode*.

**rebind**

To associate all backed-up versions of a file with a new management class name. For example, a file that has an active backup version is rebound when a later version of the file is backed up with a different management class association. See also *bind*.

**recall** In Tivoli Storage Manager, to copy a migrated file from server storage back to its originating file system using the space management client. See also *transparent recall*, *selective recall*, and *recall mode*.

**recall mode**

A mode that is assigned to a migrated file with the `dsmatrr` command that determines how the file is processed when it is recalled. It determines whether the file is stored on the local file system,

is migrated back to Tivoli Storage Manager storage when it is closed, or is read from Tivoli Storage Manager storage without storing it on the local file system.

**receiver**

A server repository that contains a log of server and client messages as events. For example, a receiver can be a file exit, a user exit, or the Tivoli Storage Manager server console and activity log. See also *event*.

**reclamation**

The process of consolidating the remaining data from many sequential-access volumes onto fewer, new sequential-access volumes.

**reclamation threshold**

The percentage of space that a sequential-access media volume must have before the server can reclaim the volume. Space becomes reclaimable when files are expired or are deleted.

**reconciliation**

The process of synchronizing a file system with the Tivoli Storage Manager server, and then removing old and obsolete objects from the Tivoli Storage Manager server.

**recovery log**

A log of updates that are about to be written to the database. The log can be used to recover from system and media failures. The recovery log consists of the active log (including the log mirror) and archive logs.

**register**

To define a client node or administrator ID that can access the server.

**registry**

A repository that contains access and configuration information for users, systems, and software.

**remote**

- (1) Pertaining to a system, program, or device that is accessed through a communication line.
- (2) For HSM products, pertaining to the origin of migrated files that are being moved.

**resident file**

On a Windows system, a complete file on

a local file system that might also be a migrated file because a migrated copy can exist in Tivoli Storage Manager storage. On a UNIX or Linux system, a complete file on a local file system that has not been migrated or premigrated, or that has been recalled from Tivoli Storage Manager storage and modified. Contrast with *stub file* and *premigrated file*. See *migrated file*.

**restore**

To copy information from its backup location to the active storage location for use. For example, to copy information from server storage to a client workstation.

**retention**

The amount of time, in days, that inactive backed-up or archived files are kept in the storage pool before they are deleted. Copy group attributes and default retention grace periods for the domain define retention.

**retrieve**

To copy archived information from the storage pool to the workstation for use. The retrieve operation does not affect the archive version in the storage pool.

**roll back**

To remove changes that were made to database files since the last commit point.

**root user**

A system user who operates without restrictions. A root user has the special rights and privileges needed to perform administrative tasks.

**S**

**SAN** See *storage area network*.

**schedule**

A database record that describes client operations or administrative commands to be processed. See *administrative command schedule* and *client schedule*.

**scheduling mode**

The type of scheduling operation for the server and client node that supports two scheduling modes: client-polling and server-prompted.

**scratch volume**

A labeled volume that is either blank or contains no valid data, that is not defined, and that is available for use.

**script** A series of commands, combined in a file, that carry out a particular function when the file is run. Scripts are interpreted as they are run. Contrast with *Tivoli Storage Manager command script*.

**Secure Sockets Layer (SSL)**

A security protocol that provides communication privacy. With SSL, client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery.

**selective backup**

The process of backing up certain files or directories from a client domain. The files that are backed up are those that are not excluded in the include-exclude list. The files must meet the requirement for serialization in the backup copy group of the management class that is assigned to each file. Contrast with *incremental backup*.

**selective migration**

The process of copying user-selected files from a local file system to Tivoli Storage Manager storage and replacing the files with stub files on the local file system. Contrast with *threshold migration* and *demand migration*.

**selective recall**

The process of copying user-selected files from Tivoli Storage Manager storage to a local file system. Contrast with *transparent recall*.

**serialization**

The process of handling files that are modified during backup or archive processing. See *dynamic serialization*, *static serialization*, *shared static serialization*, and *shared dynamic serialization*.

**server** A software program or a computer that provides services to other software programs or other computers.

**server options file**

A file that contains settings that control various server operations. These settings affect such things as communications, devices, and performance.

**server-prompted scheduling mode**

A client/server communication technique where the server contacts the client node when tasks must be done. Contrast with *client-polling scheduling mode*.

**server storage**

The primary, copy, and active-data storage pools that are used by the server to store user files such as backup versions, archive copies, and files migrated from space manager client nodes (space-managed files). See also *active-data pool*, *primary storage pool*, *copy storage pool*, *storage pool volume*, and *volume*.

**session**

A logical or virtual connection between two stations, software programs, or devices on a network that allows the two elements to communicate and exchange data.

**session resource usage**

The amount of wait time, processor time, and space that is used or retrieved during a client session.

**shadow copy**

A snapshot of a volume. The snapshot can be taken while applications on the system continue to write data to the volumes.

**shadow volume**

The data stored from a snapshot of a volume. The snapshot can be taken while applications on the system continue to write data to the volumes.

**shared dynamic serialization**

A value for serialization that specifies that a file must not be backed up or archived if it is being modified during the operation. Tivoli Storage Manager retries the backup or archive operation a number of times; if the file is being modified during each attempt, Tivoli Storage Manager will back up or archive the file on its last try. See also *serialization*. Contrast with *dynamic serialization*, *shared static serialization*, and *static serialization*.

**shared library**

A library device that is used by multiple storage manager servers.

**shared static serialization**

A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. Tivoli Storage Manager attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See also

*serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *static serialization*.

**snapshot**

An image backup type that consists of a point-in-time view of a volume.

**space-managed file**

A file that is migrated from a client node by the space manager client. The space manager client recalls the file to the client node on demand.

**space management**

The process of keeping sufficient free storage space available on a local file system for new data by migrating files to server storage. Synonymous with *hierarchical storage management*.

**space manager client**

A program that runs on a UNIX or Linux system to manage free space on the local file system by migrating files to server storage. The program can recall the files either automatically or selectively. Also called *hierarchical storage management (HSM) client*.

**space monitor daemon**

A daemon that checks space usage on all file systems for which space management is active, and automatically starts threshold migration when space usage on a file system equals or exceeds its high threshold.

**sparse file**

A file that is created with a length greater than the data it contains, leaving empty spaces for the future addition of data.

**special file**

On AIX, UNIX, or Linux systems, a file that defines devices for the system, or temporary files that are created by processes. There are three basic types of special files: first-in, first-out (FIFO); block; and character.

**SSL** See *Secure Sockets Layer*.

**stabilized file space**

A file space that exists on the server but not on the client.

**stanza** A group of lines in a file that together have a common function or define a part of the system. Each stanza is identified by a name that occurs in the first line of the

stanza. Depending on the type of file, a stanza is ended by the next occurrence of a stanza name in the file, or by an explicit end-of-stanza marker. A stanza can also be ended by the end of the file.

**startup window**

A time period during which a schedule must be initiated.

**static serialization**

A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. If the file is in use during the first attempt, the storage manager cannot back up or archive the file. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *shared static serialization*.

**storage agent**

A program that enables the backup and restoration of client data directly to and from storage attached to a storage area network (SAN).

**storage area network (SAN)**

A dedicated storage network that is tailored to a specific environment, combining servers, systems, storage products, networking products, software, and services.

**storage hierarchy**

(1) A logical order of primary storage pools, as defined by an administrator. The order is typically based on the speed and capacity of the devices that the storage pools use. The storage hierarchy is defined by identifying the next storage pool in a storage pool definition. See also *storage pool*.

(2) An arrangement of storage devices with different speeds and capacities. The levels of the storage hierarchy include: main storage, such as memory and direct-access storage device (DASD) cache; primary storage (DASD containing user-accessible data); migration level 1 (DASD containing data in a space-saving format); and migration level 2 (tape cartridges containing data in a space-saving format).

**storage pool**

A named set of storage volumes that are the destination that is used to store client

data. A storage pool contains backup versions, archive copies, and files that are migrated from space manager client nodes. A primary storage pool is backed up to a copy storage pool. See also *primary storage pool*, *copy storage pool*, and *active-data pool*.

**storage pool volume**

A volume that has been assigned to a storage pool. See also *volume*, *active-data pool*, *copy storage pool*, and *primary storage pool*.

**storage privilege class**

A privilege class that gives an administrator the authority to control how storage resources for the server are allocated and used, such as monitoring the database, the recovery log, and server storage. See also *privilege class*.

**stub**

A shortcut on the Windows file system that is generated by the hierarchical storage management (HSM) client for a migrated file that allows transparent user access. A stub is the sparse file representation of a migrated file, with a reparse point attached.

**stub file**

A file that replaces the original file on a local file system when the file is migrated to storage. A stub file contains the information that is necessary to recall a migrated file from Tivoli Storage Manager storage. It also contains additional information that can be used to eliminate the need to recall a migrated file.

**stub file size**

The size of a file that replaces the original file on a local file system when the file is migrated to Tivoli Storage Manager storage. The size that is specified for stub files determines how much leader data can be stored in the stub file. The default for stub file size is the block size defined for a file system minus 1 byte.

**subscription**

In a Tivoli environment, the process of identifying the subscribers that the profiles are distributed to. For Tivoli Storage Manager, a subscription is the process by which a managed server receives configuration information associated with a particular profile on a

configuration manager. See also *managed server*, *configuration manager*, and *profile*.

**system privilege class**

A privilege class that gives an administrator the authority to issue all server commands. See also *privilege class*.

**Systems Network Architecture (SNA)**

The description of the logical structure, formats, protocols, and operational sequences for transmitting information through and controlling the configuration and operation of networks.

**T**

**tape library**

A set of equipment and facilities that support an installation's tape environment. The tape library can include tape storage racks, mechanisms for automatic tape mounting, a set of tape drives, and a set of related tape volumes mounted on those drives.

**tape volume prefix**

The high-level-qualifier of the file name or the data set name in the standard tape label.

**target node**

A client node for which other client nodes (called agent nodes) have been granted proxy authority. The proxy authority allows the agent nodes to perform operations such as backup and restore on behalf of the target node, which owns the data.

**TCA** See *trusted communications agent*.

**TCP/IP**

See *Transmission Control Protocol/Internet Protocol*.

**threshold migration**

The process of moving files from a local file system to Tivoli Storage Manager storage based on the high and low thresholds that are defined for the file system. Contrast with *demand migration*, *selective migration*, and *migration job*.

**throughput**

In storage management, the total bytes in the workload, excluding overhead, that are backed up or restored, divided by elapsed time.

**timeout**

A time interval that is allotted for an event to occur or complete before operation is interrupted.

**timestamp control mode**

A mode that determines whether commands preserve the access time for a file or set it to the current time.

**Tivoli Storage Manager command script**

A sequence of Tivoli Storage Manager administrative commands that are stored in the database of the Tivoli Storage Manager server. The script can run from any interface to the server. The script can include substitution for command parameters and conditional logic.

**tombstone object**

A small subset of attributes of a deleted object. The tombstone object is retained for a specified period, and at the end of the specified period, the tombstone object is permanently deleted.

**Transmission Control Protocol/Internet Protocol (TCP/IP)**

An industry-standard, nonproprietary set of communication protocols that provides reliable end-to-end connections between applications over interconnected networks of different types.

**transparent recall**

The process that is used to automatically recall a file to a workstation or file server when the file is accessed. See also *recall mode*. Contrast with *selective recall*.

**trusted communications agent (TCA)**

A program that handles the sign-on password protocol when clients use password generation.

**U**

**UCS-2** A 2-byte (16-bit) encoding scheme based on ISO/IEC specification 10646-1. UCS-2 defines three levels of implementation: Level 1-No combining of encoded elements allowed; Level 2-Combining of encoded elements is allowed only for Thai, Indic, Hebrew, and Arabic; Level 3-Any combination of encoded elements are allowed.

**UNC** See *Universal Naming Convention name*.

**Unicode**

A character encoding standard that

supports the interchange, processing, and display of text that is written in the common languages around the world, plus some classical and historical texts. The Unicode standard has a 16-bit character set defined by ISO 10646.

**Unicode-enabled file space**

Unicode file space names provide support for multilingual workstations without regard for the current locale.

**Unicode transformation format 8**

Unicode Transformation Format (UTF), 8-bit encoding form, which is designed for ease of use with existing ASCII-based systems. The CCSID value for data in UTF-8 format is 1208.

**Universal Naming Convention (UNC) name**

The server name and network name combined. These names together identify the resource on the domain.

**Universally Unique Identifier (UUID)**

The 128-bit numeric identifier that is used to ensure that two components do not have the same identifier.

**UTF-8** See *Unicode transformation format 8*.

**UUID** See *Universally Unique Identifier*.

**V****validate**

To check a policy set for conditions that can cause problems if that policy set becomes the active policy set. For example, the validation process checks whether the policy set contains a default management class.

**version**

A backup copy of a file stored in server storage. The most recent backup copy of a file is the active version. Earlier copies of the same file are inactive versions. The number of versions retained by the server is determined by the copy group attributes in the management class.

**virtual file space**

A representation of a directory on a network-attached storage (NAS) file system as a path to that directory.

**virtual volume**

An archive file on a target server that represents a sequential media volume to a source server.

**volume**

A discrete unit of storage on disk, tape or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control. See also *scratch volume*, and *storage pool volume*.

**volume history file**

A file that contains information about volumes that have been used by the server for database backups and for export of administrator, node, policy, or server data. The file also has information about sequential-access storage pool volumes that have been added, reused, or deleted. The information is a copy of volume information that is recorded in the server database.

**Volume Shadow Copy Service**

A set of Microsoft application-programming interfaces (APIs) that you can use to create shadow copy backups of volumes, exact copies of files, including all open files, and so on.

**VSS** See *Volume Shadow Copy Service*.

**VSS Backup**

A backup operation that uses Microsoft Volume Shadow Copy Service (VSS) technology. The backup operation produces an online snapshot (point-in-time consistent copy) of Microsoft Exchange data. This copy can be stored on local shadow volumes or on Tivoli Storage Manager server storage.

**VSS Fast Restore**

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on local shadow volumes.

**VSS Instant Restore**

A volume-level hardware-assisted Microsoft Volume Shadow Copy Service (VSS) function where target volumes that contain the snapshot are copied back to the original source volumes.

**VSS offloaded backup**

A backup operation that uses a Microsoft Volume Shadow Copy Service (VSS) hardware provider (installed on an alternate system) to move IBM Data

Protection for Microsoft Exchange data to the Tivoli Storage Manager server. This type of backup operation shifts the backup load from the production system to another system.

**VSS Restore**

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on Tivoli Storage Manager server storage to their original location.

**W****wildcard character**

A special character such as an asterisk (\*) or a question mark (?) that can be used to represent one or more characters. Any character or set of characters can replace the wildcard character.

**workload partition (WPAR)**

A partition within a single operating system instance.

**workstation**

A terminal or personal computer at which a user can run applications and that is usually connected to a mainframe or a network.

**worldwide name**

A 64-bit, unsigned name identifier that is unique.



---

# Index

## Numerics

3490 device type (CARTRIDGE) 224  
3494SHARED option 1447  
349X  
    DEFINE LIBRARY command 249  
    UPDATE LIBRARY command 1287  
3590 device type 149, 213  
3592 device class 1255  
3592 device type 153, 218, 1197  
3592 drives and media  
    logical block protection 153, 1197  
4MM device type 159  
8MM device type 163

## A

ACCEPT DATE command 28  
ACCESS parameter 394, 1424  
accessibility features 1593  
account record, setting 1055  
ACSACCESSID option 1448  
ACSLOCKDRIVE option 1449  
ACSL 1442  
ACSL library 252, 1289  
ACSQUICKINIT option 1450  
ACSTIMEOUTX option 1451  
action, restore 123, 305, 1347  
ACTIVATE POLICYSET command 29  
activating policy sets 29  
active log 1564  
active log mirror 1564  
active-data pool  
    defining a new active-data pool 378  
    identifying with a QUERY command 917  
    restoring data  
        restoring storage pools 1025  
        restoring volumes 1029  
    specifying as backups for primary storage pools 350, 367,  
        1387, 1402  
    updating an existing active-data pool 1412  
active-data pools  
    list 724  
ACTIVELOGDIRECTORY option 1452  
ACTIVELOGSIZE option 1453  
activity log  
    querying 661  
    setting retention period 1056  
    setting size limit 1056  
ADMINCOMMTIMEOUT option 1454  
ADMINIDLETIMEOUT option 1455  
administrative client  
    batch mode 1, 3  
    console mode 1, 2  
    continuation characters 12  
    interactive mode 1, 3  
    mount mode 1, 3  
    options 5  
    privilege classes 19  
    starting 2  
    stopping 2  
administrative client (*continued*)  
    using 1  
administrative client options  
    alwaysprompt 5  
    commadelimited 5  
    consolemode 5  
    dataonly 6  
    displaymode 6  
    id 6  
    itemcommit 6  
    mountmode 6  
    newlineafterprompt 6  
    noconfirm 6  
    outfile 2, 3, 6  
    password 6  
    quiet 6  
    serveraddress 6  
    tabdelimited 7  
    tcpport 7  
    tcpserveraddress 7  
administrative command  
    background processing 15  
    components 8  
    entry rules 8  
    foreground processing 15  
    help 563  
    issued by all administrators 25  
    issued with operator privilege 24  
    issued with policy privilege 22  
    issued with storage privilege 23  
    issued with system privilege 19  
    naming conventions 13  
    parameters, entering 11  
    privilege class for 19  
    scheduling 314  
    using 27  
    using wildcard characters 14  
administrative commands  
    LOCK ADMIN 599  
    QUERY ADMIN 667  
    REGISTER ADMIN 965  
    REMOVE ADMIN 989  
    RENAME ADMIN 996  
    UNLOCK ADMIN 1159  
    UPDATE ADMIN 1170  
administrative schedule  
    associating with profile 288  
    copying 105  
    defining 314  
    deleting 444  
    distributing 294  
    updating 1342  
administrative session, canceling 74  
administrator  
    associating with profile 288  
    distributing 294  
    exporting 490  
    importing 570  
ADMINONCLIENTPORT option 1456  
    administrative sessions using TCPPOINT 1456  
    using with TCPADMINPORT 1541

- affect performance, raw partitions 341
- ALIASHALT server option 1457
- ALLOWREORGINDEX option 1458
- ALLOWREORGTABLE option 1459
- ANR9999D error message 1080
- ANS1082E message 125, 307, 1349
- ANS1102E message 125, 307, 1349
- ARCHDELETE parameter 975, 1313
- ARCHFAILOVERLOGDIRECTORY option 1460
- archive
  - defining 141
  - description of 141
  - frequency 142
  - retention grace period 239, 1275
- archive copy group
  - attributes of 141
  - defining 141
  - serialization 141
- archive file deletion by a client 975, 1313
- archive log failover directory 1564
- ARCHLOGDIRECTORY option 1461
- ASSIGN DEFMGMTCLASS command 31
- assigning
  - client node to domain 971, 1309
  - default management class 31
- ASSISTVCRRECOVERY option 1462
- association of client with schedule
  - defining 115
  - deleting 401
  - querying 679
- association, object with profile
  - defining 288
  - deleting 437
- AUDIT LDAPDIRECTORY command 34
- AUDIT LIBRARY command 36
- AUDIT LICENSES command 39
- audit occupancy
  - displaying storage pool utilization 681
  - querying 681
- AUDIT VOLUME command 40
- auditing
  - library 36
  - storage pool volume 40
- AUDITSTORAGE option 1463
- authentication
  - password 1073
  - setting open or closed 1073
- authority
  - client access 556
  - client owner 556
  - for command use 19
  - granting 556
  - QUERYAUTH option 1513
  - REQSYSAUTHOUTFILE option 1521
  - revoking 1033
- auto-update 1135
- AUTOMIGNONUSE parameter 272, 1307

## B

- BACKDELETE parameter 975, 1313
- background process
  - canceling 69
  - query process 69, 838
- backup
  - backup file deletion 975, 1313
  - client 136

- backup (*continued*)
  - database 47
  - device information 51
  - primary storage pool 58
  - scheduling for client 301
  - volume history 62
- backup copy group
  - attributes of 136
  - defining 136
  - serialization 136
  - TOCDestination parameter 139
  - updating 1183
- BACKUP DB command 47
- BACKUP DEVCONFIG command 51
- BACKUP NODE command 53
- backup set
  - defining 117
  - deleting 403
  - displaying contents of 689
  - generating 545
  - node groups
    - defining a group 274
    - defining a member 275
    - deleting a group 432
    - deleting a member 433
    - querying a group 823
    - updating a group 1325
  - querying 684
  - updating retention period 1174
- BACKUP STGPOOL command 58
- BACKUP VOLHISTORY command 62
- batch mode
  - restrictions 3
  - starting and ending 3
  - using 3
- BEGIN EVENTLOGGING command 64
- BEGINDATE parameter 662, 749
- BEGINREORGTIME option 1516, 1517
- BEGINTIME parameter 662, 749, 806
- books
  - See* publications

## C

- CACHE parameter
  - DEFINE STGPOOL, primary random access 346
  - UPDATE STGPOOL, primary random access 1384
- cancel
  - administrative session 74
  - client node session 74
  - expiration process 67
  - mount requests 72
  - process 69
  - restartable restore session 73
- CANCEL EXPIRATION command 67
- CANCEL PROCESS 16
- CANCEL PROCESS command 69
- CANCEL REPLICATION command 71
- CANCEL REQUEST command 72
- CANCEL RESTORE command 73
- CANCEL SESSION command 74
- canceling commands 16
- CAP parameter
  - MOVE DRMEDIA 83, 615, 631
- CARTRIDGE device type 224
- CD support 204

- CENTERA device type
  - concurrent read/write access to FILE volumes
    - increasing number of client mount points 971, 1309
- changing date and time on the server 28
- characters available for specifying passwords 13
- CHECKIN LIBVOLUME command 75
- checking in
  - library volume 75
  - with labeling 75
- checking out a library volume 83
- CHECKLABEL parameter
  - AUDIT LIBRARY 37
  - CHECKIN LIBVOLUME 79
  - CHECKOUT LIBVOLUME 86
- CHECKOUT LIBVOLUME command 83
- CHECKTAPEPOS option 1464
- CLASSES parameter
  - GRANT AUTHORITY 556
  - REVOKE AUTHORITY 1033
- CLEAN DRIVE command 89
- cleaning a drive, frequency 243
- client access authority 556
- client node session
  - canceling 74
  - scheduling 300
- client option 127
- client owner authority 556
- client-server options, setting 1443
- client, backing up subfiles for 1149
- CLIENTDEDUPTXNLIMIT option 1466
- COLLOCATE parameter
  - DEFINE STGPOOL
    - active-data pool 379
    - copy sequential access 370
    - primary sequential access 360
  - UPDATE STGPOOL
    - active-data pool 1413
    - copy sequential access 1406
    - primary sequential access 1396
- collocation
  - group
    - defining 131
    - deleting 410
    - querying 694
    - updating 1181
  - group member
    - defining 133
    - deleting 411
  - specifying for a storage pool
    - active-data pool 379, 1413
    - copy sequential access 370, 1406
    - primary sequential access 360, 1396
- command
  - DELETE ALERTTRIGGER 400
  - QUERY ALERTSTATUS 674
  - UPDATE ALERTTRIGGER 1165
- command line interface
  - administrative client 1
  - server console 1
  - using 1
- command line, Administration Center 7
- command output, formatting 4
- command routing
  - defining server groups for 17
  - defining servers for 17
- command scripts
  - copying 109
- command scripts (*continued*)
  - defining 323
  - deleting 447
  - querying 888
  - renaming 1003
  - running 1039
  - updating 1365
- commands 466
  - ACCEPT DATE 28
  - ACTIVATE POLICYSET 29
  - ASSIGN DEFMGMTCLASS 31
  - AUDIT LDAPDIRECTORY 34
  - AUDIT LIBRARY 36
  - AUDIT LICENSES 39
  - AUDIT VOLUME 40
  - BACKUP DB 47
  - BACKUP DEVCONFIG 51
  - BACKUP NODE 53
  - BACKUP STGPOOL 58
  - BACKUP VOLHISTORY 62
  - BEGIN EVENTLOGGING 64
  - CANCEL EXPIRATION 67
  - CANCEL EXPORT 68
  - CANCEL PROCESS 69
  - CANCEL REPLICATION 71
  - CANCEL REQUEST 72
  - CANCEL RESTORE 73
  - CANCEL SESSION 74
  - CHECKIN LIBVOLUME 75
  - CHECKOUT LIBVOLUME 83
  - CLEAN DRIVE 89
  - COMMIT 90
  - COPY ACTIVATEDATA 92
  - COPY CLOPTSET 96
  - COPY DOMAIN 97
  - COPY MGMTCLASS 99
  - COPY POLICYSET 101
  - COPY PROFILE 103
  - COPY SCHEDULE 105
  - COPY SCRIPT 109
  - COPY SERVERGROUP 110
  - DEFINE ALERTTRIGGER 112
  - DEFINE ASSOCIATION 115
  - DEFINE BACKUPSET 117
  - DEFINE CLIENTACTION 121
  - DEFINE CLIENTOPT 127
  - DEFINE CLOPTSET 130
  - DEFINE COLLOCGROUP 131
  - DEFINE COLLOCMEMBER 133
  - DEFINE COPYGROUP 135
  - DEFINE DATAMOVER 145
  - DEFINE DEVCLASS 148
  - DEFINE DOMAIN 239
  - DEFINE DRIVE 241
  - DEFINE EVENTSERVER 245
  - DEFINE GRPMEMBER 246
  - DEFINE LIBRARY 247, 249, 252, 254, 256, 257, 259, 262, 263, 266
  - DEFINE MACHINE 267
  - DEFINE MACHNODEASSOCIATION 269
  - DEFINE MGMTCLASS 271
  - DEFINE NODEGROUP 274
  - DEFINE NODEGROUPMEMBER 275
  - DEFINE PATH 276
  - DEFINE PATH - Destination is a drive 277
  - DEFINE PATH - Destination is a library 282
  - DEFINE PATH - Destination is a zosmedia library 285

commands (continued)

DEFINE POLICYSET 286  
 DEFINE PROFASSOCIATION 288  
 DEFINE PROFILE 294  
 DEFINE RECMEDMACHASSOCIATION 296  
 DEFINE RECOVERYMEDIA 298  
 DEFINE SCHEDULE 300  
 DEFINE SCRIPT 323  
 DEFINE SERVER 326  
 DEFINE SERVERGROUP 333  
 DEFINE SPACETRIGGER 334  
 DEFINE STATUSTHRESHOLD 337  
 DEFINE STGPOOL 341  
 DEFINE SUBSCRIPTION 387  
 DEFINE VIRTUALFSMAPPING 389  
 DEFINE VOLUME 391  
 DELETE ASSOCIATION 401  
 DELETE BACKUPSET 403  
 DELETE CLIENTOPT 408  
 DELETE CLOPTSET 409  
 DELETE COLLOGGROUP 410  
 DELETE COLLOCMEMBER 411  
 DELETE COPYGROUP 413  
 DELETE DATAMOVER 415  
 DELETE DEVCLASS 416  
 DELETE DOMAIN 417  
 DELETE DRIVE 418  
 DELETE EVENT 419  
 DELETE EVENTSERVER 421  
 DELETE FILESPACE 422  
 DELETE GRPMEMBER 426  
 DELETE KEYRING 427  
 DELETE LIBRARY 428  
 DELETE MACHINE 429  
 DELETE MACHNODEASSOCIATION 430  
 DELETE MGMTCLASS 431  
 DELETE NODEGROUP 432  
 DELETE NODEGROUPMEMBER 433  
 DELETE PATH 434  
 DELETE POLICYSET 436  
 DELETE PROFASSOCIATION 437  
 DELETE PROFILE 440  
 DELETE RECMEDMACHASSOCIATION 442  
 DELETE RECOVERYMEDIA 443  
 DELETE SCHEDULE 444  
 DELETE SCRIPT 447  
 DELETE SERVER 448  
 DELETE SERVERGROUP 449  
 DELETE SPACETRIGGER 450  
 DELETE STATUSTHRESHOLD 451  
 DELETE STGPOOL 453  
 DELETE SUBSCRIBER 454  
 DELETE SUBSCRIPTION 455  
 DELETE VIRTUALFSMAPPING 456  
 DELETE VOLHISTORY 457  
 DELETE VOLUME 462  
 DISABLE REPLICATION 470  
 DISABLE SESSIONS 471  
 DISMOUNT VOLUME 474  
 DISPLAY OBJNAME 475  
 DSMADMC 1  
 DSMSERV 1560  
 DSMSERV DISPLAY DBSPACE 1563  
 DSMSERV DISPLAY LOG 1564  
 DSMSERV FORMAT 1566  
 DSMSERV REMOVEDB 1573  
 DSMSERV RESTORE DB 1575, 1578

commands (continued)

DSMSERV RUNFILE 1560  
 ENABLE EVENTS 477  
 ENABLE REPLICATION 480  
 ENABLE SESSIONS 481  
 END EVENTLOGGING 483  
 EXPIRE INVENTORY 485  
 EXPORT ADMIN 490  
 EXPORT NODE 497  
     directly to another server 509  
 EXPORT POLICY 518  
 EXPORT SERVER 525  
     directly to another server 535  
     to sequential media 527  
 EXTEND DB 543  
 GENERATE 544  
 GENERATE BACKUPSET 545  
 GENERATE BACKUPSETTOC 553  
 GRANT AUTHORITY 556  
 GRANT PROXYNODE 560  
 HALT 561  
 HELP 563  
 IDENTIFY DUPLICATES 565  
 IMPORT ADMIN 570  
 IMPORT NODE 574  
 IMPORT POLICY 581  
 IMPORT SERVER 584  
 INSERT MACHINE 589  
 ISSUE MESSAGE 590  
 LABEL LIBVOLUME 592  
 LOCK NODE 601  
 LOCK PROFILE 603  
 MACRO 605  
 MIGRATE STGPOOL 607  
 MOVE DATA 611  
 MOVE DRMEDIA 615  
 MOVE GRPMEMBER 630  
 MOVE MEDIA 631  
 MOVE NODEDATA 639  
 NOTIFY SUBSCRIBERS 648  
 PERFORM LIBACTION 649  
 PING SERVER 652  
 PREPARE 653  
 QUERY ACTLOG 661  
 QUERY ALERTTRIGGER 672  
 QUERY ASSOCIATION 679  
 QUERY AUDITOCCUPANCY 681  
 QUERY BACKUPSET 684  
 QUERY BACKUPSETCONTENTS 689  
 QUERY CLOPTSET 692  
 QUERY COLLOGGROUP 694  
 QUERY CONTENT 697  
 QUERY COPYGROUP 705  
 QUERY DATAMOVER 710  
 QUERY DB 713  
 QUERY DBSPACE 716  
 QUERY DEVCLASS 718  
 QUERY DIRSPACE 723  
 QUERY DOMAIN 724  
 QUERY DRIVE 727  
 QUERY DRMEDIA 731  
 QUERY DRMSTATUS 740  
 QUERY ENABLED 743  
 QUERY EVENT 745  
 QUERY EVENTRULES 757  
 QUERY EVENTSERVER 759  
 QUERY EXPORT 760

commands (continued)

QUERY FILESPACE 767  
 QUERY LIBRARY 774  
 QUERY LIBVOLUME 777  
 QUERY LICENSE 780  
 QUERY LOG 782  
 QUERY MACHINE 784  
 QUERY MEDIA 787  
 QUERY MGMTCLASS 793  
 QUERY MONITORSETTINGS 796  
 QUERY MONITORSTATUS 799  
 QUERY MOUNT 803  
 QUERY NASBACKUP 805  
 QUERY NODE 809  
 QUERY NODEDATA 820  
 QUERY OCCUPANCY 825  
 QUERY OPTION 829  
 QUERY PATH 831  
 QUERY POLICYSET 835  
 QUERY PROCESS 838  
 QUERY PROFILE 841  
 QUERY PROXYNODE 844  
 QUERY PVESTIMATE 845  
 QUERY RECOVERYMEDIA 849  
 QUERY REPLICATION 852  
 QUERY REPLNODE 862  
 QUERY REPLRULE 866  
 QUERY REQUEST 868  
 QUERY RESTORE 869  
 QUERY RPFCONTENT 872  
 QUERY RPFFILE 874  
 QUERY SAN 877  
 QUERY SCHEDULE 880  
 QUERY SCRIPT 888  
 QUERY SERVER 891  
 QUERY SERVERGROUP 895  
 QUERY SESSION 897  
 QUERY SHREDSTATUS 901  
 QUERY SPACETRIGGER 903  
 QUERY SSLKEYRINGPW 905  
 QUERY STATUS 906  
 QUERY STATUSTHRESHOLD 914  
 QUERY STGPOOL 917  
 QUERY SUBSCRIBER 928  
 QUERY SUBSCRIPTION 930  
 QUERY SYSTEM 932  
 QUERY TAPEALERTMSG 934  
 QUERY TOC 935  
 QUERY VIRTUALFSMAPPING 938  
 QUERY VOLHISTORY 940  
 QUERY VOLUME 948  
 QUIT 956  
 RECLAIM STGPOOL 957  
 RECONCILE VOLUMES 961  
 REGISTER LICENSE 969  
 REGISTER NODE 971  
 REMOVE NODE 991  
 REMOVE REPLNODE 993  
 RENAME FILESPACE 998  
 RENAME NODE 1001  
 RENAME SCRIPT 1003  
 RENAME SERVERGROUP 1004  
 RENAME STGPOOL 1005  
 REPLICATE NODE 1006  
 RESET PASSEXP 1014  
 RESTORE NODE 1019  
 RESTORE STGPOOL 1024

commands (continued)

RESTORE VOLUME 1028  
 RESUME EXPORT 1016  
 REVOKE AUTHORITY 1033  
 REVOKE PROXYNODE 1037  
 ROLLBACK 1038  
 RUN 1039  
 SELECT 1042  
 SET ACCOUNTING 1055  
 SET ACTLOGRETENTION 1056  
 SET ALERTACTIVEDURATION 1058  
 SET ALERTCLOSEDDURATION 1059  
 SET ALERTEMAIL 1060  
 SET ALERTEMAILFROMADDR 1062  
 SET ALERTEMAILSMTPHOST 1063  
 SET ALERTEMAILSMTPPORT 1064  
 SET ALERTINACTIVEDURATION 1066  
 SET ALERTMONITOR 1067  
 SET ALERTSUMMARYTOADMINS 1065  
 SET ALERTUPDATEINTERVAL 1068  
 SET ARCHIVERETENTIONPROTECTION 1069  
 SET ARREPLRULEDEFAULT 1071  
 SET AUTHENTICATION 1073  
 SET BKREPLRULEDEFAULT 1074  
 SET CLIENTACTDURATION 1077  
 SET CONFIGMANAGER 1078  
 SET CONFIGREFRESH 1079  
 SET CONTEXTMESSAGING 1080  
 SET CPUINFOREFRESH 1081  
 SET CROSSDEFINE 1082  
 SET DBRECOVERY 1083  
 SET DBREPORTMODE 1085  
 SET DEDUPVERIFICATIONLEVEL 1086  
 SET DEFAULTAUTHENTICATION 1088  
 SET DRMACTIVEDATASTGPOOL 1090  
 SET DRMCHECKLABEL 1091  
 SET DRMCMDFILENAME 1092  
 SET DRMCOPYSTGPOOL 1093  
 SET DRMCOURIERNAME 1094  
 SET DRMDBBACKUPEXPIREDAYS 1095  
 SET DRMFILEPROCESS 1097  
 SET DRMINSTRPREFIX 1098  
 SET DRMNOTMOUNTABLENAME 1100  
 SET DRMPPLANPREFIX 1101  
 SET DRMPPLANVPOSTFIX 1103  
 SET DRMPRIMSTGPOOL 1104  
 SET DRMRPFEXPIREDAYS 1105  
 SET DRMVaultNAME 1107  
 SET EVENTRETENTION 1108  
 SET INVALIDPWLIMIT 1109  
 SET LDAPPASSWORD 1111  
 SET LDAPUSER 1112  
 SET LICENSEAUDITPERIOD 1114  
 SET MAXCMDRETRIES 1115  
 SET MAXSCHEDSESSIONS 1116  
 SET MINPWLENGTH 1118  
 SET MONITOREDSEVERGROUP 1120  
 SET MONITORINGADMIN 1119  
 SET PASSEXP 1121  
 SET QUERYSCHEDPERIOD 1123  
 SET RANDOMIZE 1124  
 SET REGISTRATION 1126  
 SET REPLRETENTION 1128  
 SET REPLSERVER 1130  
 SET RETRYPERIOD 1132  
 SET SCHEDMODES 1133  
 SET SERVERHLADDRESS 1135

commands (*continued*)

- SET SERVERLLADDRESS 1136
- SET SERVERNAME 1137
- SET SERVERPASSWORD 1138
- SET SPREPLRULEDEFAULT 1139
- SET SSLKEYRINGPW 1141
- SET STATUSATRISKINTERVAL 1142
- SET STATUSMONITOR 1144
- SET STATUSREFRESHINTERVAL 1146
- SET STATUSSKIPASFAILURE 1147
- SET SUMMARYRETENTION 1150
- SET TOCLOADRETENTION 1152
- SETOPT 1153
- SHRED DATA 1155
- SUSPEND EXPORT 1157
- UNLOCK NODE 1161
- UNLOCK PROFILE 1163
- UPDATE ALERTSTATUS 1168
- UPDATE BACKUPSET 1174
- UPDATE CLIENTOPT 1179
- UPDATE CLOPTSET 1180
- UPDATE COLLOGGROUP 1181
- UPDATE COPYGROUP 1182
- UPDATE DATAMOVER 1190
- UPDATE DEVCLASS 1192
- UPDATE DOMAIN 1275
- UPDATE DRIVE 1277
- UPDATE FILESPACE 1281
- Update LIBRARY 1298
- UPDATE LIBRARY 1286, 1287, 1289, 1291, 1292, 1293, 1295, 1299
- UPDATE LIBVOLUME 1302
- UPDATE MACHINE 1304
- UPDATE MGMTCLASS 1306
- UPDATE NODE 1309
- UPDATE NODEGROUP 1325
- UPDATE PATH 1326
- UPDATE PATH - Destination is a drive 1327
- UPDATE PATH - Destination is a library 1331
- UPDATE PATH - Destination is a zosmedia library 1334
- UPDATE POLICYSET 1335
- UPDATE PROFILE 1337
- UPDATE RECOVERYMEDIA 1338
- UPDATE REPLRULE 1340
- UPDATE SCHEDULE 1342
- UPDATE SCRIPT 1365
- UPDATE SERVER 1368
- UPDATE SERVERGROUP 1372
- UPDATE SPACETRIGGER 1373
- UPDATE STATUSTHRESHOLD 1375
- UPDATE STGPOOL 1379
- UPDATE VIRTUALFSMAPPING 1419
- UPDATE VOLHISTORY 1421
- UPDATE VOLUME 1423
- VALIDATE LANFREE 1428
- VALIDATE POLICYSET 1430
- VALIDATE REPLICATION 1432
- VARY 1436

commands in a macro

- committing 90
- rolling back 1038

commands, administrative 8

commands, canceling 16

COMMIT command 90

committing commands in a macro 90

COMMMETHOD option 1468

COMMTIMEOUT option 1470

communications, server-to-server 326

- COMMMETHOD option 1468
- shared memory between server and client 1468

complex password

- AUDIT LDAPDIRECTORY command 34
- SET DEFAULTAUTHENTICATION 1088
- SET LDAPUSR command 1111

configuration manager 387

configuration profile 294

console mode

- ending 2
- restrictions 2
- using 2

continuation characters

- for a list of values 12
- for a quoted list of values 12
- in output file 5
- using the maximum number of 12

conventions

- typographic xiv

COPIED parameter, QUERY CONTENT 699

COPY ACTIVATEDATA command 92

COPY CLOPTSET command 96

COPY DOMAIN command 97

copy group

- defining archive 141
- defining backup 136
- deleting 413
- description of 135
- querying 705
- restriction 135
- updating archive 1187
- updating backup 1183

COPY MGMTCLASS command 99

COPY POLICYSET command 101

COPY PROFILE command 103

COPY SCHEDULE command 105

COPY SCRIPT 109

COPY SERVERGROUP command 110

COPYCONTINUE parameter

- DEFINE STGPOOL, primary random access 349
- DEFINE STGPOOL, primary sequential access 366
- UPDATE STGPOOL, primary random access 1387
- UPDATE STGPOOL, primary sequential access 1401

copying

- management class 99
- policy domain 97
- policy set 101
- profile 103
- schedule 105
- script 109
- server group 110

COPYSTGPOOLS parameter

- DEFINE STGPOOL, primary random access 348
- DEFINE STGPOOL, primary sequential access 365
- UPDATE STGPOOL, primary random access 1386
- UPDATE STGPOOL, primary sequential access 1401

COUNT parameter, QUERY CONTENT 698

CRCDATA parameter

- DEFINE STGPOOL, primary random access 345
- DEFINE STGPOOL, primary sequential access 355, 376
- UPDATE STGPOOL, primary random access 1383
- UPDATE STGPOOL, primary sequential access 1393, 1410

creating

- backup set 545
- client files onto a set of media 545

customer support  
contact 1589

## D

DAMAGED parameter, QUERY CONTENT 699

data

moving 611  
removing expired 485

data deduplication 367, 368, 376, 385, 1403, 1410, 1411, 1417, 1418

data deduplication verification level, setting 1086

data mover

defining 145  
deleting 415  
querying 710  
updating 1190

data protection

logical block protection 153, 175, 186, 1197, 1227

data protection using WORM FILE volumes and SnapLock

when defining active-data pools 382  
when defining copy storage pools 373  
when defining sequential access storage pools 359

data shredding, storage pools

backing up 60  
defining 351  
moving data 612  
updating 1388

database

backup 47  
extending 543  
installing 1566  
querying 713  
removing 1573  
restoring 1575, 1578  
setting options for 1443  
transfer by data mover 145

database recoverable space 713

database recovery

back up volume history 62  
delete volume history 457  
query volume history 940

database storage space 716, 1563

DATAFORMAT parameter, define primary sequential access

storage pool 363, 375

dataonly 6

DATES parameter

IMPORT NODE 577  
IMPORT SERVER 586

DAYOFWEEK parameter

UPDATE SCHEDULE  
client 1353

DBDIAGLOGSIZE option 1471

DBDIAGPATHFSTHRESHOLD option 1472

DBMEMPERCENT option 1473

DBMTCPPORT option 1474

debug ANR9999D message 1080

deduplicating data 565

deduplication 367, 368, 376, 385, 1403, 1410, 1411, 1417, 1418

DEDUPTIER2FILESIZE option 1476

DEDUPTIER3FILESIZE option 1477

default management class, assigning 31

Define a 349X library 249

Define a file library 256

Define a manual library 257

Define a SCSI library 259

Define a shared library 262

Define a VTL library 263

Define a ZOSMEDIA library 266

DEFINE ALERTTRIGGER command 112

Define an ACSLS library 252

DEFINE ASSOCIATION command 115

DEFINE BACKUPSET command 117

DEFINE CLIENTACTION command 121

DEFINE CLIENTOPT command 127

DEFINE CLOPTSET command 130

DEFINE COLLOGGROUP command 131

DEFINE COLLOCMEMBER command 133

DEFINE COPYGROUP command 135

DEFINE DATAMOVER command 145

DEFINE DEVCLASS command 148

DEFINE DOMAIN command 239

DEFINE DRIVE command 241

DEFINE EVENTSERVER command 245

DEFINE GRPMEMBER command 246

DEFINE LIBRARY command 247, 249, 252, 254, 256, 257, 259, 262, 263, 266

DEFINE MACHINE command 267

DEFINE MACHNODEASSOCIATION command 269

DEFINE MGMTCLASS command 271

DEFINE PATH command 276, 277, 282, 285

DEFINE POLICYSET command 286

DEFINE PROFASSOCIATION command 288

DEFINE PROFILE command 294

DEFINE RECMEDMACHASSOCIATION command 296

DEFINE RECOVERYMEDIA command 298

DEFINE SCHEDULE command 300

DEFINE SCRIPT command 323

DEFINE SERVER command 326

DEFINE SERVERGROUP command 333

DEFINE SPACETRIGGER command 334

DEFINE STATUSTHRESHOLD command 337

DEFINE STGPOOL command 341

DEFINE SUBSCRIPTION command 387

DEFINE VIRTUALFSMAPPING command 389

DEFINE VOLUME command 391

defining

association 115

backup set 117

client with schedule 301

collocation group member 133

collocation groups 131

configuration manager 387

copy group 135

device class 148

domain 239

drive 241

event server 245

file library 256

group member 246

library 247, 249, 252, 254, 257, 259, 262

management class 271

node group 274

node group member 275

object with profile 288

path for NAS data mover 276

path for NDMP (NAS) connection 276

policy set 286

profile 294

profile association 288

schedule 300

script 323

server 326

server group 333

- defining *(continued)*
  - space trigger 334
  - storage pool 341
  - storage pool volume 391
  - subscription 387
  - VTL library 263
  - ZOSMEDIA library 266
- definitions 1599
- DELETE ALERTTRIGGER command 400
- DELETE ASSOCIATION command 401
- DELETE BACKUPSET 403
- DELETE CLIENTOPT command 408
- DELETE CLOPTSET command 409
- DELETE COLLOGROUP command 410
- DELETE COLLOCMEMBER command 411
- DELETE COPYGROUP command 413
- DELETE DATAMOVER 415
- DELETE DEVCLASS command 416
- DELETE DOMAIN command 417
- DELETE DRIVE command 418
- DELETE EVENT command 419
- DELETE EVENTSERVER command 421
- DELETE FILESPACE command 422
- DELETE KEYRING command 427
- DELETE LIBRARY command 428
- DELETE MACHINE command 429
- DELETE MACHNODEASSOCIATION command 430
- DELETE MGMTCLASS command 431
- DELETE PATH 434
- DELETE POLICYSET command 436
- DELETE PROFASSOCIATION command 437
- DELETE PROFILE command 440
- DELETE RECMEDMACHASSOCIATION command 442, 1092
- DELETE RECOVERYMEDIA command 443
- DELETE SCHEDULE command 444
- DELETE SCRIPT command 447
- DELETE SERVER command 448
- DELETE SERVERGROUP command 449
- DELETE SPACETRIGGER command 450
- DELETE STATUSTHRESHOLD command 451
- DELETE STGPOOL command 453
- DELETE SUBSCRIBER command 454
- DELETE SUBSCRIPTION command 455
- DELETE VIRTUALFSMAPPING command 456
- DELETE VOLHISTORY command 457
- DELETE VOLUME command 462
- deleting
  - archive file deletion by a client, allowing 975, 1313
  - backup file deletion by a client, allowing 975
  - backup file deletion by a client, ing 1313
  - backup set 403
  - collocation group 410
  - collocation group member 411
  - copy group 413
  - device class 416
  - domain 417
  - drive 418
  - event record 419
  - event server 421
  - expired data 485
  - file space 422
  - group member 426
  - library 428
  - management class 431
  - node group 432
  - node group member 433
- deleting *(continued)*
  - policy set 436
  - profile 440
  - profile association 437
  - schedule 444
  - script 447
  - server 448
  - server group 449
  - space trigger 450
  - storage pool 453
  - storage pool volume 462
  - subscriber 454
  - subscription 455
  - volume history 457
- deployment
  - automatic 1135
- determining retention periods for policy domains 724
- DEVCONFIG option 1478
- device class
  - 3590 1193
  - 3592 153, 218
  - defining 148
  - deleting 416
  - NAS 193, 1233
  - OPTICAL 1235
  - QIC 199, 1238
  - querying 718
  - updating 1192
  - VOLSAFE 79, 208, 1246
  - WORM 1235
- device configuration file 51
- device names
  - for devices connected to NAS file server 276
- device type
  - 3590 149, 213, 1193
  - 3592 153, 218, 1197
  - 4MM 159, 1202
  - 8MM 163, 1206
  - CARTRIDGE 224, 1260
  - DLT 169, 1211
  - ECARTRIDGE 175, 229, 1216, 1265
  - FILE 183, 236, 723, 1223, 1272
  - LTO 186, 187, 1227
  - NAS 193, 1233
  - OPTICAL 196, 1235
  - QIC 199, 1238
  - REMOVABLEFILE 204, 1242
  - SERVER 206, 1244
  - VOLSAFE 208, 1246
  - WORM 196, 1235
  - WORM12 196, 1235
  - WORM14 196, 1235
- directories, specifying multiple for FILE volumes 184, 1224
- disability 1593
- DISABLE EVENTS 466
- DISABLE EVENTS command 466
- DISABLE REPLICATION command 470
- DISABLE SESSIONS command 471
- DISABLESCHEDULES option 1479
- disabling
  - context messaging for ANR9999D messages 1080
  - server sessions, inbound and outbound 471
- DISCARDDATA parameter 463
- disk space, migrating data to create 346, 608, 1383
- disk volume performance 341
- disk volumes 341

- disk-only backup
  - defining FILE device classes 184
  - defining volumes 391
  - updating FILE device classes 1224
- DISMOUNT VOLUME command 474
- dismounting volumes 474
- DISPLAY OBJNAME command 475
- displaying
  - full object name 475
  - output 6
  - product version information 6
  - scheduled and completed events 745
- DISPLAYLFINFO option 1480
- DLT
  - device type 169
  - WORM media 170
- DNSLOOKUP option 1481
- documentation
  - See* publications
- domain
  - associating with profile 288
  - copy 97
  - defining 239
  - deleting 417
  - distributing 294
  - querying 724
  - updating 1275
- drive
  - cleaning 89, 243
  - defining 241, 727
  - deleting 418
  - querying 727
  - updating 1277
- drive encryption
  - 3592 device class 158, 1201
  - ECARTRIDGE device class 181, 1221
  - LTO device class 191, 1231
- DRIVEACQUIRERETRY option 1482
- DRIVEENCRYPTION parameter
  - 3592 device class 158, 1201
  - ECARTRIDGE device class 181, 1221
  - LTO device class 191, 1231
- DSMADMC command 1
- DSMADMC options 5
- DSMSERV command 1560
- DSMSERV DISPLAY DBSPACE command 1563
- DSMSERV DISPLAY LOG command 1564
- DSMSERV FORMAT command 1566
- DSMSERV REMOVEDB command 1573
- DSMSERV RESTORE DB command 1575, 1578
- DSMSERV RUNFILE command 1560
- duplicate-identification processes 368, 376, 385, 1403, 1411, 1418
- DURATION parameter
  - DEFINE SCHEDULE, administrative command 318
  - DEFINE SCHEDULE, client 310
  - SHRED DATA, administrative command 1155
  - UPDATE SCHEDULE, administrative command 1360
  - UPDATE SCHEDULE, client 1351

**E**

- ECARTRIDGE device type 175
- education
  - see* Tivoli technical training 1587
- ENABLE EVENTS command 477
- ENABLE REPLICATION command 480
- ENABLE SESSIONS command 481
- ENABLENASDEDUP option 1483
- enabling
  - context messaging for ANR9999D 1080
  - server sessions, inbound and outbound 481
- encryption
  - application
    - 3592 158, 1201
    - ECARTRIDGE 181, 1221
    - LTO 191, 1231
  - drive
    - 3592 158, 1201
    - ECARTRIDGE 181, 1221
    - LTO 191, 1231
  - library 158, 1201
  - system 158, 1201
- END EVENTLOGGING command 483
- ENDDATE parameter 663, 749, 806
- ending
  - batch mode 3
  - console mode 2
  - interactive mode 2
  - mount mode 3
- ENDTIME parameter 663, 749, 807
- error
  - ANR9999D message 1080
  - file specification 125, 307, 1349
  - In Progress status for QUERY EVENT 751
- ESTCAPACITY parameter
- DEFINE DEVCLASS
  - 3590 151
  - 3592 156
  - 4MM 161
  - 8MM 166
  - DLT 172
  - ECARTRIDGE 180
  - LTO 190
  - NAS 194
  - Optical/WORM 197
  - QIC 201
  - VOLSAFE 209
- UPDATE DEVCLASS
  - 3590 1194
  - 3592 1200
  - 4MM 1203
  - 8MM 1209
  - DLT 1214
  - ECARTRIDGE 1220
  - NAS 1234
  - Optical/WORM 1236
  - QIC 1240
  - VOLSAFE 1247
- event logging
  - send events to receivers 64
  - set logging off by receivers 483
- event record
  - description of 419, 1108
  - managing 1108
  - querying 745
  - removing from the database 419
  - setting retention period 1108
- event record retention period
  - managing 1108
  - querying 906
- event, querying 745
- EVENTSERVER option 1484
- EXCEPTIONSONLY parameter 749

- EXPINTERVAL option 1485
- EXPIRATION parameter
  - DEFINE SCHEDULE, administrative command 314
  - DEFINE SCHEDULE, client 301
  - UPDATE SCHEDULE, administrative command 1356
  - UPDATE SCHEDULE, client 1343
- EXPIRE INVENTORY command 485
- expiring inventory 485
- export
  - administrator 490
  - node 497
  - policy information 518
  - server 525
- EXPORT ADMIN
  - information directly to another server 495
  - to sequential media 492
- EXPORT ADMIN command 490
- EXPORT NODE
  - export node definitions to sequential media 500
- EXPORT NODE command 497
- EXPORT POLICY command 518
- EXPORT SERVER command 525
- EXPQUIET option 1486
- EXTEND DB command 543
- extending
  - database 543

## F

- failover directory for the archive log 1564
- FFDCLOGNAME option 1487
- FFDCMAXLOGSIZE option 1488
- FILE device type
  - concurrent read/write access to FILE volumes
    - increasing number of client mount points 971, 1309
    - visible as output in QUERY MOUNT 803
    - visible as output in QUERY SESSION 897
  - creating and formatting private volumes 391
  - defining the device class 183, 236
  - updating the device class 1223, 1272
- file organization 236
- file server, NAS
  - names for connected devices 276
  - node 145, 976
- file space
  - automatic rename for Unicode support 1314
  - deleting 422
  - FSID 999
  - hexadecimal representation 999
  - querying 767
  - renaming 998
  - specification errors, restrictions 123, 305, 1347
- FILEDATA parameter
  - EXPORT NODE 502, 511
  - IMPORT NODE 576
- FILEEXIT option 1489
- FILENAMES parameter
  - BACKUP DEVCONFIG 51
  - BACKUP VOLHISTORY 63
- files
  - collocating by client 360, 371, 379
  - moving 611
  - removing expired 485
- FILESPEC parameter
  - EXPORT NODE 501, 510
  - IMPORT NODE 576
- FILETEXTEXIT option 1490

- firewall, opening server ports
  - administrative sessions 1541
  - client sessions 1544
  - TCP/IP ports 1541, 1544
- fixes, obtaining 1589
- FORCE parameter, CHECKOUT LIBVOLUME 87
- FORMAT parameter
  - DEFINE DEVCLASS
    - 3590 149
    - 3592 155
    - 4MM 160
    - 8MM 164
    - DLT 170
    - ECARTRIDGE 178
    - LTO 188
    - Optical/WORM 197
    - QIC 200
    - VOLSAFE 209
  - UPDATE DEVCLASS
    - 3590 1193
    - 3592 1198
    - 4MM 1202
    - 8MM 1206
    - DLT 1212
    - ECARTRIDGE 1218
    - LTO 1228
    - QIC 1235, 1239
    - VOLSAFE 1246
- FORMAT, DSMSEV 1566
- FREQUENCY parameter
  - DEFINE COPYGROUP, archive 142
  - DEFINE COPYGROUP, backup 137
  - UPDATE COPYGROUP, archive 1187
  - UPDATE COPYGROUP, backup 1184
- FSUSEDTHRESHOLD option 1491
- full object name
  - displaying 475

## G

- GENERATE BACKUPSET command 545
- GENERATE BACKUPSETTOC 553
- globally unique identifier (GUID) 809
- glossary 1599
- GRANT AUTHORITY command 556
- GRANT PROXYNODE command 560
- granting
  - authority 556
  - types of privilege classes 556
- group, collocation by
  - defining a group 131
  - defining a group member 133
  - deleting a group 410
  - deleting a group member 411
  - querying a group 694
  - updating a group 1181
- group, node
  - defining 274
  - defining a member 275
  - deleting a group 432
  - deleting a member 433
  - querying a group 823
  - updating a group 1325
- GUID 809

## H

- HALT command 561
- halting the server 561
- HELP command 563
- help for administrative commands 563
- hierarchical storage management
  - DEFINE MGMTCLASS 271
  - description of 271
  - UPDATE MGMTCLASS 1306
- hierarchy of storage pools
  - considerations for backup 58
  - defining 341
- HIGHMIG parameter
  - DEFINE STGPOOL
    - primary random access 346
    - primary sequential access 356
  - UPDATE STGPOOL
    - primary random access 1383
    - primary sequential access 1394
- history information
  - backing up 62
  - deleting 457
  - querying 940
- HSM (hierarchical storage management) 272, 1306

## I

- IBM Publications Center xi, xiv
- IBM Support Assistant 1588
- IDENTIFY DUPLICATES command 565
- IDLETIMEOUT option 1492
- IMPORT ADMIN command 570
- IMPORT NODE command 574
- IMPORT POLICY command 581
- IMPORT SERVER command 584
- importing
  - administrator 570
  - node 574
  - policy information 581
  - server 584
- In Progress status for QUERY EVENT command 751
- incremental backup of database 47
- inheritance, copy storage pool lists and
  - defining random access storage pools 349
  - defining sequential access storage pools 365
  - updating random access storage pools 1386
  - updating sequential access storage pools 1401
- INSERT MACHINE command 589
- installing
  - database 1566
  - recovery log 1566
- interactive mode
  - continuation characters 12
  - ending 2
  - quitting 956
  - restrictions 3
  - using 3
- Internet, searching for problem resolution 1588, 1589
- inventory expiring 485
- ISSUE MESSAGE command 590
- issuing commands for a one-time action 121
- ITEMCOMMIT option 6

## K

- keepalive, TCP
  - enabling 1505
  - specifying connection idle time (AIX, Linux, and Windows) 1506
- keyboard 1593
- knowledge bases, searching 1587

## L

- label
  - for REMOVABLEFILE device type 241
  - library volume 592
- LAN
  - DEFINE LIBRARY command 254, 1291
- LAN-free data movement
  - validating 1428
- language limitations 1
- LANGUAGE option 1493
- LBPROTECT parameter 153, 175, 186, 1197, 1227
- LDAP-authenticated password xviii
  - LOCK NODE command 601
  - SET LDAPUSER command 1112
- LDAPCACHEDURATION server option 1495
- LDAPURL server option 1496
- library
  - auditing 36
  - defining 247, 249, 252, 254, 257, 259, 262, 266
  - defining file 256
  - deleting 428
  - querying 774
  - shared 262, 1298
  - updating 1286, 1287, 1289, 1291, 1293, 1295, 1298
  - updating file 1292
- library volume
  - checking in 75
  - checking out 83
  - labeling 592
  - querying 777
  - updating 1302
- license
  - auditing 39
  - querying 780
  - registering 969
  - setting audit period 1114
- local area network (LAN)
  - DEFINE LIBRARY command 254, 1291
- LOCK ADMIN command 599
- LOCK NODE command 601
- LOCK PROFILE command 603
- locking
  - administrator 599
  - node 601
  - profile 603
- logic flow statements in scripts 1583
- logical block protection 153, 175, 186, 1197, 1227
- logs, recovery 1564
- LOWMIG parameter
  - DEFINE STGPOOL
    - primary random access 346
    - primary sequential access 357
  - UPDATE STGPOOL
    - primary random access 1384
    - primary sequential access 1394

LTO  
 device class  
   defining 186, 187  
   updating 1227  
 WORM 188  
 LTO Ultrium drives and media  
   logical block protection 186, 1227

## M

MACRO command 605  
 macros  
   continuation characters 12  
   rolling back 1038  
   using 605  
 maintenance 1135  
 managed server 387  
 managed system for storage area network (SAN)  
   DEFINE LIBRARY command 247  
   UPDATE LIBRARY command 1286  
 management class  
   copying 99  
   defining the default 31  
   deleting 431  
   querying 793  
   updating 1306  
 manuals  
   *See* publications  
 MAXCAPACITY parameter  
   DEFINE DEVCLASS  
     FILE 184  
     REMOVABLEFILE 204  
     SERVER 206  
   UPDATE DEVCLASS  
     FILE 1223  
     REMOVABLEFILE 1242  
     SERVER 1244  
 maximum retries, setting 1115  
 MAXPROCESS parameter  
   BACKUP STGPOOL 59  
   UPDATE STGPOOL 1026  
 MAXSCRATCH parameter  
   DEFINE STGPOOL  
     active-data pool 383  
     copy sequential access 374  
     primary sequential access 361  
   UPDATE STGPOOL  
     active-data pool 1416  
     copy sequential access 1409  
     primary sequential access 1397  
 MAXSESSIONS option 1498  
 MAXSIZE parameter  
   DEFINE STGPOOL  
     primary random access 345  
     primary sequential access 355  
   UPDATE STGPOOL  
     primary random access 1382  
     primary sequential access 1392  
 media mount requests 72  
 media support, CD 204  
 member, server group 246  
 MERGEFILESPPACES parameter  
   EXPORT NODE 516  
   EXPORT SERVER 541  
   IMPORT NODE 579  
   IMPORT SERVER 584  
 message diagnosis, ANR9999D 1080

MESSAGEFORMAT option 1499  
 MIGDESTINATION parameter  
   DEFINE MGMTCLASS 272  
   UPDATE MGMTCLASS 1307  
 MIGPROCESS parameter  
   DEFINE STGPOOL, primary random access 346  
   UPDATE STGPOOL, primary random access 1384  
 MIGRATE STGPOOL command 607  
 migration  
   files from client node 272  
   specifying multiple concurrent processes  
     using the DEFINE STGPOOL command 362  
     using the UPDATE STGPOOL command 1399  
   starting manually 607  
   storage pool, random access  
     high migration threshold 346, 1383  
     low migration threshold 346, 608, 1384  
   storage pool, sequential access  
     high migration threshold 356, 1394  
     low migration threshold 357, 608, 1394  
 MIGREQUIRESBKUP parameter  
   DEFINE MGMTCLASS 272  
   UPDATE MGMTCLASS 1307  
 MIRRORLOGDIRECTORY option 1500  
 MODE parameter  
   DEFINE COPYGROUP, archive 143  
   DEFINE COPYGROUP, backup 138  
   UPDATE COPYGROUP, archive 1188  
   UPDATE COPYGROUP, backup 1185  
 monitoring  
   Tivoli Storage Manager activities 3  
 monitoring Tivoli Storage Manager activities 2  
 mount mode  
   ending 3  
   using 3  
 mount request 72  
   canceling 72  
   querying 868  
 mounted sequential access volumes  
   dismounting 474  
   querying 803  
 MOUNTLIMIT parameter  
   DEFINE DEVCLASS  
     3590 151  
     3592 157  
     4MM 162  
     8MM 163  
     DLT 173  
     ECARTRIDGE 181  
     LTO 191  
     NAS 194  
     Optical/WORM 198  
     QIC 202  
     REMOVABLEFILE 205  
     SERVER 207  
     VOLSAFE 210  
   UPDATE DEVCLASS  
     3590 1193  
     3592 1200  
     4MM 1204  
     8MM 1206  
     DLT 1214  
     ECARTRIDGE 1220  
     FILE 1223  
     LTO 1231  
     Optical/WORM 1237  
     QIC 1241

MOUNTLIMIT parameter *(continued)*  
   UPDATE DEVCLASS *(continued)*  
     REMOVABLEFILE 1243  
     VOLSAFE 1248  
 MOUNTRETENTION parameter  
   DEFINE DEVCLASS  
     3590 151  
     3592 157  
     4MM 161  
     8MM 167  
     DLT 173  
     ECARTRIDGE 180  
     LTO 190  
     NAS 193  
     Optical/WORM 198  
     QIC 202  
     REMOVABLEFILE 205  
     SERVER 207  
     VOLSAFE 210  
   UPDATE DEVCLASS  
     3590 1195  
     3592 1200  
     4MM 1204  
     8MM 1209  
     DLT 1214  
     ECARTRIDGE 1220  
     LTO 1230  
     NAS 1233  
     Optical/WORM 1236  
     QIC 1241  
     REMOVABLEFILE 1242  
     SERVER 1245  
     VOLSAFE 1247  
 MOUNTWAIT parameter  
   DEFINE DEVCLASS  
     3590 151  
     3592 157  
     4MM 161  
     8MM 167  
     DLT 173  
     ECARTRIDGE 181  
     LTO 191  
     NAS 194  
     Optical/WORM 198  
     QIC 202  
     REMOVABLEFILE 205  
     VOLSAFE 210  
   UPDATE DEVCLASS  
     3590 1195  
     3592 1200  
     4MM 1204  
     8MM 1209  
     DLT 1214  
     ECARTRIDGE 1220  
     LTO 1231  
     NAS 1233  
     Optical/WORM 1236  
     QIC 1241  
     REMOVABLEFILE 1242  
     SERVER 1245  
     VOLSAFE 1248  
 MOVE DATA command 611  
 MOVE DRMEDIA command 615  
 MOVE GRPMEMBER command 630  
 MOVE MEDIA command 631  
 MOVE NODEDATA command 639  
 MOVEBATCHSIZE option 1501

MOVESIZETHRESH option 1502  
 moving  
   data in storage pool 611  
   files 611  
   group member 630  
 MSGINTERVAL option 1503  
 MSGNO parameter, QUERY ACTLOG 664

## N

name  
   devices 276  
   for device connected to NAS file server 276  
 naming conventions  
   introduction 13  
   naming volumes 13  
   passwords 13  
   restrictions 13  
 NAS file server  
   data mover, defining 145  
   name for connected device 276  
   node name 145  
   path, defining 276  
 NDMP operations for NAS file servers  
   prevent closing of inactive connections  
     enabling TCP keepalive 1505  
     specifying connection idle time (AIX, Linux, and Windows) 1506  
 NDMPCONTROLPORT option 1504  
 NDMPENABLEKEEPALIVE server option 1505  
 NDMPKEEPIDLEMINUTES server option 1506  
 NDMPPORTRANGE option 1507  
 NDMPREFDATAINTERFACE option 1508  
 NetApp DataONTAP, for managing FILE volumes  
   when defining active-data pools 382  
   when defining copy storage pools 373  
   when defining sequential access storage pools 359  
 network-attached storage (NAS) file server  
   data mover, defining 145  
   device class 193, 1233  
   path, defining 276  
 NEXTSTGPOOL parameter  
   DEFINE STGPOOL  
     primary random access 346  
     primary sequential access 356  
   UPDATE STGPOOL  
     primary random access 1383  
     primary sequential access 356, 1393  
 node  
   associating with a schedule 115  
   exporting 497  
   finding file for 825  
   importing 574  
   locking 601  
   name for data mover 145  
   NAS file server 145, 976  
   node group  
     defining a group 274  
     defining a member 275  
     deleting a group 432  
     deleting a member 433  
     querying a group 823  
     updating a group 1325  
   querying 809  
   registering 971  
   removing 991  
   renaming 1001

- node (*continued*)
  - unlocking 1161
  - updating 1309
- node replication
  - configuration, measuring the effectiveness of 862
  - configuration, validating 1432
  - disabling 470
  - disabling and enabling
    - all client nodes, replication for 471, 481
    - data types in file spaces 1281
    - individual nodes, replication for 971, 1309
    - replication rules 1340
  - enabling 480
  - file spaces
    - data types, disabling and enabling 1281
    - purging data in 1281
    - querying replication results 767, 852
    - rules, changing 1281
  - nodes
    - adding for replication 971, 1309
    - individual, disabling and enabling replication for 971, 1309
    - removing from replication 993
    - system-wide, disabling and enabling replication for 471
    - systemwide, disabling and enabling replication for 481
  - previewing results 1006
  - process information
    - displaying 767, 852
    - record retention 906, 1128
  - records, setting retention for 1128
  - removing client node 993
  - replicating
    - canceling processes 71
    - data by command 1006
    - data by type, priority, and file space 1006
    - throughput, managing 1006
  - results, previewing 1006
  - rules
    - client node 809, 971, 1309
    - disabling and enabling 1340
    - file space 767, 1281
    - querying 767, 809, 866
    - server 1071, 1074, 1139
  - Secure Sockets Layer (SSL) 326, 891, 1368
  - synchronizing exported-imported data 971, 1309
  - target replication server
    - changing 1130
    - setting 1130
  - validating a configuration 1432
- NOPREEMPT option 1509
- NORETRIEVEDATE option 1510
- normal recovery log mode 47
- NOTIFY SUBSCRIBERS command
  - example 648
  - related commands 648
- NUMOPENVOLSALLOWED option 1511

## O

- OBJECTS parameter
  - DEFINE SCHEDULE, client 310
  - UPDATE SCHEDULE, client 1351
- occupancy
  - displaying file space information 825
  - querying 825

- offsite volumes, specifying the number to be reclaimed
  - using the DEFINE STGPOOL command 373, 382
  - using the UPDATE STGPOOL command 1408, 1415
- OPTICAL device type 196, 1235
- option set 127
- OPTIONS parameter
  - DEFINE SCHEDULE, client 310
  - UPDATE SCHEDULE, client 1351
- options, client 127
- options, server
  - importance for recovery 47
  - querying 829
- output
  - displaying 6
- output headers
  - displaying 6
- overflow location for logs 1564
- OWNER parameter, DELETE FILESPACE 424

## P

- parameters, descriptions in 15
- parameters, entering 11
- Passport Advantage 1590
- password
  - case-sensitivity 13
  - characters allowed for entering 13
  - entered with the REGISTER ADMIN command 965
  - entered with the REGISTER NODE command 971
  - entered with the UPDATE ADMIN command 1170
  - setting expiration interval 1121
  - setting the maximum length 13
- path
  - defining 276
  - defining to a drive 277
  - defining to a library 282
  - defining to a zosmedia library 285
  - deleting 434
  - querying 831
  - updating 1326
  - updating to a drive 1327
  - updating to a library 1331
  - updating to a zosmedia library 1334
- pattern-matching expression, creating 14
- PERFORM LIBACTION command 649
- performance
  - improving communications with client using shared memory 1468, 1531
  - limiting offsite volumes to be reclaimed
    - using the DEFINE STGPOOL command 373, 382
    - using the UPDATE STGPOOL command 1408, 1415
  - specifying multiple concurrent migration processes
    - using the DEFINE STGPOOL command 362
    - using the UPDATE STGPOOL command 1399
  - specifying multiple concurrent reclamation processes
    - using the DEFINE STGPOOL command 372, 381
    - using the UPDATE STGPOOL command 1408, 1415
  - storage requirements, reducing by deduplication 367, 368, 376, 385, 1403, 1410, 1411, 1417, 1418
- PERUNITS parameter
  - DEFINE SCHEDULE, administrative command 321
  - UPDATE SCHEDULE, administrative command 1363
  - UPDATE SCHEDULE, client 1353
- policy domain
  - associating with profile 288
  - copy 97
  - defining 239

- policy domain (*continued*)
  - deleting 417
  - distributing 294
  - querying 724
  - updating 1275
- policy privilege class
  - privilege class
    - policy 22
    - restricted 22
    - unrestricted 22
- policy set
  - activating 29
  - copying 101
  - defining 286
  - deleting 436
  - querying 835
  - updating 1335
  - validating 1430
- polling
  - information about 1123
  - setting frequency 1123
- ports for firewall
  - administrative sessions 1541
  - client sessions 1544
  - TCP/IP ports 1541, 1544
- preempting of operations 1509
- PREFIX parameter
  - DEFINE DEVCLASS
    - 3590 151
    - 3592 156
    - 4MM 161
    - 8MM 167
    - DLT 173
    - ECARTRIDGE 180
    - LTO 190
    - NAS 194
    - QIC 202
    - SERVER 207
    - VOLSAFE 210
  - UPDATE DEVCLASS
    - 3590 1195
    - 3592 1200
    - 4MM 1204
    - 8MM 1209
    - DLT 1214
    - ECARTRIDGE 1220
    - LTO 1230
    - NAS 1233, 1234
    - QIC 1240
    - SERVER 1244
    - VOLSAFE 1247
- PREPARE command 653
- PREVIEW parameter
  - EXPORT NODE 502, 577
- PREVIEWIMPORT parameter
  - EXPORT ADMIN 495
  - EXPORT NODE 516
  - EXPORT POLICY 523
  - EXPORT SERVER 541
- printing redirected output 4
- privilege class
  - administrator issued 25
  - none required 25
  - operator 24
  - policy 23
  - restricted 23
  - storage 23, 24
- privilege class (*continued*)
  - system 19
  - unrestricted 23
- problem determination
  - describing problem for IBM Software Support 1590
  - determining business impact for IBM Software Support 1590
  - submitting a problem to IBM Software 1591
- process
  - canceling 69
  - querying 838
- product version information
  - displaying 6
- profile 294
- profile association 288
- protecting WORM volumes with SnapLock
  - when defining active-data pools 382
  - when defining copy storage pools 373
  - when defining sequential access storage pools 359
- publications
  - download xi
  - order xi
  - related hardware xiv
  - search xi
  - Tivoli Storage FlashCopy Manager xiv
  - Tivoli Storage Manager xii

**Q**

- QIC device type 199, 1238
- QUERY ACTLOG command 661
- QUERY ADMIN command 667
- QUERY ALERTSTATUS command 674
- QUERY ALERTTRIGGER command 672
- QUERY ASSOCIATION command 679
- QUERY AUDITOCUPANCY command 681
- QUERY BACKUPSET 684
- QUERY BACKUPSETCONTENTS 689
- QUERY CLOPTSET command 692
- QUERY COLLOGGROUP command 694
- QUERY CONTENT command 697
- QUERY COPYGROUP command 705
- QUERY DATAMOVER command 710
- QUERY DB command 713
- QUERY DBSPACE command 716
- QUERY DEVCLASS command 718
- QUERY DOMAIN command 724
- QUERY DRIVE command 727
- QUERY DRMEDIA command 731
- QUERY DRMSTATUS command 740
- QUERY ENABLED command 743
- QUERY EVENT command 745
- QUERY EVENTRULES command 757, 759
- QUERY FILESPACE command 767
- QUERY LIBRARY command 774
- QUERY LIBVOLUME command 777
- QUERY LICENSE command 780
- QUERY LOG command 782
- QUERY MACHINE command 784
- QUERY MEDIA command 787
- QUERY MGMTCLASS command 793
- QUERY MONITORSETTINGS command 796
- QUERY MONITORSTATUS command 799
- QUERY MOUNT command 803
- QUERY NASBACKUP command 805
- QUERY NODE command 809
- QUERY NODEDATA command 820

- QUERY OCCUPANCY command 825
- QUERY OPTION command 829
- QUERY PATH command 831
- QUERY POLICYSET command 835
- QUERY PROCESS command 838
- QUERY PROFILE command 841
- QUERY PROXYNODE command 844
- QUERY PVUESTIMATE command 845
- QUERY RECOVERYMEDIA command 849
- QUERY REPLICATION command 852
- QUERY REPLNODE command 862
- QUERY REPLRULE command 866
- QUERY REQUEST command 868
- QUERY RESTORE command 869
- QUERY RPFCONTENT command 872
- QUERY RPFIL command 874
- QUERY SAN command 877
- QUERY SCHEDULE command 880
- QUERY SCRIPT command 888
- QUERY SERVER command 891
- QUERY SERVERGROUP command 895
- QUERY SESSION command 897
- QUERY SHREDSTATUS command 901
- QUERY SPACETRIGGER command 903
- QUERY SSLKEYRINGPW command 905
- QUERY STATUS command 906
- QUERY STATUSTHRESHOLD command 914
- QUERY STGPOOL command 917
- QUERY SUBSCRIBER command 928
- QUERY SUBSCRIPTION command 930
- QUERY SYSTEM command 932
- QUERY TAPEALERTMSG command 934
- QUERY TOC command 935
- QUERY VIRTUALFSMAPPING command 938
- QUERY VOLHISTORY command 940
- QUERY VOLUME command 948
- QUERYAUTH server option 1513
- querying
  - activity log 661
  - administrator 667
  - audit occupancy 681
  - backup sets 684
  - clients with schedules 679
  - completed events 745
  - contents of a backup set 689
  - contents of a volume 697
  - copy group 705
  - database 713
  - device class 718
  - domain 724
  - drive 727
  - event server 759
  - file space 767
  - library 774
  - library volume 777
  - license 780
  - management class 793
  - mount requests 868
  - node 809
  - node groups 823
  - occupancy 825
  - option 829
  - policy set 835
  - process 838
  - profile 841
  - recovery log 782
  - recovery plan file 874

- querying (*continued*)
  - recovery plan file content 872
  - scheduled events 745
  - schedules 880
  - script 888
  - server group 895
  - session 897
  - space trigger 903
  - SSL key database password 905, 1141
  - status 906
  - storage pool 917
  - storage pool volume 948
  - subscriber 928
  - subscription 930
  - volume history file 940
- QUIT command 956
- quitting
  - batch mode 3
  - interactive mode 2, 956
  - mount mode 3

## R

- random access volumes 341
- raw partitions, performance effect 341
- RECLAIM parameter
  - DEFINE STGPOOL
    - active-data pool 380
    - copy sequential access 371
    - primary sequential access 357
  - UPDATE STGPOOL
    - active-data pool 1414
    - copy sequential access 1407
    - primary sequential access 1394
- RECLAIM STGPOOL command 957
- RECLAIMDELAY option 1514
- RECLAIMPERIOD option 1515
- reclamation of volumes
  - specifying a threshold of reclaimable space
    - using the DEFINE STGPOOL command 357, 371, 380
    - using the UPDATE STGPOOL command 1395, 1407, 1414
  - specifying multiple concurrent processes
    - using the DEFINE STGPOOL command 372, 381
    - using the UPDATE STGPOOL command 1408, 1415
  - starting manually 957
- RECLAMATIONTYPE parameter
  - defining active-data pools 382
  - defining copy storage pools 373
  - defining sequential access storage pools 359
- RECONCILE VOLUMES command 961
- recovery log
  - installing 1566
  - setting options for 1443
- recovery logs 1564
- recovery plan file
  - display contents 872
  - query information 874
  - set expire days 1105
- recovery plan prefix 1101
- redirecting command output 4
- redirection characters
  - types of 4
  - using 4
- refresh interval
  - setting 1081

- REFRESHSTATE parameter
  - AUDIT LIBRARY 37
- REGISTER ADMIN command 965
- REGISTER LICENSE command 969
- REGISTER NODE command 971
- registering
  - administrator 965
  - license 969
  - node 971
- registration
  - closed 1126
  - description of 1126
  - setting 1126
- REMOVABLEFILE device type 204, 1242
- REMOVABLEFILE, CD support 204
- REMOVE ADMIN command 989
- REMOVE NODE command 991
- REMOVE REPLNODE command 993
- removing
  - administrator 989
  - client association 401
  - node 991
- RENAME ADMIN command 996
- RENAME FILESPACE command 998
- RENAME NODE command 1001
- RENAME SCRIPT command 1003
- RENAME SERVERGROUP command 1004
- RENAME STGPOOL command 1005
- renaming
  - administrator 996
  - file space 998
  - node 1001
  - script 1003
  - server group 1004
  - storage pool 1005
- REPLACEDFS parameter
  - EXPORT ADMIN 495
  - EXPORT NODE 516
  - EXPORT POLICY 523
  - EXPORT SERVER 541
  - IMPORT ADMIN 572
  - IMPORT NODE 578
  - IMPORT POLICY 582
- REPLBATCHSIZE option 1519
- REPLICATE NODE command 1006
- replication
  - See* node replication
- replication server, setting or removing target 1130
- REPLSIZETHRESH option 1520
- REPLY command 1013
- reporting error message for ANR9999D 1080
- REPORTRETRIEVE server option 1518
- REQSYSAUTHOUTFILE option 1521
- request, mount 72, 868
- RESET PASSEXP command 1014
- RESOURCETIMEOUT option 1522
- RESTORE DB, DSMSEV 1574
- RESTORE NODE command 1019
- RESTORE STGPOOL command 1024
- RESTORE VOLUME command 1028
- RESTOREINTERVAL option 1523
- restoring
  - as action on client command 123, 305, 1347
- retention period
  - description of 1108
  - setting 1108

- retention period, event record
  - managing 1108
  - querying 906
- RETENTIONEXTENSION option 1524
- RETONLY parameter
  - DEFINE COPYGROUP, backup 138
  - UPDATE COPYGROUP, backup 1185
- retry period
  - description of 1132
  - setting 1132
- return code checking 1583
- RETVR parameter
  - DEFINE COPYGROUP, archive 142
  - UPDATE COPYGROUP, archive 1187
- REUSEDELAY parameter
  - DEFINE STGPOOL
    - active-data pool 383
    - copy sequential access 375
    - primary sequential access 361
  - UPDATE STGPOOL
    - active-data pool 1417
    - copy sequential access 1410
    - primary sequential access 1398
- reusing volumes
  - active-data pool 380, 1414
  - copy sequential access 371, 1407
  - primary sequential access 357, 1395
- REVOKE AUTHORITY command 1033
- REVOKE PROXYNODE command 1037
- revoking
  - authority 1033
  - types of privilege classes 1033
- roll-forward recovery log mode 47
- ROLLBACK command 1038
- rolling back commands in a macro 1038
- routing commands 17
- RUN command 1039

## S

- SAN
  - DEFINE LIBRARY command 247
  - UPDATE LIBRARY command 1286
- SAN tape devices 249, 252, 256, 259, 1287, 1289, 1295
- SANDISCOVERY 1525
- SANDISCOVERYTIMEOUT option 1526
- SANREFRESHTIME server option 1527
- schedule
  - administrative command 314
  - associating with a client node 115
  - client 301, 314
  - copying 105
  - defining 271, 300
  - deleting 444
  - description of 314
  - querying 880
  - querying results of (events) 745
  - restrictions of 301
  - startup window, defining schedule 310, 318
  - startup window, updating schedule 1351, 1360
  - types of 300
  - updating 1342
- schedule event
  - querying 745
  - setting start date for displaying 746
  - setting start time for displaying 747, 748, 753, 755
  - viewing information about 745

- scheduling mode
  - information about 1133
  - setting 1133
- scratch volumes in storage pool
  - defining a storage pool 361, 374, 383
  - updating a storage pool 1397, 1409, 1416
- scripts
  - copying 109
  - defining 323
  - deleting 447
  - querying 888
  - renaming 1003
  - running 1039
  - updating 1365
- SEARCHMPQUEUE option 1528
- Secure Sockets Layer (SSL) 326, 891, 1368
- security options and licensing options 1445
- security, encryption
  - 3592 devices 158, 1201
  - LTO devices 191, 1231
  - StorageTek devices 181, 1221
- SELECT command 1042
- sequential volume history
  - backing up 62
  - deleting 457
  - display sequential volume history 940
  - querying 940
- SERIALIZATION parameter
  - DEFINE COPYGROUP, archive 143
  - DEFINE COPYGROUP, backup 138
  - UPDATE COPYGROUP, archive 1188
  - UPDATE COPYGROUP, backup 1185
- server
  - disabling sessions, inbound and outbound 471
  - enabling sessions, inbound and outbound 481
  - exporting 525
  - importing 584
  - migrating 607
  - reclaiming 957
  - setting name for 1137
- server console
  - restrictions 7
  - using 7
- server name settings 1137
- server options
  - 3494SHARED 1447
  - ACSACCESSID 1448
  - ACSLCKDRIVE 1449
  - ACSQUICKINIT 1450
  - ACSTIMEOUTX 1451
  - ACTIVELOGDIRECTORY 1452
  - ACTIVELOGSIZE 1453
  - ADMINCOMMTIMEOUT 1454
  - ADMINIDLETIMEOUT 1455
  - ADMINONCLIENTPORT 1456
  - ALIASHALT 1457
  - ALLOWREORGINDEX 1458
  - ALLOWREORGTABLE 1459
  - ARCHFAILOVERLOGDIRECTORY 1460
  - ARCHLOGDIRECTORY 1461
  - ASSISTVCRRECOVERY 1462
  - AUDITSTORAGE 1463
  - BEGINREORGTIME 1516, 1517
  - changing with SETOPT command 829
  - CHECKTAPEPOS 1464
  - CLIENTDEDUPTXNLIMIT 1466
  - COMMMETHOD 1468
- server options (*continued*)
  - COMMTIMEOUT 1470
  - DBDIAGLOGSIZE 1471
  - DBDIAGPATHFSTHRESHOLD 1472
  - DBMEMPERCENT 1473
  - DBMTCPPORT 1474
  - DEDUPTIER2FILESIZE 1476
  - DEDUPTIER3FILESIZE 1477
  - DEVCONFIG 1478
  - DISABLESCHEDS 1479
  - DISPLAYLFINFO 1480
  - DNSLOOKUP 1481
  - DRIVEACQUIRERETRY 1482
  - ENABLENASDEDUP 1483
  - EVENTSERVER 1484
  - EXPINTERVAL 1485
  - EXPQUIET 1486
  - FFDCLOGNAME 1487
  - FFDCMAXLOGSIZE 1488
  - FILEEXIT 1489
  - FILETEXTEXIT 1490
  - FSUSEDTHRESHOLD 1491
  - IDENTIFYAUTOSTART 1475
  - IDLETIMEOUT 1492
  - LANGUAGE 1493
  - LDAPCACHEDURATION 1495
  - LDAPURL 1496
  - MAXSESSIONS 1498
  - message options 1444
  - MESSAGEFORMAT 1499
  - MIRRORLOGDIRECTORY 1500
  - modifying the file 1439
  - MOVEBATCHSIZE 1501
  - MOVESIZETHRESH 1502
  - MSGINTERVAL 1503
  - NDMPCONTROLPORT 1504
  - NDMPENABLEKEEPALIVE 1505
  - NDMPKEEPIDLEMINUTES 1506
  - NDMPPORTRANGE 1507
  - NDMPREFDATAINTERFACE 1508
  - NOPREEMPT 1509
  - NORETRIEVEDATE 1510
  - NUMOPENVOLSAALLOWED 1511
  - QUERYAUTH 1513
  - querying 1439
  - RECLAIMDELAY 1514
  - RECLAIMPERIOD 1515
  - REPLBATCHSIZE 1519
  - REPLSIZETHRESH 1520
  - REPORTRETRIEVE 1518
  - REQSYSAUTHOUTFILE 1521
  - RESOURCETIMEOUT 1522
  - RESTOREINTERVAL 1523
  - RETENTIONEXTENSION 1524
  - SANDISCOVERY 1525
  - SANDISCOVERYTIMEOUT 1526
  - SANREFRESHTIME 1527
  - SEARCHMPQUEUE 1528
  - SERVERDEDUPTXNLIMIT 1529
  - SHMPORT 1531
  - SHREDDING 1532
  - SNMPHEARTBEATINTERVAL 1533
  - SNMPMESSAGECATEGORY 1534
  - SNMPSUBAGENT 1535
  - SNMPSUBAGENTHOST 1536
  - SNMPSUBAGENTPORT 1537
  - SSLFIPSMODE 1537

server options (*continued*)

- SSLTCPADMINPORT 1538
- SSLTCPSPORT 1539
- SSLTLS12 1540
- tailoring 1439
- TCPADMINPORT 1541
- TCPBUFSIZE 1542
- TCPNODELAY 1543
- TCPPORT 1544
- TCPWINDOWSIZE 1545
- TECBEGINEVENTLOGGING 1546
- TECHOST 1547
- TECPORT 1548
- TECUTF8EVENT 1549
- THROUGHPUTDATATHRESHOLD 1550
- THROUGHPUTTIMETHRESHOLD 1551
- TXNGROUPMAX 1552
- UNIQUETDPTCEVENTS 1553
- UNIQUETECEVENTS 1554
- VOLUMEHISTORY 1557

server scripts

- copying 109
- defining 323
- deleting 447
- querying 888
- renaming 1003
- running 1039
- updating 1365

server storage

- setting options for 1442

server-to-server communications 326

- COMMMETHOD option 1468
- shared memory between server and client 1468

SERVERDEDUPTXNLIMIT option 1529

session

- maximum number scheduled 1116
- querying 897

SET ACCOUNTING command 1055, 1149

SET ACTLOGRETENTION command 1056

SET ALERTACTIVEDURATION 1058

SET ALERTEMAIL 1060

SET ALERTEMAILFROMADDR 1062

SET ALERTEMAILSMTPHOST 1063

SET ALERTEMAILSMTPPORT 1064

SET ALERTINACTIVEDURATION 1059, 1066

SET ALERTMONITOR 1067

SET ALERTSUMMARYTOADMINS 1065

SET ALERTUPDATEINTERVAL 1068

SET ARREPLRULEDEFAULT command 1071

SET AUTHENTICATION command 1073

SET BKREPLRULEDEFAULT command 1074

SET CLIENTACTDURATION command 1077

SET commands 1052

SET CONFIGMANAGER command 1078

SET CONFIGREFRESH command 1079

SET CPUINFOREFRESH command 1081

SET CROSSDEFINE command 1082

SET DBRECOVERY command 1083

SET DBREPORTMODE command 1085

SET DEDUPVERIFICATIONLEVEL command 1086

SET DEFAULTAUTHENTICATION command 1088

SET DRMACTIVEDATASTGPOOL command 1090

SET DRMCHECKLABEL command 1091

SET DRMCMDFILENAME command 1092

SET DRMCOPYSTGPOOL command 1093

SET DRMCOURIERNAME command 1094

SET DRMDBBACKUPEXPIREDAYS command 1095

SET DRMFILPROCESS command 1097

SET DRMINSTRPREFIX command 1098

SET DRMNOTMOUNTABLENAME command 1100

SET DRMPPLANPREFIX command 1101

SET DRMPPLANVPOSTFIX command 1103

SET DRMPRIMSTGPOOL command 1104

SET DRMRPFEXPIREDAYS command 1105

SET DRMVaultNAME command 1107

SET EVENTRETENTION command 1108

SET INVALIDPWLIMIT command 1109

SET LDAPPASSWORD command 1111

SET LDAPUSER command 1112

SET LICENSEAUDITPERIOD command 1114

SET MAXCMDRETRIES command 1115

SET MAXSCHEDESESSIONS command 1116

SET MINPWLENGTH command 1118

SET MONITOREDSEVERGROUP command 1120

SET MONITORINGADMIN command 1119

SET PASSEXP command 1121

SET QUERYSCHEDPERIOD 1123

SET RANDOMIZE command 1124

SET REGISTRATION command 1126

SET REPLRETENTION command 1128

SET REPLSERVER command 1130

SET RETRYPERIOD command 1132

SET SCHEDMODES command 1133

SET SERVERHLADDRESS command 1135

SET SERVERLLADDRESS command 1136

SET SERVERNAME command 1137

SET SERVERPASSWORD command 1138

SET SPREPLRULEDEFAULT command 1139

SET SSLKEYRINGPW command 1141

SET STATUSATRISKINTERVAL command 1142

SET STATUSMONITOR command 1144

SET STATUSREFRESHINTERVAL command 1146

SET STATUSSKIPFAILURE command 1147

SET SUBFILE command 1149

SET SUMMARYRETENTION command 1150

SET TAPEALERTMSG command 1151

SET TOCLOADRETENTION command 1152

SETARCHIVERETENTIONPROTECTION command 1069

SETOPT command 1153

setting

- accounting record 1055
- authentication 1073
- configuration manager 1078
- configuration refresh 1079
- cross define of a server 1082
- data deduplication verification level 1086
- frequency for client-polling 1123
- frequency for one-time client action 1077
- high level address of a server 1135
- license audit period 1114
- low level address of a server 1136
- maximum retries of a command 1115
- maximum scheduled sessions 1116
- password for a server 1138
- password for expiration date 1121
- randomization of start times 1124
- registration 1126
- retention period for activity log 1056
- retention period for event records 1108
- retry period 1132
- scheduling mode 1133
- server name 1137
- summary retention days 1150
- setting communication options 1440

- shared library type 262, 1298
- SHMPORT option
  - specifying a TCP/IP port address 1531
  - using shared memory 1531
- SHRED DATA command 1155
- shredding data, storage pools
  - backing up 60
  - defining 351
  - moving data 612
  - updating 1388
- SHREDDING option 1532
- simultaneous write to copy storage pools
  - defining random access storage pools 349
  - defining sequential access storage pools 365
  - updating random access storage pools 1386
  - updating sequential access storage pools 1401
- single-instance store (data deduplication) 367, 368, 376, 385, 1403, 1410, 1411, 1417, 1418
- SKIPPARTIAL parameter, AUDIT VOLUME 43
- SnapLock for managing WORM FILE volumes
  - when defining active-data pools 382
  - when defining copy storage pools 373
  - when defining sequential access storage pools 359
- SNMPHEARTBEATINTERVAL option 1533
- SNMPMESSAGECATEGORY option 1534
- SNMPSUBAGENT option 1535
- SNMPSUBAGENTHOST option 1536
- SNMPSUBAGENTPORT option 1537
- software support
  - describing problem for IBM Software Support 1590
  - determining business impact for IBM Software Support 1590
  - submitting a problem 1591
- Software Support
  - contact 1589
- Sony Blue Laser RW and WORM media 197, 1236
- space management
  - DEFINE MGMTCLASS 271
  - description of 271
  - UPDATE MGMTCLASS 1306
- space trigger commands
  - DEFINE SPACETRIGGER 334
  - DELETE SPACETRIGGER 450
  - QUERY SPACETRIGGER 903
  - UPDATE SPACETRIGGER 1373
- SPACEMGTECHNIQUE parameter
  - DEFINE MGMTCLASS 272
  - UPDATE MGMTCLASS 1306
- SSL (Secure Sockets Layer) 326, 891, 1368
- SSL key database password
  - querying 905
  - setting 1141
- SSLFIPSMODE option 1537
- SSLTCPADMINPORT server option 1538
- SSLTCPPOINT server option 1539
- SSLTLS12 option 1540
- start time
  - information about 300
  - setting 1124
- STARTDATE parameter
  - DEFINE SCHEDULE, administrative command 318
  - DEFINE SCHEDULE, client 310
  - UPDATE SCHEDULE, administrative command 1360
  - UPDATE SCHEDULE, client 1351
- STARTTIME parameter
  - DEFINE SCHEDULE, administrative command 318
  - DEFINE SCHEDULE, client 310
- STARTTIME parameter (*continued*)
  - UPDATE SCHEDULE, administrative command 1360
  - UPDATE SCHEDULE, client 1351
- status
  - event completion 751
  - information about 906
  - querying 906
- STGPOOL parameter, MOVE DATA 612
- STGPOOLS parameter
  - GRANT AUTHORITY 558
  - REVOKE AUTHORITY 1035
- storage 1442
- storage area network (SAN)
  - DEFINE LIBRARY command 247
  - UPDATE LIBRARY command 1286
- storage pool
  - collocating
    - active-data pool 379, 1413
    - copy sequential access 370, 1406
    - primary sequential access 360, 1396
  - defining 341
  - deleting 453
  - migrating 607
  - querying 917
  - reclaiming 957
  - restoring 1024
  - updating 1379
- storage pool volume
  - auditing 40
  - defining 391
  - deleting 462
  - querying 697, 948
  - scratch, active-data pool 383, 1416
  - scratch, copy sequential access 374, 1409
  - scratch, primary sequential access 361, 1397
  - updating 1423
  - varying 1436
- storage pools 334
- storage volumes
  - naming 13
  - performance increase with raw partitions 341
- StorageTek drives
  - logical block protection 175
- support contract 1590
- support information 1587
- support subscription 1590
- SWAP parameter, CHECKIN LIBVOLUME 80
- syntax diagram
  - abbreviations 8
  - default value 9
  - fragments 11
  - optional choice 9
  - repeatable choice 10
  - repeating values 10
  - required parameters 9
  - symbols 10
  - using 8
  - variables 10
- SYSCONFIG command (on NAS file server) 276
- system privilege class
  - administrative commands 19

## T

- tape
  - AUDIT LIBRARY command 36
  - AUDIT VOLUME command 40

tape (*continued*)

- CHECKIN LIBVOLUME command 75
- CHECKOUT LIBVOLUME command 83
- DEFINE DEVCLASS command 148
- DEFINE LIBRARY command 247, 249, 252, 254, 257, 259, 262
- DEFINE STGPOOL command 341, 352, 369
- DEFINE VOLUME command 391
- DELETE DEVCLASS command 416
- DELETE DRIVE command 418
- DELETE LIBRARY command 428
- DELETE VOLUME command 462
- LABEL LIBVOLUME command 592
- QUERY DEVCLASS command 718
- QUERY DRIVE command 727
- QUERY LIBRARY command 774
- QUERY LIBVOLUME command 777
- QUERY VOLUME command 948
- storage area network (SAN) 249, 252, 256, 259, 1287, 1289, 1295
- UPDATE DEVCLASS command 1192
- UPDATE LIBRARY command 1286, 1287, 1289, 1291, 1293, 1295, 1298
- UPDATE LIBVOLUME command 1302
- UPDATE VOLUME command 1423

tape-device encryption

- 3592 158, 1201
- ECARTRIDGE 181, 1221
- LTO 191, 1231

target replication server, setting or removing 1130

TCP keepalive

- enabling 1505
- specifying connection idle time (AIX, Linux, and Windows) 1506

TCPADMINPORT server option 1541

TCPBUFSIZE option 1542

TCPNODELAY option 1543

TCPPORT option 1544

TCPWINDOWSIZE option 1545

TECBEGINEVENTLOGGING option 1546

TECHOST option 1547

TECPORT option 1548

TECUTF8EVENT option 1549

threshold, migration

- random-access storage pools
  - high threshold 346, 1383
  - low threshold 346, 1384
- sequential-access storage pools
  - high threshold 356, 1394
  - low threshold 357, 1394

THROUGHPUTDATATHRESHOLD option 1550

THROUGHPUTTIMETHRESHOLD option 1551

time interval, setting for checking in volumes 1214, 1231, 1236

Tivoli Monitoring V6.3.3 updates xix

Tivoli Monitoring V6.3.4 updates xvii

Tivoli technical training 1587

TOCDestination parameter 139, 1186

TOSERVER parameter

- EXPORT ADMIN 495
- EXPORT NODE 516
- EXPORT POLICY 523
- EXPORT SERVER 541

training, Tivoli technical 1587

transfer data by data mover 145

troubleshooting LAN-free data movement 1428

TXNGROUPMAX option 1552

type, device

- 3590 149, 1193
- 3592 153, 218, 1197
- 4MM 159, 1202
- 8MM 163, 1206
- CARTRIDGE 224, 1260
- DLT 169, 1211
- ECARTRIDGE 175, 1216
- FILE 183, 206, 236, 1223, 1244, 1272
- LTO 186, 187, 1227
- OPTICAL 196, 1235
- QIC 199, 1238
- REMOVABLEFILE 204, 1242
- SERVER 206, 1244
- VOLSAFE 208, 1246
- WORM 196, 1235
- WORM12 196, 1235

typographic conventions xiv

## U

Ultrium, LTO device type

- logical block protection 186, 1227

UNIQUEDPTECEVENTS option 1553

UNIQUETECEVENTS option 1554

UNLOCK ADMIN command 1159

UNLOCK NODE command 1161

UNLOCK PROFILE command 1163

unlocking

- administrator 1159
- node 1161
- profile 1163

Update a 349X library 1287

Update a file library 1292

Update a manual library 1293

Update a SCSI library 1295

Update a shared library 1298

Update a VTL library 1299

UPDATE ADMIN command 1170

UPDATE ALERTSTATUS 1168

UPDATE ALERTTRIGGER command 1165

Update an ACSLS library 1289

UPDATE BACKUPSET 1174

UPDATE CLIENTOPT command 1179

UPDATE CLOPTSET command 1180

UPDATE COLLOGGROUP command 1181

UPDATE COPYGROUP command 1182

UPDATE DATAMOVER command 1190

UPDATE DEVCLASS command 1192

UPDATE DOMAIN command 1275

UPDATE DRIVE command 1277

UPDATE FILESPACE command 1281

UPDATE LIBRARY command 1286, 1287, 1289, 1291, 1292, 1293, 1295, 1298, 1299

UPDATE LIBVOLUME command 1302

UPDATE MACHINE command 1304

UPDATE MGMTCLASS command 1306

UPDATE NODE command 1309

UPDATE PATH command 1326, 1327, 1331, 1334

UPDATE POLICYSET command 1335

UPDATE PROFILE command 1337

UPDATE RECOVERYMEDIA command 1338

UPDATE REPLRULE command 1340

UPDATE SCHEDULE command 1342

UPDATE SCRIPT command 1365

UPDATE SERVER command 1368

UPDATE SERVERGROUP command 1372

UPDATE SPACETRIGGER command 1373  
UPDATE STATUSTHRESHOLD command 1375  
UPDATE STGPOOL command 1379  
UPDATE VIRTUALFSMAPPING command 1419  
UPDATE VOLHISTORY command 1421  
UPDATE VOLUME command 1423

#### updating

- administrator 1170
- archive copy group 1187
- backup copy group 1183
- backup set 1174
- copy group 1182
- device class 1192
- domain 1275
- drive 1277
- file library 1292
- library 1286, 1287, 1289, 1291, 1293, 1295, 1298
- library volume 1302
- management class 1306
- node 1309
- node group 1325
- policy set 1335
- profile 1337
- retention period for a backup set 1174
- schedule 1342
- script 1365
- server group 1372
- storage pool 1379
- storage pool space triggers 1373
- storage pool volume 1423
- TOCDestination copy group parameter 1186
- volume history 1421
- VTL library 1299

USEREXIT option 1555

#### using

- administrative client options 5
- command-line interface 1
- continuation characters 12
- macros 605
- redirection 4
- redirection characters 4
- syntax diagram 8
- system date on the server 28

## V

VALIDATE LANFREE command 1428  
VALIDATE POLICYSET command 1430  
VALIDATE REPLICATION command 1432

#### validating

- LAN-free 1428
- policy set 1430

VARY command 1436

varying volumes 1436

#### VERDELETED parameter

- DEFINE COPYGROUP, backup 137
- UPDATE COPYGROUP, backup 1184

#### VEREXISTS parameter

- DEFINE COPYGROUP, backup 137
- UPDATE COPYGROUP, backup 1184

verify automated library inventory 36

virtual tape library 263

#### VOLSAFE

- device type 208, 1246
- labeling volumes 596

#### volume

- labeling library 592

#### volume history

- deleting 457
- querying 940

VOLUMEHISTORY option 1557

#### volumes, storage

- defining 391
- limiting offsite volumes to be reclaimed
  - using the DEFINE STGPOOL command 373, 382
  - using the UPDATE STGPOOL command 1408, 1415
- naming 13
- reclaiming
  - from active-data pool 380, 1414
  - from copy sequential access 371, 1407
  - from primary sequential access 357, 1395
- restoring 1028
- scratch, maximum in active-data pool 383, 1416
- scratch, maximum in copy sequential access 374, 1409
- scratch, maximum in primary sequential access 361, 1397

#### VTL

- DEFINE LIBRARY command 263
- UPDATE LIBRARY command 1299

#### VTL library

- defining 263
- updating 1299

## W

#### WAITTIME parameter

- CHECKIN LIBVOLUME command 80
- LABEL LIBVOLUME command 596

WHEREACCESS parameter, UPDATE VOLUME 1425

WHERESTATUS parameter, UPDATE VOLUME 1425

WHERESTGPOOL parameter, UPDATE VOLUME 1425

#### wildcard

- match-any examples 14
- match-exactly-one examples 14
- using 14

#### WORM devices and media

- 8mm devices 164
  - DLT 170
  - IBM 3592 154
  - LTO 188
  - SnapLock for managing WORM FILE volumes 359, 373, 382
  - Sony AIT50/AIT100, supported in 8mm class
    - definition 164
  - Sony Blue Laser 197
  - StorageTek VolSafe 178
- WORM optical devices
  - defining the class 196
  - updating the class 1235

## Z

#### z/OS media server

- 3590 device class 213, 1250
- 3592 device class 1255
- 3592 device type 218
- CARTRIDGE device type 1260
- ECARTRIDGE device class 229
- ECARTRIDGE device type 1265
- FILE device type 236
- ZOSMEDIA 249, 252, 254, 262, 1287, 1289, 1291, 1298





Product Number: 5608-E01  
5608-E02  
5608-E03

Printed in USA

SC23-9777-05

