



Performance Tuning Guide



Performance Tuning Guide

Note

Before using this information and the product it supports, read the information in "Notices" on page 59.

This edition applies to Version 6.1 of IBM Tivoli Storage Manager and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters.

© **Copyright International Business Machines Corporation 1996, 2009.**

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

Contents

Preface	v
Who should read this guide	v
Publications	v
Tivoli Storage Manager publications	v
Support information	vii
Getting technical training	vii
Searching knowledge bases	vii
Contacting IBM Software Support	ix

New for IBM Tivoli Storage Manager Version 6.1 xi

New for the server in Version 6.1.2.	xii
Enabled functions	xii
Licensing changes	xii
ACSLS functionality for 64-bit Windows systems	xii
TSMDLST utility for AIX.	xii
Device support for AIX	xii
PREVIEW parameter for DSMSERV INSERTDB	xii
Replacement passthru device driver for HP-UX 11i 3 (IA64 only)	xii
SAN discovery support for Linux on zSeries	xiii
New server option SANDISCOVERYTIMEOUT	xiii
New for the server in Version 6.1.0	xiii
Disabled functions in 6.1.0 and 6.1.1	xiii
Changes to the Version 6.1 Administration Center	xiv
Data deduplication	xvi
Storage devices	xvi
Disaster recovery manager support for active-data pools	xvii
EXPIRE INVENTORY command enhancements	xvii
No-query restore changes	xviii
Server database	xviii
Support for NetApp SnapMirror to Tape feature	xviii
Reporting and monitoring feature	xviii
ODBC driver support.	xix
Backup sets and client node enhancements	xix

Chapter 1. Overview of IBM Tivoli Storage Manager tuning. 1

Chapter 2. IBM Tivoli Storage Manager server performance tuning. 3

Tuning server options	3
DBMEMPERCENT	3
DISKSTGPOOLMEMSIZE	4
EXPINTERVAL	4
MAXSESSIONS	5
MOVEBATCHSIZE and MOVESIZETHRESH	5
RESTOREINTERVAL	6
TCPNODELAY	6
TCPWINDOWSIZE	6
TXNGROUPMAX	7

Server hardware recommendations	8
Database manager for IBM Tivoli Storage Manager	9
Database and log performance	9
Database manager tuning.	10
Backup performance	11
Tuning inventory expiration	11
Disaster recovery performance	12
Searching the server activity log	12
Scheduling sessions and processes.	12
LAN-free backup	12
Maximum number of mount points for a node	13
Managing storage pools and volumes	13
Cached disk storage pools	13
Tuning storage pool migration	14
Improving storage agent performance	15
Modifying the IBM Tivoli Monitoring environment file for reporting performance	15
Performance improvement activities for server platforms	16
Actions for better performance on all server platforms	17
AIX server performance	17
AIX: vmo and ioo commands	17
UNIX file systems	19
HP-UX server	19
Linux server	19
Sun Solaris server	19
Windows server	20
Estimating throughput in untested environments.	21
Tuning tape drive performance.	21
Using collocation with tape drives.	22
Tape drive transfer rate	23
Tuning disk performance.	24
Busses	24

Chapter 3. IBM Tivoli Storage Manager client performance tuning 25

COMMRESTARTDURATION and COMMRESTARTINTERVAL.	25
COMPRESSALWAYS	25
COMPRESSION	26
DISKBUFFSIZE	27
MEMORYEFFICIENTBACKUP	27
PROCESSORUTILIZATION	28
QUIET	29
RESOURCEUTILIZATION	29
TAPEPROMPT	31
TCPBUFFSIZE	31
TCPNODELAY	31
TCPWINDOWSIZE.	32
TXNBYTELIMIT.	32
Multiple session backup and restore	33
Client command line options	35
Performance recommendations by client platform	35
Macintosh client.	35

Windows client	35
Client performance recommendations for all platforms	36
Hierarchical Storage Manager tuning	37
Data Protection for Domino for z/OS.	37
Chapter 4. Administration Center performance tuning	39
Administration Center capacity planning	39
Maximum number of active administrators.	39
Processing capacity.	39
I/O throughput	40
Processing memory.	40
Java heap memory size	40
Administration Center setup recommendations	42
Installation requirements	42
Locating the Administration Center	42
Minimizing memory usage	43
Optimizing Windows Server 2003 memory	43
Using the default action	43
Tuning Administration Center performance.	44
Tuning processor performance	44
Tuning network performance	44
Tuning memory performance	44

Chapter 5. Network protocol tuning . . 47

TCP/IP communication concepts and tuning	47
Sliding window	49
Networks	49
Limiting network traffic	50
AIX network settings	50
MTU and MSS settings	51
NetWare client cache tuning.	53
Maximum transmission unit settings	53
Sun Solaris network settings.	54
z/OS network settings.	55
USS client with IBM TCP/IP for z/OS	55
TCP/IP and z/OS UNIX system services performance tuning.	56

Appendix. Accessibility features for Tivoli Storage Manager. 57

Notices	59
Trademarks	61

Glossary 63

Index 65

Preface

This publication helps you tune the performance of the servers and clients in your IBM® Tivoli® Storage Manager environment.

Before using this publication, you should be familiar with the following areas:

- The operating systems on which your IBM Tivoli Storage Manager servers and clients reside
- The communication protocols installed on your client and server machines

Who should read this guide

The audience for this publication is anyone who wants to improve the performance of the Tivoli Storage Manager server and client, and the network and attached hardware as they apply to Tivoli Storage Manager.

Publications

Tivoli Storage Manager publications and other related publications are available online.

You can search all publications in the Tivoli Storage Manager Information Center: <http://publib.boulder.ibm.com/infocenter/tsminfo/v6>.

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/>.

You can also order some related publications from the IBM Publications Center Web site. The Web site provides information for ordering publications from countries other than the United States. In the United States, you can order publications by calling 800-879-2755.

Tivoli Storage Manager publications

Publications are available for the server, storage agent, client, and Data Protection.

Table 1. Tivoli Storage Manager server publications

Publication title	Order number
<i>IBM Tivoli Storage Manager Messages</i>	GC23-9787
<i>IBM Tivoli Storage Manager Performance Tuning Guide</i>	GC23-9788
<i>IBM Tivoli Storage Manager Problem Determination Guide</i>	GC23-9789
<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

Table 1. Tivoli Storage Manager server publications (continued)

Publication title	Order number
<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 Sun Solaris Installation Guide</i>	GC23-9784
<i>IBM Tivoli Storage Manager for Sun Solaris Administrator's Guide</i>	SC23-9772
<i>IBM Tivoli Storage Manager for Sun 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 Server Upgrade Guide</i>	SC23-9554
<i>IBM Tivoli Storage Manager for System Backup and Recovery Installation and User's Guide</i>	SC32-6543

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 Sun 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
<i>IBM Tivoli Storage Manager for Windows: Backup-Archive Clients Installation and User's Guide</i>	SC23-9792
<i>IBM Tivoli Storage Manager for Space Management for UNIX and Linux: User's Guide</i>	SC23-9794
<i>IBM Tivoli Storage Manager for HSM for Windows Administration Guide</i>	SC23-9795
<i>IBM Tivoli Storage Manager Using the Application Program Interface</i>	SC23-9793
<i>Program Directory for IBM Tivoli Storage Manager z/OS Edition Backup-Archive Client</i>	GI11-8912
<i>Program Directory for IBM Tivoli Storage Manager z/OS Edition Application Program Interface</i>	GI11-8911

Table 4. Tivoli Storage Manager Data Protection publications

Publication title	Order number
<i>IBM Tivoli Storage Manager for Advanced Copy Services: Data Protection for Snapshot Devices Installation and User's Guide</i>	SC33-8331
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Installation and User's Guide</i>	SC32-9059
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for UNIX and Linux Installation and User's Guide</i>	SC32-9064
<i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for Windows Installation and User's Guide</i>	SC32-9065
<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 Mail: Data Protection for Lotus Domino for UNIX, Linux, and OS/400® Installation and User's Guide</i>	SC32-9056
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino for Windows Installation and User's Guide</i>	SC32-9057
<i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Installation and User's Guide</i>	SC23-9796
<i>Program Directory for IBM Tivoli Storage Manager for Mail (Data Protection for Lotus Domino)</i>	GI11-8909

Support information

You can find support information for IBM products from a variety of sources.

Getting technical training

Information about Tivoli technical training courses is available online.

Go to <http://www.ibm.com/software/tivoli/education/>.

Searching knowledge bases

If you have a problem with Tivoli Storage Manager, there are several knowledge bases that you can search.

You can begin with the Tivoli Storage Manager Information Center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6>. From this Web site, you can search all Tivoli Storage Manager publications.

Searching the Internet

If you cannot find an answer to your question in the Tivoli Storage Manager information center, search the Internet for the latest, most complete information that might help you resolve your problem.

To search multiple Internet resources, go to the support Web site for Tivoli Storage Manager at <http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>. From there, you can search a variety of resources including:

- IBM technotes
- IBM downloads

- IBM Redbooks®

If you still cannot find the solution to the problem, you can search forums and newsgroups on the Internet for the latest information that might help you resolve your problem. To share your experiences and learn from others in the user community, go to the Tivoli Storage Manager wiki at <http://www.ibm.com/developerworks/wikis/display/tivolistoragemanager/Home>.

Using IBM Support Assistant

At no additional cost, you can install on any workstation the IBM Support Assistant, a stand-alone application. You can then enhance the application by installing product-specific plug-in modules for the IBM products that you use.

The IBM Support Assistant helps you gather support information when you need to 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

For more information, see the IBM Support Assistant Web site at <http://www.ibm.com/software/support/isa/>.

Finding product fixes

A product fix to resolve your problem might be available from the IBM Software Support Web site.

You can determine what fixes are available by checking the Web site:

1. Go to the IBM Software Support Web site at <http://www.ibm.com/software/tivoli/products/storage-mgr/product-links.html>.
2. Click the **Support Pages** link for your Tivoli Storage Manager product.
3. Click **Download**, and then click **Fixes by version**.

Getting e-mail notification of product fixes

You can get notifications about fixes and other news about IBM products.

To receive weekly e-mail notifications about fixes and other news about IBM products, follow these steps:

1. From the support page for any IBM product, click **My support** in the upper-right corner of the page.
2. If you have already registered, skip to the next step. If you have not registered, click **Register** in the upper-right corner of the support page to establish your user ID and password.
3. Sign in to **My support**.
4. On the My support page, click **Edit profiles** in the left navigation pane, and scroll to **Select Mail Preferences**. Select a product family and check the appropriate boxes for the type of information you want.
5. Click **Submit**.
6. For e-mail notification for other products, repeat steps 4 and 5.

Contacting IBM Software Support

You can contact IBM Software Support if you have an active IBM software maintenance contract and if you are authorized to submit problems to IBM.

Before you contact IBM Software Support, follow these steps:

1. Set up a software maintenance contract.
2. Determine the business impact of your problem.
3. Describe your problem and gather background information.

Then see “Submit the problem to IBM Software Support” on page x for information on contacting IBM Software Support.

Setting up a software maintenance contract

Set up a software maintenance 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, Tivoli, Lotus[®], and Rational[®] products, as well as IBM DB2[®] and IBM WebSphere[®] products that run on Microsoft[®] Windows[®] or UNIX[®] operating systems), enroll in IBM Passport Advantage[®] in one of the following ways:
 - **Online:** Go to the Passport Advantage Web page at <http://www.ibm.com/software/lotus/passportadvantage/>, click **How to enroll**, and follow the instructions.
 - **By Phone:** For the phone 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**.
- For server software products, you can purchase a software maintenance agreement by working directly with an IBM sales representative or an IBM Business Partner. For more information about support for server software products, go to the IBM Technical support advantage Web page at <http://www.ibm.com/servers/>.

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States. For a list of telephone numbers of people who provide support for your location, go to the Software Support Handbook page at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html>.

Determine the business impact

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you need to understand and assess the business impact of the problem you are reporting.

Severity 1	Critical business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.
Severity 2	Significant business impact: The program is usable but is severely limited.
Severity 3	Some business impact: The program is usable with less significant features (not critical to operations) unavailable.
Severity 4	Minimal business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem has been implemented.

Describe the problem and gather 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 recreated? 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 currently using a workaround for this problem? If so, be prepared to explain it when you report the problem.

Submit the problem to IBM Software Support

You can submit the problem to IBM Software Support online or by phone.

Online

Go to the IBM Software Support Web site at <http://www.ibm.com/software/support/probsub.html>. Enter your information into the appropriate problem submission tool.

By phone

For the phone number to call in your country, go to the contacts page of the IBM Software Support Handbook at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html>.

If the problem that you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. If a workaround is possible, IBM Software Support provides one for you to implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the Tivoli Storage Manager product support Web site at <http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliStorageManager.html>, so that users who experience the same problem can benefit from the same resolutions.

New for IBM Tivoli Storage Manager Version 6.1

Many features in the Tivoli Storage Manager Version 6.1 server are new for previous Tivoli Storage Manager users.

New for the server in Version 6.1.2

Server fix pack 6.1.2 contains several new features, in addition to fixes for problems.

Enabled functions

Functions that were disabled in Tivoli Storage Manager V6.1.0 and V6.1.1 are now enabled in Version 6.1.2.

Until Tivoli Storage Manager V6.1.2, a database that contained backup sets or tables of contents (TOCs) could not be upgraded to V6. These restrictions no longer exist.

In addition, the following commands have been enabled in Version 6.1.2:

- BACKUP NAS client command if the TOC parameter specifies PREFERRED or YES
- BACKUP NODE if the TOC parameter specifies PREFERRED or YES
- DEFINE BACKUPSET
- GENERATE BACKUPSET
- GENERATE BACKUPSETTOC

Licensing changes

Following the release of Tivoli Storage Manager Version 6.1.2, Tivoli Storage Manager Version 6.1.0 will no longer be available for download or purchase. Due to this unique circumstance, certain 6.1.2 packages will be available with a license module. See the following information for details on how this situation affects your environment.

Existing Version 6.1.0 and 6.1.1 users

If you have installed version 6.1.0 and are using a version 6.1.0 license, you can download the 6.1.2 package from the Service FTP site. You can install the 6.1.2 package using the instructions in Installing a Tivoli Storage Manager fix pack.

Version 5 users

If you have not yet installed a version of the V6.1 server, when you upgrade, you must upgrade directly to version 6.1.2. Version 6.1.2 is available with a license module from Passport Advantage or from your Tivoli Storage Manager sales representative. You can upgrade from V5 to V6.1.2 using the instructions in Upgrading the server.

New users

Version 6.1.2 is available from Passport Advantage or from your Tivoli Storage Manager sales representative. You can install version 6.1.2 using the instructions in Installing Tivoli Storage Manager.

ACSLS functionality for 64-bit Windows systems

Windows

Tivoli Storage Manager Version 6.1.0 requires the installation of StorageTek Library Attach software to utilize Sun StorageTek Automated Cartridge System Library Software (ACSLS) functions for the Windows operating system.

Support for ACSLS library functions is now available for both 32-bit and 64-bit Windows operating systems in fix pack level 6.1.2.

TSMDLST utility for AIX

AIX

The `tsmdlst` utility is now available for AIX® beginning in Tivoli Storage Manager fix pack level 6.1.2.

Use the `tsmdlst` utility to obtain information for AIX operating systems about medium-changer, tape, and optical devices controlled by the Tivoli Storage Manager device driver. With this utility you can obtain usage information, device names, serial numbers, and other device information.

Device support for AIX

AIX

The number of devices per driver is extended for AIX.

The Tivoli Storage Manager device driver can configure 1024 devices for each driver.

PREVIEW parameter for DSMSERV INSERTDB

A **PREVIEW** parameter is available for the `DSMSERV INSERTDB` utility in Tivoli Storage Manager fix pack level 6.1.2. The `DSMSERV INSERTDB` utility is used only as part of the process for upgrading a V5 Tivoli Storage Manager server to V6.1.

When you use the **PREVIEW=YES** parameter, the operation includes all the 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 before you run the actual upgrade process for your server.

Replacement passthru device driver for HP-UX 11i 3 (IA64 only)

HP-UX

The V6.1.2 passthru driver for HP-UX 11i v3 IA64 replaces the V6.1 passthru driver for HP-UX 11i v3 IA64. The passthru driver in V6.1.2 lets you configure 32 LUNs for each port.

If you are using the passthru driver for the V6.1 version level, reconfigure existing devices using the `autoconf` configuration script after installing the Tivoli Storage Manager server. You must also load the `esctl`, `estape`, `eschgr`, and `esdisk` drivers

into the HP-UX kernel. The passthru driver is packaged as part of the Tivoli Storage Manager server.

SAN discovery support for Linux on zSeries

Linux

Support for SAN discovery functions is now available for the Linux on zSeries operating system beginning in Tivoli Storage Manager fix pack level 6.1.2.

New server option SANDISCOVERYTIMEOUT

AIX

HP-UX

Linux

Solaris

A new server option is available for SAN discovery functions beginning in Tivoli Storage Manager fix pack level 6.1.2.

The SANDISCOVERYTIMEOUT option specifies the amount of time allowed for host bus adapters to respond when they are queried by the SAN discovery process.

New for the server in Version 6.1.0

Tivoli Storage Manager server version 6.1.0 contains many new features and changes.

Disabled functions in 6.1.0 and 6.1.1

Some functions have been disabled in Tivoli Storage Manager 6.1.0 and 6.1.1.

Note: The restrictions described here have been removed in Tivoli Storage Manager V6.1.2. If your server is at the V6.1.0 or V6.1.1 level, migrate to V6.1.2 to enable these functions.

A database containing backup sets or tables of contents (TOCs) cannot be upgraded to V6.1.0 or 6.1.1. The database upgrade utilities check for defined backup sets and existing TOCs. If either exists, the upgrade stops and a message is issued saying that the upgrade is not possible at the time. In addition, any operation on a V6.1 server that tries to create or load a TOC fails.

When support is restored by a future V6.1 fix pack, the database upgrade and all backup set and TOC operations will be fully enabled.

In the meantime, the following commands have been disabled:

- BACKUP NAS client command if the TOC parameter specifies PREFERRED or YES
- BACKUP NODE if the TOC parameter specifies PREFERRED or YES
- DEFINE BACKUPSET
- GENERATE BACKUPSET
- GENERATE BACKUPSETTOC

Changes to the Version 6.1 Administration Center

Many features in the Tivoli Storage Manager Administration Center Version 6.1 are new for previous users.

Updated Integrated Solutions Console

In V6.1, the Administration Center is hosted by the IBM Integrated Solutions Console (ISC) Advanced Edition Version 7.1. After installation of the Integrated Solutions Console installation completes, open a Web browser and enter the following URL, which will display the logon screen for the Integrated Solutions Console: `https://local_host:9043/ibm/console`. This screen indicates a successful installation of the Integrated Solutions Console.

To learn about console updates:

1. Start the ISC.
2. Click **Help** in the ISC banner.
3. In the Help navigation tree, click **Console Updates**.

Windows

WebSphere Windows service

In V6.1, the WebSphere Windows service is named TSM Administration Center - TsmAC.

Identify managing servers

The table of servers that is the hub of the enterprise-management work page has a column that identifies the managing server, if one exists, for each listed server. By sorting or filtering on the column, you can display the set of servers that are managed by a given server.

Hover help for table links

The Administration Center typically displays Tivoli Storage Manager objects in a table. In V6.1, when the cursor hovers over an object image, hover-help text is displayed. The hover help identifies the default action that results when you click the link that is associated with the object.

Links to information about server messages and Administration Center messages

When a problem or issue occurs with the server or Administration Center, you are immediately notified and provided with a brief message about the problem or issue. The message number is also provided. In V6.1, you can obtain detailed information about a message by clicking the link that is associated with the message number. The information is displayed in a new browser window.

Maintenance script enhancements

Tivoli Storage Manager utilizes a maintenance script to perform scheduled maintenance tasks. In V6.1, you can generate a maintenance script in one of two styles: predefined and custom.

A *predefined* maintenance script is one that is generated through a wizard. This script contains standard commands that cannot be altered. A predefined script can only be modified in the wizard.

A *custom* maintenance script is created using the Administration Center maintenance script editor. To have more control of your maintenance tasks, you can modify the commands that you specify. You can also use the editor to update your custom maintenance script.

Client nodes and backup sets enhancements

The redesigned Administration Center displays information about backup sets, client nodes, and client-node groups in one portlet. The design includes search functions that you can use to find and display information more quickly. When you select a client node, a summary panel is displayed with the current operation status, server actions, and client-node actions.

The work item **Client nodes and backup sets** appears in the ISC navigation tree.

Session and process information available in the health monitor

The Administration Center health monitor now includes information about server processes and sessions. The information is also available in the properties notebooks for servers.

Centralized server-connection management

In V6.1, server-connection tasks, such as adding a server connection, changing a password, and creating a server instance, are consolidated in a single location: the **Manage Servers** work item, located in the ISC navigation tree.

With actions available in this work item, you can quickly upload server-connection information to the Administration Center using an XML file. This file can optionally include a set of server credentials for multiple servers. To help create an XML file, you can download a list of server connections, without the credential information.

Changes to management-class activation

In the V6.1, Tivoli Storage Manager no longer activates changes to existing management classes automatically. You must activate the changes manually. Before the changes take effect, they are validated. Results of the validation are displayed. You or another administrator can review them, and then either confirm or cancel the activation.

Because changes are manually activated, you can prepare the management class in advance and activate the changes at an appropriate time.

Data deduplication

Data deduplication is a method of eliminating redundant data in sequential-access disk (FILE) primary, copy, and active-data storage pools. One unique instance of the data is retained on storage media, and redundant data is replaced with a pointer to the unique data copy. The goal of deduplication is to reduce the overall amount of time that is required to retrieve data by letting you store more data on disk, rather than on tape.

Data deduplication in Tivoli Storage Manager is a two-phase process. In the first phase, duplicate data is identified. During the second phase, duplicate data is removed by certain server processes, such as reclamation processing of storage-pool volumes. By default, a duplicate-identification process begins automatically after you define a storage pool for deduplication. (If you specify a duplicate-identification process when you update a storage pool, it also starts automatically.) Because duplication identification requires extra disk I/O and CPU resources, Tivoli Storage Manager lets you control when identification begins as well as the number and duration of processes.

You can deduplicate any type of data except encrypted data. You can deduplicate client backup and archive data, Tivoli Data Protection data, and so on. Tivoli Storage Manager can deduplicate whole files as well as files that are members of an aggregate. You can deduplicate data that has already been stored. No additional backup, archive, or migration is required.

For optimal efficiency when deduplicating, upgrade to the version 6.1 backup-archive client.

Restriction: You can use the data-deduplication feature with Tivoli Storage Manager Extended Edition only.

Storage devices

New device support and other changes to storage devices are available in Tivoli Storage Manager Version 6.1.

ACSLS functionality for Windows systems

Windows

Tivoli Storage Manager Version 6.1.0 requires the installation of StorageTek Library Attach software to utilize Sun StorageTek Automated Cartridge System Library Software (ACSLS) functions for the Windows operating system.

Support for ACSLS library functions is only available on 32-bit Windows operating systems in version 6.1.0.

Passthru device driver for HP-UX 11i v2 and v3 on the IA64 architecture

HP-UX

The HP-UX passthru device driver replaces the Tivoli Storage Manager device driver tsm SCSI and is packaged as part of the Tivoli Storage Manager server. The passthru driver can be used with versions 2 and 3 of the HP-UX 11i operating system.

If you are running either of these versions, reconfigure existing devices using the autoconf configuration script after installing Tivoli Storage Manager.

Tips:

- The passthru driver, Tivoli Storage Manager server, and storage agent packages are available in 64-bit mode only.
- In V6.1, the passthru driver lets you configure eight LUNs for each port. The passthru driver in V6.1.2 lets you configure 32 LUNs for each port.

Support for HP and Quantum DAT160 drives and media

AIX

HP-UX

Linux

Solaris

Windows

With Tivoli Storage Manager, you can now use HP and Quantum DAT160 (DDS6) tape drives and media. New recording formats are available for the 4MM device type.

Support for Sun StorageTek T10000 drives, T10000B drives, and T10000 media

With Tivoli Storage Manager, you can now use Sun StorageTek T10000 drives, T10000B drives, and T10000 media. New recording formats are available for the ECARTRIDGE device type. Tivoli Storage Manager supports Volsafe media with the Sun StorageTek T10000 and T10000B drives.

Disaster recovery manager support for active-data pools

To restore your client systems more quickly and efficiently, you can now use active-data pools in your recovery plans and procedures.

Active-data pools are storage pools that contain only active versions of client backup data. Like copy storage pool volumes, disaster recovery manager lets you:

- Specify the names of active-data pool volumes to be managed by the disaster recovery manager.
- Recycle on-site and off-site active-data pool volumes according to server policies and processes.
- Include active-data pool volumes in the scripts, macros, and documentation that is part of the recovery plan file.
- Track and manage active-data pool media as required by your operations.

By default, active-data pools are not eligible for processing at the time of installation. Copy storage pools, on the other hand, are processed at installation time even if you have not explicitly specified a copy storage pool or pools to be managed.

EXPIRE INVENTORY command enhancements

The EXPIRE INVENTORY command is now enhanced with new functionality.

The additional parameters that you can now use are NODE, DOMAIN, TYPE, DURATION, AND RESOURCE. You can use these parameters to target specific client nodes and domains, and also to determine the type of data to be processed. You can use the RESOURCE parameter to specify the number of parallel processes that you want to run within the single EXPIRE INVENTORY process. You can run up to ten threads at one time, but if you are processing one node, only one thread is utilized.

No-query restore changes

The no-query restore (NQR) function and the internal algorithms responsible for NQR were changed to take advantage of DB2 capabilities and to improve performance.

The NQR function has been rewritten to resolve a performance problem encountered when restoring a small number of objects for a client file system with a large number of backup objects spread across a large number of Tivoli Storage Manager server storage pool volumes. NQR performance is now comparable to that of the classic restore under these conditions. NQR now performs a volume determination phase that must be completed before any objects are restored from DISK, FILE, or tape storage volumes.

Server database

Tivoli Storage Manager version 6.1 provides a new server database. Advantages include automatic statistics collection and database reorganization, full-function SQL queries, and elimination of the need for offline audits of the database.

Upgrading to V6.1 requires that data in a current Tivoli Storage Manager server database be extracted and then inserted into the new database structure. Tivoli Storage Manager provides utilities to perform the process.

Support for NetApp SnapMirror to Tape feature

AIX

HP-UX

Linux

Solaris

Windows

With Tivoli Storage Manager you can create SnapMirror to Tape images of file systems on NetApp file servers.

SnapMirror to Tape provides an alternative method for backing up very large NetApp file systems. Because this backup method has limitations, use this method when copying very large NetApp file systems to secondary storage for disaster recovery purposes.

Reporting and monitoring feature

The reporting and monitoring feature uses a combination of the Tivoli Common Reporting tool, IBM Tivoli Monitoring, and the IBM Tivoli Data Warehouse to offer you reports and real time monitoring information about Tivoli Storage Manager servers and client activity.

HP-UX

Installing the Tivoli Storage Manager reporting and monitoring feature directly on a Tivoli Storage Manager HP-UX server is not supported. You can monitor and report on HP-UX Tivoli Storage Manager servers by creating a monitoring agent instance for these servers on an AIX, Linux®, or Windows IBM Tivoli Monitoring server.

Solaris

Installing the Tivoli Storage Manager reporting and monitoring feature directly on a Tivoli Storage Manager Sun Solaris server is not supported. You can monitor and report on Sun Solaris Tivoli Storage Manager servers by creating a monitoring agent instance for these servers on an AIX, Linux, or Windows IBM Tivoli Monitoring server.

ODBC driver support

Tivoli Storage Manager Version 6.1 uses the DB2[®] open database connectivity (ODBC) driver to query the database and display the results.

The Tivoli Storage Manager ODBC driver is no longer supported with the server.

Backup sets and client node enhancements

The Administration Center now displays the backup sets, client nodes, and client node groups from one portlet.

You can view all of the client nodes, view them by server, or search for a node from the three available client node tabs. The All Client Nodes tab lists all of the nodes and has a Filter feature to help in your search. The filter works differently than the other table filters in the Administration Center, in that here you do not have to press the Enter key to get results. The search is initiated when you enter a text character in the filter field. As you add characters, the results are filtered even more.

When you select a client node, a summary panel is displayed with the current operation status, server actions, and client node actions. You can also access the Server Actions by right-clicking on the selected row.

The Search tab lets you refine your search parameters to include the server name, the client node name, policy domain name, and other fields that are available.

In the Client Node Groups section, you can find a client node group from the All Client Node Groups tab or from the By Server tab. You can use the filter and the right-click menu on these pages also.

Backup sets are found in the Backup Set Collections section. Search a server by selecting it and clicking **Update Table**.

Chapter 1. Overview of IBM Tivoli Storage Manager tuning

Tivoli Storage Manager performance can be influenced by various factors. Tuning for optimal performance requires care and expertise.

Tuning Tivoli Storage Manager can be complex because of the many operating systems, network configurations, and storage devices that Tivoli Storage Manager supports. Performance tuning even for a single platform function is complex. The factors that can affect performance and have significant effects include:

- Average client file size
- Percentage of files changed since last incremental backup
- Percentage of bytes changed since last incremental backup
- Server hardware (CPUs, RAM, disk drives, network adapters)
- Server storage pool devices (disk, tape, optical)
- Server operating system
- Server activity (non-Tivoli Storage Manager workload)
- Client hardware (CPUs, RAM, disk drives, network adapters)
- Client operating system
- Client activity (non-Tivoli Storage Manager workload)
- Network hardware and configuration
- Network utilization
- Network reliability
- Communication protocols
- Communication protocol tuning
- Final output repository type (disk, tape, optical)

Chapter 2. IBM Tivoli Storage Manager server performance tuning

You can tune the performance of Tivoli Storage Manager servers through server options, commands, and other configuration settings.

See the *Administrator's Reference* to determine the commands and options that are available for your platform. The options are tunable on all Tivoli Storage Manager servers. You can change any option setting in the server options file (dsmserv.opt). If you change the server options file, you must stop and restart the server for the changes to take effect. You can change some settings with the server SETOPT command.

Tuning server options

Some server options can be tuned to improve Tivoli Storage Manager performance.

The following options are tunable on all Tivoli Storage Manager servers. You can change any option setting in server options file (dsmserv.opt). You can change some options with the server SETOPT command. If you change the server options file, you must halt and restart the server for the changes to take effect.

DBMEMPERCENT

DBMEMPERCENT Sets a limit on the percentage of the system memory that is used for the database manager.

By default, the percentage of the virtual address space that is dedicated to the database manager processes is set to 70 to 80 % of system RAM. To change this setting to a value from 10 to 99 %, modify the DBMEMPERCENT server option. Ensure that the value allows adequate memory for both the Tivoli Storage Manager server and any other applications that are running on the system. The default value is 'AUTO'.

It is generally not necessary to change this setting on a system that is dedicated to a single Tivoli Storage Manager server. If there are other applications that require significant amounts of memory on a system, changing this setting to an appropriate amount might reduce paging and improve system performance. For systems with multiple Tivoli Storage Manager servers, changing this setting for each server is recommended. For example, this could be set to 25% for each of three servers on a system. Each server could also have a different value for this setting, as appropriate for the workload on that server.

DISKSTGPOOLMEMSIZE

The DISKSTGPOOLMEMSIZE server option specifies the size of the cache that the server can use to manage operations for storage pools with the device type of DISK.

The more memory available, the less disk storage pool metadata must be retrieved from the database server. Performance might be improved during operations that store data into or delete data from disk storage pools.

The DISKSTGPOOLMEMSIZE server option specifies, in megabytes, the size of the memory available to manage disk storage pools. Each megabyte can manage 32 gigabytes of disk storage. This option should be large enough to accommodate the maximum amount of data expected to be stored in or deleted from disk storage pools per second.

For example, if a maximum of 96 gigabytes of data per second is expected to be stored in or deleted from disk storage pools, a size of 3 is recommended. If this option is not specified, it defaults to 80, which can manage 2560 gigabytes of disk storage. For 32-bit servers, it defaults to 20 which can manage 640 gigabytes of disk storage.

EXPINTERVAL

Inventory expiration removes client backup and archive file copies from the server. EXPINTERVAL specifies the interval, in hours, between automatic inventory expirations run by the Tivoli Storage Manager server. The default is 24.

Backup and archive copy groups can specify the criteria that make copies of files eligible for deletion from data storage. However, even when a file becomes eligible for deletion, the file is not deleted until expiration processing occurs. If expiration processing does not occur periodically, storage pool space is not reclaimed from expired client files, and the Tivoli Storage Manager server requires increased disk storage space.

Expiration processing is CPU and database I/O intensive. If possible, it should be run when other Tivoli Storage Manager processes are not occurring. To enable this, set EXPINTERVAL to 0 and either schedule expiration to occur once each day, or manually start the process with the EXPIRE INVENTORY server command. Expiration processing can be scheduled by defining an administrative schedule.

When using the DURATION parameter on an administrative schedule, periodically check that expiration is actually completing within the specified time.

This is the recommended setting:

```
EXPINTERVAL 0
```

This setting specifies that there is no expiration processing. Use an administrative schedule to run expiration at an appropriate time each day.

MAXSESSIONS

The MAXSESSIONS option specifies the maximum number of simultaneous client sessions that can connect with the Tivoli Storage Manager server.

The default value is 25 client sessions. The minimum value is 2 client sessions. The maximum value is limited only by available virtual memory or communication resources. By limiting the number of clients, server performance can be improved, but the availability of Tivoli Storage Manager services to the clients is reduced.

A typical production Tivoli Storage Manager server could have the MAXSESSIONS parameter set to 100 or greater.

MOVEBATCHSIZE and MOVESIZETHRESH

The MOVEBATCHSIZE and MOVESIZETHRESH options tune the performance of server processes that involve the movement of data between storage media. These processes include storage pool backup and restore, migration, reclamation, and move data operations.

MOVEBATCHSIZE specifies the number of files to be moved and grouped together in a batch, within the same server transaction. The default value for MOVEBATCHSIZE is 1000; and the maximum value is also 1000. 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. The default value for MOVESIZETHRESH is 2048; and the maximum value is also 2048.

The number of client files moved for each server database transaction during a server storage pool backup/restore, migration, reclamation or move data operation are determined by the number and size of the files in the batch. If the number of files in the batch equals the MOVEBATCHSIZE before the cumulative size of the files becomes greater than the MOVESIZETHRESH, then the MOVEBATCHSIZE is used to determine the number of files moved or copied in the transaction. If the cumulative size of files being gathered for a move or copy operation exceeds the MOVESIZETHRESH value before the number of files becomes equivalent to the MOVEBATCHSIZE, then the MOVESIZETHRESH value is used to determine the number of files moved or copied in the transaction.

When the MOVEBATCHSIZE or MOVESIZETHRESH parameters are increased from their default values, the server requires more space in the recovery log. The recovery log might require an allocation of space two or more times larger than a recovery log size which uses the defaults. In addition, the server requires a longer initialization time at startup. The impact of a larger recovery log size is felt while running the server with the log mode set to NORMAL (the default value). If you choose to increase these values for performance reasons, be sure to monitor recovery log usage during the first few storage pool backup/restore, migration, reclamation, or move data executions to ensure sufficient recovery log space is available.

These are the recommended settings:

```
MOVEBATCHSIZE 1000
MOVESIZETHRESH 2048
```

RESTOREINTERVAL

The RESTOREINTERVAL option specifies how long, in minutes, that a restartable restore session can be in the database before it can be expired. Restartable restores allow restores to continue after an interruption without starting at the beginning.

Restartable restores reduce duplicate effort or manual determination of where a restore process was terminated. The RESTOREINTERVAL option defines the amount of time an interrupted restore can remain in the restartable state.

The minimum value is 0. The maximum is 10080 (one week). The default is 1440 (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. Restartable restore sessions consume resources on the Tivoli Storage Manager server. You should not keep these sessions any longer than they are needed.

It is recommended that you tune the RESTOREINTERVAL option to your environment.

TCPNODELAY

The TCPNODELAY server option specifies whether the server allows data packets that are less than the maximum transmission unit (MTU) size to be sent out immediately over the network.

When TCPNODELAY is set to NO, the server buffers data packets that are less than the MTU size:

- Buffering can improve network utilization.
- Buffering requires a delay that can impact session throughput greatly.

When set to YES, it disables the TCP/IP Nagle algorithm, which allows data packets less than the MTU size to be sent out immediately. Setting this option to YES might improve performance in higher speed networks. The default is YES.

This is the recommended setting:

```
TCPNODELAY YES
```

Note: This option also exists on the Tivoli Storage Manager client.

TCPWINDOWSIZE

The TCPWINDOWSIZE server option specifies the amount of receive data in kilobytes that can be on a TCP/IP connection at one time. The TCPWINDOWSIZE server option applies to backups and archives. The TCPWINDOWSIZE client option applies to restores and retrieves.

The sending host cannot send more data until an acknowledgement and TCP receive window update are received. Each TCP packet contains the advertised TCP receive window on the connection. A larger window allows the sender to continue sending data, and might improve communication performance, especially on fast networks with high latency.

The TCPWINDOWSIZE option overrides the operating system's TCP send and receive spaces. In AIX, for instance, these parameters are `tcp_sendspace` and `tcp_recvspace` and are set as "no" options. For Tivoli Storage Manager, the default is 63 KB, and the maximum is 2048 KB. Specifying TCPWINDOWSIZE 0, results in Tivoli Storage Manager using the operating system default. This is not

recommended because the optimal setting for Tivoli Storage Manager might not be same as the optimal setting for other applications.

The TCPWINDOWSIZE option specifies the size of the TCP sliding window for all clients and all servers. On the server this applies to all sessions. Therefore, raising TCPWINDOWSIZE can increase memory significantly when there are multiple, concurrent sessions. A larger window size can improve communication performance, but uses more memory. It enables multiple frames to be sent before an acknowledgment is obtained from the receiver. If long transmission delays are being observed, increasing the TCPWINDOWSIZE might improve throughput.

For all platforms, rfc1323 must be set to have window sizes larger than 64 KB-1:

- **AIX:** Use `no -o rfc1323=1`
- **HP-UX:** Using a window size greater than 64 KB-1 automatically enables large window support.
- **Sun Solaris 10:** Use `"ndd "set /dev/tcp tcp_wscales_always 1"` This should be enabled by default.
- **Linux:** Should be on by default for recent kernel levels. Check with `"cat /proc/sys/net/ipv4/tcp_window_scaling"`. Recent Linux kernels use autotuning, and changing TCP values might have a negative effect on autotuning. Make changes with caution.
- **Windows XP and 2003:** Add or modify, with `regedit`, the following registry name/value pair under [HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters] `Tcp1323Opts, REG_DWORD, 3`

Attention: Before modifying the registry name and value pair, you should back up the entire registry.

These are the recommended settings:

```
TCPWINDOWSIZE 63
```

Or for a Gigabit Ethernet with Jumbo Frames – 9000 MTU:

```
TCPWINDOWSIZE 128
```

Note: This option is also on the Tivoli Storage Manager client.

TXNGROUPMAX

The TXNGROUPMAX server option specifies the number of objects that are transferred between a client and server in a single transaction. The minimum value is 4 objects, and the maximum value is 65,000 objects. The default value has now been set 4096 objects. An object is a file or directory.

It is possible to affect the performance of client backup, archive, restore, and retrieve operations by using a larger value for this option:

1. 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.
2. 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. For more information on managing the recovery log, see the *Administrator's Guide*.

3. A larger value of the TXNGROUPMAX option can also increase the number of objects that must be resent if the transaction is aborted 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 that must be resent.
4. Increasing the TXNGROUPMAX value affects the responsiveness of stopping the operation, and the client might 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 the REGISTER NODE and UPDATE NODE commands.

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.

Set TXNGROUPMAX to 256 in your server options file. Settings higher than 4096 typically provide no benefit. If some clients have small files and go straight to a tape storage pool, then raising TXNGROUPMAX to the higher value can benefit those clients.

Server hardware recommendations

The Tivoli Storage Manager Version 6.1 server has greater scalability and makes much greater use of additional server hardware than previous versions.

Additional processors and memory can be especially useful for reducing operational windows for incremental backups, archiving, inventory expiration, and storage pool reclamation. New functions such as deduplication can require additional processing capacity. Here are some recommended minimums:

Processors

At least one processor core per concurrent inventory expiration or deduplication identify process.

Memory

At least 4 GB of memory for each production Tivoli Storage Manager server on a system. However, 8 GB or more might be useful for improving the performance of certain operations, particularly incremental backups. For systems with multiple Tivoli Storage Manager servers, see “DBMEMPERCENT” on page 3 for tuning each server’s memory limits.

I/O adapters

Server performance depends on the system I/O throughput capacity for both local area networks (LAN) and storage (SCSI, fibre channel).

Database manager for IBM Tivoli Storage Manager

The Tivoli Storage Manager Version 6.1 introduces a new database manager that automatically manages the space available to the directories as database space.

You create and designate directories that the server can use for the database. Space allocation is done automatically in the directory locations specified for the database. If the server needs additional space, you can add directories by using the EXTEND DBSPACE command. You cannot adjust the assigned capacity of the database. Space allocation is done automatically in the directory locations specified for the database. You can specify a log mirror to protect the active log.

Locate the database on fast, reliable storage, such as Redundant Array of Independent Disks hardware. Ensure that the directories are not placed on file systems that might run out of space. During the successful installation of a Tivoli Storage Manager server, these steps are completed:

- The database is allocated and initialized.
- The active log is allocated and initialized.
- The archive log is allocated and initialized.

In addition, it is strongly recommended that an active log mirror and an archive log be allocated and initialized. Mirroring of the database is not supported by Tivoli Storage Manager.

The database manager determines when database backups are required and automatically initiates them. The device class used for the database backups is set by using the SET DBRECOVERY command.

The server does not track individual database volumes. The server uses its own fixed-size recovery log buffer pool and adjusts the value of the buffer pool size dynamically.

Note: Raw logical volumes are no longer supported for the database and logs.

Database and log performance

The Tivoli Storage Manager database, active log, and archive log have different I/O operations and must be considered separately for tuning.

For most operations the database access pattern is random. Therefore, it is best to use the fastest available disk to support database operations. Also, write caching should be enabled on the disk volumes that the database and logs reside on, but only if the cache is nonvolatile and can survive unexpected power outages and other failure modes.

It is useful to gauge the performance of a database volume by measuring the I/O per second (IOPS) rate. This works best when only one database volume is allocated per physical volume. A typical fiber channel attached disk can handle roughly 150 IOPS before queuing occurs. Some disks, such as SATA attached disk, can handle fewer IOPS, and some disks might be able to handle more. If a physical disk approaches the 150 IOPS mark, add additional physical disks and additional database volumes to the database configuration. Operating system tools, such as iostat or filemon, are available to measure IOPS for physical volumes. The Tivoli Storage Manager server instrumentation can be used to measure IOPS for a database volume.

It is best to use multiple directories for the database, with up to 4 or 8 directories for a large Tivoli Storage Manager database. Each database directory should be located on a disk volume that uses separate physical disks from other database directories. The Tivoli Storage Manager server database I/O workload is spread over all directories, thus increasing the read and write I/O performance. A large number of small capacity physical disks is better than a small number of large capacity physical disks with the same rotation speed.

The access pattern for the active log is always sequential. Physical placement on the underlying disk is very important. It is best to isolate the log from the database and from the disk storage pools. If this cannot be done, then place the log with storage pools and not with the database.

Active log mirroring provides higher reliability, but comes at a cost in performance. Locate the mirror log directory on a disk volume that uses separate physical disks from the active log by using the `MIRRORLOGDIR` parameter in the `DSMSERV FORMAT` command. After installation, change the mirror log directory location by changing the value of the `MIRRORLOGDIR` option in the server option file and restarting the server.

Database manager tuning

Proper tuning of the database for Tivoli Storage Manager is essential for optimal performance. The Tivoli Storage Manager V6.1 database manager enables self-tuning memory.

Key areas to address when determining your Tivoli Storage Manager database manager configuration are:

- The database manager (instance-level) configuration parameters
- Database configuration parameters that have the largest impact on the performance of the system
- Maintenance operations that you must perform to ensure that the system continues to perform optimally once you have it running

Configuration parameters

The primary considerations for the Tivoli Storage Manager database are: enough memory for the database manager server, configuration of enough physical disks to handle the I/O requirements, and enough CPUs to handle the workload.

Self tuning memory

For the best database manager operation remember, that the data and indexes are manipulated in the database buffer pools allocated in memory. Performance is stifled by a lot of paging operations required when there is more buffer pool space defined than real memory. Beside the buffer pools which use the most memory, the sort list, the lock list and the package cache are other memory segments allocated by the database manager.

The database manager supported with Tivoli Storage Manager version 6.1 enables self tuning memory by default. This automatically samples the workload and performance characteristics of the database. Using this feature the database manager adapts the sort heap, lock list, package cache, buffer pool and total database memory improving performance and through put as your environment requires. Tivoli Storage Manager mixes the database workload between transactions recording backups to heavy query usage during restore operations.

Enough disk to handle I/O

Efficient performance relies on having enough physical disk drives to service the throughput required for the workload. The correct relationship of physical disks to CPU in the database manager server helps to maintain good performance. One of the most CPU and database intensive parts of the Tivoli Storage Manager server workload is inventory expiration. There should be a ratio of one database directory, array, or LUN for each inventory expiration process.

CPUs to handle the workload

The power of the database manager system depends on the number and speed of its CPUs. For a balanced system, under normal operations only consume approximately 80% of the available CPUs.

Backup performance

When possible, limit the number of versions of any backup file to the minimum required.

File backup performance is degraded when there are many versions of an object. Use the `DEFINE COPYGROUP` command and modify the `VEREXISTS` parameter to control the number of versions, or use the `UPDATE COPYGROUP` command. The default number of backup versions is 2.

If the retention requirements in your environment differ among client machines, use different copy groups rather than taking the lowest common denominator. For example, if your accounting machines require records to be kept for seven years, but other machines need data kept for only two years, do not specify seven for all. Instead, create two separate copy groups. Not only are backups potentially faster, but you also consume less storage because you are not keeping data that you do not need.

Tuning inventory expiration

You can improve inventory expiration by limiting the scope of the data to be expired and by specifying the number of threads that can be run in parallel.

In the `EXPIRE INVENTORY` command you can specify parameters to target specific client nodes and domains, and also to determine the type of data to be processed. For example, you can run expiration on a specific set of nodes if policy changes had been made recently that would cause a lot of the data for these nodes to be eligible to be expired.

In addition, you can use the `RESOURCE` parameter to specify the number of parallel processes that you want to run during a single `EXPIRE INVENTORY` process. You can run up to ten threads at one time, but if you are processing one node, only one thread is utilized. The workload is distributed across specified number of threads. You can run up to 10 expiration processes in parallel. The resources represent parallel work by the server within the single expiration process; expiration still runs as a single process. For example, if you specify 3 nodes and a 3 threads or more, then expiration processing for the three client nodes runs in parallel. If you specify 3 nodes and 5 threads, then expiration processing for the three nodes runs in parallel, and the two extra resources are ignored.

See “EXPINTERVAL” on page 4 for information about the EXPINTERVAL server option.

Disaster recovery performance

Avoid use of export and import operations for disaster recovery.

Tivoli Storage Manager provides procedures for backing up and restoring the storage pools for disaster recovery. The performance of storage pool backup and recovery for disaster recovery is superior to the performance of export and import.

Searching the server activity log

When performance problems occur, an abnormal system condition is often the cause.

You might determine the cause of these problems by examining server activity logs, client error files, or the appropriate system logs for your operating system.

Scheduling sessions and processes

If possible, schedule server processes such as expiration, storage pool backup, move data, export and import operations, and so on, when client backups are not active.

Tivoli Storage Manager throughput can degrade if all client backups are started simultaneously. It is better to spread out backup session start-ups over time. Creating several client schedules with staggered start times and assign nodes to those schedules appropriately. For nodes that use the client polling method of scheduling, use the SET RANDOMIZE command to spread out the nodes startup times.

LAN-free backup

Using LAN-free backup can improve performance. To do so requires the Tivoli Storage Manager storage agent on the client for LAN-free backups to SAN attached tape, and Tivoli SANergy[®] if backups are to be sent to FILE volumes on SAN-attached disk.

- Back up and restore to tape or disk using the SAN. The advantages are:
 - Metadata is sent to the server using the LAN while client data is sent over the SAN.
 - Frees the Tivoli Storage Manager server from handling data leading to better scalability.
 - Potentially faster than LAN backup and restore.
 - Better for large file workloads, databases (Data Protection).
 - Small file workloads have bottlenecks other than data movement.
- Ensure that there are sufficient data paths to tape drives.
- Do not use LAN-free backup if you bundle more than 20 separate dsmc commands in a script.
 - dsmc start/stop overhead is higher due to tape mounts.
 - Use the new file list feature to back up a list of files.

Maximum number of mount points for a node

On the REGISTER and UPDATE NODE commands, the MAXNUMMP parameter specifies the maximum number of mount points a node is allowed to use on the server.

The MAXNUMMP parameter value can be from 0-999. Zero specifies that the node cannot acquire any mount point for a backup or archive operation. However, the server still allows the node to use a mount point for a restore or retrieve operation. If the client stores its data in a storage pool that has copy storage pools defined for simultaneous backup, the client might require additional mount points.

As a general rule, assign one mount point for each copy storage pool of sequential device type. If the primary storage pool is of sequential device type, assign a mount point for the primary storage pool also.

Managing storage pools and volumes

Logical storage pools and storage volumes are the principal components in the Tivoli Storage Manager model of data storage. By manipulating the properties of these objects, you can optimize the usage of storage devices.

When you configure devices so that the server can use them to store client data, you create storage pools and storage volumes. The procedures for configuring devices use the set of defaults that provides for storage pools and volumes. The defaults can work well. However, you might have specific requirements not met by the defaults. There are three common reasons to change the defaults:

1. Optimize and control storage device usage by arranging the storage hierarchy and tuning migration through the hierarchy (next storage pool, migration thresholds).
2. Reuse tape volumes through reclamation.
3. Reuse is also related to policy and expiration. Keep a client's files on a minimal number of volumes (collocation).

Cached disk storage pools

Using cached disk storage pools can increase restore performance by avoiding tape mounts. Cached disk storage pool does not refer to cache in the disk hardware, or file system cache in the operating system.

The benefit of using cached disk storage pools is seen for restoring files that were recently backed up. If the disk pool is large enough to hold a day's worth of data, then caching is a good option. Use the DEFINE STGPOOL or UPDATE STGPOOL command with the CACHE=YES parameter to enable caching. However, when the storage pool is large and CACHE is set to YES, the storage pool might become fragmented and response will suffer. If this condition is suspected, our recommendation is to turn disk storage pool caching off. Also, disk caching can affect backup throughput because database updates are required to delete cached files in order to create space for the backup files.

Tuning storage pool migration

You can improve performance by tuning storage pool migration processes and thresholds.

Tuning migration processes

Use the `DEFINE STGPOOL` command and modify the `MIGPROCESS` parameter to control the number of migration processes that are used for migrating files from a disk storage pool to a tape storage pool.

When data is migrated from disk to tape, multiple processes can be used if multiple tape drives are available. In some cases, this can improve the time to empty the disk storage volumes, since each migration process works on data for different client nodes. For the `MIGPROCESS` option, you can specify an integer from 1-999, inclusive, but it should not exceed the number of tape drives available. The default value is 1. You can also use the `UPDATE STGPOOL` command to modify the number of migration processes.

Restoration of deduplicated data

Restore operations from a sequential-access disk (`FILE`) storage pool that is set up for deduplication have different performance characteristics than restore operations from a `FILE` storage pool that is not set up for deduplication.

In a `FILE` storage pool that is not set up for deduplication, files are typically restored in a mainly sequentially process. In a `FILE` storage pool that is set up for deduplication, however, data is distributed throughout the storage pool. As a result, the I/O is more random, which can lead to slower restore times. This behavior occurs more often with small (less than 100 KB) files. In addition, more server processor resources are consumed when restoring from a deduplicated storage pool. This occurs because the data is checked to ensure that it has been reassembled properly using MD5 algorithms.

Although restore operations of small files from a deduplicated storage pool might be relatively slow, these operations are still typically faster than restore operations of small files from tape because of the added tape mount and locate time.

Tuning migration thresholds

Use the `DEFINE STGPOOL` command and modify the `HIGHMIG` and `LOWMIG` parameters to tune migration thresholds. If the thresholds are set too high, migration is delayed.

This can cause the disk storage pool to fill, and when a client attempts to send data to the disk storage pool, it sees the full condition and attempts to go to the volume indicated at the next level in the storage hierarchy. If this is a tape pool, then all drives might be in use by a migration process, in which case the client session waits on the tape media to be freed by the migration process. The client then sits idle. In this case, the migration thresholds should be lowered so migration starts earlier, or more disk space should be allocated to the disk storage pool. You can also use the `UPDATE STGPOOL` command to modify the migration thresholds.

Improving storage agent performance

There are steps you can take to get the best performance from your Tivoli Storage Manager storage agents.

To get the best performance consider the following items when you set options for storage agents:

- The storage agent has its own configuration file, `dsmsta.opt`, containing many of the same options as the server `dsmserv.opt`. In general, use the same settings as recommended for the server.
- You can use the `DSMSTA SETSTORAGESERVER` command for some options.
- Ensure `TCPNODELAY` is set to `YES` (the default) on both the storage agent and server. Use the option `LANFREECOMMMETHOD SHAREDMMEM` in the client options file on client platforms that support it to obtain the lowest CPU usage.

Modifying the IBM Tivoli Monitoring environment file for reporting performance

Using the environment file that was automatically created for you when you added a Tivoli Storage Manager monitoring agent instance, you can improve reporting performance by modifying the environment variables.

The Windows environment file is named `KSKENV_xxx`, where `xxx` is the instance name of the monitoring agent you created. This file is located in the IBM® Tivoli Monitoring installation directory. (for example: `\IBM\ITM\TMAITM6`)

The AIX and Linux environment files are named `sk_xxx.config`, where `xxx` is the instance name of the monitoring agent you created. This file is located in the `/opt/tivoli/tsm/reporting/itm/config` directory on both Linux and AIX systems

The following list contains the environment variables that you can change to modify the performance for your monitoring agent. Use any text editor to edit the file.

***KSK_PREFETCH_MINUTES*, Default value=30**

The delays in minutes that the Tivoli Storage Manager server is queried for certain attribute groups. You can modify this variable in the following ways:

- Make this value larger to reduce the number of times that the agent queries the Tivoli Storage Manager server over a 24 hour period
- Reduce the value to increase the frequency of Tivoli Storage Manager server queries.

This value is for all attribute groups collected by the agent.

***KSK_PREFETCH_RETENTION_DAYS*, Default Value=2**

The number of days that the pre-fetch data is stored in the pre-fetch cache. The pre-fetch cache is for short-term storage of Tivoli Storage Manager data that is transferred later to the data warehouse. Two days is normally a sufficient amount of time for this variable.

***KSK_MAXIMUM_ROWS_RETURNED*, Default Value=2500**

The maximum number of rows that are returned at any one time to the IBM Tivoli Monitoring server. Changing this value could cause your Tivoli Enterprise Portal to receive so many rows of data that it will not be able to

display them. The speed of the processor and the amount of memory installed in the IBM Tivoli Monitoring server dictates the value of this variable.

Important: Do not increase this variable to a value greater than 3500 rows to prevent data overflow.

***KSK_APIHRLIMIT*, Default value=1**

The age, in hours, of the data that is collected by the Tivoli Storage Manager and common reporting agent. Do not increase this value unless you are running the agent on a very high performing server.

***KSK_APIBUFFER*, Default value=50 000**

The maximum number of rows that are retrieved from the Tivoli Storage Manager database at any one time. This value should be set to 50 000 rows or less. If the total number of rows defined by this value exceeds the total number of rows in the Tivoli Storage Manager database, no data is returned.

***KSK_APITIMEOUT*, Default value=480**

The amount of time, in minutes, before the Tivoli Storage Manager Administrators API times out.

In addition to the agent variables that you can change in the environment variable file, there are also two other values that are important when tuning the reporting servers. These values are modified in the Tivoli Enterprise Portal history configuration panels and they are:

***Warehouse interval*, Default value=daily**

Specifies how often the data that has been collected from IBM Tivoli Monitoring is sent to the Tivoli Data Warehouse for historical data storage. Possible values are hourly, daily, and off.

***Collection interval*, Default value=15**

Specifies the length of time between requests from IBM Tivoli Monitoring to the Tivoli Storage Manager Tivoli Common Reporting data collection agent. This value should be twice the value of the *KSK_PREFETCH_MINUTES* variable. Possible values are 1, 5, 15, 30 minutes, hourly, daily.

Tip: Keep the *KSK_MAXIMUM_ROWS_RETURNED*, and the *KSK_APIBUFFER* variables as low as possible to prevent data overflows on the Tivoli Storage Manager Tivoli common Reporting agent or the Tivoli Enterprise Portal.

Performance improvement activities for server platforms

The values for some server options vary depending on the platform.

Actions for better performance on all server platforms

There are some actions that you can take on most or all platforms to improve server performance.

The following actions will help you optimize Tivoli Storage Manager server performance for your environment.

- Format disk storage pool volumes sequentially, one at a time, if they are placed into the same file system. This creates files that have only a few gaps and improves sequential read and write performance.
- Optimize the transaction sizes for the Tivoli Storage Manager client/server and server maintenance functions. Refer to the appropriate options TXNGROUPMAX, TXNBYTELIMIT, MOVEBATCHSIZE, and MOVESIZETHRESH. A larger transaction size increases the size of the server file aggregates. File aggregation provides throughput improvements for many server data movement and inventory functions, such as storage pool migration, storage pool backup, and inventory expiration. A larger transaction size when using backup directly to tape reduces the number of tape buffer flushes and therefore can improve throughput.

AIX server performance

There are a number of actions that can improve performance for a Tivoli Storage Manager server running in an AIX environment.

- Use raw partitions for disk storage pool volumes on the AIX platform. Customer experience and measurements in the lab show that raw logical volumes offer better client backup/restore throughput and server administrative process performance.
- If you choose not to use raw logical volumes, always use JFS2 file systems. For JFS file systems, direct I/O might cause degradation, especially with large-file enabled file systems. JFS2 file systems generally provide better performance than JFS file systems.
- If you have new generation SP™ nodes, set the TCPWINDOWSIZE to 128 for both SP client and SP AIX server. This is especially true if you have ATM card in the SP machine. On the newer and faster SP machines (and faster HPS), TCPWINDOWSIZE 128 works well.

AIX: vmo and ioo commands

You can use the vmo and ioo commands to improve the performance of an AIX server.

The AIX Virtual Address space is managed by the Virtual Memory Manager (VMM). VMM is administered by the vmo AIX command; I/O that can be tuned is controlled by the ioo command. In AIX 5.3 and later, the vmo and ioo commands replace vmtune.

- The vmo command is used to tune the AIX virtual memory system.
- The vmo and ioo options can be displayed using the vmo -a and ioo -a commands.
- You can change options by specifying the appropriate option and value.
- Changes to vmo parameters do not survive reboot unless the -p option is specified.

Read ahead (iio maxpgahead)

When AIX detects sequential file reading is occurring, it can read ahead even though the application has not yet requested the data.

- Read ahead improves sequential read performance on JFS and JFS2 file systems.
- Tivoli Storage Manager client - Improves large file backup throughput.
- Tivoli Storage Manager server - Improves storage pool migration throughput on JFS volumes only (does not apply to JFS2 or raw logical volumes).
- The recommended setting of maxpgahead is 256 for both JFS and JFS2:

```
iio -p -o maxpgahead=256 -o j2_maxPageReadAhead=256
```

- When altering the read ahead parameter, you must also alter the maxfree parameter so that there is enough free memory to store the read ahead data.
- The following equation must be true:

$$\text{minfree} + \text{maxpgahead} \leq \text{maxfree}$$

To calculate minfree and maxfree use these formulae:

- minfree = 120 multiplied by the number of processors (or default if larger)
- maxfree = 120 + maxpgahead (or j2_maxPageReadAhead) multiplied by the number of processors (or the default if larger)

This does not improve read performance on raw logical volumes or JFS2 volumes on the Tivoli Storage Manager server. The server uses direct I/O on JFS2 file systems.

- Using raw logical volumes for the server can cut CPU consumption, but doing so might be slower during storage pool migration due to lack of read ahead.

AIX file cache (vmo minperm/maxperm)

You can determine how much memory AIX sets aside for file system cache.

- AIX can page out application memory (for example the Tivoli Storage Manager application) in favor of caching file system data. This can cause paging of the database buffer pool leading to slow database performance.
- Paging of the database buffer pool can cause database cache hit statistics to be overly optimistic.
- Tivoli Storage Manager server does not benefit greatly from file system cache.
- Lowering maxperm causes AIX to retain more application memory.
- Stop Tivoli Storage Manager server virtual memory paging by modifying the minperm/maxperm parameters. There are two exceptions: RAM-constrained systems and when the database buffer pool size is too large. See "DBMEMPERCENT" on page 3 for the proper settings.
- A good starting point is setting aside a maximum of 50% (vmo -p -o maxperm%=50) for file system caching, instead of the default of 80%. Lower file system caching further if 50% is not effective (change realtime). As maxperm approaches minperm, consider lowering minperm as well. Watch vmstat for progress, if pageouts go to zero, pageins eventually lower as well.
- The maxclient parameter might need to be lowered along with the maxperm parameter.
- Lowering the maxperm parameter can improve the upgrading of a V5.5 database to a V6.1 database.

UNIX file systems

When raw logical volumes are used, I/O is not processed through the virtual memory layer of the operating system.

This can result in lower CPU utilization for each disk I/O. On the other hand, when reading from server disk volumes, raw logical volumes do not use the read ahead mechanism of the operating systems. This can result in poorer performance on restores and server move operations from disk to tape. However, because many disk subsystems have their own read-ahead mechanisms, this might be of little concern.

If you choose to use raw logical volumes instead of file systems, consider the following warnings:

- Do not change the size of raw logical volumes except by using Tivoli Storage Manager facilities.
- Be careful not to start multiple instances of the server that might use the same raw logical volumes. Tivoli Storage Manager implements a locking mechanism that is designed to prevent the overwrite of volumes by another server instance. However, it is best to name raw logical volumes in order to prevent overwrite.

Solaris direct I/O should be used if not using raw logical volumes. For AIX, if you choose not to use raw logical volumes, direct I/O should always be used for JFS2 file systems. For JFS, direct I/O might cause degradation, especially with large-file-enabled file systems.

HP-UX server

Use a raw partition for disk storage pools on HP-UX Tivoli Storage Manager server.

Using a raw partition can improve performance. Raw partition volumes offer better backup and restore throughput than when VXFS volumes are used on HP-UX.

Linux server

Disable any unneeded daemons (services).

Most enterprise distributions come with many features. However, most of the time only a small subset of these features are used. For example, the TCP/IP data movement could be blocked or slowed down significantly by the internal firewall in SUSE 9 x86_64. It can be ended by `/etc/init.d/SuSEfirewall2_setup stop`.

Sun Solaris server

There are a number of actions that can improve performance for a Tivoli Storage Manager server running in a Sun Solaris environment.

- Use raw partitions for disk storage pool volumes on the Solaris platform. Raw logical volumes offer better client backup and restore throughput than UFS or Veritas file system volumes.
- The VxFS file system with direct I/O offers I/O characteristics similar to raw I/O on raw partitions. By switching to raw or direct I/O and by giving enough memory, a much larger working set of data can be cached, and a much higher cache hit rate can be sustained with obvious performance benefits. When physical I/O is required, the CPU cost of performing that I/O is reduced

because the data is not first copied to file system buffers. We recommend using VxFS file systems mounted with the direct I/O option (`mincache=direct`).

- When UFS file system volumes are used, mount these file systems using the `forcedirectio` flag. If the file system is mounted using `forcedirectio`, data is transferred directly between user address space and the disk. If the file system is mounted using `noforcedirectio`, data is buffered in kernel address space when data is transferred between user address space and the disk. The `forcedirectio` performance option benefits only from large sequential data transfers. The default behavior is `noforcedirectio`.

Windows server

There are a number of actions that can improve performance for a Tivoli Storage Manager server running in a Windows environment.

- Use a 64-bit version of the Windows Server 2003 or Windows Server 2008 to realize the following benefits: a larger virtual memory-address space, support for a larger physical RAM, and improved performance.
- Use the NTFS file system for the disk storage required by the Tivoli Storage Manager server, including the database directories, active log directory, archive log directory, and storage pool volumes. NTFS has the following advantages:
 - Support for larger disk partitions
 - Better data recovery
 - Better file security
 - Formatting storage pool volumes on NTFS is much faster
- NTFS file compression should not be used on disk volumes that are used by the Tivoli Storage Manager server, because of the potential for performance degradation.
- For optimal Tivoli Storage Manager for Windows server performance with respect to Windows real memory usage, use the server property setting for Maximize Throughput for Network Applications. This setting gives priority to application requests for memory over requests from the Cache Manager for file system cache. This setting will make the most difference in performance on systems that are memory constrained.
- For optimal backup and restore performance when using a local client on a Windows system, use the shared memory communication method. This is done by including the `COMMMETHOD SHAREDMEM` option in the server options file and in the client options file.
- Here are other actions that can affect Tivoli Storage Manager client and server performance:
 - Be aware that antivirus software can negatively affect backup performance.
 - Disable or do not install unused services.
 - Disable or do not install unused network protocols.
 - Give preference to background application performance.
 - Avoid screen savers.
 - Ensure that the paging file is not fragmented.
 - Ensure that device drivers are current, especially for new hardware.

Estimating throughput in untested environments

You can estimate the Tivoli Storage Manager throughput rate for untested environments.

Estimating throughput for environments that have not been directly tested can be difficult. However, the following important observations can be made:

- Throughput over a network can be expected to reach saturation at around 80 percent of its rated capacity. Efficiency indicates the percentage of maximum throughput rate that can realistically be achieved. This leads to the following maximum throughputs that can be obtained for given networks:

Network	Mbit per second	MB per second	GB per hour	% Efficiency
Ethernet	100	10	17.6	90
ATM	155	15.5	34.1	50
SPSwitch		120	265	50
T3	45	4.48	15.8	80
T1	1.54	0.16	0.56	80
Gigabit Ethernet	1 GB	100	350	80

- Throughput for backup and restore of small file workloads is basically independent of network type, as long as the network remains unsaturated, and propagation delays are not excessive due to intervening routers or switches.
- Gigabit Ethernet performance is highly dependent on the quality of the Ethernet chipset and the type of bus used. In addition, taking advantage of certain chipset features, such as jumbo frames and other TCP offload features, can have a large impact on performance. Therefore, performance can vary widely. On some chipsets and machines only 25% efficiency may be possible while on others 90% is easily reached.

Tuning tape drive performance

There are a few basic procedures for maintaining the performance of your tape drives.

Configuring enough tape drives

You must configure enough tape drives for your environment:

- The maximum number of Tivoli Storage Manager client sessions backing up directly to tape at any time during the peak backup window.
- Additional tape drives for other functions that run during the backup window. For example, storage pool migration, storage pool backup, and reclamation.

Cleaning tape drives

Cleaning the tape drive according to the manufacturer's specifications is very important to ensure maximum tape drive performance. Failure to clean tape drives can cause read and write errors, drive failures, and generally poor performance.

Enabling tape compression

In many cases, enabling compression at the tape drive improves Tivoli Storage Manager throughput.

You can use the `FORMAT` option of the `DEFINE DEVCLASS` command to specify the appropriate recording format to be used when writing data to sequential access media. The default is `DRIVE`, which specifies that Tivoli Storage Manager selects the highest format that can be supported by the sequential access drive on which a volume is mounted. This setting usually allows the tape control unit to perform compression.

Tip: Avoid specifying the `DRIVE` value when a mixture of devices are used in the same library. For example, if you have drives that support recording formats superior to other drives in a library, do not specify the `FORMAT=DRIVE` option. Refer to the appropriate Tivoli Storage Manager *Administrator's Guide* for more information

If you do not use compression at the client and your data is compressible, you should achieve higher system throughput if you use compression at the tape control unit. Refer to the appropriate Tivoli Storage Manager *Administrator's Guide* for more information concerning your specific tape drive. If you compress the data at the client, we recommend that you not use compression at the tape drive. In this case, you might lose up to 10-12% of the tape capacity at the tape media.

Using collocation with tape drives

Using collocation can significantly improve the performance of restores for large amounts of data with many files, because fewer tapes are searched for the necessary files. Collocation also decreases the chance for media contention with other clients. The trade-off is that more tapes are needed. Use the `COLLOCATE` option on the `DEFINE STGPOOL` command or `UPDATE STGPOOL` command to enable collocation.

The default is collocation by group. Until node groups are defined, however, no collocation occurs. When node groups nodes are defined, the server can collocate data based on these groups. Collocation by group can yield the following performance benefits:

- Reduce unused tape capacity by allowing more collocated data on individual tapes
- Minimize mounts of target volumes
- Minimize database scanning and reduce tape passes for sequential-to-sequential transfer

Collocation by group gives the best balance of restore performance versus tape volume efficiency. For those nodes where collocation is needed to improve restore performance, use collocation by group. Manage the number of nodes in the groups so that backup data for the entire group is spread over a manageable number of volumes. Where practical, a node can be moved from one collocation group to another by first changing the group affinity with the `DEFINE NODEGROUPMEMBER` command then using the `MOVE NODEDATA` command.

Tape drive transfer rate

There are many factors that affect the sustained transfer rate of Tivoli Storage Manager when using tape drives. The sustained transfer rate takes into account the net effect of all these factors.

These factors include:

- Native transfer rate
- Compression ratio
- File size
- Server attachment
- Server attachment HBA type
- Disk transfer rate
- Network bandwidth
- Server utilization
- Start/stop performance
- Application control file activity
- Bus bandwidth
- Quality of the media

Tape drive streaming rate performance

Streaming rate is the rate at which a tape drive can read and write, not including any start and stop operations. Most uses of tape do include some start and stop operations, which slow down the sustained rate at which the drive operates.

When writing to a tape drive, normally the drive returns control to the application when the data is in the tape drive buffer but before the data has been written to tape. This mode of operation provides all tape drives a significant performance improvement. However, the drive's buffer is volatile. For the application to ensure that the write makes it to tape, the application must flush the buffer. Flushing the buffer causes the tape drive to back hitch (start/stop). The Tivoli Storage Manager parameters TXNBYTELIMIT and TXNGROUPMAX control how frequently Tivoli Storage Manager issues this buffer flush command.

When writing to a tape drive, you must consider network bandwidth. For example, 100BaseT Ethernet LAN can sustain 5 to 6 MB per second. Therefore, you cannot backup to a tape drive faster than that.

High performance tape drives

When you use high performance tape drives with Tivoli Storage Manager, it is important to use the appropriate server and client options to enhance performance.

These are the recommended settings:

Server options

```
TXNGROUPMAX 4096
MOVESIZETHRESH 2048
MOVEBATCHSIZE 1000
```

Client options

```
TXNBYTELIMIT 2097152
```

If on average Tivoli Storage Manager clients have files smaller than 100 KB, it is recommended that these clients back up to a disk storage pool for later migration

to tape. This allows more efficient data movement to tape.

Tuning disk performance

You can configure Tivoli Storage Manager disk storage to optimize performance.

- **Server configuration – disk:** Use as many independent physical disks as you can afford to minimize I/O contention. Configure one Tivoli Storage Manager volume per physical disk, or at most two. Separate the disk storage pool volumes, the recovery logs, and the database.

Using raw logical volumes on UNIX systems can cut CPU consumption but might be slower during storage pool migrations due to lack of read-ahead. However, many disk subsystems have read-ahead built in, which negates this concern.

- **Server configuration - disk write cache:** Use disk subsystem/adaptor write cache for all RAID 5 arrays and physical disks with Tivoli Storage Manager database volumes (random I/O).
- **Choosing physical disks (JBOD) or RAID arrays:** RAID requires more physical disks for equivalent performance. Be sure to consider the write penalty of RAID5 arrays. Write throughput is important during backup and archive. Tivoli Storage Manager active log mirroring provides better recoverability than hardware redundancy.

Busses

If your server has multiple PCI busses, distribute high-throughput adaptors among the different busses. For systems with busses that have different speeds, match the adaptor to the appropriate buss based on the speed.

For example, if you are going to do a lot of backups to disk, you probably do not want your network card and disk adaptor on the same PCI bus. Theoretical limits of busses are just that, theoretical, though you should be able to get close in most cases. As a general rule, it is best to have only one or two tape drives per SCSI bus and one to four tape drives per fibre host bus adapter (HBA). Mixing tape and disk on the same fibre HBA is not recommended.

Even if a given tape drive is slower than the fibre channel SAN being used, tape drive performance is usually better on the faster interfaces. This is because the individual blocks are transferred with lower latency. This allows Tivoli Storage Manager and the operating system to send the next block more quickly. For example, an LTO 4 drive perform better on a 4 Gbit SAN than a 2 Gbit SAN even though the drive is only capable of speeds under 2 Gbit.

Chapter 3. IBM Tivoli Storage Manager client performance tuning

You can tune the performance of Tivoli Storage Manager clients.

Some Tivoli Storage Manager client options can be tuned to improve performance. Tivoli Storage Manager client options are specified in either the `dsm.sys` file or the `dsm.opt` file.

COMMRESTARTDURATION and COMMRESTARTINTERVAL

The `COMMRESTARTDURATION` and `COMMRESTARTINTERVAL` options control the restart window of time and interval between restarts.

To make clients more tolerant of network connectivity interruptions, use the options `COMMRESTARTDURATION` and `COMMRESTARTINTERVAL` to control the restart window of time and interval between restarts. These options assist in environments which are subject to heavy network congestion or frequent interruptions, and eases the manageability of large numbers of clients by reducing intervention on error conditions. In a sense, performance is improved if consideration is given to account for time to detect, correct, and restart as a result of an error condition.

Note: A scheduled event continues if the client reconnects with the server before the `COMMRESTARTDURATION` value elapses, even if the event's startup window has elapsed.

The `COMMRESTARTDURATION` option sets the maximum time for the client to try to reconnect with a server after a communication failure. The value is specified in minutes from 0 to 9999. The default is 60.

The `COMMRESTARTINTERVAL` option sets the time for the client to wait between attempts to reconnect with a server after a communication failure. The value is specified in seconds from 0 to 65535. The default is 15.

It is recommended that you tune the `COMMRESTARTDURATION` and `COMMRESTARTINTERVAL` options to your environment.

COMPRESSALWAYS

The `COMPRESSALWAYS` option specifies whether to continue compressing an object if it grows during compression, or resend the object, uncompressed. This option is used with the compression option.

The `COMPRESSALWAYS` option is used with the archive, incremental, and selective commands. This option can also be defined on the server. If `COMPRESSALWAYS YES` (the default) is specified, files continue compression even if the file size increases. To stop compression if the file size grows, and resend the file, uncompressed, specify `COMPRESSALWAYS NO`. This option controls compression only if your administrator specifies that your client node determines the selection. To reduce the impact of retries, use `COMPRESSALWAYS YES`.

It is better to identify common types of files that do not compress well and list these on one or more client option EXCLUDE.COMPRESSION statements. Files that contain large amounts of graphics, audio, or video files and files that are already encrypted do not compress well. Even files that seem to be mostly text data (for example, Microsoft Word documents) can contain a significant amount of graphic data that might cause the files to not compress well.

Using Tivoli Storage Manager client compression and encryption for the same files is valid. The client first compresses the file data and then encrypts it, so that there is no loss in compression effectiveness due to the encryption, and encryption is faster if there is less data to encrypt. For example, to exclude objects that are already compressed or encrypted, enter the following statements:

```
exclude.compression ?:\...\*.gif
exclude.compression ?:\...\*.jpg
exclude.compression ?:\...\*.zip
exclude.compression ?:\...\*.mp3
exclude.compression ?:\...\*.cab
```

This is the recommended setting:
COMPRESSALWAYS YES

COMPRESSION

The COMPRESSION client option specifies if compression is enabled on the Tivoli Storage Manager client. For optimal backup and restore performance with a large number of clients, consider using client compression.

Compressing the data on the client reduces demand on the network and the Tivoli Storage Manager server. The reduced amount of data on the server continues to provide performance benefits whenever this data is moved, such as for storage pool migration and storage pool backup. However, client compression significantly reduces the performance of each client, and the reduction is more pronounced on the slowest client systems.

For optimal backup and restore performance when using fast clients and heavily loaded network or server, use client compression. For optimal backup and restore performance when using a slow client, or a lightly loaded network or server, do not use compression. However, be sure to consider the trade-off of greater storage requirements on the server when not using client compression. The default for the COMPRESSION option is NO.

For maximum performance with a single fast client, fast network, and fast server, turn compression off.

Two alternatives exist to using compression:

- If you are backing up to tape, and the tape drive supports its own compression, use the tape drive compression. See “Tuning tape drive performance” on page 21 for more information.
- Do not use compression if a client has built-in file compression support. Compression on these clients does reduce the amount of data backed up to the server. NetWare and Windows have optional built-in file compression.

Compression can cause severe performance degradation when there are many retries due to failed compression attempts. Compression fails when the compressed file is larger than the original. The client detects this and will stop the compression, fail the transaction and resend the entire transaction uncompressed. This occurs

because the file type is not suitable for compression or the file is already compressed (zip files, tar files, and so on). Short of turning off compression, there are two options you can use to reduce or eliminate retries due to compression:

- Use the COMPRESSALWAYS option. This option eliminates retries due to compression.
- Use the EXCLUDE.COMPRESSION option in the client options file. This option disables compression for specific files or sets of files (for example, zip files or jpg files). Look in the client output (dsmsched.log) for files that are causing compression retries and then filter those file types.

These are the recommended settings:

- For a single fast client, fast network, fast server
COMPRESSION NO
- For multiple clients, slow network, slow server
COMPRESSION YES

DISKBUFSIZE

The DISKBUFSIZE client option specifies the maximum disk I/O buffer size (in kilobytes) that the client may use when reading files.

Optimal backup, archive, or HSM migration client performance may be achieved if the value for this option is equal to or smaller than the amount of file read ahead provided by the client file system. A larger buffer will require more memory and might not improve performance.

The default value is 32 for all clients except AIX. For AIX, the default value is 256 except when ENABLELANFREE YES is specified. When ENABLELANFREE YES is specified on AIX, the default value is 32. API client applications have a default value of 1023, except for Windows API client applications (version 5.3.7 and later), which have a default value of 32.

The recommended setting is to use the default value for the client platform.

MEMORYEFFICIENTBACKUP

The MEMORYEFFICIENTBACKUP client option specifies the method Tivoli Storage Manager uses during database backups to determine which objects are new or changed and must be backed up, and which objects are deleted and must be expired.

The memory required by the client depends on the method used and the number of objects in the client file systems. An object is a file or a directory.

To choose a value for the MEMORYEFFICIENTBACKUP option, begin by determining the number of objects in the client file systems and rounding that number up to the nearest million. For example, if your client file systems have 4,200,000 objects, round up to 5,000,000. You would then use 5 as the value of numobjs.

Follow the steps below in sequence, and use the option parameter (YES, NO, or DISKCACHEMETHOD) for the first step that applies. For example, if a 64 bit backup-archive client has 2 GB of real memory available for use by the client process and the value of numobjs is 5, you would specify NO.

1. If the client system is using the 32-bit backup-archive client, and numobjs is less than or equal to 5, and at least numobjs x 300 MB of real memory is available for use by the client process, then specify NO, which is the default.
2. If the client system is using the 64-bit backup-archive client, and at least numobjs x 300 MB of real memory is available for use by the client process, then specify NO, which is the default.
3. If the client system has at least the following amount of fast temporary disk storage available for use by the client process, specify DISKCACHEMETHOD:
 - UNIX or Linux: numobjs x 300 MB
 - Windows: numobjs x 600 MB
 - Mac OS X: numobjs x 1200 MB

By default, this disk storage space must be available on the volumes being backed up, or you must specify the DISKCACHELOCATION option with the path to the available space.

4. If none of the above apply, specify YES.

Note: Using MEMORYEFFICIENTBACKUP YES can increase the work required on the Tivoli Storage Manager server. The result can be a significant increase in the incremental backup elapsed time, particularly if this option is used for a large number of clients each with a large number of directories.

Here are some alternatives to using the MEMORYEFFICIENTBACKUP option to reduce client memory consumption:

- Use the client include/exclude options to back up only what is necessary.
- Use journal based incremental backup for Windows (NTFS) and AIX (JFS2) clients only.
- Use the VIRTUALMOUNTPOINT option (UNIX only) to define multiple virtual mount points within a single file system, and back up these mount points sequentially.
- Spread the data across multiple file systems and back up these file systems sequentially
- Use the image backup function to backup the entire volume. This might require less time and resources than using incremental backup on some file systems with a large number of small files.

The recommended setting is to use the default:

```
MEMORYEFFICIENTBACKUP NO
```

PROCESSORUTILIZATION

The PROCESSORUTILIZATION option (only on the Novell client) specifies, in hundredths of a second, the amount of CPU time the length of time Tivoli Storage Manager controls the CPU. Because this option can affect other applications on your client node, use it only when speed is a high priority.

The default is 1. The recommended values are from 1 to 20. If set to less than 1, this parameter could have a negative impact on performance. Increasing this value increases the priority of Tivoli Storage Manager to the CPU, lessening the priorities other process. Setting PROCESSORUTILIZATION greater than 20 might prevent other scheduled processes or NetWare requestors from accessing the file server.

QUIET

The QUIET client option prevents messages from being displayed during Tivoli Storage Manager backups.

The default is the VERBOSE option, and Tivoli Storage Manager displays information about each file it backs up. To prevent this, use the QUIET option. However, messages and summary information are still written to the log files. There are two main benefits to using the QUIET option:

- For tape backup, the first transaction group of data is always resent. To avoid this, use the QUIET option to reduce retransmissions at the client.
- If you are using the client scheduler to schedule backups, using the QUIET option dramatically reduces disk I/O overhead to the schedule log and improves throughput.

RESOURCEUTILIZATION

The RESOURCEUTILIZATION client option regulates the number of concurrent sessions that the Tivoli Storage Manager client and server can use during processing. Multiple sessions can be initiated automatically through a Tivoli Storage Manager backup, restore, archive, or retrieve command. Although the multiple session function is transparent to the user, there are parameters that enable the user to customize it.

The RESOURCEUTILIZATION option increases or decreases the ability of the client to create multiple sessions. For backup or archive, the value of RESOURCEUTILIZATION does not directly specify the number of sessions created by the client. However, the setting does specify the level of resources the server and client can use during backup or archive processing. The higher the value, the more sessions the client can start.

The range for the parameter is from 1 to 10. If the option is not set, by default only two sessions can be started, one for querying the server and one for sending file data. A setting of 5 permits up to four sessions: two for queries and two for sending data. A setting of 10 permits up to eight sessions: four for queries and four for sending data. The relationship between RESOURCEUTILIZATION and the maximum number of sessions created is part of an internalized algorithm and, as such, is subject to change. This table lists the relationships between RESOURCEUTILIZATION values and the maximum sessions created. Producer sessions scan the client system for eligible files. The remaining sessions are consumer sessions and are used for data transfer. The threshold value affects how quickly new sessions are created.

RESOURCEUTILIZATION value	Recommendations		
	Maximum number of sessions	Unique number of producer sessions	Threshold (seconds)
1	1	0	45
2	2	1	45
3	3	1	45
4	3	1	30
5	4	2	30
6	4	2	20

Recommendations			
RESOURCEUTILIZATION value	Maximum number of sessions	Unique number of producer sessions	Threshold (seconds)
7	5	2	20
8	6	2	20
9	7	3	20
10	8	4	10
0 (default)	2	1	30

Backup throughput improvements that can be achieved by increasing the RESOURCEUTILIZATION level vary from client node to client node. Factors that affect throughputs of multiple sessions include the configuration of the client storage subsystem (the layout of file systems on physical disks), the client's ability to drive multiple sessions (sufficient CPU, memory), the server's ability to handle multiple client sessions (CPU, memory, number of storage pool volumes), and sufficient bandwidth in the network to handle the increased traffic.

The MAXSESSIONS parameter controls the maximum number of simultaneous client sessions with the Tivoli Storage Manager server. The total number of parallel sessions for a client is counted for the maximum number of sessions allowed with the server. You need to decide whether to increase the value of the MAXSESSIONS parameter in the server option file.

When using the RESOURCEUTILIZATION option to enable multiple client/server sessions for backup direct to tape, the client node maximum mount points allowed parameter, MAXNUMMP, must also be updated at the server (using the UPDATE NODE command).

If the client file system is spread across multiple disks (RAID 0 or RAID 5), or multiple large file systems, the recommended RESOURCEUTILIZATION setting is a value of 5 or 10. This enables multiple sessions with the server during backup or archive and can result in substantial throughput improvements in some cases. It is not likely to improve incremental backup of a single large file system with a small percentage of changed data. If a backup is direct to tape, the client node maximum mount points allowed parameter, MAXNUMMP, must also be updated at the server using the update node command.

RESOURCEUTILIZATION can be set to value other than default if a client backup involves many files and they span or reside on multiple physical disks. A setting of 5 or greater is recommended. However, for optimal utilization of the Tivoli Storage Manager environment, you need to evaluate the load of server, network bandwidth, client CPU and I/O configuration and take that into consideration before changing the option.

When a restore is requested, the default is to use a maximum of two sessions, based on how many tapes the requested data is stored on, how many tape drives are available, and the maximum number of mount points allowed for the node.

The default value for the RESOURCEUTILIZATION option is 1, and the maximum value is 10. For example, if the data to be restored are on five different tape volumes, and the maximum number of mount points for the node requesting the restore is five, and RESOURCEUTILIZATION is set to 3, then three sessions are

used for the restore. If the RESOURCEUTILIZATION setting is increased to 5, then five sessions are used for the restore. There is a one-to-one relationship to the number of restore sessions allowed and the RESOURCEUTILIZATION setting.

Here are the recommended settings:

For workstations

RESOURCEUTILIZATION 1

For a small server

RESOURCEUTILIZATION 5

For a large server

RESOURCEUTILIZATION 10

Note: Non-root UNIX is limited to one session.

TAPEPROMPT

The TAPEPROMPT client option specifies whether to prompt the user for tape mounts.

The TAPEPROMPT option specifies if you want Tivoli Storage Manager to wait for a tape to be mounted for a backup, archive, restore or retrieve operation, or to prompt you for your choice.

The recommended setting is:

TAPEPROMPT NO

TCPBUFFSIZE

The TCPBUFFSIZE option specifies the size of the internal TCP communication buffer, that is used to transfer data between the client node and the server. A large buffer can improve communication performance, but requires more memory.

The default is 32 KB, and the maximum is now 512 KB. The recommended setting is:

TCPBUFFSIZE 32

TCPNODELAY

Use the TCPNODELAY option to disable the TCP/IP Nagle algorithm, which allows data packets of less than the Maximum Transmission Unit (MTU) size to be sent out immediately.

The default is YES. This generally results in better performance for Tivoli Storage Manager client/server communications.

The recommended setting is:

TCPNODELAY YES

Note: TCPNODELAY is also available as a server option.

TCPWINDOWSIZE

The TCPWINDOWSIZE client option specifies the size of the TCP sliding window in kilobytes.

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. The TCPWINDOWSIZE option is valid for all Tivoli Storage Manager clients and servers.

The TCPWINDOWSIZE option overrides the operating system's default TCP/IP session send and receive window sizes. For AIX, the defaults are set as no options `tcp_sendspace` and `tcp_recvspace`. For Solaris, the defaults are set via the `tcp_xmit_hiwat` and `tcp_rcv_hiwat` tunable parameters. Specifying TCPWINDOWSIZE 0 causes Tivoli Storage Manager to use the operating system default. This is not recommended because the optimal setting for Tivoli Storage Manager might not be same as the optimal setting for other applications. The default is 63 KB, and the maximum is 2048 KB.

The recommended setting is:

```
TCPWINDOWSIZE 63
```

Note: This option is also on the Tivoli Storage Manager server.

TXNBYTELIMIT

The TXNBYTELIMIT client option specifies the maximum transaction size in kilobytes for data transferred between the client and server.

The range of values is 300 KB through 2097152 KB (2 GB); the default is 25600. A transaction is the unit of work exchanged between the client and server. Because the client program can transfer more than one file or directory between the client and server before it commits the data to server storage, a transaction can contain more than one file or directory. This is called a transaction group. This option permits you to control the amount of data sent between the client and server before the server commits the data and changes to the server database, thus affecting the speed with which the client performs work. The amount of data sent applies when files are batched together during backup or when receiving files from the server during a restore procedure. The server administrator can limit the number of files or directories contained within a group transaction using the TXNGROUPMAX option. The actual size of a transaction can be less than your limit. Once this number is reached, the client sends the files to the server even if the transaction byte limit is not reached.

There are several items to consider when setting this parameter:

- Increasing the amount of data per transaction increases recovery log requirements on the server. Check log and log pool space to ensure there is enough space. Also note that a larger log might result in longer server start-up times.
- Increasing the amount of data per transaction might result in more data being retransmitted if a retry occurs. This might negatively affect performance.

- The benefits of changing this parameter are dependent on configuration and workload characteristics. In particular, this parameter benefits tape storage pool backup more so than disk storage pool backup, especially if many small files are in the workload.

When setting the size of transactions consider setting a smaller size if you are suffering many resends due to files changing during backup when using static, shared static, or shared dynamic. This would apply to static as well as to shared because when the client realizes a file has changed during a backup and decides to not send it, the file that is, it would still have to re-send the other files in that transaction. To enhance performance, set TXNBYTELIMIT to the max 2097152, and on the server, raise TXNGROUPMAX to 256. Additionally, for small file workloads, first stage the backups to a disk storage pool and then migrate to tape.

The recommended settings are:

```
TXNBYTELIMIT 25600
```

Or for backup directly to tape:

```
TXNBYTELIMIT 2097152
```

Multiple session backup and restore

Multiple session restore allows the backup-archive clients to perform multiple restore sessions for no-query restore operations, thus increasing the speed of restores. Multiple session restore is similar to multiple session backup support.

Multiple session restores can be used under the following conditions:

- The data to be restored resides on several tapes.
- There are sufficient mount points available.
- The restore is done using the no-query restore protocol. For details about no-query restore, refer to the *Backup-Archive Clients Installation and User's Guide*.

For backup or archive operations, the MAXNUMMP parameter on the UPDATE NODE or REGISTER NODE command controls the number of mount points allowed to a client. The RESOURCEUTILIZATION option affects how many sessions the client can use. Set RESOURCEUTILIZATION to one greater than the number of required sessions (equal to the number of drives that the client to use will use). Issue the restore command so that it results in a no-query-restore process. For backup or archive operations, if the MAXNUMMP setting is too low and if there are not enough mount points for each of the sessions, it might not be possible to use the number of sessions allowed by the RESOURCEUTILIZATION option.

The number of sessions used depends on the following settings:

- The number of mount points available to the client. This number is controlled by the MOUNTLIMIT setting in the DEFINE DEVCLASS and UPDATE DEVCLASS commands and by the MAXNUMMP setting in REGISTER NODE and UPDATE NODE server commands. The MAXNUMMP setting is not enforced for restore or retrieve operations. See "Maximum number of mount points for a node" on page 13 for details.
- The RESOURCEUTILIZATION client option setting. Because the number of sessions increases for a multiple session restore, set the MAXSESSIONS server option accordingly. See "RESOURCEUTILIZATION" on page 29 and "MAXSESSIONS" on page 5 for details.

If all the files are on random disk, only one session is used. There is no multi-session restore for a random disk-only storage pool restore. However, if you are performing a restore in which the files reside on four tapes or four sequential disk volumes and some on random disk, you can use up to five sessions during the restore. You can use the MAXNUMMP parameter to set the maximum number of mount points a node can use on the server. If the RESOURCEUTILIZATION option value exceeds the value of the MAXNUMMP on the server for a node, you are limited to the number of session specified in MAXNUMMP.

If the data you want to restore is on five different tape volumes, the maximum number of mount points is 5 for your node, and RESOURCEUTILIZATION is set to three, then three sessions are used for the restore. If you increase the RESOURCEUTILIZATION setting to 5, then 5 sessions are used for the restore. There is a one to one relationship to the number of restore sessions allowed for the RESOURCEUTILIZATION setting. Multiple restore sessions are only allowed for no-query-restore operations.

The server sends the MAXNUMMP value to the client during sign-on. During an no-query restore, if the client receives a notification from the server that another volume has been found and another session can be started to restore the data, the client checks the MAXNUMMP value. If another session would exceed that value, the client will not start the session.

Some backup considerations:

- Only one session per file system compares attributes for incremental backup. Incremental backup throughput does not improve for a single file system with a small amount of changed data.
- Data transfer sessions do not have file system affinity; each session could send files from multiple file systems. This is good for workload balancing. This is not so good if you are backing up directly to a tape storage pool collocated by filesystem. Do not use multiple sessions to back up directly to a storage pool collocated by filesystem. Use multiple commands, one per filesystem.
- Multiple sessions might not start if there are not enough entries on the transaction queue.
- For backup operations directly to tape, you can prevent multiple sessions so that data is not spread across multiple volumes by setting RESOURCEUTILIZATION to 2.

Some restore considerations:

- Only one session is used when restoring from random access disk storage pools.
- Only one file system can be restored at a time with the command line; multiple sessions may still be used for a single file system.
- Even small clients can gain throughput for restores requiring many tape mounts or locates.
- Tape cartridge contention might occur, especially when restoring from a collocated node.

Client command line options

Two options can be used only on the command line and only with specific commands. When specifying options with a command, always precede the option with a dash (-). In general, the command line interface is faster than the GUI and requires less overhead.

Two command line options that might improve Tivoli Storage Manager performance are:

IFNEWER

This option is used in conjunction with the restore command and restores files only if the server date is newer than the date of the local file. This option might result in lower network utilization if less data travels across the network.

INCRBYDATE

In a regular incremental backup, the server reads the attributes of all the files in the file system and passes this information to the client. The client then compares the server list to a list of its current file system. This comparison can be very time-consuming, especially for clients on NetWare, Macintosh, and Windows. These clients sometimes have a limited amount of memory.

With an incremental-by-date backup, the server only passes the date of the last successful backup. It is not necessary to query every active file on the Tivoli Storage Manager server. The time savings are significant. However, regular, periodic incremental backups are still needed to back up files that have only had their attributes changed. For example, if a new file in your file system has a creation date previous to the last successful backup date, future incremental-by-date backups will not back up this file. This is because the client sees it as already backed up. Also, files that have been deleted are not detected by an incremental-by-date backup. These deleted files will be restored if you perform a full system restore.

Performance recommendations by client platform

Some client performance recommendations vary by platform.

Macintosh client

Limit the use of Extended Attributes. When Extended Attributes are used, limit their length. Antivirus software can negatively affect backup and restore performance

Windows client

Performance recommendations for Windows clients include the shared memory communication method and the use of antivirus products.

- For optimal backup and restore performance when using a local client on a Windows system, use the shared memory communication method. Specify `COMMMETHOD SHAREDMEM` in both the server options file and the client options file.
- Antivirus products and backup and restore products can use significant amounts of system resources and therefore impact application and file system performance. They may also interact with each other to seriously degrade the performance of either product. For optimal performance of backup and restore:

- Schedule antivirus file system scans and incremental backups for non-overlapping times.
- If the antivirus program allows, change the antivirus program properties so that files are not scanned when opened by the client processes. Some antivirus products can automatically recognize file reads by backup products and do not need to be configured. Check the IBM support site for additional details.

Windows journal-based backup

Instead of comparing the current state of files with the Tivoli Storage Manager database, you can back up those files indicated as changed in the change journal.

Journal-based backup uses a real-time determination of changed files and directories and avoids file system scan and attribute comparison. Here are some advantages of journal-based backup:

- Much faster than classic incremental, but improvement depends on the amount of changed data.
- Requires less memory usage and less client disk I/O than full incremental backup. The amount of disk space required for the journal database depends on the number of changed files and directories between successive backups.
- Good for large file systems with many files that do not change often.

For journal-based backup, you must install a Tivoli Journal Engine Service, which monitors file system activity for file changes. This affects the performance of the file system slightly (approximately 5% during a Netbench test run). Journal options are specified in `tsmjbbd.ini`. The defaults work well; you just add the file systems to be monitored.

Client performance recommendations for all platforms

There are a number of actions you can take to improve client performance.

Run concurrent sessions on a single client

Running two or more client program instances at the same time on the same system might provide better overall throughput, depending on the available resources. Scheduling backups for multiple file systems concurrently on one Tivoli Storage Manager client system can be done with any of the following methods:

- Using one node name, running one client scheduler, and setting the `RESOURCEUTILIZATION` client option to 5 or greater with multiple file systems included in the schedule or domain specification.
- Using one node name, running one client scheduler, and scheduling a command that runs a script on the client system that includes multiple command line client statements (using `dsmc`).
- Using multiple node names and running one client scheduler for each node name, in which each scheduler uses a unique client options file, etc.

Reduce data flow from the client

Using `INCLUDE/EXCLUDE` options appropriately can reduce the amount of data that is backed up and therefore reduce the overall backup and restore elapsed time and system resource utilization.

Minimize client processing overhead

Using the -FROMNODE option creates additional overhead on all clients. Consider using the VIRTUALNODENAME instead of the -FROMNODE option.

Hierarchical Storage Manager tuning

If you must migrate a group of small files to the server, performance is better if the data moves to disk rather than to tape. After the files are migrated to disk, you can use storage pool migration to move the files to tape.

Performance tuning of Hierarchical Storage Manager (HSM) migration is poor for very small files that are grouped together with wildcard invocation of DSMIGRATE command as an example. HSM works on one file at a time, unlike archive, retrieve, restore and backup which group files at a transaction boundary. There is one transaction for each file when you use HSM migration and recall. For a group of small files, it is better to use archive or backup to store them to the server.

Data Protection for Domino for z/OS

DFSMS™ HFS HIPER APARs OW51210 and OW51732, and their resolving PTFs, have identified and fixed a performance problem with TDP for Domino®, Domino for z/OS®, and HFS. Ensure that these PTFs are installed on your system.

Chapter 4. Administration Center performance tuning

You can tune the performance of the Administration Center by careful capacity planning and installation, and by memory, processor, and network settings.

Administration Center capacity planning

You can estimate how many Tivoli Storage Manager Administration Center servers are required to support a given number of administrators and servers in a large enterprise. You can also size the hardware required by the Administration Center servers.

With this information you can allocate sufficient resources to meet the peak application demands with little or no performance degradation and loss of function. If the Administration Center server shares a system with other applications, add the processing and memory requirements of those applications to the Administration Center requirements to determine the total system requirements.

Note: This applies only to the version 6.1 Administration Center using the Integrated Solutions Console Advanced Edition 7.1.

Maximum number of active administrators

The number of administrators who are active at a given time is limited.

Many administrators can be defined to one Administration Center server without concern. However, the maximum number of administrators who are logged in and performing administrative tasks at a given time for one Administration Center server instance is 16 for a Windows or Linux server. Exceeding this limit is likely to result in an unacceptable rate of task failures and lost administrator productivity. More administrators can likely be logged in at any given time if they run only a few tasks, for example 3 or 4 per hour, and if those tasks are not concurrent. These limits do not appear to be related to processing power or memory constraints. If the required number of active administrators is greater than these limits, then plan for multiple Administration Center instances.

Processing capacity

You can meet an Administration Center environment with high performance requirements with a dual-processor system having a speed of 3 GHz or faster.

The Administration Center server can use additional processors, but there might not be a noticeable effect on the application performance.

Estimate the Administration Center processor utilization by using the following equation:

$$\text{CpuUtilization (\%)} = 0.15 + \text{TasksCompleted (per Hour)} * 0.006$$

The tasks completed per hour rate is the highest total number of tasks per hour expected to be run using the Administration Center server. It includes tasks run by all administrators logged in at the time. A task is best thought of as the minimum amount of interaction within the administration interface that produces some usable information or completes a desired operation. The number of tasks run per

hour by a single administrator could be between 20 and 100 but not to exceed 2850 tasks per hour. Adjust the CPU utilization further by multiplying by the ratio of 3.4 GHz relative to the planned processor speed, and multiplying by the ratio of two processors to the planned number of processors.

I/O throughput

Administration Center disk and network I/O requirements are not particularly demanding, and there is no need for sustained high I/O throughput.

However, application response time suffers if network delays or disk I/O delays occur. A low latency network provides the best administrator response time. Networks that are poorly tuned, networks that are already saturated by other applications, or networks that have significantly higher latency (WANs) could significantly affect Administration Center performance.

Processing memory

The most important resource requirement for the Administration Center is memory.

The maximum Java heap size is the value specified for the Administration Center server; the default is 512 MB. The largest value that can be configured for the maximum Java heap size is 1536 MB for all platforms except AIX, which allows up to 2048 MB. Thus, the Administration Center process working set memory requirement is determined by the amount of Java heap memory specified.

Add the additional memory required by the operating system to any other applications on the Administration Center server. Configure the server configured with at least this much real memory. It is important that the required real memory be available. Without adequate real memory, significant response time degradation can occur as the result of system memory paging.

Java heap memory size

The Administration Center maintains information for each active administrative session as a set of objects in the configured Java heap memory. The memory requirements of an administrative session depend on the activities that the administrator runs during the session.

Each open work page requires memory, and each open table requires additional memory that depends on the number of rows and columns in the table.

Estimate the Administration Center Java heap size by using the following equation:

$$\text{JavaHeapAllocated (MB)} = 206 + \text{ActiveAdmins} * 30$$

ActiveAdmins is the maximum number of administrators logged in at a given time. Additional administrators can be defined in the Integrated Solutions Console, but as long as they are not logged in, no additional memory is required. The number of Tivoli Storage Manager server connections defined by an administrator in the Administration Center is not an important variable in determining the Java heap size requirements, except in the sense that more servers imply that more actual work may be required. A larger maximum Java heap size provides additional memory in the case of unexpected administration activity or workload growth. However, more real memory would be required.

Using a maximum Java heap size that is too small for the amount of work being run in the Administration Center causes the Java Virtual Machine (JVM) to

perform garbage collection more frequently. This, in turn, causes higher processor utilization and slower application response time. In excessive conditions, memory allocation failures can result in the application being unable to perform the requested action until memory is freed up by closing work pages or logging out sessions.

Tips for reducing administrator session memory requirements:

- Close work pages as soon as you are finished with them.
- Logout if you are not using any administrative functions for more than 30 minutes.
- Do not configure the session timeout period for more than 30 minutes.

Configure both the Administration Center session timeout period and the maximum memory size (Java heap size) by using the Administration Center Support Utility. This utility is located in the AC/products/tsm/bin subdirectory below the Tivoli Storage Manager installation directory. Here is an example of its usage:

```
/opt/tivoli/tsm/AC/products/tsm/bin # ./supportUtil.sh
Administration Center Support Utility - Main Menu
=====
1. Manage Administration Center tracing
2. Manage the maximum memory size the Administration Center can use
3. Manage the Administration Center session timeout setting
4. Collect trace files, logs and system information to send to support
5. Generate a heap dump of the Java virtual machine
6. Generate a Java core dump of the Java virtual machine
7. View the log file for this utility
9. Exit
Enter Selection: 3
Administration Center Support Utility - Manage the Session
=====
1. Update the Administration Center session timeout setting
2. View the Administration Center session timeout setting
99. Return to main menu
Enter Selection: 1
The session timeout setting determines how long a session can be idle before it
times out. After a timeout occurs the user must log in again. The default
timeout setting is 30 minutes. The minimum timeout setting is 10 minutes. To
cancel this operation enter an empty value.
Enter the new session timeout (minutes): 30
Updating the session timeout to 30 minutes.....
Session timeout successfully updated. Restart ISC for changes to take effect.
Administration Center Support Utility - Main Menu
=====
1. Manage Administration Center tracing
2. Manage the maximum memory size the Administration Center can use
3. Manage the Administration Center session timeout setting
4. Collect trace files, logs and system information to send to support
5. Generate a heap dump of the Java virtual machine
6. Generate a Java core dump of the Java virtual machine
7. View the log file for this utility
9. Exit
Enter Selection: 2
Administration Center Support Utility - Manage the JVM
=====
1. Update the maximum memory size the Administration Center can use
2. View the maximum memory size the Administration Center can use
99. Return to main menu
Enter Selection: 1
The maximum memory size determines the largest amount of memory that can be
used by the Administration Center. A minimum heap size of 512 MB is recommended.
When used by 10 or more users, the recommendation is at least 1024 MB. To cancel
```

this operation, enter an empty value.
Enter the new JVM max memory size (MB): 1536
Updating the maximum memory size to 1536 MB.....
Maximum memory size successfully updated.

Remember: Do not configure the maximum memory size (Java heap size) to be greater than the available real system memory, or significant performance degradation may occur.

The Tivoli Storage Manager Administration Center Capacity Planner tool can simplify using the equations and can provide recommendations for Administration Center hardware sizings. See your IBM representative to obtain this tool.

Administration Center setup recommendations

There are a number of measures you can take that can improve the performance of the Administration Center.

Installation requirements

The Administration Center installation must meet the minimum hardware requirements.

You can find these requirements at: <http://www.ibm.com/support/docview.wss?uid=swg21328445>. If the Administration Center server is installed in an environment in which the workload is light (that is, a single administrator), then the memory requirements calculated using the information in “Administration Center capacity planning” on page 39 would indicate a smaller memory requirement than the minimums provided on the Web. Administration Center server performance might be acceptable when using this smaller memory amount if the workload is light. Additional memory and processing power can provide significant performance benefits in the case of unexpected demand or workload growth.

If you plan to upgrade an existing Tivoli Storage Manager server, and the existing hardware cannot meet the additional Administration Center requirements, then consider upgraded hardware or an additional system for the Administration Center function.

Locating the Administration Center

Where you install your Administration Center can affect its performance.

You can use a single Administration Center installation to administer multiple Tivoli Storage Manager servers. You can install the Administration Center on the same machine as the Tivoli Storage Manager server if the machine meets the memory requirements described above and if it is installed on an operating system that the Administration Center supports. If additional applications are being deployed to the same Integrated Solutions Console (ISC) system, or if the Administration Center is to be used with multiple Tivoli Storage Manager servers and administrators, then install the Administration Center on a system other than one with a Tivoli Storage Manager server. For availability reasons, you might prefer to install multiple Administration Center images.

During administration activities, more network traffic occurs between the Administration Center system and the administrator’s browser than between the Administration Center system and a Tivoli Storage Manager server. Therefore,

install the Administration Center close (in network topology) to the administrators, rather than close to the Tivoli Storage Manager servers. For example, if you are in Chicago and administer Tivoli Storage Manager servers in Los Angeles, Paris, and Tokyo, the Administration Center should be installed in Chicago.

Minimizing memory usage

You can take steps to keep the Administration Center server memory available for those users and tasks that need it the most.

For the best results, take all of the following steps:

- Turn all Administration Center server tracing off, except as needed for problem determination.
- Do not increase the ISC session timeout above 30 minutes. Idle sessions can hold on to large amounts of memory until the session is canceled due to the timeout.
- Close portlet pages when you are finished using them. Closing the pages frees up any memory held by those pages.
- Use available filtering options where provided when displaying data in large tables. The additional search criteria reduces the amount of data that the Administration Center receives from the Tivoli Storage Manager server and reduces the amount of data sent to the browser and reduces the amount of memory required.

Optimizing Windows Server 2003 memory

For a memory-constrained Administration Center server on Windows, optimize the memory usage for programs.

Follow this procedure:

1. Click **Start** → **Settings** → **Control Panel**.
2. In the Control Panel, double click **System**.
3. In the System Properties panel, click the **Advanced** tab then click **Performance Settings**.
4. In the Performance Options panel, click the **Advanced** tab.
5. In the Memory Usage section, click **Programs**. The default is System cache.
6. Click **OK** and reboot the server.

Using the default action

There is a shortcut in the interface for selecting a storage object and using the default action.

From a list of selectable objects (servers, policy domains, and so on), the default action in the “Select action” drop-down box is run when the underlined object is clicked. Usually the default action is to display the properties for an object. Using this shortcut can improve performance by not having to click the radio button next to the object, then click the drop-down box, and then select the requested action.

Tuning Administration Center performance

If you experience performance problems with the Administration Center, you can check a number of possible areas that can be tuned.

Tuning processor performance

Problems with excessive processor utilization when the Administration Center is installed are likely due to other applications using too much of the processors.

Use the available platform-specific performance tools on the Administration Center server to investigate the problem.

If total processor utilization is over 90% for extended periods, either reduce the workload on the system, or add more processing capacity.

Tuning network performance

If you suspect that your network performance is not optimal, there are steps that you can take.

Check that network response time is acceptable. Use the ping command to determine the response time between the Administration Center server and the Tivoli Storage Manager server, and between the system used for the browser and the Administration Center server. Acceptable round-trip response times might be 10 ms for a LAN, or 100 ms for a WAN. If either network response time is consistently higher than expected, contact your network administrator.

Even if the ping network response time is good, bulk data transfer rates can be poor. Check the FTP data transfer rate of a 10 MB file between the Administration Center server and the Tivoli Storage Manager server, and between the system used for the browser and the Administration Center server. An FTP put or FTP get for a 10 MB file should take no more than a few seconds in a properly configured network. For 100 Mb Ethernet networks, the media speed setting can cause performance problems. This setting can be set to "100 Mb full-duplex", "100 Mb half-duplex", or "Auto". Although newer hardware works well with the "Auto" setting, some combinations of switches, hubs, and adapters exhibit poor performance if incorrectly configured.

Tuning memory performance

There are three areas that could be the source of memory performance problems.

Perform these check to determine

- Check that the maximum memory size that the Administration Center server can use, also known as the maximum Java heap, is correctly configured for the expected workload in terms of the number of concurrently logged in administrators and task execution rate. Refer to the procedure described in "Administration Center capacity planning" on page 39. If the maximum Java heap is set too low, use the support utility to increase this setting, and stop and restart the Administration Center server. Stopping the server frees memory for idle sessions without waiting for the session timeout to expire.
- Check that the Administration Center server is not paging system memory. Use the performance tools available on the Administration Center server to investigate this. For example, use the vmstat command on UNIX, or check the Windows Performance Monitor Memory object, Pages Output/sec counter. If the

| system is paging memory, then reduce the amount of memory in use by the
| active processes, or add more real memory.

- | • Check the amount of real memory currently in use by the Administration Center
| server process. Use the ps command on UNIX, or check the Windows Task
| Manager (Processes tab, Mem Usage column), or Windows Performance Monitor
| Process object, Java instance, Working Set counter. The following commands can
| be used on AIX to find the Administration Center server process ID and then the
| resident set memory (RSS) for this process:

```
| ps -ef | grep [I]SC_Portal | awk '{ print $2 }'  
| ps avxw PID
```

| In addition, check the IBM support site for updates to the Integrated Solutions
| Console and Tivoli Storage Manager Administration Center, and for information
| that may describe your problem.

Chapter 5. Network protocol tuning

Tuning network protocols can improve the performance of IBM Tivoli Storage Manager operations.

TCP/IP communication concepts and tuning

To tune TCP/IP, the most significant performance improvements are found by modifying parameters that affect the data transfer block size, window values, and connection availability.

The following tasks require system resources:

- Keeping communication connections available
- Keeping user data until it is acknowledged (on the transmit side)
- Managing the communications layers

These resources include memory, CPU, communications adapters, link utilizations, and involve the limitations of various communication layer implementations. Data sizes and flow control are the two main factors that cause resource over-commitment, which results in system performance degradation.

TCP/IP protocols and functions

The TCP/IP protocols and functions can be categorized by their functional groups.

The groups are: the network layer, internetwork layer, transport layer, and application layer. The following table shows the functional groups and their related protocols.

Group	Protocols and Functions
Network layer	<ul style="list-style-type: none">• Token-Ring• Ethernet• Others
Internetwork layer	<ul style="list-style-type: none">• Internet Protocol (IP)• Internet Control Message Protocol (ICMP)• Address Resolution Protocol (ARP)
Transport layer	<ul style="list-style-type: none">• Transmission Control Protocol (TCP)• User Datagram Protocol (UDP)
Application layer	<ul style="list-style-type: none">• Telnet File Transfer Protocol (FTP)• Remote Procedure Call (RPC)• Socket Interfaces• Others

Protocol functions

The protocol functions can be categorized as the following:

- Reliable delivery

- Packet assembly and disassembly
- Connection control
- Flow control
- Error control

Reliable Delivery

Reliable delivery services guarantee to deliver a stream of data sent from one machine to another without duplication or loss of data. The reliable protocols use a technique called acknowledgment with retransmission, which requires the recipient to communicate with the source, sending back an acknowledgment after it receives data.

Packet assembly and disassembly

Each layer of a communications protocol can potentially perform some sort of assembly or disassembly function. If the source and destination nodes do not lie on the same physical network, then the TCP/IP software has to fragment the packets traveling from one network to another if the maximum transmission units (MTUs) on the networks do not match. The TCP/IP software at the receiving station then reassembles the fragments.

There are advantages to assembly and disassembly:

- A communications network may only accept data blocks up to a certain size, hence requiring that larger blocks be broken down. For example, an Ethernet LAN has an MTU size of 1500 bytes, whereas a Token-Ring LAN has an MTU size of up to 16000 bytes.
- Error control might be more efficient for smaller blocks.
- More equitable access, with shorter delay, may be provided to shared transmission facilities. For example, if the line is slow, allowing too big a block to be transmitted could cause a monopolization of the line.

There are also disadvantages to assembly and disassembly:

- Each transmitted unit of data requires some fixed amount of overhead. Hence the smaller the block, the larger the percentage of overhead.
- More blocks have to be processed for both sending and receiving sides in order to transmit equal amounts of user data, which can take more time.

Flow control and error control

- Flow control is a function provided by a receiving system that limits the amount or rate of data that is sent by a transmitting system. The aim is to regulate traffic to avoid exceeding the receivers system resources.
- Error control is needed to guard against loss or damage of data and control information. Most techniques involve error detection and retransmission.

Sliding window

The sliding window allows TCP/IP to use communication channels efficiently, in terms of both flow control and error control. The sliding window is controlled in Tivoli Storage Manager through the TCPWINDOWSIZE option.

To achieve reliability of communication, the sender sends a packet and waits until it receives an acknowledgment before transmitting another. The sliding window protocol enables the sender to transmit multiple packets before waiting for an acknowledgment. The advantages are:

- Simultaneous communication in both directions.
- Better utilization of network bandwidth, especially if there are large transmission delays.
- Traffic flow with reverse traffic data, known as piggybacking. This reverse traffic might or might not have anything to do with the acknowledgment that is riding on it.
- Variable window size over time. Each acknowledgment specifies how many octets have been received and contains a window advertisement that specifies how many additional octets of data the receiver is prepared to accept, that is, the receiver's current buffer size. In response to decreasing window size, the sender decreases the size of its window. Advantages of using variable window sizes are flow control and reliable transfers.

Tip: A client continually shrinking its window size is an indication that the client cannot handle the load. Increasing the window size does not improve performance.

Networks

Tuning your networks can provide significant performance improvements.

There is a variety of actions you can take to tune your networks.

- Use dedicated networks for backup (LAN or SAN).
- Keep device drivers updated.
- Using Ethernet adapter auto detect to set the speed and duplex generally works well with newer adapters and switches. If your network hardware is more than three years old and backup and restore network performance is not as expected, set the speed and duplex to explicit values (for example, 100 MB full-duplex, 100 MB half-duplex, and so on). Make sure that all connections to the same switch are set to the same values.
- Gb Ethernet jumbo frames (9000 bytes) can give improved throughput and lower host CPU usage especially for larger files. Jumbo frames are only available if they are supported on client, server, and switch. Not all Gb Ethernet hardware supports jumbo frames.
- In networks with mixed frame-size capabilities (for example, standard Ethernet frames of 1500 bytes and jumbo Ethernet frames of 9000 bytes) it can be advantageous to enable path maximum transmission unit (PMTU) discovery on the systems. Doing so means that each system segments the data sent into frames appropriate to the session partners. Those that are fully capable of jumbo frames use jumbo frames. Those that have lower capabilities automatically use the largest frames that do not cause frame fragmentation and re-assembly somewhere in the network path. Avoiding fragmentation is important in optimizing the network.

Limiting network traffic

There are several Tivoli Storage Manager server SET commands that can limit the amount of network traffic due to client sessions.

The SET commands are:

- SET QUERYSCHEDPERIOD sets the frequency that a client can contact the server to obtain scheduled work (polling mode). This overrides the client setting. A shorter frequency means more network traffic due to polling. Use longer settings (6 to 12 hours) to reduce network traffic. Alternately, use Server Prompted schedule mode to eliminate network traffic due to polling.
- SET MAXCMDRETRIES sets a global limit on number of times a client scheduled command retries. This overrides the client setting. A smaller number reduces network traffic due to retries.
- SET RETRYPERIOD specifies the number of minutes between a retry of a scheduler after a failed attempt to contact the server. This overrides the client setting. A larger value will reduce network traffic due to retries and will make successful retry more likely. Be sure to consider your schedule start-up windows when setting the MAXCMDRETRIES and RETRYPERIOD. If a retry is attempted outside of the start-up window, it fails.

AIX network settings

It is important to minimize all performance constraints on AIX to achieve maximum throughput on the server. This is accomplished by tuning the network option parameters on AIX.

Tivoli Storage Manager uses TCP/IP communication protocol over the network. It is important to tune the TCP protocols to obtain maximum throughput. This requires changing the network parameters that control the behavior of TCP/IP protocols and the system in general.

In AIX, an application using the TCP/IP communication protocol opens a TCP socket and writes data to this socket. The data is copied from the user space into the socket send buffer, called the `tcp_sndspace` in kernel space. The receive buffers are called `tcp_rcvspace`. The send and receive buffers are made up of smaller buffers called mbufs.

An mbuf is a kernel buffer that uses pinned memory and comes in two sizes, 256 bytes and 4096 bytes called mbuf clusters or simply clusters. The maximum socket buffer size limit is determined by the `sb_max` kernel variable. Because mbufs are primarily used to store data for incoming and outgoing network traffic, they must be configured to have a positive effect on network performance. To enable efficient mbuf allocation at all times, a minimum number of mbuf buffers are always kept in the free buffer pool. The minimum number of mbufs is determined by `lowmbuf`, whereas the minimum number of clusters is determined by the `lowclust` option. The `mb_cl_hiwat` option controls the maximum number of free buffers the cluster pool can contain.

The `thewall` network option controls the maximum RAM that can be allocated from the Virtual Memory Manager (VMM) to the mbuf management routines. The `netstat -m` can be used to obtain detailed information on the mbufs. The `netstat -I interface-id` command can be used to determine if there are errors in packet transmissions.

If the number is greater than 0, overflows have occurred. At the device driver layer, the mbuf chain containing the data is put on the transmit queue, and the adapter is signaled to start the transmission operation. On the receive side, packets are received by the adapter and then are queued on the driver-managed receive queue. The adapter transmit and receive queue sizes can be configured using the System Management Interface Tool (SMIT).

At the device driver layer, both the transmit and receive queues are configurable. It is possible to overrun these queues. To determine this use `netstat -v` command, which shows Max Transmits Queued and Max receives Queued.

MTU and MSS settings

The maximum transmission unit (MTU) and maximum segment size (MSS) setting are important factors in tuning AIX for throughput.

For best throughput for systems on the same type of network, it is advisable to use a large MTU. In multi-network environments, if data travels from a network with a large MTU to a smaller MTU, the IP layer has to fragment the packet into smaller packets (to facilitate transmission on a smaller MTU network), which costs the receiving system CPU time to reassemble the fragment packets. When the data travels to a remote network, TCP in AIX defaults to a maximum segment size (MSS) of 512 bytes. This conservative value is based on a requirement that all IP routers support an MTU of at least 576 bytes.

Network type	MTU	MSS (RFC1323 0)	MSS (RFC1323 1)
FDDI	4352	4312	4300
Token ring	4096	4056	4044
Ethernet	1500	1460	1448

Note: Jumbo frames can be enabled on Gigabit Ethernet and 10 Gigabit Ethernet adapters. Doing so raises the MTU to 9000 bytes. Because there is less overhead per packet, jumbo frames typically provide better performance, or CPU consumption, or both. Consider jumbo frames especially if you have a network dedicated to backup tasks. Jumbo frames should only be considered if all equipment between most of your Tivoli Storage Manager clients and server supports jumbo frames, including routers and switches.

You can override the default MSS in the following three ways:

1. Specify a static route to a specific remote network and use the `-mtu` option of the route command to specify the MTU to that network. Disadvantages of this approach are:
 - It does not work with dynamic routing.
 - It is impractical when the number of remote networks increases.
 - Routes must be set at both ends to negotiate a value larger than a default MSS.
2. Use the `tcp_mssdflt` option of the `no` command to change the default value of MSS. This is a system wide change. In a multi-network environment with multiple MTUs, the value specified to override the MSS default should be the minimum MTU value (of all specified MTUs) less 40. In an environment with a large default MTU, this approach has the advantage that MSS does not need to be set on a per-network basis. The disadvantages are:

- Increasing the default can lead to IP router fragmentation if the destination is on a remote network, and the MTUs of intervening networks is not known.
 - The `tcp_mssdflt` parameter must be set to the same value on the destination host.
3. Subnet and set the `subnetsarelocal` option of the `no` command. Several physical networks can be made to share the same network number by subnetting. The `subnetsarelocal` option specifies, on a system-wide basis, whether subnets are to be considered local or remote networks. With `subnetsarelocal=1` (the default), Host A on subnet 1 considers Host B on subnet 2 to be on the same physical network. The consequence of this is that when Host A and Host B establish connection, they negotiate the MSS assuming they are on the same network. This approach has the following advantages:
- It does not require any static bindings MSS is automatically negotiated.
 - It does not disable or override the TCP MSS negotiation so that small differences in the MTU between adjacent subnets can be handled appropriately.

The disadvantages are:

- Potential IP router fragmentation when two high-MTU networks are linked through a lower-MTU network.
- Source and destination networks must both consider subnets to be local.

In an SP2 environment with a high speed switch, use an MTU of 64 KB

AIX - no (network options) - You can configure the network option parameters by using the `no` command.

- Use `no -a` to view current settings.
- When using TCP window sizes ≥ 64 , set `rfc1323` to 1.
- If you see non-zero "No mbuf errors" in `entstat`, `fddistat`, or `atmstat`, raise `thewall`.
- Set `thewall` to at least 131072 and `sb_max` to at least 1310720. Newer versions of AIX have larger defaults.
- Because the settings for the `no` command do not survive reboot, use the `-p` option.
- Recommended changes: `no -o rfc1323=1`

Here are the recommended values for the parameters described in this section.

```
lowclust = 200
lowmbuf = 400
thewall = 131072
mb_cl_hiwat = 1200
sb_max = 1310720
rfc1323 = 1
```

With the exception of `rfc1323`, use the current values if it is greater.

NetWare client cache tuning

You can improve performance by tuning the NetWare client cache.

The following table provides recommendations for tunable parameters in NetWare 5.1 and 6.5 (through the SERVMAN utility).

NetWare Parameter	Default	Range	Recommendation
MAXIMUM NUMBER OF INTERNAL DIRECTORY HANDLES	100	40-1000	1000
DIRECTORY CACHE ALLOCATION WAIT TIME	2.2	0.5 - 120	0.5
MAXIMUM PHYSICAL receive PACKET SIZE (TR)	tba	tba	4202
MAXIMUM PHYSICAL receive PACKET SIZE (EN)	tba	tba	1442
MINIMUM DIRECTORY CACHE BUFFERS	tba	tba	2000
MAXIMUM DIRECTORY CACHE BUFFERS	100	20-2000	4000
MAXIMUM CONCURRENT DISK CACHE WRITES	tba	tba	1000
MAXIMUM CONCURRENT DIRECTORY CACHE WRITES	tba	tba	25
MINIMUM PACKET receive BUFFERS	tba	tba	100

Maximum transmission unit settings

The maximum transmission unit (MTU) in a NetWare environment must be tuned to the appropriate size for the best performance.

In multi-network environments, data traveling from a network with a larger MTU to a smaller MTU, the IP layer must fragment the packet into smaller packets. This costs the receiving system CPU time to reassemble the fragment packets. The MTU size is configured by editing STARTUP.NCF (NetWare V3.x) or AUTOEXEC.NCF file (NetWare V4.x) as shown below.

TcpMSSinternetlimit

When data travels to a remote network or a different subnet, the TCPIP.NLM sets the MTU size to the default maximum segment size (MSS) value of 536 bytes.

The TcpMSSinternetlimit parameter can be used to override the default MSS value and to set a larger MTU. For NetWare v4.x with TCP/IP v3.0, setting TcpMSSinternetlimit off in SYS:\ETC\TCPIP.CFG causes the TCPIP. NLM to use the MTU value specified in STARTUP.NCF file (maximum physical receive packet size).

Important: The TcpMSSinternetlimit parameter is case sensitive. If this parameter is not specified correctly, it will be dropped automatically from the tcpip.cfg file by NetWare.

TCPIP.CFG
TcpMSSinternetlimit off

For NetWare v3.x, Novell patch TCP31A.EXE (for TCP/IP v2.75) can provide the same option.

Sun Solaris network settings

Tuning TCP/IP settings for Sun Solaris servers and clients can improve performance.

Here are some settings that can improve performance.

- TCPWINDOWSIZE 32K, which is set in the client dsm.sys file, is recommended for the Solaris client in the FDDI and fast (100 Mbit) Ethernet network environment.
- TCPWINDOWSIZE 63K or higher is recommended for Gigabit Ethernet network environment. One good way to find the optimal TCPWINDOWSIZE value in your specific network environment is to run the TTCP program multiple times, with different TCPWINDOWSIZE set for each run. The raw network throughput number reported by TTCP can be used as a guide for selecting the best TCPWINDOWSIZE for your Tivoli Storage Manager server and client. TTCP is freeware which can be downloaded from many Sun freeware web sites. The default values for TCP xmit and rcv buffers are only 8 KB for Solaris. The default value for tcp_xmit_hiwat and tcp_rcv_hiwat must be changed to the value of TCPWINDOWSIZE to avoid any TCP buffer overrun problem. You can use the Solaris "ndd -set" command to change the value of these two TCP buffers.
- On SunOS, the TCP/IP software parameters can be changed by editing the min_inetnet/in_proto.c file in the release kernel build directory (usually /usr/sys). After changing the parameters, you need to rebuild the kernel. The parameters that can affect performance are:

tcp_default_mss

Specifies the default maximum segment size (MSS) for TCP in bytes. The MSS is based on the maximum transmission unit (MTU) size of the network if the destination is on the same network. To avoid fragmentation, the conservative value of 512 is used. For improved performance on Ethernet or Token-Ring, larger MSS values are recommended. (For example, settings of 1024, 1500, 2048, or 4096 can be used.) On Ethernet LANs, the largest MTU value is 1500.

tcp_sendspace

Specifies the number of bytes that the user can send to a TCP socket buffer before being blocked. The default values can be changed on a given socket with the SO_SNDBUF ioctl. The default value is 4096

Recommendation: Set the tcp_sendspace parameter to 16 KB or 32 KB.

tcp_recvspace

Specifies the number of bytes that the remote TCP can send to a TCP socket buffer before being blocked. The default values can be changed on a given socket with the SO_RCVBUF ioctl. The default value is 4096.

Recommendation: Set the parameter to tcp_recvspace 16 KB or 32 KB.

z/OS network settings

You can configure the TCP/IP address space for IBM TCP/IP for z/OS and tune TCP/IP and z/OS UNIX system services.

USS client with IBM TCP/IP for z/OS

You can configure the TCP/IP address space for IBM TCP/IP for z/OS.

During initialization of the TCP/IP address space, system operation and configuration parameters are read from a configuration dataset. The program searches for the data set `job_name.node_name.TCPIP`, where `node` is the node name of the system as specified on the VMCF initialization record. VMCF is a subsystem defined by a line in the IEFSSNxx member that causes the VMCF address space to be created and initialized. If this dataset is not found, the program uses the first of the following data sets it finds:

- `tcPIP.node_name.TCPIP`
- `job_name.PROFILE.TCPI`
- `tcPIP.PROFILE.TCPIP`

We discuss only the configuration parameters that affect overall system performance. The various free pool sizes can be configured depending on the user environment and are discussed in a later section. In our lab environment, default values were used, except as noted. These settings are recommended values, but they might need to be altered depending on system capacity requirements.

TCPIP.DATA

TCPIP.DATA contains `hostname`, `domainorigin`, `nsinteraddr`, and so on.

The content of TCPIP.DATA is the same as for previous releases of TCP/IP for z/OS. For a sample TCPIP.DATA, see the *IP Configuration* manual or see the sample provided with the product.

One important recommendation is to keep the statement "TRACE RESOLVER" commented out to avoid complete tracing of all name queries. This trace should be used for debugging purposes only.

PROFILE.TCPIP

During initialization of the TCPIP stack, configuration parameters for the stack are read from the PROFILE.TCPIP configuration data set. Reference the *z/OS IP Configuration* manual for additional information on the parameters that are used in this file.

The PROFILE.TCPIP contains TCP buffer sizes, LAN controller definitions, ports, home IP addresses, gateway statements, VTAM® LUs for Telnet use, and so on.

The TCPWINDOWSIZE client option allows you to set the TCP/IP send and receive buffers independently from TCP/IP. The default size is 63 KB. Therefore, you only need to set the TCP/IP profile TCPMAXRCVBUFRSIZE parameter to a value equal to or larger than the value you want for the client TCPWINDOWSIZE option. You can set the TCPSENBFRSIZE and TCPRCVBUFRSIZE parameters to values appropriate for the non-Tivoli Storage Manager network workloads on the system, because these parameters are overridden by the client TCPWINDOWSIZE

option. When send/rcv buffer sizes are not specified in the PROFILE, a default size of 16 KB is used for send/rcv buffers.

```
IPCONFIG  PATHMTUDISCOVERY
TCPCONFIG TCPMAXRCVBUFRSIZE 524288
          TCPSENDBFRSIZE    65535
          TCPRCVBUFRSIZE    65535
```

Note: The FTP server and client application override the default settings and use 64 KB-1 as the TCP window size and 180 KB bytes for send/rcv buffers. Therefore, there is no change required in the TCPCONFIG statement for FTP server and client.

TCP/IP and z/OS UNIX system services performance tuning

There is a variety of actions you can take to tune TCP/IP and z/OS UNIX system services.

- Set the client/server TCP window size to the allowed maximum.

Tip: Set the TCP window size on z/OS to the allowed maximum by setting TCPRCVBUFRSIZE to 32K or larger. If client workstation permits, set the client window size to 65535. However, if the installation is storage constrained, use the default TCPRCVBUFRSIZE of 16K.

- Ensure that client and server MTU/packet size are equal. Follow the recommendations given in the PROFILE.TCPIP section.
- Ensure that TCP/IP and all other traces are turned off for optimal performance. Trace activity does create an extra processing overhead.
- Follow the z/OS UNIX System Services performance tuning guidelines in the *z/OS UNIX System Services Planning* manual or at this URL <http://www.ibm.com/servers/eserver/zseries/zos/unix/bpxa1tun.html>.
- Region sizes and dispatching priority: It is highly recommended that you set the region size to 0K or 0M for the TCPIP stack address space and for started tasks such as the FTP server, the SMTP/NJE server, and so on.
- If your environment permits, set the dispatching priority for TCPIP and VTAM equivalent and keep servers slightly lower than TCPIP and VTAM. For other started tasks, such as FTP, keep them slightly lower than the TCPIP task.
- If you are using Work Load Manager, follow the above recommendations when your installation defines performance goals in a service policy. Service policies are defined through an ISPF application and they set goals for all types of z/OS managed work.
- If you are using TCP/IP V3R2, see the *MVS™ TCP/IP V3R2 Performance Tuning Guide*. This tuning guide also includes step by step process for tuning other TCP/IP platforms such as AIX.
- Estimate how many z/OS UNIX System Services users, processes, ptys, sockets and threads would be needed for your z/OS UNIX installation. Update your BPXPRMxx member in SYS1.PARMLIB
- Spread z/OS UNIX user HFS datasets over more DASD volumes for optimal performance.
- Monitor your z/OS UNIX resources with RMF™ or system commands (DISPLAY ACTIVE, and DISPLAY OMVS, and so on.)

Appendix. Accessibility features for Tivoli Storage Manager

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features of Tivoli Storage Manager are described in this topic.

Accessibility features

The following list includes the major accessibility features in Tivoli Storage Manager:

- 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
- User documentation provided in HTML and PDF format. Descriptive text is provided for all documentation images.

The Tivoli Storage Manager Information Center, and its related publications, are accessibility-enabled.

Keyboard navigation

Windows The Tivoli Storage Manager for Windows Console follows Microsoft conventions for all keyboard navigation and access. Drag and Drop support is managed using the Microsoft Windows Accessibility option known as MouseKeys. For more information about MouseKeys and other Windows accessibility options, please refer to the Windows Online Help (keyword: MouseKeys).

AIX Tivoli Storage Manager follows AIX operating system conventions for keyboard navigation and access.

HP-UX Tivoli Storage Manager follows HP-UX operating-system conventions for keyboard navigation and access.

Linux Tivoli Storage Manager follows Linux operating-system conventions for keyboard navigation and access.

Mac OS X Tivoli Storage Manager follows Macintosh operating-system conventions for keyboard navigation and access.

Solaris Tivoli Storage Manager follows Sun Solaris operating-system conventions for keyboard navigation and access.

Vendor software

Tivoli Storage Manager 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.

Related accessibility information

You can view the publications for Tivoli Storage Manager in Adobe® Portable Document Format (PDF) using the Adobe Acrobat Reader. You can access these or any of the other documentation PDFs at the IBM Publications Center at <http://www.ibm.com/shop/publications/order/>.

IBM and accessibility

For more information about the commitment that IBM has to accessibility, see the IBM Human Ability and Accessibility Center at <http://www.ibm.com/able>.

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:

*IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, 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 wish 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 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.

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)[®] are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol ([®] or [™]), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe is either a registered trademark or trademark of Adobe Systems Incorporated in the United States, other countries, or both.

Intel[®] and Itanium[®] are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States 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 company, product, or service names may be trademarks or service marks of others.

Glossary

A glossary is available with terms and definitions for the IBM Tivoli Storage Manager server and related products.

The glossary is located in the Tivoli Storage Manager Version 6.1 information center:
<http://publib.boulder.ibm.com/infocenter/tsminfo/v6>

Index

Numerics

64-bit ACSLS xii

A

accessibility features 57
ACSLS xvi
active log
 mirroring 24
adapters per fiber HBA 24
Administration Center
 estimating requirements 39
 I/O throughput 40
 installing 42
 Java heap memory size 40
 location 42
 memory usage 43
 number of administrators 39
 optimizing Windows Server 2003
 memory 43
 performance tuning 39
 processing capacity 39
 processing memory 40
 setup 42
 tuning 44
 memory performance 44
 network performance 44
 processor performance 44
AIX
 ioo command 17
 performance recommendations 17
 server and client TCP/IP tuning 50
 Virtual Address space 17
 vmo command 17
antivirus software 35

B

backup
 LAN-free 12
 operations 11
 performance 11
 throughput 29
BACKUP DB server command 9
BEGINROUTES/ENDROUTES block 55
busses
 multiple PCI 24

C

cache size 4
cached disk storage pools 13
client
 tuning options 25
client commands
 DSMMIGRATE 37
client options 35
 command line only
 IFNEWER 35

client options (*continued*)
 command line only (*continued*)
 INCRBYDATE 35
 COMMETHOD SHARED MEM 20,
 35
 COMMRESTARTDURATION 25
 COMMRESTARTINTERVAL 25
 COMPRESSALWAYS 25, 26
 COMPRESSION 26
 DISKBUFSIZE 27
 ENABLELANFREE 27
 MEMORYEFFICIENTBACKUP 27
 PROCESSORUTILIZATION 28
 QUIET 29
 RESOURCEUTILIZATION 29, 33, 36
 TAPEPROMPT 31
 TCPBUFSIZE 31
 TCPNODELAY 31
 TCPWINDOWSIZE 17, 20, 32, 54
 TXNBYTELIMIT 17, 23, 32
 VIRTUALMOUNTPOINT 27
 VIRTUALNODENAME 36
 Windows 35
collocation 22
COMMETHOD SHARED MEM client
 option 20, 35
COMMETHOD SHARED MEM server
 option 20, 35
COMMRESTARTDURATION client
 option 25
COMMRESTARTINTERVAL client
 option 25
COMPRESSALWAYS client option 25,
 26
compression
 enabling on tape drives 21
COMPRESSION client option 26
configuration parameters
 self-tuning memory 10
customer support
 contact ix

D

Data Protection for Domino for z/OS 37
database
 performance 9
database manager 10
database performance 10
DBMEMPERCENT server option 3
deduplication
 in FILE storage pools 14
DEFINE COPYGROUP server
 command 11
DEFINE DEVCLASS server
 command 21
DEFINE STGPOOL server command 13,
 14, 22
device drivers xii, 49
direct I/O
 AIX 19

direct I/O (*continued*)
 Sun Solaris 19
disaster recovery 12
disk
 performance considerations 24
 write cache 24
disk I/O
 CPU 10
DISKBUFSIZE client option 27
DISKSTGPOOLMEMSIZE server
 option 4
dsm.opt file 25
dsm.sys file 25
DSMMIGRATE client command 37
DSMSERV INSERTDB
 preview xii

E

education
 see Tivoli technical training vii
ENABLELANFREE client option 27
environment file
 modifying reporting performance 15
Ethernet adapters 49
EXPINTERVAL server option 4, 11
EXPIRE INVENTORY server
 command 11
export 12

F

fixes, obtaining viii
FROMNODE option 36

G

Gb Ethernet jumbo frames 49

H

hardware
 server 8
Hierarchical Storage Manager
 migration 37

I

IBM Software Support
 submitting a problem x
IBM Support Assistant viii
import 12
INCLUDE/EXCLUDE lists 36
installing
 Administration Center 42
Internet, searching for problem
 resolution vii, viii
inventory expiration 11
ioo command 17

J

- Journal File System 17
- journal-based backup
 - Windows 35

K

- knowledge bases, searching vii

L

- LAN-free backup 12
- licensing xi
- Linux servers
 - performance recommendations 19
- log
 - performance 9
- logical storage pools 13

M

- Macintosh client
 - antivirus software 35
 - Extended Attributes 35
- manage operations 4
- maximum segment size (MSS) 51
- maximum transmission unit (MTU) 51
 - NetWare 53
- Maximum Transmission Unit (MTU) 31
- MAXNUMMP server option 13, 29, 33
- MAXSESSIONS server option 5, 29, 33
- MEMORYEFFICIENTBACKUP client
 - option 27
- migration
 - Hierarchical Storage Manager 37
 - processes 14
 - thresholds 14
- mount points, virtual 36
- MOVE DATA command 5
- MOVEBATCHSIZE server option 5, 23
- MOVESIZETHRESH server option 5, 23
- multi-client backups and restores 36
- multiple client sessions 29
- multiple session backup and restore 33

N

- NetWare
 - client cache tuning 53
- networks
 - dedicated 49
 - for backup 49
 - protocol tuning 47
 - settings
 - AIX 50
 - Sun Solaris 54
 - z/OS 55
 - traffic 50
- NTFS file compression 20
- NTFS file system 20

P

- problem determination
 - describing problem for IBM Software Support x
 - determining business impact for IBM Software Support ix
 - submitting a problem to IBM Software x
- PROCESSORUTILIZATION client
 - option 28
- PROFILE.TCPIP configuration data
 - set 55
- publications
 - download v
 - order v
 - search v
 - Tivoli Storage Manager v

Q

- QUIET client option 29

R

- RAID arrays 9, 24
- raw logical volumes 17
 - advantages and disadvantages 19
- raw partitions 17, 19
- recommended values by platform 35
- REGISTER NODE server command 33
- RESOURCEUTILIZATION client
 - option 29, 33, 36
- RESTOREINTERVAL server option 6

S

- SAN discovery
 - zSeries xiii
- SANDISCOVERYTIMEOUT server
 - option xiii
- scheduling
 - processes 12
 - sessions 12
- server
 - hardware 8
- server activity log
- searching 12
- server commands 13
 - BACKUP DB 9
 - database manager 3
 - DEFINE COPYGROUP 11
 - DEFINE DEVCLASS 21
 - DEFINE STGPOOL 13, 14, 22
 - EXPIRE INVENTORY 11
 - REGISTER NODE 33
 - SET MAXCMDRETRIES 50
 - SET QUERYSCHEDPERIOD 50
 - SET RETRYPERIOD 50
 - storage pools 4
 - UPDATE COPYGROUP 11
 - UPDATE NODE 29, 33
 - UPDATE STGPOOL 14, 22
- server options 3
 - best performance settings by platform 16

server options (*continued*)

- COMMETHOD SHAREDMMEM 20, 35
- DBMEMPERCENT 3
- DISKSTGPOOLMEMSIZE 4
- EXPINTERVAL 4, 11
- MAXNUMMP 13, 29, 33
- MAXSESSIONS 5, 29, 33
- MOVEBATCHSIZE 5, 23
- MOVESIZETHRESH 5, 23
- RESTOREINTERVAL 6
- SANDISCOVERYTIMEOUT xiii
- TCPNODELAY 6, 15
- TCPWINDOWSIZE 6, 17, 20
- TXNBYTELIMIT 7, 23
- TXNGROUPMAX 7, 17, 23, 32
- server tuning overview 1
- SET MAXCMDRETRIES server
 - command 50
- SET QUERYSCHEDPERIOD server
 - command 50
- SET RETRYPERIOD server command 50
- sliding window 32
- Software Support
 - contact ix
 - describing problem for IBM Software Support x
 - determining business impact for IBM Software Support ix
- storage
 - volumes 13
- Storage Agent 15
- storage pool
 - backup and restore 12
 - migrating files 14
 - migration 14
- storage pools
 - cached disk 13
- Sun Solaris
 - server and client TCP/IP tuning 54
 - server performance
 - recommendations 19
 - TCPWINDOWSIZE client option 54
- support information vii
- system memory 3

T

- tape drives
 - cleaning 21
 - compression 21
 - on a SCSI bus 24
 - required number 21
 - streaming rate 23
 - transfer rate 23
- TAPEPROMPT client option 31
- TCP communication buffer 31
- TCP/IP
 - AIX server and client tuning 50
 - concepts 47
 - connection availability 47
 - data transfer block size 47
 - error control 47
 - flow control 47
 - functional groups
 - application layer 47
 - internetwork layer 47

- TCP/IP (*continued*)
 - functional groups (*continued*)
 - network layer 47
 - transport layer 47
 - HP-UX server and client tuning 19
 - maximum segment size (MSS) 51
 - maximum transmission unit (MTU) 51
 - NetWare
 - client cache tuning 53
 - maximum transmission unit (MTU) 53
 - packet assembly and disassembly 47
 - sliding window 32, 49
 - Sun Solaris server and client
 - tuning 54
 - tuning 47
 - window values 47
 - z/OS tuning 55
- TCP/IP and z/OS UNIX system services
 - performance tuning 56
- TCPBUFFSIZE client option 31
- TCPIP.DATA 55
- TCPNODELAY client option 31
- TCPNODELAY server option 6, 15
- TCPWINDOWSIZE client option 17, 20, 32, 54
- TCPWINDOWSIZE server option 6, 17, 20
- thresholds
 - migration 14
- throughput
 - estimating
 - in untested environments 21
- Tivoli technical training vii
- training, Tivoli technical vii
- transaction size 32
- tsmdlst utility xii
- tuning 3
- TXNBYTELIMIT client option 17, 23, 32
- TXNBYTELIMIT server option 7, 23
- TXNGROUPMAX server option 7, 17, 23, 32

U

- UFS file system volumes 19
- UNIX file systems
 - advantages and disadvantages 19
- UPDATE COPYGROUP server
 - command 11
- UPDATE NODE server command 29, 33
- UPDATE STGPOOL server
 - command 14, 22

V

- Virtual Memory Manager 17, 50
- VIRTUALMOUNTPOINT client
 - option 27
- VIRTUALNODENAME client option 36
- vmo command 17
- VxFS file system 19

W

- Windows
 - journal-based backup 35
 - performance recommendations 20
- workload requirements 10

Z

- z/OS client
 - TCP/IP tuning 55



Program Number: 5608-E01
5608-E02
5608-E03
5608-E07
5608-E12

Printed in USA

GC23-9788-01

