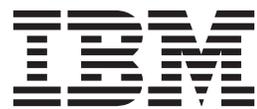


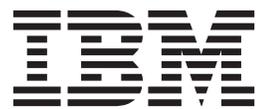
IBM Tivoli Storage Manager for Enterprise Resource
Planning
Version 6.3

*Data Protection for SAP[®]
Installation and User's Guide for DB2*



IBM Tivoli Storage Manager for Enterprise Resource
Planning
Version 6.3

*Data Protection for SAP[®]
Installation and User's Guide for DB2*



Note:

Before using this information and the product it supports, read the information in "Notices" on page 163.

This edition applies to Version 6.3 IBM Tivoli Storage Manager for Enterprise Resource Planning (product number 5608-E05), available as a licensed program product, and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters. This edition replaces SC33-6341-11.

© **Copyright IBM Corporation 1995, 2011.**

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

Contents

Figures	vii
--------------------------	------------

Tables	ix
-------------------------	-----------

About this publication xi

Who Should Read This Publication	xi
Note on Advanced Copy Services and FlashCopy Manager	xi
Publications	xi
Tivoli Storage Manager publications	xii
Tivoli Storage FlashCopy Manager publications	xiv
Support information	xiv
Getting technical training	xiv
Searching knowledge bases	xv
Contacting IBM Software Support	xvi

New for Tivoli Storage Manager for ERP Version 6.3 xix

Chapter 1. Protection for SAP database servers 1

Data Protection for SAP for DB2 overview	1
Data Protection for SAP integration with SAP	2
DB2 command line processor	3
DB2 Backup Object Manager utility	4
DB2 Log Manager	6
Backup objects and types of failures	8
Administration Assistant function for Data Protection for SAP	9
Administration Assistant function for Data Protection for SAP: Features	12
Minimizing backup and restore processing with Tivoli Storage Manager for ERP	12

Chapter 2. Planning for Data Protection for SAP for DB2 operations 15

Database Server Considerations	15
Network Considerations	16
Backup Server Considerations	16
Storing data on a Tivoli Storage Manager server	17
Alternate or parallel backup paths and backup servers	18
Archiving Inactive Data	18
Restore versus Backup	19
Planning for using IBM HACMP for AIX	19
HACMP impact on Data Protection for SAP for DB2	20
Digital Signing of Executable Files for Windows Systems	21

Chapter 3. Installing Data Protection for SAP for DB2 for V6.3 23

Required installation tasks	23
---------------------------------------	----

Installing the Data Protection for SAP for DB2 V6.3 base product	24
Prerequisites	24
Installing Tivoli Storage Manager for ERP for DB2 in silent mode	25
Installing Tivoli Storage Manager for ERP for DB2 on UNIX or Linux	26
Installing Tivoli Storage Manager for ERP for DB2 on Windows	28
Uninstalling the Old Version of Tivoli Storage Manager for ERP for DB2 under UNIX or Linux	30
Uninstalling the Old Version of Data Protection for SAP for DB2 under Windows	30
Verifying the Initial and Upgrade Installation	30
Installing the Administration Assistant function for Data Protection for SAP V6.3	31
Prerequisites for Installing the Administration Assistant function for Data Protection for SAP	31
Installing the Administration Assistant function for Data Protection for SAP Server-Level Components	32

Chapter 4. Upgrading to Data Protection for SAP for DB2 for V6.3 35

Upgrade the Data Protection for SAP for DB2 V6.3 base product	35
Migrate the Data Protection for SAP for DB2 profile	36
Upgrade the Administration Assistant function for Data Protection for SAP V6.3	36
Migrate Administration Assistant function for Data Protection for SAP data from a previous release	37

Chapter 5. Configuring Data Protection for SAP for DB2 39

Configuration tasks for the Data Protection for SAP for DB2 base product	39
Verification tasks	39
Profile tasks	39
DB2 tasks	43
Administration Assistant function for Data Protection for SAP tasks	46
Distributed file system tasks	50
HACMP tasks	52
Configuration tasks for Tivoli Storage Manager	54
Tivoli Storage Manager client tasks	54
Tivoli Storage Manager server tasks	58

Chapter 6. Protecting SAP data with Data Protection for SAP for DB2 V6.3 67

Backing up SAP data	67
Implementing the Strategy by Scheduling Automated Backup Runs	67

Windows Scheduling Example	68
Starting Backups in a Non-Partitioned Database Environment	69
Using DB2 Single System View for Backup	69
Multiple DB2 Log File Copies	69
Schedule Batch Sample	70
Full Offline Backup Batch File Sample	70
Full Offline Backup Shell Script Sample	71
Segmenting large backup objects	71
Restoring SAP data	73
Starting Restores in a Non-Partitioned Database Environment	73
Redirected Restore in Automatic Mode	74
Tablespace Definition Information (TDI)	75
Redirected Restore in Batch Mode	78
Redirected Restore in Interactive Mode	79
Sample Work Flow for Redirected Restore	80
Redirected Restore Plausibility Checks	83
DB2 Redirected Restore Using Backup Object Manager	84
Create Tablespace Definition Information	85
Redirected Restore Prerequisites	85
Tablespace Definition Information	86
Protecting SAP data with the Administration Assistant function for Data Protection for SAP	87
Administering User IDs	87
Specifying a new Administration Assistant function for Tivoli Storage Manager for ERP	87
Generating Reports Using Report Templates	88
Requesting a Report from the Administration Assistant function for Data Protection for SAP Client	89
Starting and Stopping the Administration Assistant function for Data Protection for SAP Manually	89
Changing the Password for the Administration Assistant function for Data Protection for SAP Database User ID	90

Chapter 7. Performance tuning for Data Protection for SAP for DB2 91

Overview of a balanced system.	91
Example of a disk bottleneck	92
Example of a network or Tivoli Storage Manager bottleneck	93
Viewing performance data	93
Drilling Down on Special Situations	94
Using reports.	95
Reporting on the Performance of Backup Operations	95
Reporting on Backup Status	97
Creating a Report	98
Reporting on Failed Actions	98
Modifying Report Output	98
Reporting on Operations Details	99
Reporting on Backup Operation Trends.	100
Reporting on Data Protection for SAP for DB2 Activities.	101
Working with Report Templates	102
Server-related tuning	102
Managing Data on the Backup Server	102

Alternate Network Paths and Servers	102
Options	102
Performance Options of Data Protection for SAP for DB2	103
Buffer Copies	104
Buffer Size	104
Compression	104
Automation Options for Data Protection for SAP for DB2.	105
Data transfer	106
Observations on the Data Protection for SAP for DB2 Data Throughput	106
Data Protection for SAP for DB2 Performance Sensors	107
General Performance Considerations	107
Multiple Servers	108
Multiple Sessions	109
Multiplexing	109
Multiple Network Paths	110
Storage space	110
Automated Tablespace Adaptations	111
Tablespace Normalizing	111
Tablespace Scaling.	112

Chapter 8. Troubleshooting IBM Tivoli Storage Manager for Enterprise Resource Planning 115

Troubleshooting IBM Tivoli Storage Manager for Enterprise Resource Planning common problems	115
Random problems	115
Reproducible (repeatable) problems	115
Internet Protocol version 6 (IPv6) support	116
Understanding the Setup	116
Providing information to IBM or Tivoli support	118
Troubleshooting Data Protection for SAP for DB2 problems	118
General problem resolution.	118
Location of log files	121
How to find files containing message output (log files)	121
DB2 vendor reason codes	121

Chapter 9. Data Protection for SAP for DB2 reference information. 123

Commands used with Data Protection for SAP for DB2	123
Versioning	123
Backups and Restores in Partitioned Database Environments	123
Managing Backup Objects	125
Backup Object Manager Command Options	126
Backup Command (Backup Database Data)	128
Delete Commands (Remove Backup Objects from Tivoli Storage Manager)	128
Password Command (Verify and Save Tivoli Storage Manager Password)	129
Query Commands (List Backup Objects)	129
Restore Commands (Restore Backup Objects)	130
BACKOM command examples	131
UNIX or Linux Crontab Example.	132

Crontab File Sample	132
The Data Protection for SAP for DB2 Profile . . .	133
Tivoli Storage Manager for ERP for DB2 profile parameter descriptions	134
Sample Tivoli Storage Manager for ERP for DB2 Profile for UNIX or Linux	139
Sample Data Protection for SAP for DB2 Profile for Windows	142
Defining the Custom SQL file	145
Defining Thresholds Using the Custom SQL File	147
Sample Custom SQL File	147
Data Protection for SAP for DB2 files and samples	147
Sample Shell Script for Scheduling a Report from a UNIX Scheduling Client	148
Sample Command File for Scheduling a Report from a Windows Scheduling Client	148
Client User Options File Sample (dsm.opt)	
UNIX and Linux	149
Client User Options File Sample (dsm.opt)	
Windows	149
Client System Options File Sample (dsm.sys)	149
Include/Exclude List Sample (UNIX and Linux)	150
Include/Exclude List Sample (Windows) . . .	151
Client Options Files Sample (<i>server.opt</i>). . . .	151

Sample DB2 Vendor Environment File	152
Data Protection for SAP for DB2 planning sheets	152
Data Protection for SAP for DB2 (base product) planning sheet	152
Administration Assistant function for Data Protection for SAP planning sheet	154
Tips for network settings	158
Network Settings of the Tivoli Storage Manager	158
Networks with Large Bandwidth-Delay Product	159
SP Switch (RISC 6000)	159

Appendix. Accessibility features for the Tivoli Storage Manager product family. 161

Notices 163
Trademarks 165

Glossary 167

Index 189

Figures

1.	Scope of Data Protection for SAP for DB2	1	17.	Performance Report - Tabular Presentation	
2.	Integration of Data Protection for SAP with DB2	2		Section	97
3.	DB2 Backup Architecture	3	18.	Status Report	97
4.	Data Protection for SAP Backup Object Manager	5	19.	Operations - Failure Report	98
5.	Log Management with DB2 Log Manager and		20.	Operations – Detailed Report	100
	Data Protection for SAP.	7	21.	Operations Daily Report	101
6.	DB2 Backup Objects	8	22.	Data Transfer for a Backup / Restore	103
7.	Administration Assistant function for Data		23.	Null Block Compression	104
	Protection for SAP Components (with Default		24.	High-level View of the Data Flow During	
	Port Numbers)	10		Backup	106
8.	Example of an SAP Landscape	11	25.	Performance Optimizing by Using Sensors	107
9.	Sample Environment for HACMP Takeover	20	26.	Data Protection for SAP Data Transfer	108
10.	Production Backup Example	68	27.	Multiple Servers	108
11.	Redirected Restore Overview.	84	28.	Parallel (Multiple) Sessions	109
12.	Indicating a Balanced Configuration	91	29.	Multiplexing	110
13.	Indicating a Disk Bottleneck	92	30.	Parallel (Multiple) Paths	110
14.	Indicating a Network or Tivoli Storage		31.	Tablespace Normalizing	111
	Manager Bottleneck.	93	32.	Tablespace Scaling	112
15.	Showing Data Throughput and I/O Utilization	94	33.	SAP and Data Protection for SAP	
16.	Performance Report - Graphical Presentation			configuration files on UNIX or Linux	117
	Section	96	34.	General Problem Isolation	119

Tables

1. Tivoli Storage Manager server publications	xii	10. Password Handling for UNIX or Linux	63
2. Tivoli Storage Manager storage agent publications	xii	11. Password Handling for Windows	64
3. Tivoli Storage Manager client publications	xii	12. DB2 Vendor Reason Codes	121
4. Tivoli Storage Manager data protection publications	xiii	13. Contents of the Custom SQL File	145
5. IBM Tivoli Storage Manager troubleshooting and tuning publications	xiii	14. Tags for Defining Thresholds in the Custom SQL File	147
6. Tivoli Storage FlashCopy Manager publications	xiv	15. Installation Parameters for Data Protection for SAP.	153
7. File Extensions for Shared Libraries	26	16. Installation Parameters for the Administration Assistant function for Data Protection for SAP	154
8. SERVER Statement and Appropriate Profile and Option File Settings.	40	17. Tuning Tivoli Storage Manager Configuration File Attributes	158
9. Configuration parameters for DB2 database backup and restore, and log archive and retrieve	44	18. Tuning of Network Settings.	159
		19. Tuning of SP Switch Buffer Pools	159

About this publication

This publication documents how to use IBM® Tivoli® Storage Manager for Enterprise Resource Planning Data Protection for SAP® Version 6.3. It describes the procedures needed to install and customize IBM Tivoli Storage Manager for Enterprise Resource Planning which is the interface between SAP® and Tivoli Storage Manager.

Who Should Read This Publication

This publication (or topic collection) is intended for system programmers and administrators who are responsible for implementing a backup solution in an SAP environment using the Tivoli Storage Manager. It describes the procedures needed to install and customize Data Protection for SAP, the interface between SAP and the Tivoli Storage Manager. The reader should be familiar with the documentation for SAP and Tivoli Storage Manager.

Note on Advanced Copy Services and FlashCopy Manager

The product *IBM Tivoli Storage Manager for Advanced Copy Services (TSM for ACS)* was replaced by *IBM Tivoli FlashCopy Manager*. References in this publication, and in error messages, to *TSM for ACS* are also applicable to the *Tivoli FlashCopy Manager*.

Publications

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

To search all publications, go to the Tivoli Storage Manager information center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3>.

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

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

You can also order some related publications from the IBM Publications Center website. The website provides information about ordering publications from countries other than the United States. In the United States, you can order publications by calling 1-800-879-2755.

Tivoli Storage Manager publications

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

Table 1. Tivoli Storage Manager server publications

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

Table 2. Tivoli Storage Manager storage agent publications

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

Table 3. Tivoli Storage Manager client publications

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

Table 3. Tivoli Storage Manager client publications (continued)

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

Table 4. Tivoli Storage Manager data protection publications

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

Table 5. IBM Tivoli Storage Manager troubleshooting and tuning publications

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

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

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

Note: You can find information about IBM System Storage® Archive Manager at http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/c_complydataretention_ovr.html.

Tivoli Storage FlashCopy Manager publications

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

Table 6. Tivoli Storage FlashCopy Manager publications

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

Support information

You can find support information for IBM products from various sources.

Start at the IBM Support Portal: <http://www.ibm.com/support/entry/portal/>. You can select the products that you are interested in and search for a wide variety of relevant information.

Getting technical training

Information about Tivoli technical training courses is available online.

Go to the following websites to sign up for training, ask questions, and interact with others who use IBM storage products.

Tivoli software training and certification

Choose from instructor led, online classroom training, self-paced Web classes, Tivoli certification preparation, and other training options at <http://www.ibm.com/software/tivoli/education/>

Tivoli Support Technical Exchange

Technical experts share their knowledge and answer your questions in webcasts at http://www.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html.

Storage Management community

Interact with others who use IBM storage management products at <http://www.ibm.com/developerworks/servicemanagement/sm/index.html>

Global Tivoli User Community

Share information and learn from other Tivoli users throughout the world at <http://www.tivoli-ug.org/>.

IBM Education Assistant

View short "how to" recordings designed to help you use IBM software products more effectively at <http://publib.boulder.ibm.com/infocenter/ieduasst/tivv1r0/index.jsp>

Searching knowledge bases

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

Begin by searching the Tivoli Storage Manager Information Center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3>. From this website, you can search the current Tivoli Storage Manager documentation.

Searching the Internet

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

To search multiple Internet resources, go to the IBM support website at <http://www.ibm.com/support/entry/portal/>.

You can search for information without signing in. Sign in using your IBM ID and password if you want to customize the site based on your product usage and information needs. If you do not already have an IBM ID and password, click **Sign in** at the top of the page and follow the instructions to register.

From the support website, you can search various resources including:

- IBM technotes
- IBM downloads
- IBM Redbooks® publications
- IBM Authorized Program Analysis Reports (APARs)

Select the product and click **Downloads** to search the APAR list.

If you still cannot find a solution to the problem, you can search forums and newsgroups on the Internet for the latest information that might help you find problem resolution.

An independent user discussion list, ADSM-L, is hosted by Marist College. You can subscribe by sending an email to listserv@vm.marist.edu. The body of the message must contain the following text: `SUBSCRIBE ADSM-L your_first_name your_family_name`.

To share your experiences and learn from others in the Tivoli Storage Manager and Tivoli Storage FlashCopy Manager user communities, go to the following wikis:

Tivoli Storage Manager wiki

<http://www.ibm.com/developerworks/wikis/display/tivolistoragemanager>

Tivoli Storage FlashCopy Manager wiki

[https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli Storage FlashCopy Manager](https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli%20Storage%20FlashCopy%20Manager)

Finding product fixes

A product fix to resolve your problem might be available from the IBM software support website.

You can determine what fixes are available by checking the IBM software support website at <http://www.ibm.com/support/entry/portal/>.

- If you previously customized the site based on your product usage:
 1. Click the link for your product, or a component for which you want to find a fix.
 2. Click **Downloads**, and then click **Fixes by version**.
- If you have not customized the site based on your product usage, click **Downloads** and search for your product.

Receiving notification of product fixes

You can receive notifications about fixes, flashes, upgrades, and other news about IBM products.

To sign up to receive notifications about IBM products, follow these steps:

1. From the support page at <http://www.ibm.com/support/entry/portal/>, click **Sign in to create, manage, or view your subscriptions** in the **Notifications** pane.
2. Sign in using your IBM ID and password. If you do not have an ID and password, click **register now** and complete the registration process.
3. Click **Manage all my subscriptions** in the **Notifications** pane.
4. Click the **Subscribe** tab and then click **Tivoli**.
5. Select the products for which you want to receive notifications and click **Continue**.
6. Specify your notification preferences and click **Submit**.

Contacting IBM Software Support

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

To obtain help from IBM Software Support, complete the following steps:

1. Ensure that you have completed the following prerequisites:
 - a. Set up a subscription and support contract.
 - b. Determine the business impact of your problem.
 - c. Describe your problem and gather background information.
2. Follow the instructions in “Submitting the problem to IBM Software Support” on page xviii.

Setting up a subscription and support contract

Set up a subscription and support contract. The type of contract that you need depends on the type of product you have.

For IBM distributed software products (including, but not limited to, IBM Tivoli, Lotus®, and Rational® products, as well as IBM DB2® and IBM WebSphere® products that run on Microsoft Windows or on operating systems such as AIX or Linux), enroll in IBM Passport Advantage® in one of the following ways:

- **Online:** Go to the Passport Advantage website at <http://www.ibm.com/software/lotus/passportadvantage/>, click **How to enroll**, and follow the instructions.
- **By telephone:** You can call 1-800-IBMSERV (1-800-426-7378) in the United States. For the telephone number to call in your country, go to the IBM Software Support Handbook web page at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.

Determining the business impact

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

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.

Describing the problem and gathering background information

When explaining a problem to IBM, it is helpful to be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently.

To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be re-created? If so, what steps led to the failure?
- Have any changes been made to the system? For example, hardware, operating system, networking software, and so on.
- Are you using a workaround for this problem? If so, be prepared to explain it when you report the problem.

Submitting the problem to IBM Software Support

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

Online

Go to the IBM Software Support website at [http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_\(general\)](http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_(general)). Sign in to access IBM Service Requests and enter your information into the problem submission tool.

By telephone

For the telephone number to call in your country, go to the IBM Software Support Handbook at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.

New for Tivoli Storage Manager for ERP Version 6.3

There are new features for Tivoli Storage Manager for ERP for Oracle.

There are no new features for Tivoli Storage Manager for ERP for DB2 Version 6.3.

Chapter 1. Protection for SAP database servers

Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP for DB2 protects SAP® system data and is integrated with the database-specific utilities of IBM DB2. Data Protection for SAP improves the availability of SAP database servers and reduces administration workload with automated data protection features that are designed specifically for SAP environments.

Data Protection for SAP provides these features and functions.

Data Protection for SAP for DB2 overview

Data Protection for SAP for DB2 architecture and product features are discussed.

Data Protection for SAP and Tivoli Storage Manager provide a reliable, high performance, and production-oriented solution that enables back up and restore of DB2-based SAP® systems. It is integrated with DB2 backup and recovery facilities and applies SAP backup and recovery procedures. Data Protection for SAP is optimized for SAP databases and therefore provides efficient management of large data volumes.

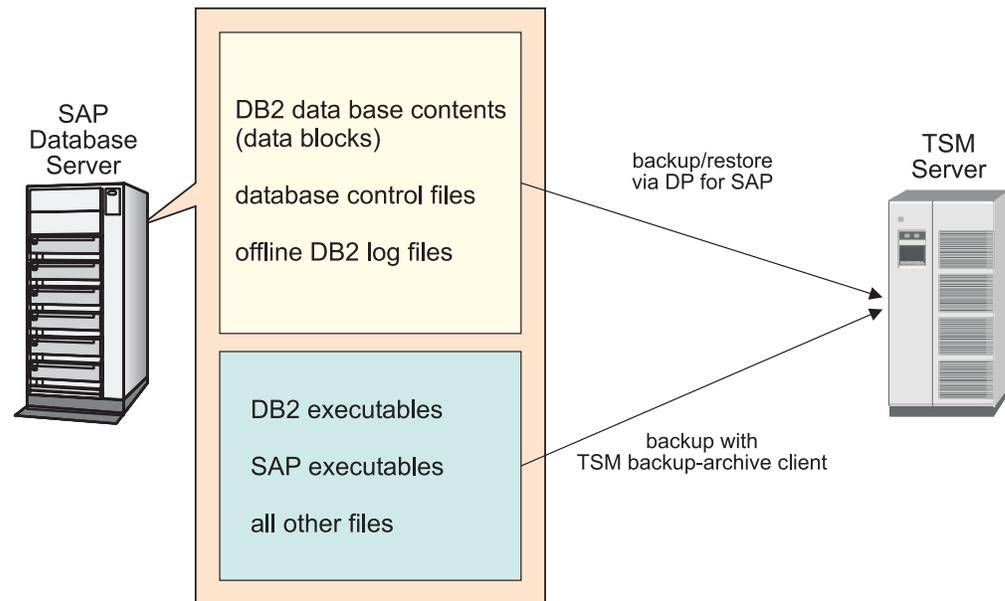


Figure 1. Scope of Data Protection for SAP for DB2

As demonstrated in this graphic, SAP backup and recovery utilities center on database objects where more than 90% of the data resides on an SAP database server. As a result, Data Protection for SAP backs up and restores database contents, database specific control files, e.g. the database configuration, the history and the log file header, and offline DB2 log files.

Other files (such as SAP and DB2 executable files) can be backed up using the IBM Tivoli Storage Manager Backup-Archive Client. This is important for disaster recovery purposes, as all SAP and DB2 executable files must be available before using Data Protection for SAP to restore and recover the database.

Data Protection for SAP integration with SAP®

Data Protection for SAP operates as a transparent link between DB2 and the Tivoli Storage Manager. A shared library is dynamically linked by DB2 backup/archive processes.

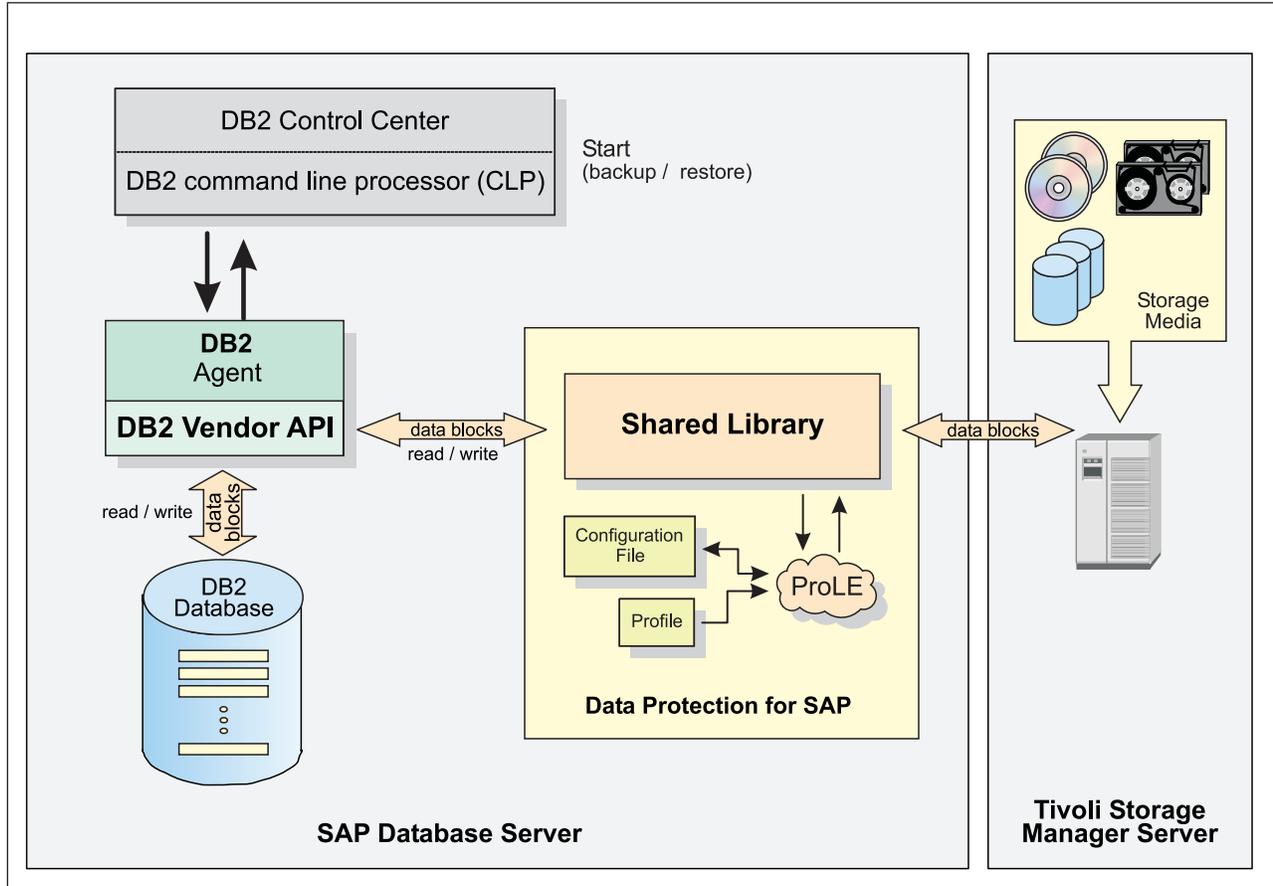


Figure 2. Integration of Data Protection for SAP with DB2

Data Protection for SAP also provides the Administration Assistant function for Data Protection for SAP which is used to increase administrator productivity. The Administration Assistant function for Data Protection for SAP can control multiple instances of Data Protection for SAP, communicates with Data Protection for SAP through TCP/IP, and typically resides on a different server. It is used to configure a Data Protection for SAP instance, monitor data transfer performance, backup status of all SAP systems backed up by Data Protection for SAP, and Tivoli Storage Manager server activity related to SAP. In addition, the Administration Assistant can remotely monitor and administer all Data Protection for SAP instances through an applet running on a Web browser. Information regarding how to use the Administration Assistant to register an SAP instance during installation or at a later time is available in "Specifying a new Administration Assistant function for Tivoli Storage Manager for ERP" on page 87.

DB2 command line processor

The DB2 Command Line Processor (CLP) interprets commands for the DB2 database and passes control to a DB2 Server Process. In the case of Data Protection for SAP for DB2, the "LOAD *libraryname*" option instructs DB2 to invoke the Data Protection for SAP shared library. This process launches the backup or restore operation, dynamically loads the library, and communicates with Data Protection for SAP through the Vendor API.

For starting a backup or restore, the DB2 CLP communicates with the DB2 Server Process and provides information to the Server Process for processing the database.

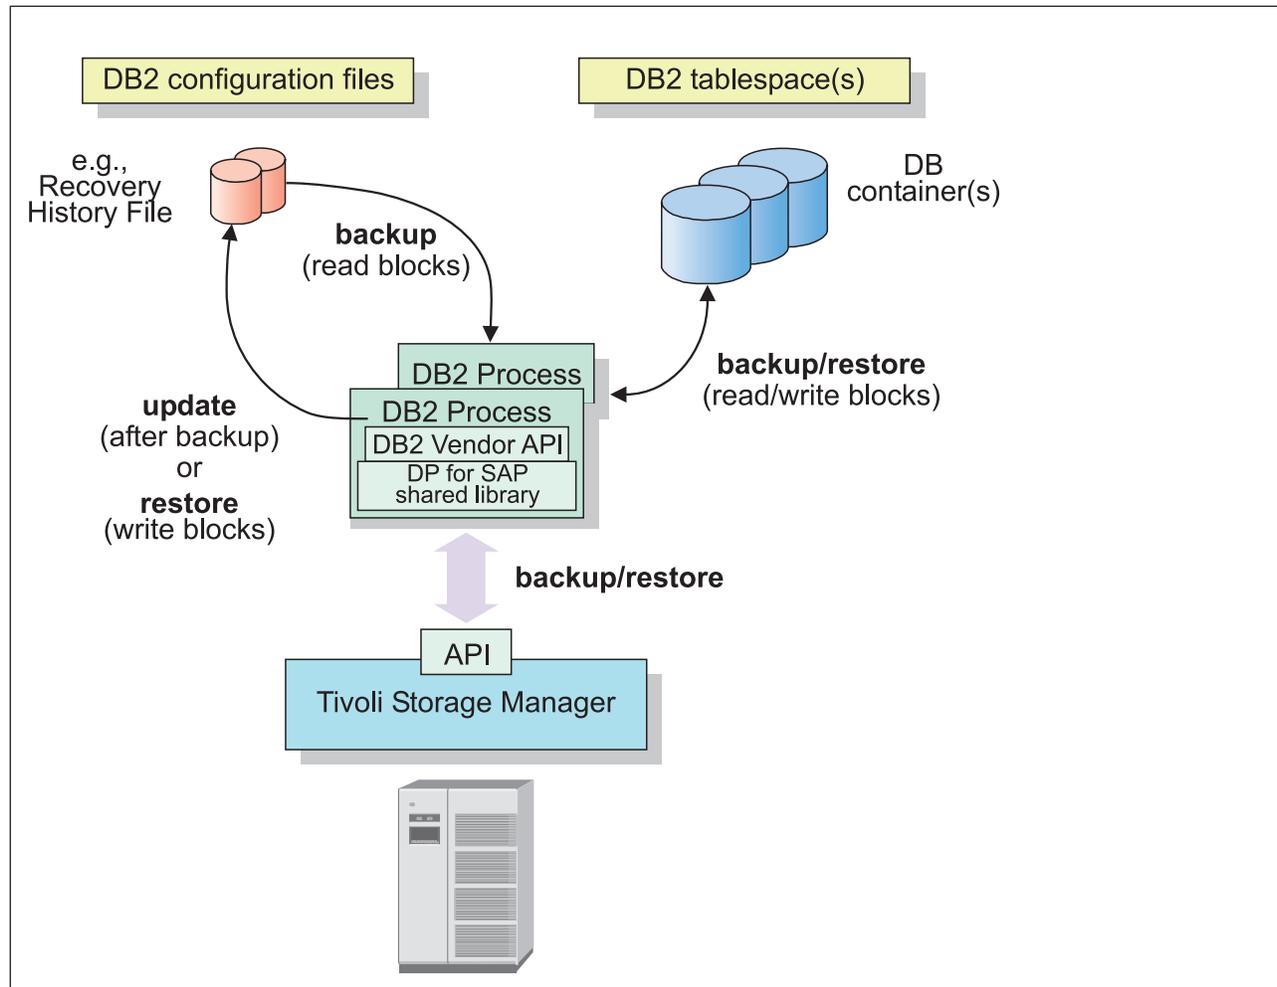


Figure 3. DB2 Backup Architecture

The DB2 **BACKUP DATABASE** command performs this DB2 Server process:

- creates a unique timestamp to identify the backup
- loads Data Protection for SAP dynamically as a shared library
- reads the data from the database containers
- reads the DB2 configuration files
- creates data blocks containing the backup image and passes these blocks to the data mover part of Data Protection for SAP

The Data Protection for SAP shared library sends the data to the Tivoli Storage Manager server storage (tape or disk). At the end of the backup process, the DB2 Server process logs the backup in the Recovery History File.

The DB2 **RESTORE DATABASE** command performs this DB2 Server process:

- loads Data Protection for SAP dynamically as a shared library
- requests the backup data from the shared library

The Data Protection for SAP shared library

- checks with the Tivoli Storage Manager if the backup image is available
- retrieves the data blocks from TSM
- passes the data blocks to the DB2 Server Process

The DB2 Server Process

- restores the DB2 data to the database containers
- logs the restore in the Recovery History File

DB2 Backup Object Manager utility

Backup objects, such as database or tablespace backups and DB2 log files, can be managed with the Data Protection for SAP for DB2 Backup Object Manager. Information about Backup Object Manager commands and options is provided.

The Backup Object Manager is a utility that performs these tasks:

- Verify and store a Tivoli Storage Manager password.
- Find backup objects in Tivoli Storage Manager.
- Check the properties of the backup objects in Tivoli Storage Manager.
- Remove any backup object from Tivoli Storage Manager.
- Backup database and selected tablespaces.
- Restore database and tablespace backups to the corresponding database.
- Retrieve files from Tivoli Storage Manager and restore them to the file system.
- Perform a redirected restore of databases (cloning).

The Backup Object Manager is designed to handle DB2 log files archived with Data Protection for SAP, the SAP[®] tool BRARCHIVE, and those files archived with Data Protection for SAP and the DB2 Log Manager. No special Backup Object Manager customization or configuration is necessary. Due to the log chain concept used by the DB2 Log Manager, all log files archived on Tivoli Storage Manager with Data Protection for SAP will be associated to one of these chains by the Backup Object Manager. However, the SAP-DB2 Administration Tools BRARCHIVE and BRRESTORE do not support the log chain concept. Therefore, log files archived with BRARCHIVE and Data Protection for SAP will be associated with a default value. For example, the first log chain is '0' or 'C0000000'. However, log files archived with the DB2 Log Manager and Data Protection for SAP will be associated with the appropriate log chain number and handled by the Backup Object Manager accordingly. Detailed information regarding the DB2 Log Manager and the log chain concept is available in your DB2 *Administration Guide* documentation.

This graphic displays how the Backup Object Manager interacts with the Tivoli Storage Manager server and the SAP database Server:

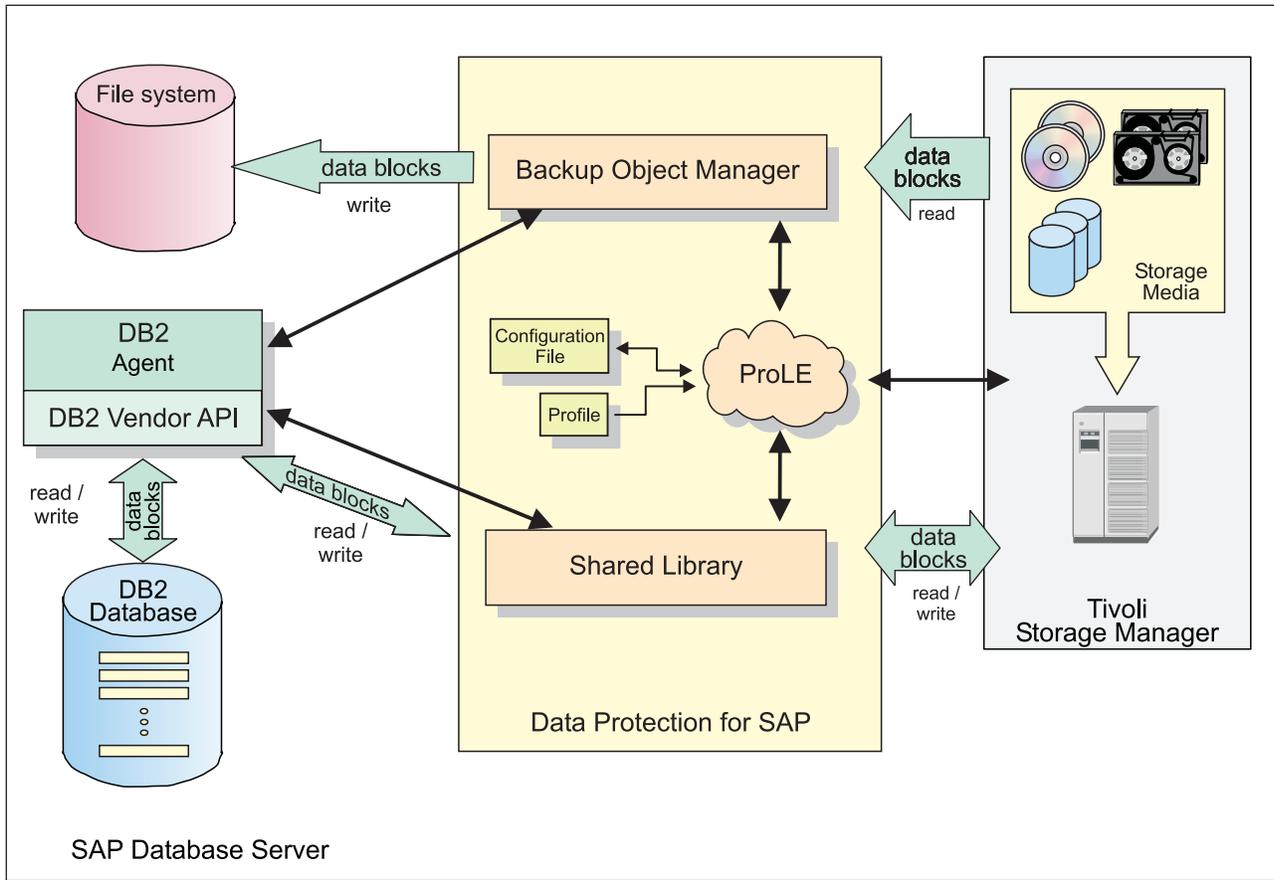


Figure 4. Data Protection for SAP Backup Object Manager

The Backup Object Manager works with database backups, DB2 log files, and raw files that might comprise any files of the file system. The tasks that can be performed with the Backup Object Manager are processed in different ways:

- Requests to verify the Tivoli Storage Manager password are passed directly to Tivoli Storage Manager.
- Requests to display or delete any data are answered by accessing the Tivoli Storage Manager server directly, thus working with the data actually available on Tivoli Storage Manager.
- Requests to restore DB2 log files and raw files are also processed using the Tivoli Storage Manager client.
- Requests to backup or restore any DB2 database data are routed to the DB2 agent. The DB2 agent employs the Data Protection for SAP shared library.

The Backup Object Manager is available for use upon successful installation and setup of Data Protection for SAP. Since the Backup Object Manager utilizes the settings in the Data Protection for SAP profile and configuration file and the settings of the XINT_PROFILE, TDP_DIR, and DB2_VENDOR_LIB environment variables, no additional installation and setup steps are required.

DB2 Log Manager

Data Protection for SAP for DB2 is integrated with the built-in DB2 Log Manager. As a result, when Data Protection for SAP is registered within the DB2 database configuration, the DB2 Log Manager uses Data Protection for SAP for archiving and retrieving log files.

Log files used in an SAP® environment are in one of these four states:

Online active

The log file is used by DB2 for current logging transactions.

Online retained

The log file is not used by DB2 for current logging transactions. However, it contains transactions with unwritten data pages. An unwritten data page is a page that has not received data from the buffer pool to disk. As a result, the log file is needed by DB2 to perform a crash recovery or roll-back operation. The DB2 Log Manager copies a filled online log file to a possible archive location. Do not use operating system commands for copying online log files.

Offline retained

The log file is not used by DB2 for current logging transactions and it does not contain transactions with unwritten data pages. In addition, it is not needed to perform a crash recovery or a roll-back operation. The log file is archived to a location specified by the database configuration. When archived successfully, DB2 deletes the log from the database log directory.

Archived

Filled or closed log files that were archived to Tivoli Storage Manager storage.

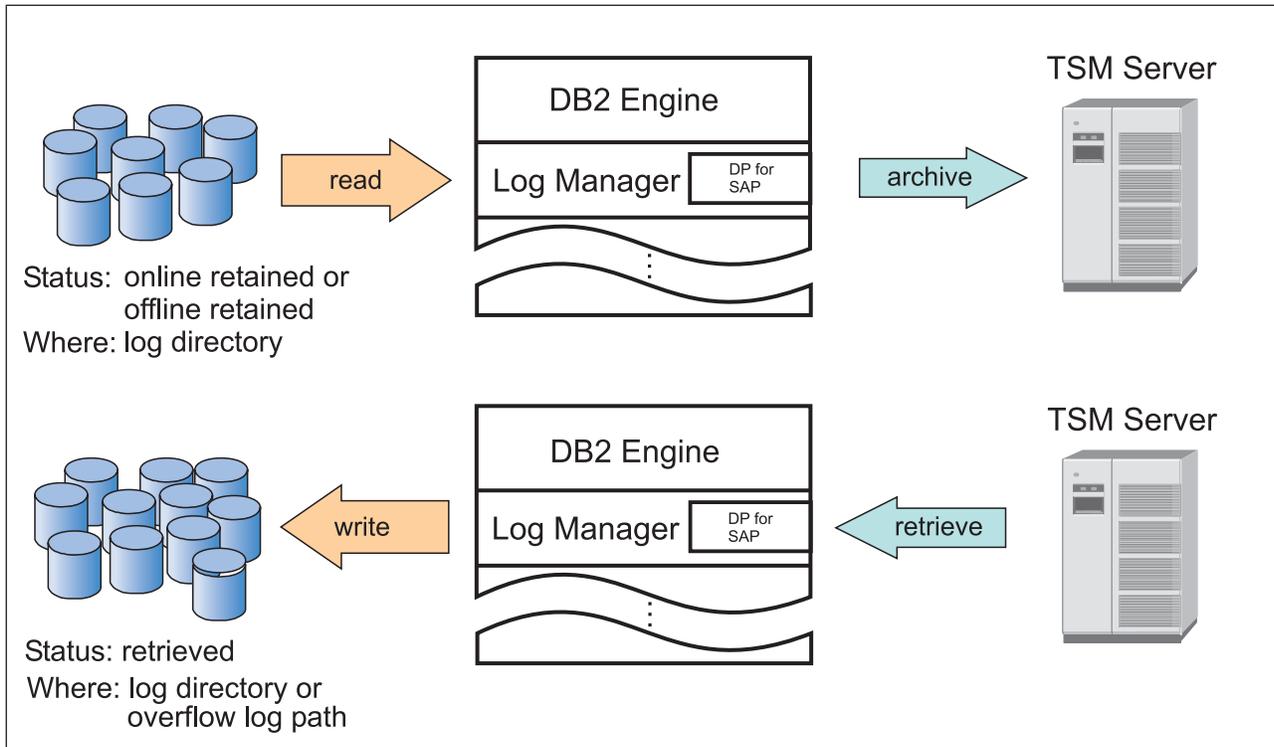


Figure 5. Log Management with DB2 Log Manager and Data Protection for SAP

Data Protection for SAP for DB2 is loaded dynamically by the DB2 Log Manager as a shared library on UNIX or Linux, or as a dynamic link library (DLL) on Windows, and runs as part of the DB2 engine. When a log file is ready to be archived (online/offline retained), the DB2 Log Manager starts the archive process by passing the file as blocks to Data Protection for SAP. The data is then sent (by Data Protection for SAP) to Tivoli Storage Manager storage.

When a database rollforward recovery is issued, the DB2 Log Manager first checks if the corresponding log files are located either in the log path or in an overflow log path as specified in the DB2 rollforward command invocation. If the log files are not found at one of these locations, the DB2 Log Manager accesses Data Protection for SAP to determine if the log images are available on Tivoli Storage Manager. If available, Data Protection for SAP retrieves the data from Tivoli Storage Manager and sends them as blocks to the DB2 Log Manager which writes the log files to the file system. The log files are then applied to the database using DB2 processes.

Detailed information about the DB2 Log Manager is available in your DB2 *Administration Guide*.

Backup objects and types of failures

Data Protection for SAP for DB2 backs up and restores SAP® database objects only as shown in Figure 6.

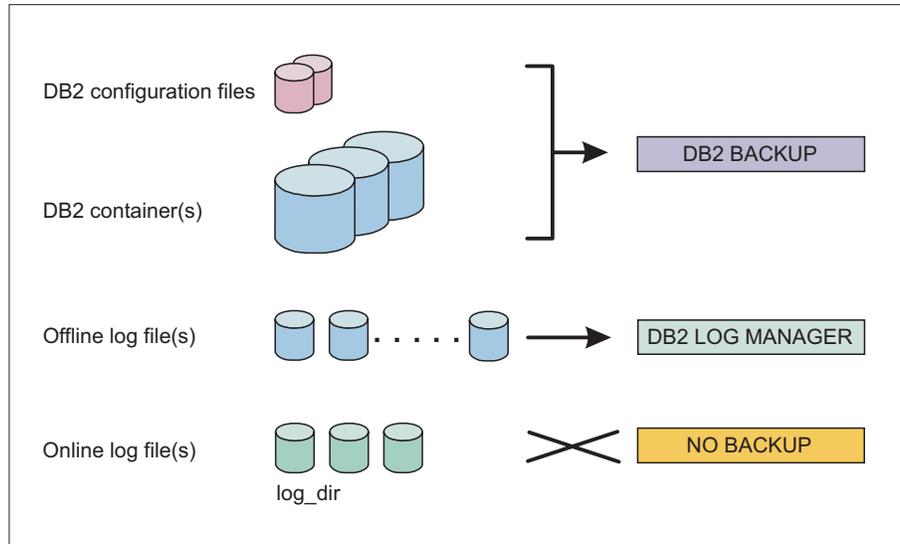


Figure 6. DB2 Backup Objects

Corrupt database

In case of a corrupted database (caused by user errors or transaction failures), the database can be restored to a specific point in time. Restoring only the database and configuration files should be sufficient for a specific point in time operation. As a result, a backup image of the database and the corresponding DB2 log files are required.

Hardware failure

In the event of a storage hardware failure, the database is typically restored to the most recent point in time. Thus, the most recent database image and DB2 log files are restored. However, the database executable files, SAP system data, and user data might also need to be restored in the event of a hardware failure. In order to protect the system against the loss of SAP executable files, user data, or even operating system data, use the Tivoli Storage Manager backup-archive client incremental backup feature. You can use the client to define an include-exclude list of files that to be backed up during incremental backup operation. The include-exclude list should exclude database container files and offline log files that have been backed up or archived by Data Protection for SAP. See "Include/Exclude List Sample (UNIX and Linux)" on page 150 or "Include/Exclude List Sample (Windows)" on page 151 for example include-exclude lists. Example include-exclude list files are also provided in the Data Protection for SAP installation directory.

Disaster recovery

For a complete disaster recovery operation, all operating system data must be restored along with the database image, DB2 log files, database executable files, SAP system data, and user data. To help prevent a complete loss of the operating system, use operating system utilities (such as mksysb for AIX®) to perform system backups. Such backups should be performed after installing, updating, or upgrading the operating system.

This will allow you to start your system from the backup medium. A configured TCP/IP environment and Tivoli Storage Manager Backup-Archive client installation should be included in a base backup in order to be able to restore all data. Since there is no provision for backing up online DB2 log files that are required for disaster recovery, place the DB2 log directory on a mirrored disk.

Administration Assistant function for Data Protection for SAP

The Administration Assistant comprises the client component and three server-level components (Server, Database Agent, and Database). Operations data is maintained in an internal database which helps prevent an insufficient memory problem in SAP® environments where a large number of Data Protection for SAP for DB2 instances are active. The internal database used by the Administration Assistant is managed by either the open-source database product Apache Derby or IBM DB2 data server. Apache Derby is bundled with, and installed by, the Administration Assistant. For more information on Apache Derby, see <http://db.apache.org/derby/>

If you prefer using the IBM DB2 data server, an existing DB2 installation must be present. It will be configured by the Administration Assistant. For more information on DB2, see <http://www.ibm.com/software/data/db2/>

The server-level components are installed together on one system (standard installation) or distributed across multiple systems (distributed installation). An example of a multiple system installation could be when the Server component resides on one system and the database components reside on a second system; or, each component is installed on a separate system. This type of distributed installation helps alleviate CPU loads on a single-system configuration (in large-scale environments) by distributing this load over two or three separate systems. If CPU load is not an issue, the single-system installation is typically used. The distributed installation requires that all connecting Data Protection for SAP instances be version 5.4 or higher. If a single-system installation is selected, earlier Data Protection for SAP versions can also connect to the Administration Assistant.

Each system hosting an Administration Assistant component can be running UNIX, Linux, or Windows. Separate configuration files are maintained by the Server (assist.cfg) and Database Agent (dbagent.cfg) component. User profiles ensure that a client user can access the data of only those SAP database servers for which permission has been granted.

This figure shows the communication relationships of the Administration Assistant components (port numbers shown are defaults).

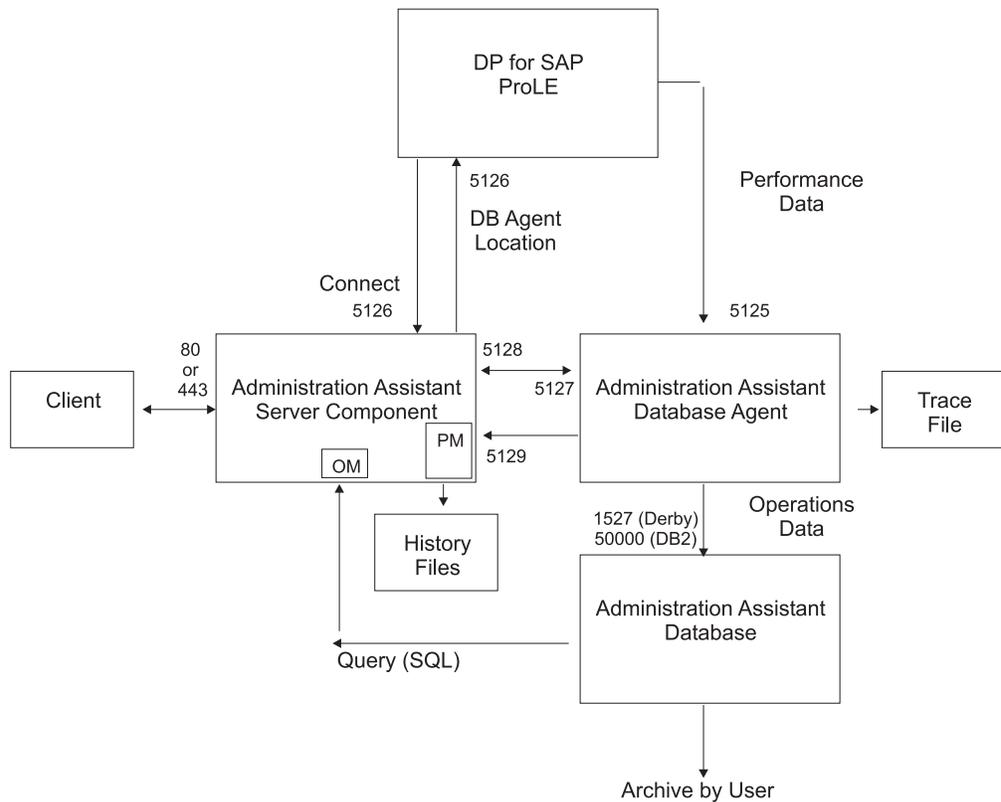


Figure 7. Administration Assistant function for Data Protection for SAP Components (with Default Port Numbers)

The Server component waits for the client requests for connections using either the HTTP or HTTPS protocols and also for connect requests (through TCP/IP) from the Data Protection for SAP ProLE service. After connecting to the Server component, the Data Protection for SAP ProLE service connects and communicates directly with the Database Agent to send backup and restore data requested through the Data Protection for SAP instance. The Database Agent collects this data and stores all information related to the Operations Monitor in the Administration Assistant database through the Database component. The Database Agent forwards performance data to the Administration Assistant Server component, which records it in history files. The retention time for this data is definable at installation time (default 14 days). This data is accessed when the clients request any of the Administration Assistant monitoring or analysis functions. The Administration Assistant server-level components must be running and connected to the Data Protection for SAP ProLE service during the backup and restore operations in order to receive and store the history data. The existence of the database-related components is transparent to the client user.

An SAP system landscape contains several SAP systems, such as production, development, test, and education systems. A single Administration Assistant Server component can monitor many SAP database servers. A typical example is shown in Figure 8 on page 11.

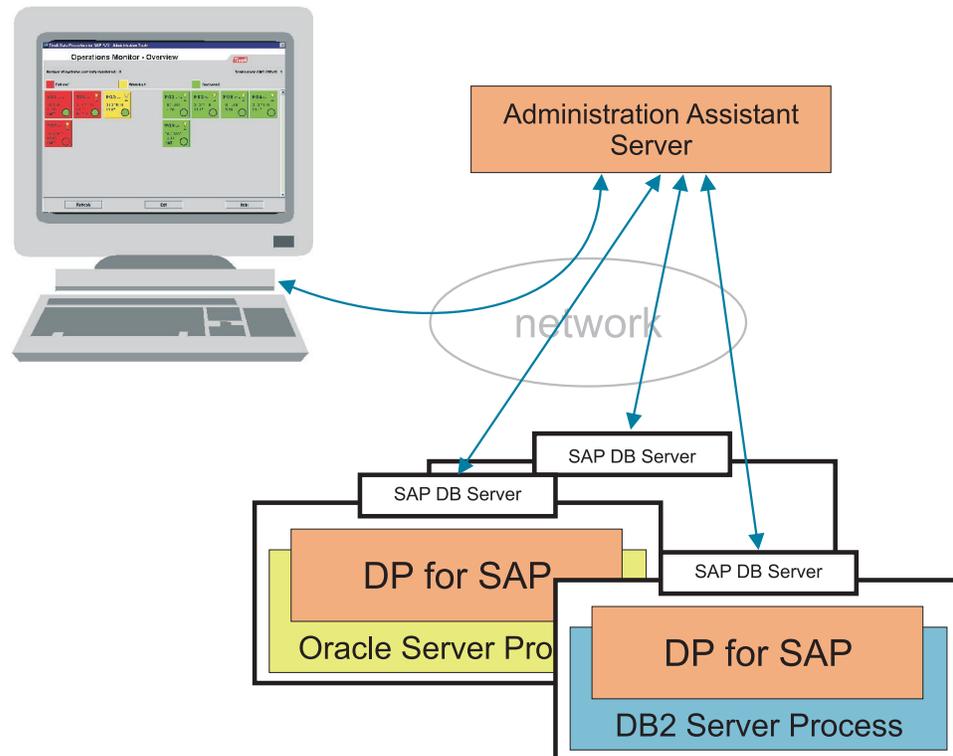


Figure 8. Example of an SAP Landscape

The Administration Assistant client is started from a browser by invoking the URL of the Server component host. The client is implemented as a Javaapplet that communicates with the Server component through a remote method invocation (RMI) connection.

- When the Administration Assistant Server component is started in non-secure mode (keyword `nonsecure` defined in the Server configuration file `assist.cfg`), it accepts connect requests from a client to its HTTP port using the HTTP protocol. In this case, further communication between the client and server is via TCP/IP.
- When the Server component is started in secure mode (keyword `nonsecure` omitted from the Server configuration file), it accepts connect requests from a client to its HTTPS port through the HTTPS secure protocol. In this case, the Secure Sockets Layer (SSL) protocol is employed for all communication between the Administration Assistant clients and the Server component. The latest SSL protocol (Version 3) can be found at <http://wp.netscape.com/eng/ssl3/>.

The latest information on PKI with X.509 certificate can be found on the Web page of the IETF Working Group 'Public Key Infrastructure (X.509) (pkix)' at: <http://www.ietf.org/html.charters/pkix-charter.html>. XML- or HTML-format reports can be created by the Administration Assistant graphical user interface (or through a command-line interface from a scheduling client). The scheduling client is an application that is based on Oracle Java™ technology. It communicates with the Administration Assistant Server component through an RMI connection.

Administration Assistant function for Data Protection for SAP: Features

The Administration Assistant provides a number of features that help you configure, manage, and monitor the data protection environment.

Monitor Operations

A centralized view of the backup status information for all SAP® systems registered with the Administration Assistant server is provided. Summaries of the backup status of all or a selection of SAP systems are available as well as detailed information on all backup runs of a specific SAP system. Thresholds can be defined to enable alerting under certain conditions.

View Performance Data

Performance information during Data Protection for SAP for DB2 backup or restore operations is displayed. The Administration Assistant also saves this performance data and provides a graphical presentation for later analysis.

Configure systems

Configuration of the SAP backup profiles, the Data Protection for SAP profile, and the IBM Tivoli Storage Manager files for each of the SAP systems registered with the Administration Assistant server is provided. Online information also supports configuration. Additionally, profiles can be copied from one system to another system. When configuration changes are performed using the Administration Assistant, a configuration history is maintained so that a previous configuration can be reused.

Request problem support

This feature sends support requests directly to IBM. Although support requests contain user-specified problems the Administration Assistant automatically collects and forwards additional information, such as profiles and error logs.

Manage report templates

This allows the generation and maintenance of templates for producing reports.

Administer users

This feature defines user IDs and permissions in order to access the server component from the Administration Assistant client.

The primary documentation for the Administration Assistant is the integrated online help. The Administration Assistant also provides administrator-created reports in XML or HTML format that are generated from the output of *Monitor operations* and *View performance data*.

Minimizing backup and restore processing with Tivoli Storage Manager for ERP

Although Tivoli Storage Manager for ERP for DB2 provides extensive storage capabilities, business-critical databases might demand even faster recovery operations. Tivoli Storage Manager for ERP and the product *IBM Tivoli Storage FlashCopy Manager* (formerly known as *IBM Tivoli Storage Manager for Advanced Copy Services IBM Tivoli Storage Manager for Advanced Copy Services*) provide backup and restore capabilities for the SAP® database on IBM FlashCopy devices (such as IBM System Storage DS8000(R), IBM System Storage SAN Volume

| Controller, IBM Storwize V7000, and IBM XIV Storage System). These products can
| minimize downtime of the production systems by exploiting point-in-time copy
| functions exploited by these products.

Starting with DB2 9.5, DB2 offers a functionally restricted version of TSM for ACS or FlashCopy Manager known as DB2 Advanced Copy Services, which can be upgraded to an unrestricted level by installing the full FlashCopy Manager version (V2.1 or higher). In this environment, the software based on TSM for ACS is also called by using the DB2 BACKUP DATABASE and RESTORE DATABASE commands, but they contain the keywords "USE SNAPSHOT" rather than the "LOAD *library*" phrase used to call the shared library for Tivoli Storage Manager for ERP. "USE SNAPSHOT" causes DB2 to load and interact with the FlashCopy Manager Snapshot Backup Library.

FlashCopy Manager product information is available at this Web site:
<http://www-01.ibm.com/software/tivoli/products/storage-flashcopy-mgr/>

Chapter 2. Planning for Data Protection for SAP for DB2 operations

Planning information regarding various component considerations is provided.

Database Server Considerations

In general, the production (SAP® database) server is the most critical component for data transfer. This is especially when parallelism is applied as described in “Performance Options of Data Protection for SAP for DB2” on page 103. As a result, special attention should be given to these items:

CPU power

Data transfer, data compression, local, or LAN-free backup operations can cause significant demands on the database server CPU. These demands are in addition to the application load caused by online backups. In many environments, the CPU is the most critical constraint. The CPU load for LAN-free backups (Managed System for SAN) can be significantly reduced by managing the buffers as described in “Buffer Copies” on page 104.

I/O paths

Fast disk attachments with internal busses (like a peripheral component interface) and file system features (like caching or reading ahead) can improve data transfer rates. These attachments and features can be especially useful for backup and restore operations that contain a significant number of files and large data volumes.

Volume Manager settings

Volume Manager provides volume mirroring options that can significantly reduce the data transfer rate during restore operations. As a result, not using volume mirroring options during restore operations can improve the data transfer rate.

Disk layout

The manner in which the database files are laid out can affect data transfer rates. Since the DB2 backup utility allows parallel access to tablespaces during backup and restore operations, distribute data across several disks in order to take advantage of this feature.

Database size

The size of a database can be reduced by offloading inactive data to an external archive. For archive support, refer to the companion product *DB2 CommonStore for SAP*. See “Archiving Inactive Data” on page 18 for additional information.

Network Considerations

Consider these items when setting up the network:

LAN-free backup

LAN-free backup can reduce the load on the network and on the Tivoli Storage Manager server, thus improving data transfer rates. When using LAN-free backup, make sure fiber channel adapter capacity to the SAN can accommodate the data transfer rates of the disk reads and tape writes.

Network bandwidth

Experience reveals that the effective throughput capacity is approximately half of the theoretical network bandwidth. For high-speed networks (such as Gigabit Ethernet LAN), the network adapters limit the throughput rather than the network itself.

Network topology

A dedicated backbone network that is used only for backup and restore operations can improve the data transfer rate.

TCP options

Use TCP options that are the most beneficial for your environment.

Multiple Paths

Data Protection for SAP for DB2 allows you to increase the overall throughput rate to the backup server by specifying multiple network paths. Details are provided in "Multiple Network Paths" on page 110.

Backup Server Considerations

Consider these items when setting up the Tivoli Storage Manager server. Note that Data Protection for SAP for DB2 uses the Tivoli Storage Manager archive function for all backup activities:

Dedicated backup server

A dedicated backup server allows sharing of resources and provides an efficient resource utilization.

CPU power

Observations show that for a given data throughput, the CPU load on the backup server is approximately 60% of that on the database server. Therefore, backup server CPU power is not quite as critical as the CPU power of the database server. However, demands on the Tivoli Storage Manager server CPU do increase when several clients access a single Tivoli Storage Manager server.

Storage hierarchy

Backup of large data files should be directed to tape in order to achieve the highest transfer rates. If disks must be used, it is recommended to use one disk pool per session. Small files (such as log files) should be directed to disk storage first and then be migrated to tape collectively to avoid excessive tape mounts.

Parallel sessions

The Tivoli Storage Manager server allows using several tape drives in parallel to store data. This can increase overall data throughput. In order to exploit this feature, the corresponding Tivoli Storage Manager node must be allowed the appropriate number of mount points and the device class must be allowed the appropriate mount limits.

Detailed information on how to set up Tivoli Storage Manager for use with Data Protection for SAP can be found in “Alternate or parallel backup paths and backup servers” on page 18 and “Configure the Tivoli Storage Manager server” on page 58.

Storing data on a Tivoli Storage Manager server

Data Protection for SAP transfers data to and from the backup server through single or multiple (parallel) sessions to the Tivoli Storage Manager server. Each session must have a storage device associated with it. The SAP backup ID is persistently linked with each backup file. This backup ID can be used later to determine all files required for a complete restore.

In SAP terminology 'backup' means backup of database contents, 'archive' means the backup of offline DB2 log files. Data Protection for SAP for DB2 uses the Tivoli Storage Manager archive function for both backup types.

Tape storage is the preferred media for storing the database contents as this is proven to provide the best data throughput for backup and restore. A disk-tape storage hierarchy is recommended for backing up log files, each DB2 log file should be backed up immediately after it is placed in the archive directory. This provides the best protection against data loss and eliminates the need to mount a tape for each DB2 log file.

Collocation is a Tivoli Storage Manager function that ensures client data is maintained together on one tape. Collocation should be deactivated in these situations:

- Deactivate collocation for Data Protection for SAP backups when enabling parallel sessions for use with multiple tape drives in parallel.
- Deactivate collocation when using the multiple log copy function as described in “Multiple DB2 Log File Copies” on page 69.

To improve availability (alternate servers) or performance (multiple servers), configure Data Protection for SAP to use multiple Tivoli Storage Manager servers. Consider the location of all backup data before removing a Tivoli Storage Manager server from the Data Protection for SAP profile. Since Data Protection for SAP only accesses those servers defined in its profile, be cautious when removing a Tivoli Storage Manager server if it contains valid backup data.

Database backups are typically retained for a specified period and then become obsolete. In order to manage backup storage space efficiently, delete obsolete backups so that the tape storage space can be reclaimed. There are two ways to perform this deletion:

- Set an appropriate archive retention period with Tivoli Storage Manager options.
- Use the Data Protection for SAP backup version control function. When the number of backup versions (specified by this function) is exceeded, entire backup generations (such as full backups and all related partial and log file backups, are automatically deleted.

Alternate or parallel backup paths and backup servers

In Data Protection for SAP[®] terminology, path denotes a connection between a Tivoli Storage Manager client (Tivoli Storage Manager node) and a Tivoli Storage Manager server. A set of communication parameters are also set for each defined communication path. A Tivoli Storage Manager server network address is an example of a communication path. This set of communication parameters is called client option data and is collected under a logical server name. The logical server name is determined by the user. On UNIX or Linux systems, all client option data can be stored in a single file. This file is the client system option `dsm.sys` file. On Windows systems, the client option data for each logical server must be stored in separate client option files that have the file name `servername.opt`. For example, if there are two logical Tivoli Storage Manager servers *fast* and *slow*, then two client option files `fast.opt` and `slow.opt` are required. Windows also requires an additional client user option file, `dsm.opt`. All option files must reside in the same directory.

Data Protection for SAP for DB2 can use several communication links between Tivoli Storage Manager clients in order to control alternate backup paths and alternate backup servers. This feature can increase throughput by transferring data over multiple paths simultaneously or to and from several servers in parallel. It can improve the availability of the Tivoli Storage Manager client-to-server communication and enable disaster recovery backup to a special (remote) Tivoli Storage Manager server.

Each path in the `initSID.utl` profile is defined by a server statement and the corresponding definitions in the Tivoli Storage Manager client system option file `dsm.sys` (UNIX and Linux) or `server.opt` (Windows). The `SERVER <server 1..n>` statement denotes Tivoli Storage Manager servers defined in the Data Protection for SAP profile. This corresponds to the statement `SERVERNAME server 1..n` in the Tivoli Storage Manager client option file(s). These servers are identified by their `TCPSEVERADDRESS` and can be located on one system (multiple paths) or several systems (multiple servers). `SESSIONS` denotes the number of parallel sessions that Data Protection for SAP schedules for the given path. If only one path is used, `SESSIONS` must be equal to `MAX_SESSIONS`, which specifies the total number of parallel sessions to be used (equivalent to number of tape drives/management classes). Data Protection for SAP attempts to communicate with the Tivoli Storage Manager server using the first path in the profile. If this proves successful, Data Protection for SAP starts the number of parallel sessions as specified for this path. If the attempt was unsuccessful, this path is skipped and Data Protection for SAP continues to the next path. This process continues until as many sessions are active as were specified in the total session number (`MAX_SESSIONS`). If this number is never reached (for example, because several paths were inactive), Data Protection for SAP terminates the backup job.

Archiving Inactive Data

Data Protection for SAP for DB2 creates a database image that is stored at the bit level and therefore, is designed for routine backup operations. Outdated backups must be restored into the same exact environment they were originally taken from in order to access the data from within SAP[®] applications. This requires maintaining older versions of SAP, operating system, database, and Tivoli Storage Manager data to rebuild this original environment. SAP provides archiving functions that can display business documents that are designated with long term retention requirements. These business documents are format-independent and can

be used for auditing and other legal purposes. Archived data can then be removed from the operational database to reduce the database size and improve backup and restore processing time.

Long term archive requirements can be achieved with the IBM DB2 CommonStore for SAP product. This product accesses the SAP ArchiveLink interface and uses Tivoli Storage Manager to archive the following document types:

- inactive data (data retention)
- printlists (e.g. reports)
- outgoing documents (e.g. printed output like invoices, bills)
- incoming documents (e.g. digitized fax, scanned letters, audio)
- local documents (e.g. text, spreadsheets, pictures, graphics)
- inactive data

This demonstrates how Tivoli Storage Manager is used as an integrated repository for backup and archive tasks. DB2 CommonStore for SAP product information is available at this Web site: <http://www.ibm.com/software/data/commonstore/sap/>.

Restore versus Backup

The majority of this section has addressed issues related to optimizing backups. In most cases, configuration changes and infrastructure problems affect both backup and restore operations similarly. Therefore, modifications supporting a fast backup while also exploiting resources can also be considered applicable to the restore operation. Generally, it is recommended to tune the backup and then run a restore test to verify that restore still works in a satisfactory manner.

During a restore operation, the values of these parameters are determined by their settings during the corresponding backup:

Compression

If compression is used during the backup, data needs to be decompressed.

Multiple servers

When a backup is performed using multiple servers, the same servers must be online and available during the restore operation.

Planning for using IBM HACMP™ for AIX

This section provides information about Data Protection for SAP for DB2 that is useful when planning for HACMP fail-over configurations. This example uses the mutual takeover configuration (each node can take over the other node). If the application server and database server are installed on different hosts, the described actions need to be taken on the database servers only.

This figure illustrates the takeover environment:

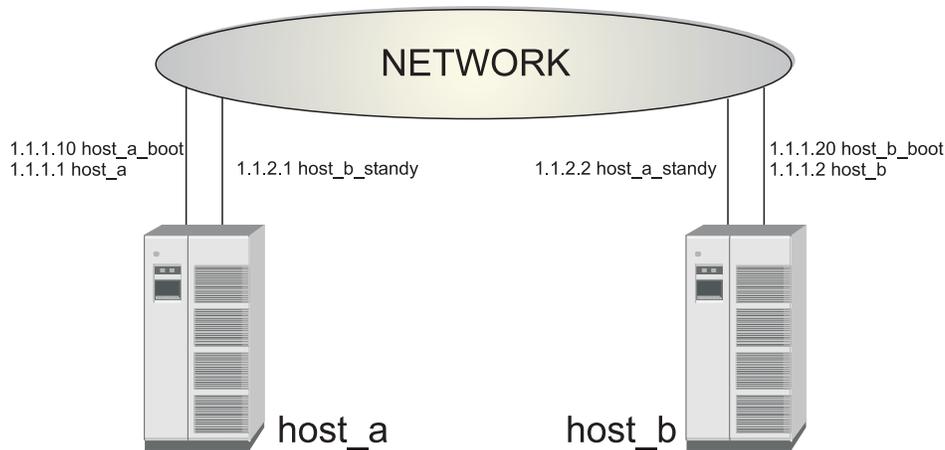


Figure 9. Sample Environment for HACMP Takeover

HACMP impact on Data Protection for SAP for DB2

A list of Data Protection for SAP for DB2 components that are impacted by HACMP are provided.

Files

- The installation directory is `/usr/tivoli/TSM/tdp_r3`.
- Lock files are located in `/var/tdp_r3`.
- There is only one ProLE running on each host (even after takeover).
- Each SAP® system has its own Data Protection for SAP configuration files (`initSID.utl`, `initSID.bki`). These files are located in a directory specified during the installation process.

Dependencies

- Both hosts should have the same level of Tivoli Storage Manager API installed.
- Both hosts must be Data Protection for SAP.
- On both hosts, the `dsm.sys` file (in `/usr/Tivoli/Tivoli Storage Manager/client/api/bin/dsm.sys`) must contain all server names required for takeover.

Communication

The Data Protection for SAP dynamic library connects to prole using the following procedure:

- Retrieves the IP address for localhost (should be 127.0.0.1 for IPv4).
- Retrieves the `tdpr3db264` service (should be 57324).
- Connects to `127.0.0.1: tdpr3db264 service>`.

Digital Signing of Executable Files for Windows Systems

Data Protection for SAP for DB2 executable files (except .jar files) for Windows systems have a digital signature. The following files are affected:

- Passport Advantage package for Windows
- Data Protection for SAP installation files
 - *version*-TIV-TSMERPDB2-WinIA64.exe
 - *version*-TIV-TSMERPDB2-WinX64.exe
- The Data Protection for SAP application executable files
 - backom.exe
 - createinfo.exe
 - prole.exe
 - tdpdb2.dll

Code signing employs digital IDs, also known as certificates.

Having a valid digital signature ensures the authenticity and integrity of an executable file. It identifies the software publisher as IBM Corporation to the person who downloads or executes it. However, it does not mean that the end-user or a system administrator implicitly trusts the publisher. A user or administrator must make the decision to install or run an application on a case-by-case basis, based on their knowledge of the software publisher and application. By default, a publisher is trusted only if its certificate is installed in the Trusted Publishers certificate store.

The customer can see the digital signature for any .EXE, .DLL, or installation wizard of Data Protection for SAP using one of the following methods:

1. The digital signature can be viewed from the Digital Signature tab of Properties of the signed file. If you select the IBM Corporation item and click Details, you will see more information about the IBM Certificate and the entire chain of trusted Certificate Authority signatures.
2. In the case of the installation wizard, there is also the possibility to see the IBM digital signature from the software publisher link displayed in the Security Warning window.

A warning is issued if the installation executable file is downloaded from a site that is not listed as a trusted site. The security warning is not related to the fact that executable files contain digital certificates. It is related to the security zone policy of the site you download the file from. There is also another condition to be met: the executable must be stored on an NTFS disk. Windows Server 2008 includes Internet Explorer 7, and its default security configurations are set according to the Internet Explorer Enhanced Security Configuration on four different security zones: Internet, local intranet, trusted, and restricted sites. The Internet Explorer Enhanced Security Configuration component (also known as Microsoft Internet Explorer hardening) reduces a server's vulnerability to attacks from Web content by applying more restrictive Internet Explorer security settings. As a consequence, Internet Explorer Enhanced Security Configuration may prevent some Web sites from displaying properly or performing as expected. It may also prevent users and administrators from accessing resources with Universal Naming Convention (UNC) paths on a corporate intranet. Refer to this document for more information on managing Internet Explorer Enhanced Security Configuration: <http://www.microsoft.com/downloads/details.aspx?FamilyID=d41b036c-e2e1-4960-99bb-9757f7e9e31b&DisplayLang;=en> You might get a security warning

displayed whenever you run an executable file downloaded using the Internet Explorer from a URL or UNC that is not a member of the trusted security zone.

When a downloaded file is saved to a disk formatted with NTFS, it will update the meta data for the file with the zone (Internet or restricted-) it was downloaded from. The meta data is saved as an Alternate Data Stream (ADS), which is a feature of NTFS with which the same filename can be used to cover multiple data streams. When opening a file which includes an ADS that identifies it as being from another zone, the Attachment Execution Services (AES) software is activated, which reacts to the following file categories as described:

- **High risk:** Blocks the file from being opened when the file is from the restricted-zone: The following security warning is issued:

Windows Security Warning:
Windows found that this file is potentially harmful.
To help protect your computer, Windows has blocked access to this file.

- **Moderate risk:** Prompts with a warning before the file is opened when the file is from the Internet zone:

Open File - Security Warning:
The publisher could not be verified. Are you sure you want to run this software?

- **Low risk:** Opens the file with no warnings.

Warning messages do not prevent the file from being used.

Note: This is different from configuring the Web Server with a digital certificate. During the installation of the Administration Assistant, the customer has the option to generate a self-signed certificate for the AA server and to use it to configure the security communication over HTTPS between the Administration Assistant server component and the clients. Alternatively there is the possibility to configure the security communication later after the installation completes using the instructions provided under "Configuring for Secure Communication".

Chapter 3. Installing Data Protection for SAP for DB2 for V6.3

Information needed to install the various IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP for DB2 components is provided.

Review the appropriate prerequisite information before attempting to perform any installation tasks.

Note: Data Protection for SAP and the Administration Assistant function for Data Protection for SAP are installed via InstallAnywhere rather than InstallShield. Slightly modified procedures are required to employ console mode (non-graphical user interface) or perform a silent installation. See “Installing Tivoli Storage Manager for ERP for DB2 in silent mode” on page 25.

Furthermore, Windows executable files (except Java) contain a digital signature to certify that the software originated by IBM. For more information, see “Digital Signing of Executable Files for Windows Systems” on page 21.

Required installation tasks

Data Protection for SAP for DB2 must be installed on all SAP® database servers. The following tasks are required to set up Data Protection for SAP:

1. Verify the Data Protection for SAP for DB2 package is complete. See the README.1ST file on each installation disc (or disc image) for a description of the contents.
2. Verify that the prerequisites are met as described in “Prerequisites” on page 24.
3. Review planning sheet information as described in “Data Protection for SAP for DB2 (base product) planning sheet” on page 152.
4. (Optional) Install the Administration Assistant function for Data Protection for SAP prior to installing Data Protection for SAP. Data Protection for SAP can automatically connect to the Administration Assistant as part of its installation procedure. Details are available in “Administration Assistant function for Data Protection for SAP: Features” on page 12.
5. Install Data Protection for SAP as described in “Installing Tivoli Storage Manager for ERP for DB2 on UNIX or Linux” on page 26 or “Installing Tivoli Storage Manager for ERP for DB2 on Windows” on page 28. See “Upgrade the Data Protection for SAP for DB2 V6.3 base product” on page 35 when upgrading a previous version of Data Protection for SAP.
6. Perform post-installation tasks as such as “Configure the Tivoli Storage Manager client options” on page 54 and “Configure the Tivoli Storage Manager server” on page 58.
7. Verify the installation completed successfully as described in “Verifying the Initial and Upgrade Installation” on page 30.

Installing the Data Protection for SAP for DB2 V6.3 base product

Information needed to install the Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP for DB2 base product is provided.

Perform the installation tasks for the appropriate operating system.

Prerequisites

The installation packages are located on the Data Protection for SAP for DB2 product installation disk, disk image (from Passport Advantage), and occasionally on the IBM public FTP server. Initial installations must always be done from the disc or image. Refer to the file README.1ST in the root path for information about where to find documents on the disc or image, and follow the appropriate installation description below. See the README.1ST file in the root directory of the disc or image for a list of its contents.

If you are going to upgrade from an earlier version of Tivoli Data Protection for R/3 or Data Protection for SAP in your environment, you have the option of either upgrading from the product disc or image, or downloading the latest version from the IBM FTP server. For the specific procedure for upgrading from an earlier version, refer to “Upgrade the Data Protection for SAP for DB2 V6.3 base product” on page 35.

These products must be installed before installing Data Protection for SAP:

- DB2
- SAP® R/3 or SAP e-business Solution, based on DB2
The SAP Service Marketplace provides current information relating to SAP features, product versions, and maintenance levels that are compatible with your version of SAP R/3 or SAP.
- Tivoli Storage Manager Backup-Archive Client
For information about configuring the Tivoli Storage Manager API client, see “Configure the Tivoli Storage Manager client options” on page 54. TCP/IP must be ready for communication between the Tivoli Storage Manager server and the Tivoli Storage Manager client.
- An operating system level supported by SAP and the Tivoli Storage Manager client

The Release Notes® file on the Tivoli Information Center contains the most current information related to Data Protection for SAP hardware, software, operating system, and maintenance levels.

In case Data Protection for SAP is to be installed on a distributed file system, the root user needs read and write access to the file system for the duration of the installation. For more information on the installation in a distributed file system, refer to: “Configuring Tivoli Storage Manager for ERP for DB2 in a Distributed File System” on page 50.

Installation planning forms for Data Protection for SAP and the Administration Assistant are available in the planning_sheet (UNIX and Linux) or planning_sheet.txt (Windows) files located in the installation directory. They are also available for printing in “Data Protection for SAP for DB2 (base product) planning sheet” on page 152. Once prerequisites are met and installation planning information is completed, Data Protection for SAP is ready to be installed.

Installing Tivoli Storage Manager for ERP for DB2 in silent mode

Information about installing the product without using a graphical user interface.

To support target systems without a window manager, the setup program supports deploying an installation in console mode. An installation running in console mode suppresses the graphical wizard panel display available with a GUI installation. Instead, user data entry and status messages are displayed on the console or in the command prompt window.

In order to perform a silent or unattended installation follow this procedure:

1. Create a response file during an installation in either graphic or console mode by using option "-DRECORDFILE" denoting the response file name:

```
./version-TIV-TSMERPDB2-platform.bin [-i console] -DRECORDFILE=properties file
```

Note: This is a UNIX command. For Windows, use the corresponding .exe file with the same options.

2. Start the executable file with the "-i silent" option (silent mode) and the "-f option" denoting the file name of the response file:

```
./version-TIV-TSMERPDB2-platform.bin -i silent -f properties file
```

Note: This is a UNIX command. For Windows, use the corresponding .exe file with the same options.

The *properties file* specification must contain a full path.

Sample properties file:

```
USER_INSTALL_DIR=//opt//tivoli//tsm//tdp_r3//db264
NAMEPORTAA_ADRESSE=AAServer
NAMEPORTAA_PORT=5126
RMANYES=
MANNO=
TSMUTL_SERVERADRESSE=TSMServer
TSMUTL_NODE=R3NODE
TSMUTL_BACKUPMGM=MDB
TSMUTL_ARCHIVEMGM=MLOG1 MLOG2
TSMUTL_YES=1
TSMUTL_NO=0
TSMAPI_DSMI_DIR=
TSMAPI_DSMI_CONFIG=
TSMAPI_DSMI_LOG=
TSMAPI_YES=
TSMAPI_NO=
SAP_BR_TOOL=
SAP_CFG_FILE=
TSM_CFG_FILE=
DBGSCRIPTS2=//dev//null
SID=SID
DB2_INSTANCE_NAME=DB2_INSTANCE_NAME
USER_MAGIC_FOLDER_1=//db2//DB2ERE//tdp_r3
LOGGING_NONE=0
LOGGING_LOGARCHMETH1=0
LOGGING_LOGARCHMETH2=0
LOGGING_BOTH=1
LOGGING_NR=12
```

Lines starting with '#' are treated as comments.

Note: This is a UNIX properties file. When installing Tivoli Storage Manager for ERP in silent mode for Windows, use the corresponding Windows properties file.

Installing Tivoli Storage Manager for ERP for DB2 on UNIX or Linux

IBM Tivoli Storage Manager for Enterprise Resource Planning for DB2 for these operating systems is delivered as a single executable file for each platform. Packages on the FTP server contain 'FTP' before the platform designation.

- For a disc or disc image, the name has the format:

version-TIV-TSMERPDB2-platform

When the file is launched, Tivoli Storage Manager for ERP guides you through the installation procedure. Read the descriptions carefully and follow the guidelines that are displayed on the panels.

Shared libraries have different file extensions on different UNIX or Linux platforms. Within the following the section, the file extensions of shared libraries are represented as '*ext*'. Replace this text with the extension applying to your platform:

Table 7. File Extensions for Shared Libraries

Operating System	Extension
AIX	a
HP-UX	sl
Linux	so
Solaris	so

Perform the following tasks to install Tivoli Storage Manager for ERP on a UNIX or Linux system:

1. Log in as the root user on the SAP database server machine.
2. Verify that the *DISPLAY* variable is set to view the installation prompts through a graphical X-Window.
3. Start the DB2 instance. The installation program makes the necessary updates to the DB2 configuration.
4. Invoke the Tivoli Storage Manager for ERP executable file and follow the installation prompts.
5. View the summary in the last page of the installation wizard. The Tivoli Storage Manager for ERP installation path is displayed in the summary where the installation log file (*log.txt*) is located.

These modifications are automatically performed to your system during installation:

- An entry is created in */etc/inittab* that automatically starts the "ProLE" daemon on UNIX systems. If *upstart* is configured, */etc/init/prole_db2.conf* is created and *upstart* starts the "ProLE" daemon.
- An entry is created in */etc/services* for the service *tdpr3db264*.
- The environment variable *XINT_PROFILE* specifies the Tivoli Storage Manager for ERP profile that is located in the path specified for configuration files during installation. The file name is *initSID.utl* where *SID* is the DB2 database *SID* specified during installation.

- The environment variable TDP_DIR points to the path where Tivoli Storage Manager for ERP configuration files and process logs reside. The default path is *profile path*/tdplog where *profile path* is the path specified for the Tivoli Storage Manager for ERP profile during installation.
- The environment variable XINT-NLS_CATALOG_PATH points to the installation path of Tivoli Storage Manager for ERP. The message catalog is located under *DP for SAP install path/lang* where *DP for SAP install path* is the installation path /usr/tivoli/tsm/tdp_r3/db264.
- If the DB2 instance is running, the installation program sets the DB2 database configuration parameter *VENDOROPT* to the Tivoli Storage Manager for ERP vendor environment file. If *VENDOROPT* is already set (for example because of the installation of a previous version), the program will use its value and not set *VENDOROPT*. If DB2 log archiving is not to be managed by Tivoli Storage Manager for ERP, the corresponding database configuration settings are not modified. If DB2 log archiving is to be managed by Tivoli Storage Manager for ERP, the corresponding DB2 database configuration values are set based on the method selected during installation:

```
LOGARCHMETHn  VENDOR:/path/library
LOGARCHOPTn   /path/vendor.env
```

If the DB2 instance was not running, you must complete these tasks manually, as described in “Specifying the *VENDOROPT* parameter” on page 44 and “Configuring the DB2 Log Manager” on page 44.

The EN_US folder is created, which contains the message catalog file tsmerp.cat. The _uninst folder is also created, which contains additional files.

These files are installed in the Tivoli Storage Manager for ERP installation directory:

```
README
README_TSMERPversionlanguage.html
TIPHINTS
libtdpdb264.a (AIX)
libtdpdb264.so (Linux or Solaris)
ProLE
backom
createinfo
initSID.utl
SanFSsetupFS.sh (AIX only)
agent.lic (only after installation from disc or disc image)
```

The folder EN_US is created and it contains the message catalog tsmerp.cat. The _uninst folder is also created, which contains sample files.

These files are installed in the directory where the Tivoli Storage Manager for ERP configuration files are located:

```
initSID.utl
vendor.env
agent.lic (copy of file in installation directory)
```

Installing Tivoli Storage Manager for ERP for DB2 on Windows

Tivoli Storage Manager for ERP for Windows is delivered as a single executable file (.exe) for each platform. Packages on the FTP server contain 'FTP' before the platform designation.

Tivoli Storage Manager for ERP for these operating systems is delivered as a single executable file for each platform. The packages are named as follows:

- The package name on the disc (or disc image):
version-TIV-TSMERPDB2-platform

Complete these tasks to install Tivoli Storage Manager for ERP on a Windows system:

1. Log in as a user with Administrator authority on the SAP database server machine.
2. If you want the installation program to make updates to the DB2 configuration, start the DB2 instance.
3. Start the Tivoli Storage Manager for ERP executable file, and follow the instructions of the installation dialog.
4. View the summary on the last page of installation wizard. The Tivoli Storage Manager for ERP installation path is displayed in the summary where the installation log file (log.txt) is located.

The following modifications are performed on your system during installation:

- The ProLE service is installed and started.
- An entry is created in %windir%\system32\drivers\etc\services (tdpr3db264) .
- (Optional) The DSMI_DIR, DSMI_CONFIG, and DSMI_LOG environment variables are modified.
- The XINT_PROFILE environment variable specifies the Tivoli Storage Manager for ERP profile located in the path specified during installation. The file name is *initSID.ut1* where *SID* is the DB2 database SID specified during installation.
- The TDP_DIR environment variable specifies the directory where Tivoli Storage Manager for ERP saves the configuration file and creates its process logs. Initially, this path is set to *profile path\tdplog* where <profile path> is the path for Tivoli Storage Manager for ERP profile specified during installation.
- The environment variable XINT-NLS_CATALOG_PATH points to the installation path of Tivoli Storage Manager for ERP. The message catalog is located under *DP for SAP install path\lang* where *DP for SAP install path* is the installation path specified by the user during the installation.
- If the DB2 instance is running, the installation program sets the DB2 database configuration parameter *VENDOROPT* to the Tivoli Storage Manager for ERP vendor environment file. If *VENDOROPT* is already set (for example because of the installation of a previous version), the program will use its value and not set *VENDOROPT*. If DB2 log archiving is not to be managed by Tivoli Storage Manager for ERP, the corresponding database configuration settings are not modified. If DB2 log archiving is to be managed by Tivoli Storage Manager for ERP, the corresponding DB2 database configuration values are set based on the method selected during installation:

```
LOGARCHMETHn  VENDOR:path\tdpdb2.d11
LOGARCHOPTn   drive:\path\vendor.env
```

If the DB2 instance is not running, these tasks are not performed and must be performed manually at a later time as described in “Specifying the VENDOROPT parameter” on page 44 and “Configuring the DB2 Log Manager” on page 44.

The following files are installed in the Tivoli Storage Manager for ERP installation directory:

README.txt
README_TSMERPversionlanguage.html
TIPHINTS.txt
tdpdb2.dll
ProLE.exe
backom.exe
createinfo.exe
initSID.utl
agent.lic (only after installation from disc or disc image)

The _uninst folder is also created, which contains sample files.

These files are installed in the directory where the Tivoli Storage Manager for ERP profile is located:

initSID.utl ('SID' is replaced by the DB2 database SID provided during installation)
vendor.env
agent.lic (copy of file in installation directory)

Enable ProLE on Windows to access configuration files on a remote share

When ProLE is started as a regular service, it operates (by default) under the ID of the local system account with Administrator privileges. However, a session opened on a remote system does not have credentials or permissions. Microsoft knowledge base article 132679 provides information about this situation:

<http://support.microsoft.com/kb/132679>

This situation prevents the ProLE service from accessing files that reside on a remote share. This is true even when the share is mapped to a local drive letter or is accessed as a Uniform Naming Convention (UNC) notation (\\server\path\). Data Protection for SAP for DB2 version 5.4 (or later) accepts UNC notation for the profile but not for all the files specified within the profile. These files will be opened by ProLE, which by default has no permission to access remote shares, as explained above.

Perform these tasks to enable ProLE to access such files on a remote share:

1. Map the share where the configuration files reside to a local drive letter.
2. Modify the profile (.utl) to refer to the path names on the mapped drive.
3. Modify the ProLE service so that it runs as an account with permissions to access the mapped drive (and not as a local system account). Note that this might have other implications when using a regular account. For example, when the password for this account expires or is changed, the service will no longer be able to start.
4. Restart the ProLE service to activate the changes.

Uninstalling the Old Version of Tivoli Storage Manager for ERP for DB2 under UNIX or Linux

Perform these tasks to uninstall a previous version of IBM Tivoli Storage Manager for Enterprise Resource Planning:

1. Log in as root user.
2. Make sure that the DISPLAY variable is set correctly as the uninstall procedure requires a graphical X-Window.
3. Make sure the previous version of Tivoli Storage Manager for ERP for DB2 is not running.
4. Invoke the uninstall executable file as shown here:

AIX 64-bit:

for version prior to 6.1:

```
/usr/tivoli/tsm/tdp_r3/db264/_uninst/uninstaller.bin [-silent | -console]
```

for Version 6.1 or later:

```
/usr/tivoli/tsm/tdp_r3/db264/Uninstall_TIV-TSMERPDB2/  
Uninstall_TIV-TSMERPDB2 [-i silent | -i console]
```

Other UNIX 64-bit or Linux 64-bit:

for version prior to 6.1:

```
/opt/tivoli/tsm/tdp_r3/db264/_uninst/uninstaller.bin [-silent | -console]
```

for Version 6.1 or later:

```
/opt/tivoli/tsm/tdp_r3/db264/Uninstall_TIV-TSMERPDB2/  
Uninstall_TIV-TSMERPDB2 [-i silent | -i console]
```

Follow the instructions of the uninstall dialog.

Uninstalling the Old Version of Data Protection for SAP for DB2 under Windows

Perform these tasks to uninstall a previous version of Data Protection for SAP for DB2 on a Windows NT, Windows 2000, or Windows 2003 machine:

1. Log on as user with administrator authority on the SAP® database server machine.
2. Ensure that the previous version of Data Protection for SAP is not running.
3. Select **Start** → **Settings** → **Control** panel.
4. Click on **Add/Remove Programs**.
5. Select the old version of **Data Protection for SAP** and click on **Remove**.
6. Follow the instructions of the uninstall procedure.

Verifying the Initial and Upgrade Installation

In order to verify the installation of Data Protection for SAP for DB2, perform a full DB2 database backup and restore with the DB2 Control Center or DB2 command line processor (CLP). A complete restore or recovery of the entire SAP® database is also recommended. However, a complete offline backup should be performed first.

Installing the Administration Assistant function for Data Protection for SAP V6.3

The Administration Assistant function for Data Protection for SAP is a Web-browser based graphical interface that provides customization, simulation, and analysis of SAP® database backup, restore, and configuration operations. Information needed to install the Administration Assistant function for Data Protection for SAP V6.3 is provided.

Perform these tasks to install the Administration Assistant function for Data Protection for SAP.

Prerequisites for Installing the Administration Assistant function for Data Protection for SAP

Prerequisites: Server-Level Components

The following products must be installed before setting up the Administration Assistant function for Data Protection for SAP server-level components:

- Java Runtime Environment (JRE) or Java Development Kit (JDK)
- Java Beans Activation Framework (JAF)
- Java Mail
- IBM DB2 data server (optional DBMS for Administration Assistant database if you do not want to use the Apache Derby database already bundled with the Administration Assistant install package). If you elect to use DB2, make sure DB2 is running. In addition, UNIX and Linux systems require that a dedicated system user (for which the DB2 instance should be installed) be created.
- For software, hardware, and maintenance levels required by the current version of the Administration Assistant, refer to the Data Protection for SAP for DB2 release notes.
- TCP/IP must be ready for communication before starting up the Administration Assistant server-level components.

Prerequisites: Client Components

These requirements must be met before starting the Administration Assistant function for Data Protection for SAP client:

- A fully Java-capable Web browser with Java plugin. The applet loaded from the Administration Assistant server must be granted these permissions:
 - Permission to establish a connection to the Administration Assistant server through RMI. For example:

```
permission java.net.SocketPermission "Server component hostname:1024-", "connect";
```
 - Permission to switch to a different language. For example:

```
permission java.util.PropertyPermission "user.language", "write";
```
- In order to view report graphics, a browser that supports Scalable Vector Graphics (SVG), like Adobe SVG Viewer, must be available.
- (UNIX or Linux): An X Window system is required for the Administration Assistant client.

- A minimum screen resolution of 1024x768 pixels (1280x1024 or higher is recommended).
- For software and maintenance levels required by the current version of the Administration Assistant, refer to the Data Protection for SAP release notes.
- TCP/IP must be ready for communication before starting up the Administration Assistant server-level components.

Prerequisites: Scheduling Client

These requirements must be met when selecting the scheduling client:

- A TCP/IP connection can be established to the Administration Assistant Server component.
- A Java VM is available.
- In order to view report graphics, a browser that supports Scalable Vector Graphics (SVG), like Adobe SVG Viewer, must be available.

Installation Planning for Server-Level Components

See Table 16 on page 154 for a list of planning requirements in table form. This information is also available in the `planning_sheet_aa` (UNIX or Linux) and `planning_sheet_aa.txt`. (Windows) files in the Data Protection for SAP installation directory.

Installing the Administration Assistant function for Data Protection for SAP Server-Level Components

Initial installations must be performed from the installation disc or disc image. Refer to the README.1ST file in the root path of the disc or disc image for the most current information. The Administration Assistant installation packages reside on each of the Data Protection for SAP for DB2 discs or disc images, and can be downloaded from the IBM FTP server. The Administration Assistant installation package is a single, platform-independent .jar file with this name convention:

`version-TIV-TSMERPAABASE-MULTI.jar`

When upgrading from an earlier version of the Administration Assistant function for Data Protection for SAP, the latest version is available for download from the IBM FTP server. Additional upgrade information is available in “Upgrade the Administration Assistant function for Data Protection for SAP V6.3” on page 36.

A setup assistant is included in the Administration Assistant package that helps guide the installation process in English or multi-language version. Be aware of the considerations before installing the Administration Assistant:

- System administrator privileges are required to install the Administration Assistant.
- If a multi-host installation (which distributes the server-level components over two or more hosts) is to be performed, copy the package file to each target host. Then perform a custom installation so that components are selected for that host.
- The CLASSPATH environment variable is not required. However, if this variable is set, you must specify the directory in which the package file resides.
- After installation, in order to switch the language (specified during installation), the Administration Assistant must be uninstalled and install again with the preferred language.

Specify this command to start the installation:

```
java -jar package file name
```

After the first component is installed, an overview panel displays the installation status and records user entries.

During installation, the following modifications are made to your system automatically:

- All necessary paths (installation, history, OnDoc, log paths) are created. Corresponding files are copied into the installation and OnDoc directories.
- These Administration Assistant startup files are created and added to the installation directory:

Component	UNIX or Linux	Windows
Server	sadma.sh	sadma.cmd
Database Agent	sdba.sh	sdba.cmd
Database	sdb.sh	sdb.cmd

- The configuration file `assist.cfg`, containing all relevant configuration parameters specified during the installation, is created and added to the installation directory.
- The configuration file `dbagent.cfg` containing all relevant configuration parameters specified during installation of the Database Agent component is created and added to the installation directory.
- On Windows systems, up to three services are installed and automatically started. These services start the Administration Assistant components: server, dbagent, and database. The database component only runs if the Apache Derby database is used.
- On UNIX or Linux systems, a new `/etc/init.d` entry is created for each Administration Assistant server-level component and the components are started automatically. Note that an administrator must create appropriate run level entries for these components in order for automatic start and stop features to function:

Component	Entry in <code>/etc/init.d</code>
Server	adminAssistant, with parameters start, stop, and status
Database Agent	databaseAgent, with parameters start, stop, and status
Database (Derby) (optional, as alternative to DB2)	apacheDerby, with parameters start and stop
Database (DB2) (optional, as alternative to Derby)	Not applicable

For an installation using IBM DB2:

- On Windows systems, the database tables are created and no other changes are made.
- On UNIX and Linux systems, a DB2 instance for the specified user (`$USERNAME`) is created. These changes are also made to the system:
 - An entry in `/etc/services` is added:

```
$USERNAME $PORT/tcp # used for Data Protection for SAP - Administration Assistant with DB2 support
```

- Changes to the created DB2 instance:
 - Set DB2 profile registry variable: DB2COMM=tcPIP
 - Set DB2 database manager parameter: SVCENAME=\$USERNAME
 - Set DB2 database manager parameter: SPM_NAME=NULL

For an installation using secure communication:

- A keystore is created on request.
- An X.509 v1 self-signed certificate containing a key pair with the hostname as an alias is created in the keystore on request.
- The server's self-signed certificate is imported into the truststore on request.
- The server's self-signed certificate is exported to a certificate file on request.
- A Certificate Signing Request is created if desired.

Consider these items before uninstalling the Administration Assistant server-level components:

- The Administration Assistant client component is not physically installed. It operates as a Java applet when the URL of the host running the Server component is called. No action needs to be taken at the client level when uninstalling the Administration Assistant server-level components.
- The public key infrastructure will not be modified when uninstalling the Administration Assistant components, even if it was originally set up during its installation process.

To uninstall the Administration Assistant server-level components, change to the uninstall directory (in the Administration Assistant installation directory) on each system on which one of the components was installed and issue this following command:

```
java -jar uninstall.jar
```

The command files open an uninstall assistant which guides you through the process.

Chapter 4. Upgrading to Data Protection for SAP for DB2 for V6.3

Information needed to upgrade to Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP for DB2 V6.3 is provided.

Perform these tasks to upgrade to Data Protection for SAP for DB2 V6.3.

Upgrade the Data Protection for SAP for DB2 V6.3 base product

Note: The format of the configuration file (.bki) was changed with version 5.4. The software accepts the previous format and converts it automatically.

If it is necessary to use a version earlier than 5.4, the old format can be recovered by overwriting the new file with the empty file (provided with the previous version). The file must then be initialized by setting the Tivoli Storage Manager password. However, the information about the current backup number will be lost. As a result, more backup versions must be retained for a longer period of time than is specified by the MAX_VERSIONS parameter.

Perform these tasks to upgrade Data Protection for SAP from an earlier version:

1. Verify that the Data Protection for SAP for DB2 package is complete. The installation packages are provided on a disc, disc image (downloadable from Passport Advantage), or the IBM FTP server. See the release notes file in the Tivoli Information Center for the most current release information.
2. Check the readme files and release notes for incompatibilities between the installed version and the new version. Make sure that data backed up with an older version of Tivoli Storage Manager for Enterprise Resource Planning can still be restored with the version to be installed.
3. Make sure that the requirements for the new version of Data Protection for SAP are met as described in “Prerequisites” on page 24.
4. Make sure planning information is available as described in “Prerequisites” on page 24.
5. A full backup of the SAP® database should be performed before upgrading to the new version.
6. Uninstall the old version as described in “Uninstalling the Old Version of Tivoli Storage Manager for ERP for DB2 under UNIX or Linux” on page 30 or “Uninstalling the Old Version of Data Protection for SAP for DB2 under Windows” on page 30.
7. Install the new version of Data Protection for SAP as described in “Prerequisites” on page 24.
8. Update the Data Protection for SAP profile as described in “Migrate the Data Protection for SAP for DB2 profile” on page 36.
9. Create the configuration file(s) as described in “Creating the configuration files” on page 46.
10. Perform the necessary configuration tasks as described in “Configure the Tivoli Storage Manager client options” on page 54.
11. Verify the installation as described in “Verifying the Initial and Upgrade Installation” on page 30.

12. A full backup should be performed after upgrading to the new version.

Migrate the Data Protection for SAP for DB2 profile

The license file, the profile, and the configuration files are not deleted when Data Protection for SAP for DB2 is uninstalled. These files can be used by the new version of Data Protection for SAP. To reuse the existing configuration and connection to the Tivoli Storage Manager server, choose not to change the profile when you are prompted during installation.

Upgrade the Administration Assistant function for Data Protection for SAP V6.3

Perform these tasks to upgrade the Administration Assistant function for Data Protection for SAP server to a new version:

1. Verify that the Administration Assistant package is complete. The Administration Assistant is provided on each of the Data Protection for SAP installation discs or disc images, or downloaded from the IBM FTP server.
2. Verify that the new Administration Assistant requirements are met as described in “Prerequisites for Installing the Administration Assistant function for Data Protection for SAP” on page 31. Note that the Data Protection for SAP for DB2 release notes contain the latest requirement information.
3. Review planning information as described in “Prerequisites for Installing the Administration Assistant function for Data Protection for SAP” on page 31.
4. If you plan to migrate existing data to the new version, perform the tasks described in “Migrate Administration Assistant function for Data Protection for SAP data from a previous release” on page 37.
5. Uninstall the old version of the Administration Assistant as described in “Installing the Administration Assistant function for Data Protection for SAP Server-Level Components” on page 32.
6. Install the new version of the Administration Assistant server-level components as described in “Installing the Administration Assistant function for Data Protection for SAP Server-Level Components” on page 32.
7. Perform the configuration tasks beginning with “1. Preparing a secure connection” on page 47.
8. Set up the Administration Assistant client as described in “2. Configuring the Administration Assistant function for Data Protection for SAP Client” on page 48.
9. Verify the installation as described in “3. Verifying the Administration Assistant function for Data Protection for SAP installation” on page 48.

Note: It is possible to use the Administration Assistant in conjunction with supported Data Protection for SAP versions prior to version 5.4, provided the Administration Assistant is installed on a single host.

Migrate Administration Assistant function for Data Protection for SAP data from a previous release

Note: The following procedure must be performed before uninstalling the Administration Assistant and installing the new version. In addition, Data Protection for SAP for DB2 does not provide support for transferring data from an installation of the Administration Assistant prior to version 5.4. If desired, the report function can be used to capture data from the prior version before the new version is installed.

Migrating Database Data

Information on transferring data from the database of a previous version of the product.

Note: It is recommended that you make a backup of the current Administration Assistant function for Data Protection for SAP database before starting the migration process.

1. From Administration Assistant function for Data Protection for SAP 5.4

- a. The export tool is provided on each Data Protection for SAP for DB2 installation disc (or disc image) in the migration directory. This directory contains:

- aaDerbyAdaption.jar
- prepareExport.sql
- export.cmd (for use with Windows systems)
- export.sh and export ksh (for use with UNIX/Linux systems)

Copy these files from the installation disc (or disc image) for the new version of the Administration Assistant to your system.

- b. If you are using Apache Derby, get information on how to connect to the Apache Derby database. These settings are provided in file assist.cfg and are listed below:

- Location of your previous installation of the Administration Assistant
- Username to connect to the Apache Derby database
- Password to connect to the Apache Derby database
- Port to connect to the Apache Derby database
- Hostname of your system
- Name of the database
- Path to file aaDerbyAdaption.jar
- Directory where the data will be exported

- c.

Start the export script. The script guides you through the export process.

The directory that you specify in this step is the same directory from which you can later import data to the latest version of the product during the installation process.

2. From Administration Assistant function for Data Protection for SAP v5.5 or higher:

If you want to be able to access data from the currently running Administration Assistant database in a newer version of the Administration Assistant database of the same type, ensure that you do not uninstall the currently running

database during the Administration Assistant uninstallation process. When you are asked which components to uninstall, specify only Administration Assistant server and Database Agent.

When you install the newer version of the Administrative Assistant database, you are asked if you want to update an existing database. If you choose this option, and are using the Apache Derby database, specify the directory that contains the existing database. (The default directory is *AA_install_dir/aaDBSupport*.) If you are using DB2, you do not need to specify an import directory.

If you want to keep performance data that is not kept in the database, back up the complete history directory, including its subdirectories, before uninstalling the old version. After installing the new version, copy the performance data into the new installation directory.

As a result, the export directory contains several *.aa files.

During the installation process, you will be asked if you want to import old data. Within this dialog box you can enter the export directory you selected during the export.

Migrating Styles and Report Templates

Information on using existing styles and templates from a previous version of the product.

If you would like to reuse your styles and reports, save these directories from the installation directory to another directory.

Note: During the installation of the Administration Assistant function for Data Protection for SAP, all data in the installation directory will be removed.

After the installation process, you can copy these directories back to the installation directory of the Administration Assistant.

Chapter 5. Configuring Data Protection for SAP for DB2

Instructions about how to configure Data Protection for SAP for DB2 are provided.

Data Protection for SAP for DB2 requires certain configuration tasks to be performed for these applications:

- Data Protection for SAP base product
- Administration Assistant
- DB2 Log Manager and related DB2 files
- HACMP
- Distributed File System
- Tivoli Storage Manager backup-archive client
- Tivoli Storage Manager server

Configuration tasks for the Data Protection for SAP for DB2 base product

Instructions about how to configure the Data Protection for SAP for DB2 base product are provided.

Data Protection for SAP for DB2 requires that you complete certain configuration tasks before it performs a backup operation. Optional configuration tasks are identified in their description.

Verification tasks

Data Protection for SAP for DB2 requires these verification tasks to be performed as part of the product configuration.

Profile tasks

Data Protection for SAP for DB2 requires these tasks to be performed in the Data Protection for SAP profile as part of the product configuration.

Setting the SERVER statement in the Data Protection for SAP for DB2 profile

The SERVER statement is specified in the Data Protection for SAP for DB2 profile and there are corresponding keywords in the Tivoli Storage Manager client option file. Depending on the choice of password handling, some parameters are ignored. The corresponding sections in the Data Protection for SAP profile and the Tivoli Storage Manager client option file are established using the logical server name. This logical server name is defined by the keywords SERVER or SERVERNAME. The logical server names are also used by the "View TSM Server Utilization" function of the Administration Assistant. This function generates a separate entry for each logical server name found in the system landscape. Identical logical server names are considered to represent the same server.

2. Run `dsmc -servername=stanza1` and use the new password to generate a proper entry.
 3. Run `dsmc -servername=stanza2` and use the new password to generate the proper entry.
- [4] If you are using Tivoli Storage Manager API 5.2 (or later), you can use the `TCP_ADDRESS` parameter in the Data Protection for SAP profile. This parameter eliminates the need to set multiple stanzas in the Tivoli Storage Manager client option file for multiple paths and eliminates the problem when updating the password (see [3]).

Example of SERVER statement with alternate paths:

This example assumes that the Tivoli Storage Manager server is configured with two tape drives and two LAN connections. A backup is typically performed through network path 1 (SERVER statement 1). If network path 1 is unavailable, the backup is performed using network path 2 (SERVER statement 2). If path 1 is active, Data Protection for SAP for DB2 begins the two sessions as defined in the SERVER statement for path 1. Since `MAX_SESSIONS` also specifies 2, no more sessions are started. If path 1 is inactive, Data Protection for SAP starts 2 sessions on path 2. Since `MAX_SESSIONS` specifies 2, the backup is performed using path 2.

This is an example of the Data Protection for SAP profile used in this alternate path configuration:

```

MAX_SESSIONS    2          # 2 tape drives
.
.
SERVER    server_a      # via network path 1
ADSMNODE    C21
SESSIONS    2
PASSWORDREQUIRED    YES
BRBACKUPMGTCCLASS    mdb
BRARCHIVEMGTCCLASS    mlog1 mlog2
# USE_AT    0 1 2 3 4 5 6

SERVER    server_b      # via network path 2
ADSMNODE    C21
SESSIONS    2
PASSWORDREQUIRED    YES
BRBACKUPMGTCCLASS    mdb
BRARCHIVEMGTCCLASS    mlog1 mlog2
# USE_AT    0 1 2 3 4 5 6

```

Note that even if the logical names `server_a` and `server_b` actually point to the same Tivoli Storage Manager server, the Administration Assistant still considers them to be two different servers.

Example of SERVER statement with parallel servers:

This example assumes the following configuration:

- Two Tivoli Storage Manager servers (each with two tape drives) with connections through two network paths:
 - `server_a` uses TCP/IP address `xxx.xxx.xxx.xxx`
 - `server_b` uses TCP/IP address `yyy.yyy.yyy.yyy`
- An SAP® database server connected to two networks.
- Daily backups are performed on both systems.

This is an example of the Data Protection for SAP for DB2 profile used in this parallel configuration:

```

MAX_SESSIONS    4          # 4 tape drives
.
.
SERVER          server_a    # via network path 1
ADSMNODE        C21
SESSIONS        2
PASSWORDREQUIRED YES
BRBACKUPMGTCCLASS MDB
BRARCHIVEMGTCLASS MLOG1 MLOG2 MLOG3 MLOG4
# USE_AT        1 2 3 4 5 6 7

SERVER          server_b    # via network path 2  ADSMNODE        C21
SESSIONS        2
PASSWORDREQUIRED YES
BRBACKUPMGTCCLASS MDB
BRARCHIVEMGTCLASS MLOG1 MLOG2 MLOG3 MLOG4
# USE_AT        1 2 3 4 5 6 7

```

Example of SERVER statement with alternate servers:

This example assumes the following configuration:

- Two Tivoli Storage Manager servers:
 - server_a uses TCP/IP address xxx.xxx.xxx.xxx and uses four tape drives (MAX_SESSIONS 4)
 - server_b uses TCP/IP address yyy.yyy.yyy.yyy and uses four tape drives (MAX_SESSIONS 4)
- An SAP[®] database server connected to this FDDI network.
- Normal backups are performed with server a, which is local to the SAP database server.
- A disaster recovery backup is stored on remote server b every Friday.

This is an example of the Data Protection for SAP for DB2 profile used in this disaster recovery configuration:

```

MAX_SESSIONS    4          # 4 tape drives
.
.
SERVER          server_a    # via network path 1
ADSMNODE        C21
SESSIONS        4
PASSWORDREQUIRED YES
BRBACKUPMGTCCLASS MDB
BRARCHIVEMGTCLASS MLOG1 MLOG2 MLOG3 MLOG4
USE_AT          1 2 3 4

SERVER          server_b    # via network path 2
ADSMNODE        C21
SESSIONS        4
PASSWORDREQUIRED YES
BRBACKUPMGTCCLASS MDB
BRARCHIVEMGTCLASS MLOG1 MLOG2 MLOG3 MLOG4
USE_AT          5          # for Disaster Recovery

```

DB2 tasks

Data Protection for SAP for DB2 requires these DB2 tasks to be performed as part of the product configuration.

Reviewing DB2 and Data Protection for SAP for DB2 configuration guidelines

Data Protection for SAP for DB2 data transfer functions are implemented in a shared library that is accessed by DB2 whenever a backup or restore, and a log archive or log retrieve command, are issued. The shared library requires information on the path of the Data Protection for SAP profile and the path of the log files that are written by Data Protection for SAP. If an action is initiated using the DB2 commands `BACKUP DATABASE` or `RESTORE DATABASE`, the information required must be specified in a vendor environment file. The name of the vendor environment file is sent to DB2 through either the `OPTIONS` parameter of the `BACKUP DATABASE` or `RESTORE DATABASE` commands, or it can be stored persistently in the database configuration parameter `VENDOROPT` (for log archive or log retrieve, this can be stored either in the database configuration parameter `LOGARCHOPT1` or in `LOGARCHOPT2`). In the case of `BACKUP DATABASE` or `RESTORE DATABASE`, use of the `OPTIONS` keyword for this purpose is no longer necessary. It is strongly recommended that you keep the settings in the vendor environment file and in the system variables synchronised at all times. For an example of a Data Protection for SAP vendor environment file, see “Sample DB2 Vendor Environment File” on page 152. If `BACKUP DATABASE` or `RESTORE DATABASE` is triggered using the `backom` utility, the information required must be specified in the environment.

Consider these additional adjustment rules for Data Protection for SAP:

- To select a different set of Data Protection for SAP environment settings for a DB2 backup or restore, specify the full path of the vendor environment file in the `OPTIONS` parameter of the `BACKUP DATABASE` or `RESTORE DATABASE` commands. For details, refer to *DB2 Command Reference*.
- To select a different Data Protection for SAP profile, modify the environment variable `XINT_PROFILE` to denote the new profile in the vendor environment file.
- To select a different Data Protection for SAP profile for future calls to the `backom` utility, modify the environment variable `XINT_PROFILE` to denote the new profile.
- To select a different Data Protection for SAP profile for a call to the `backom` utility, specify the path of the new profile in option `-e` of the `backom` command.
- To change the path for Data Protection for SAP process log files for a call to DB2 commands `BACKUP DATABASE` or `RESTORE DATABASE`, modify the environment variable `TDP_DIR` in the vendor environment file and specify the file path in the `OPTIONS` parameter of the `BACKUP DATABASE` or `RESTORE DATABASE` commands.
- To change the path for Data Protection for SAP process log files for future calls to the `backom` utility, modify the environment variable `TDP_DIR` to denote the new profile.

Specifying the VENDOROPT parameter

In order to select a default set of Data Protection for SAP for DB2 environment settings for DB2 commands BACKUP DATABASE, RESTORE DATABASE and for the DB2 Log Manager, modify the DB2 database configuration to denote a file containing the settings:

```
db2 update db cfg for SID using LOGARCHOPT1|2 vendor environment file
```

where <vendor environment file> is the fully qualified path of the file containing Data Protection for SAP environment settings for DB2. Make sure that the environment settings of your system match the settings in this file.

This command can be used as an alternative to the db2set command and provides these advantages:

- There is no need to restart the DB2 instance.
- You can define default values for the OPTIONS parameter of the BACKUP DATABASE and RESTORE DATABASE commands in the DB2 configuration, thus making the OPTIONS parameter of these commands optional. (You can still override the default setting of the database configuration by specifying the OPTIONS parameter.)
- The same settings apply to database backup/restore and to log file archive/retrieve.

When using the BACKUP DATABASE and RESTORE DATABASE commands with the USE SNAPSHOT option for snapshot-based backup and restore by DB2 ACS or FlashCopy Manager, the VENDOROPT parameter is ignored. In this case, any options other than the default values must be set using the OPTIONS keyword.

Configuring the DB2 Log Manager

The following database configuration parameters are applicable to DB2 database backup and restore, and log archive and retrieve with Data Protection for SAP:

Table 9. Configuration parameters for DB2 database backup and restore, and log archive and retrieve

Parameter	Description	Default
LOGARCHMETH1	Media type of the primary destination for archived log files	Off
LOGARCHOPT1	Options field for the primary destination for archived log files (if required).	NULL
LOGARCHMETH2	Media type of the secondary destination for archived log files. If this path is specified, log files will be archived to both this destination and the destination specified by LOGARCHMETH1.	Off
LOGARCHOPT2	Options field for the secondary destination for archived log files (if required).	NULL

Table 9. Configuration parameters for DB2 database backup and restore, and log archive and retrieve (continued)

Parameter	Description	Default
FAILARCHPATH	If DB2 is unable to archive log files to both the primary and secondary (if set) archive destinations due to a media problem, then DB2 will try to archive log files to this path. This path must be a disk.	NULL
NUMARCHRETRY	Number of retries to archive a log file to the primary or secondary archive destination before trying to archive log files to a failover directory. This is only used if FAILARCHPATH is set. If NUMARCHRETRY is not set, DB2 will continuously retry archiving to the primary or secondary log archive destination.	5
ARCHRETRYDELAY	Number of seconds to wait after a failed archive attempt before trying to archive the log file again. Subsequent retries will only take affect if NUMARCHRETRY is at least set to 1.	2

To activate log archival or retrieval with the DB2 Log Manager facility, modify the DB2 database configuration during the installation. The following two changes to the database configuration are the minimum changes necessary to use the DB2 Log Manager with Data Protection for SAP:

1. Update one of the LOGARCHMETH database configuration parameters (this example uses LOGARCHMETH1):

- (UNIX and Linux):

```
db2 update db cfg for SID using LOGARCHMETH1 VENDOR:/path/shared library
```

- (Windows):

```
db2 update db cfg for SID using LOGARCHMETH1 VENDOR:drive:\path\tdpdb2.d11
```

2. Update the Data Protection for SAP environment. A file that contains the environment settings must be made available to DB2 to allow DB2 to provide this environment for Data Protection for SAP archive or retrieve requests. This file is an additional requirement. This example shows the setup needed by Data Protection for SAP for LOGARCHMETH1:

- (UNIX and Linux):

```
db2 update db cfg for SID using LOGARCHOPT1 /path/vendor.env
```

- (Windows):

```
db2 update db cfg for SID using LOGARCHOPT1 drive:\path\vendor.env
```

The update to LOGARCHMETH takes effect during the next log file archive.

The database configuration parameters *LOGRETAIN* and *USEREXIT* are still available but are mapped to the parameter *LOGARCHMETH1*. For further description of the DB2 Log Manager, see the *DB2 Administration Guide*. Configure Data Protection for SAP so that at least one Tivoli Storage Manager session (one for the database backup and one for the log archives) is available for each of these operations.

Creating the configuration files

When setting the Tivoli Storage Manager password with the *backom* utility, the configuration files for all DB2 partitions are automatically created in the paths *path/%DB2NODE/*, where *path* is the directory denoted by the value of keyword *CONFIG_FILE* in the profile and the string *%DB2NODE* is replaced automatically by a DB2 partition name referenced in the DB2 configuration file *db2nodes.cfg*. If the directory denoted by the value of keyword *CONFIG_FILE* is not located in the same network file system where the DB2 configuration file *db2nodes.cfg* is located, this procedure must be repeated for each machine where a partition of the database resides. If the database is not partitioned, *NODE0000* is used as the only DB2 partition name.

Optional: Setting backup object segmentation

Environments that contain large databases that rapidly increase in size might encounter problems when transferring data to the Tivoli Storage Manager server. For example, you might encounter the following problems when backing up or restoring large databases:

- Canceling a running backup session takes an unacceptably long time. This behavior is due to multiple internal processing activities on the Tivoli Storage Manager server.
- The recovery log for the Tivoli Storage Manager internal database might become unavailable when processing large databases. This unavailability prevents immediate access to important recovery data.

To avoid potential problems related to transferring large objects, use the Data Protection for SAP *SEGMENTSIZ*E profile keyword. This keyword specifies the upper bound of the segments that are split from large backup objects during backup and restore processing. For more information about the *SEGMENTSIZ*E keyword, see “Tivoli Storage Manager for ERP for DB2 profile parameter descriptions” on page 134.

Administration Assistant function for Data Protection for SAP tasks

Data Protection for SAP for DB2 requires these Administration Assistant function for Data Protection for SAP tasks to be performed as part of the product configuration.

1. Preparing a secure connection

By default, the Administration Assistant function for Data Protection for SAP is set up to accept unsecure (HTTP) client requests. If the Administration Assistant was set up for secure (HTTPS) connection during installation, then proceed to the next step.

The secure communication between the Administration Assistant Server component and its clients is realized with the Secure Socket Layer (SSL) protocol. This protocol requires that both the server and client be integrated in a public key infrastructure (PKI). The Server component requires these settings:

- An HTTPS port to listen on for HTTPS connect requests.
- A keystore containing a key pair it uses to identify itself to the clients and when connecting internally to the RMI registry. The server hostname is used as an alias for this key pair. Since the keystore contains the server private key, precautions must be taken that prevent access by unauthorized persons.
- A truststore containing trusted certificates that allow verifying the server's signature. If the server certificate was digitally signed by an official certificate authority whose root certificate is available in the truststore by default, there is nothing to be done. If however, the server identifies itself with a self-signed certificate, this certificate must be imported into the truststore as well.
- Be sure to remove this trusted certificate from the truststore as soon as the officially signed server certificate is available and employed. A setup using self-signed certificates is not recommended for production environments.
- Both the keystore and truststore can be modified with your keystore management tool. This tool varies by platform and provider.

Perform these tasks to set up the Administration Assistant Server component for secure communication:

1. Remove the keyword `nonsecure` from the Server configuration file (`assist.cfg`).
2. Specify the appropriate HTTPS port number in the Server configuration file:

```
httpsport=https port number
```

The default HTTPS port number is 443.

3. Add the keystore, keystore password, and truststore to the appropriate Java call. The Java calls are shown in **bold** text:

```
-Djavax.net.ssl.keyStore=keystore  
-Djavax.net.ssl.keyStorePassword=password for keystore  
-Djavax.net.ssl.trustStore=truststore
```

- (UNIX and Linux): add the parameters to `sadma.sh`
- (Windows): add the parameters to `sadmt.cmd` and to the registry. The Windows registry key is:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\...  
...AdminAssistant\Parameters\AppParameters
```

If you do not specify one or more of these parameters, the defaults of your Java virtual machine will be used.

4. Make sure the required certificates are contained in the keystore and trust store.
5. Restart the Administration Assistant Server component.

When changing the Administration Assistant server from nonsecure to secure mode using a self-signed certificate, remember to also prepare the Administration Assistant clients as described in “2. Configuring the Administration Assistant function for Data Protection for SAP Client” and “4. Configuring a scheduling client to create reports” on page 49.

2. Configuring the Administration Assistant function for Data Protection for SAP Client

The Administration Assistant function for Data Protection for SAP client invokes a Java applet when connecting to the Administration Assistant function for Data Protection for SAP Server component. Make sure these requirements are met when setting up the Administration Assistant client:

- Make sure all Administration Assistant Client prerequisites are met as described in “Prerequisites for Installing the Administration Assistant function for Data Protection for SAP” on page 31.
- The browser must be enabled to accept cookies.
- Advertisements and pop-up panels must not be blocked unless `index.html` is used in the address.
- A secure connection requires that the client Java plugin must be able to verify the certificate presented by the Administration Assistant Server component. In a production environment, this is typically performed at the server level as the server certificate is signed by an official certificate authority whose root certificate is contained in the plugin truststore. If the server identifies itself with a self-signed certificate, this certificate must be imported into the plugin truststore. If you did not use the using the Java Plugin Control Panel to replace the plugin truststore, the file `cacerts` (located in the Java security path) is used as the truststore. The file is modified with the keystore management tool. This tool varies by platform and provider. For example, the Sun Microsystems **keytool** is modified with this command:

```
keytool -import -alias Server component hostname -file cert_file  
-keystore trustore
```

- Be sure to remove the self-signed trusted certificate from the truststore as soon as the officially signed server certificate is available and activated. A setup with self-signed certificates is not recommended for production environments.

3. Verifying the Administration Assistant function for Data Protection for SAP installation

Perform this task to verify the installation of the Administration Assistant function for Data Protection for SAP. Make sure to use the ADMIN userid (with password 'admin') for the initial login:

- (Nonsecure connection): If the Server component was started with the keyword `nonsecure` in the Server configuration file, connect to the Administration Assistant Server component from a client machine with this command:

```
http://Server component host name:http port
```

Optionally, you can make the connection without opening a new browser window by issuing this command:

```
http://Server component host name:http port/index.html
```

- (Secure connection): If the Server component was started with the keyword `secure` in the Server configuration file, connect to the Administration Assistant Server component from a client machine with this command:

```
https://Server component host name:https port
```

Optionally, you can make the connection without opening a new browser window by issuing this command:

```
https://Server component host name:https port/index.html
```

Use the client function *Administer Users* to change the default password immediately after establishing a connection. As soon as an instance of Data Protection for SAP for DB2 connects to your Administration Assistant Server component, the instance will be displayed in the list of Data Protection for SAP servers. For details on how to set up your instance of Data Protection for SAP to connect to a specific Server component, refer to “Specifying a new Administration Assistant function for Tivoli Storage Manager for ERP” on page 87.

4. Configuring a scheduling client to create reports

A scheduling client server must be set up in order to create reports with templates. Perform these tasks to set up a scheduling client server:

1. Select a system that meets the requirements as described in “Prerequisites for Installing the Administration Assistant function for Data Protection for SAP” on page 31.
2. Copy files `Admt.jar` and `NLS.jar` from the installation directory of the Administration Assistant Server component to the scheduling client system. Before generating a report, make sure that these files are specified in the `CLASSPATH` and that the JVM is included in the `PATH`. See “Sample Shell Script for Scheduling a Report from a UNIX Scheduling Client” on page 148 or “Sample Command File for Scheduling a Report from a Windows Scheduling Client” on page 148 for a sample script.
3. In case the Administration Assistant Server component is started in secure mode, set up a public key infrastructure between the scheduling client and the Server component. In a production environment, this is typically performed at the server level as the server certificate is signed by an official certificate authority whose root certificate is contained in the plugin truststore. If the server identifies itself with a self-signed certificate, this certificate must be imported into the plugin truststore. If you did not use the Java Plugin Control Panel to replace the plugin truststore, the file `cacerts` (located in the Java security path) is used as the truststore. The file is modified with the keystore management tool. This tool varies by platform and provider. For example, the Sun Microsystems **keytool** is modified with this command: `keytool -import -alias <Server component hostname> -file cert_file -keystore trustore`

Defining thresholds

You can define limits (or thresholds) for various states pertaining to the Administration Assistant function for Data Protection for SAP environment. The threshold status is shown in the "Monitor Backup States" and "Backup State - Detailed View" panels. These are predefined threshold types:

- Backup duration (in minutes or hours)
- Backup size (in MB or GB)
- Throughput rate (in GB per hour or MB per second)
- Time since the last complete backup (in hours or days)
- Size of all log file backups since the last complete backup (in MB or GB)
- Recovery point objective (maximum time permitted since the last backup, in minutes or hours)

When a threshold is exceeded, this is reported in the "Threshold Status" column of the "Monitor Backup States" panel, and an e-mail describing the exception in more detail is sent to any e-mail addresses defined for the threshold. A *lifetime* parameter associated with each threshold defines the length of time between e-mail notifications, provided the threshold remains in alert status. The Administration Assistant Online Help provides information about threshold definitions.

Distributed file system tasks

Data Protection for SAP for DB2 requires these tasks to be performed to configure Data Protection for SAP in a distributed file system.

Configuring Tivoli Storage Manager for ERP for DB2 in a Distributed File System

This set up task is not required if the following conditions exist:

- All SAP® systems to are statically assigned to specific hosts. For example, the instances are not moved between hosts.
- The root user is granted read/write access permission to the distributed file system.

If these conditions exist, the standard installation process can be used as described in "Required installation tasks" on page 23.

For a single SID located on a host, Tivoli Storage Manager for ERP sets the ProLE service to run with the *db2SID* user ID instead of root. Perform these tasks to set up the ProLE service to run with the *db2SID* user ID:

1. Enable root access to the distributed file system.
2. Install Tivoli Storage Manager for ERP using the procedure described in "Required installation tasks" on page 23.
3. On a UNIX system, replace the following entry in the */etc/inittab* file:

```
pd64:345:respawn:/usr/tivoli/tsm/tdp_r3/db264/prole -p profile
```

with this entry:

```
pd64:345:respawn:su - db2SID -c /usr/tivoli/tsm/tdp_r3/db264/prole -p profile
```

If upstart is configured, the *init* script */etc/init/prole_db2.conf* must be used. *SID* must be the actual SID.

4. Refresh the `/etc/inittab` processes.
5. Disable root access to the distributed file system.

For multiple SIDs on a host system, run the ProLE service by root with permanent read/write permission to the distributed file system.

Configuring Data Protection for SAP for DB2 in a Distributed File System in an Adaptive Computing Environment

Certain setup tasks must be performed when Data Protection for SAP for DB2 is used in an Adaptive Computing Environment. Since the Adaptive Computing Environment currently does not allow more than one SID per host, the root user does not require additional permissions for the distributed file system. Perform these tasks to prepare installation in the distributed file system:

1. Log in as root user and perform a regular installation of Data Protection for SAP on one of the systems participating in the distributed file system. During the installation procedure, make sure the configuration files reside in a directory that is not located in the distributed file system. These files will not be used and can be deleted after installation.
2. After installation completes successfully, copy the contents of the installation directory to a temporary directory in the distributed file system. For example:

```
mkdir /san/SanFS/tivoli/tdp_r3
cp -r /usr/tivoli/tsm/tdp_r3/db264 /san/SanFS/tivoli/tdp_r3
```

3. Each of the SAP® environments can now be set up to use Data Protection for SAP for backup and recovery. In the Adaptive Computing Environment, Data Protection for SAP backup and recovery tasks can be performed from the same host for all participating SIDs. For each SID, log in as the database instance owner and run the 'SanFSsetupSID.sh' script from the installation path in the distributed file system. For example:

```
/san/sanFS/tivoli/tdp_r3/db264/SanFSsetupSID.sh
```

The following information must be provided to the script:

- a. The SID for the SAP system to be backed up.
 - b. The path for the Data Protection for SAP profile and configuration file (`initSID.utl`, `initSID.bki`).
 - c. To connect to an Administration Assistant server, specify the hostname or IP address and server port for the Administration Assistant server.
4. The script `SanFSsetupSID.sh` creates scripts `prepareTDPSAP_SID.sh`. On each host, log in as root user and run the `prepareTDPSAP_SID.sh` script with the appropriate `SID`. If this script is placed in the distributed file system, make sure the root users have the appropriate permissions to run it.
 5. Whenever a SID is moved to a different host, the '`prepareTDPSAP_SID.sh`' script must be run by the root user of the new host.

HACMP tasks

Data Protection for SAP for DB2 requires these tasks to be performed to use Data Protection for SAP in a High Availability Cluster Multi-Processing environment.

Configuring Data Protection for SAP for DB2 as an HACMP Application

A prerequisite for installation is a correct setup of the Tivoli Storage Manager client. The installation steps for the Tivoli Storage Manager Backup/Archive Client for AIX can be found in the documentation *Tivoli Storage Manager Installing the Clients*.

Data Protection for SAP for DB2 must be defined as an application to HACMP. Although the *HACMP for AIX Installation Guides* should be reviewed for detailed directions, a high-level summary is provided here. Note that Data Protection for SAP must be in a resource group having a cascading or rotating takeover relationship. It does not support a concurrent access resource group. Perform these tasks to configure Data Protection for SAP an application for HACMP:

1. Enter this command start HACMP for AIX system management:

```
smit hacmp
```

2. Select Cluster Configuration > Cluster Resources > Define Application Servers > Add an Application Server.
3. Enter field values as follows:

Server Name

Enter an ASCII text string that identifies the server (for example, `tdpclientgrpA`). You use this name to refer to the application server when you define it as a resource during node configuration. The server name can include alphabetic and numeric characters and underscores. Do not use more than 31 characters.

Stop Script

Enter the full pathname of the script that stops the server (for example, `/usr/sbin/cluster/events/utls/stop_tdpr3.sh`). This script is called by the cluster event scripts. This script must be in the same location on each cluster node that might stop the server.

4. Press Enter to add this information to the HACMP for AIX ODM.
5. Press F10 after the command completes to leave SMIT and return to the command line.

Refer also to the *HACMP for AIX Planning Guide V4.4* for further information about selecting the HACMP node topology and takeover relationships.

Adding Data Protection for SAP for DB2 to an HACMP Resource Group:

A final step in enabling Data Protection for SAP for DB2 for HACMP failover is to define it to a cluster resource group. Although the *HACMP for AIX Installation Guides* should be reviewed for detailed directions, a high-level summary is provided here. Perform these tasks to define the resources that will be part of a resource group:

1. From the Cluster Resources SMIT screen, select the Change/Show Resources/Attributes for a Resource Group option and press Enter. SMIT displays a picklist of defined resource groups.
2. Pick the desired resource group.
3. Press Enter and SMIT displays the Configure a Resource Group screen.

4. Enter values that define all the resources you want to add to this resource group.
5. After entering field values, synchronize cluster resources.
6. Press F10 to exit SMIT or F3 to return to previous SMIT screens to perform other configuration tasks or synchronize the changes you just made. To synchronize the cluster definition, go to the Cluster Resources SMIT screen and select the Synchronize Cluster Resources option.

The Tivoli Storage Manager client application should be added to the same resource group that contains the file systems it will back up. The file systems defined in the resource group should also be the ones specified in the domain for this client instance in the client user options file. Note that both JFS and NFS file systems can be defined as cluster resources, although NFS supports only 2 node clusters in a cascading takeover relationship.

HACMP stop script example:

This section illustrates a stop script in an HACMP environment.

Depending on the installation environment, the sample stop script might need to ensure that any backup or restore operation in progress can be stopped.

The stop script is used in the following situations:

- HACMP is stopped.
- A failover occurs due to a failure of one component of the resource groups. The other members are stopped so that the entire group can be restarted on the target node in the failover.
- A fallback occurs and the resource group is stopped on the node currently hosting it to allow transfer back to the node re-entering the cluster.

The stop script will be called by HACMP as the root user.

Note: This script is not in its final form. It should be considered pseudo code that indicates the functions it will perform.

```

#!/bin/ksh
#####
# This sample script is provided for use with
Data Protection for SAP in an HACMP #
# environment #
# It should be reviewed and customized to meet your specific environment #
# #
# #
# Name: stop_tdpr3.sh #
# #
# #
#####

if ["$VERBOSE_LOGGING"="high"]
then
    set -x
fi

# Function to update all disk information for Data Protection for SAP

STOP_BACKUP()
{
# You may want to cancel all backups currently running
# Note that this will generate errors in the current backup logs and it will also
# cancel the connection to the Admin Assistant.
# *** Note that if you are using Data Protection for Snapshot Devices for SAP,
# this may leave your FlashCopy device in an
# inconsistent state.
# kill -9 `cat /var/tdp_r3/prole.pid`

# This stops any running backup or archive process.

STOP_BACKUP
Exit 0

```

Configuration tasks for Tivoli Storage Manager

Instructions about how to configure the Tivoli Storage manager client and server for Data Protection for SAP for DB2 operation are provided.

Data Protection for SAP for DB2 requires that you complete certain configuration tasks for the Tivoli Storage Manager backup-archive client and server.

Tivoli Storage Manager client tasks

Data Protection for SAP for DB2 requires these tasks to be performed for the Tivoli Storage Manager client as part of the product configuration.

Configure the Tivoli Storage Manager client options

The Tivoli Storage Manager clients must be configured after the Tivoli Storage Manager server is configured. These clients include the *backup-archive client* for the file system backups and the *Application Programming Interface (API) client* for interface programs. The API client allows users to enhance existing applications with backup, archive, restore, and retrieve services. An installed and confirmed API client is a prerequisite for Data Protection for SAP for DB2.

The clients must be installed on all nodes that will interface with the Tivoli Storage Manager server. In an SAP® system landscape, this means that the backup/archive client must be installed on every machine scheduled for a file system backup, such as SAP application servers and the SAP database server. The Tivoli Storage

Manager API client only needs to be installed on the SAP database server machine to enable backup and restore operations of the SAP database using Data Protection for SAP. The Administration Assistant uses the logical Tivoli Storage Manager server names in its "View TSM Server Utilization" function. Identical logical names are considered to represent the same Tivoli Storage Manager server, but different entries are generated for each logical server name found in the system landscape. Therefore, use identical logical server names when pointing to the same Tivoli Storage Manager server throughout the system landscape and use different logical server names when different Tivoli Storage Manager servers are addressed.

Set Tivoli Storage Manager client options on UNIX or Linux:

Tivoli Storage Manager clients on UNIX or Linux are configured by setting options in the `dsm.opt` and `dsm.sys` files. The `include/exclude` file is also used to define which files are included or excluded during backup, archive, or hierarchical storage processing. Examples of an `include/exclude` file for UNIX or Linux can be found in "Include/Exclude List Sample (UNIX and Linux)" on page 150. Perform these tasks to configure the Tivoli Storage Manager backup/archive clients to operate in an SAP® environment:

1. Install the Tivoli Storage Manager client software on the SAP database server machine.
2. Edit the client system options file `dsm.sys` and set these values as appropriate for your installation:

```
Servername      server_a
TCPSPort        1500
TCPSPServeraddress xxx.xxx.xxx.xxx or servername
InclExcl        /usr/tivoli/tsm/client/ba/bin/inclexcl.list
Compression     OFF
```

3. Specify `TCPSPServeraddress 127.0.0.1` or loopback if the server and client are on the same machine. This improves TCP/IP communication speed.
4. Specify `InclExcl` if you want Tivoli Storage Manager to include or exclude the files listed in `inclexcl.list`. This is optional. You may want to exclude all database files that are processed by the DB2 database backup.
5. Throughput improves when tape drives attached to the Tivoli Storage Manager server provide hardware compression. However, combining hardware compression and Tivoli Storage Manager client software compression (Compression ON) is not recommended.
6. Edit the client user options file `dsm.opt` and set these values as appropriate for your installation:

```
LANGUAGE        AMENG      (this is the default value)
NUMBERFormat    1          (this is the default value)
TAPEPROMPT      NO
TIMEFORMAT      1          (this is the default value)
```

When the Tivoli Storage Manager API client is installed on a UNIX or Linux system, make sure there is a softlink `/usr/lib/libApiDS.a` that points to the `libApiDS.a` file in the Tivoli Storage Manager API installation directory (`/usr/tivoli/tsm/client/api/bin64`).

TSM provides two features that allow specifying the location of the TSM API Client error log: the environment variable `DSMI_LOG` and the TSM system client

option `ERRORLOGName` in `dsm.sys`. `DSMI_LOG` specifies a directory to which a file named `dsierror.log` will be written, while `ERRORLOGName` sets a path and user-defined file name.

In order to achieve conclusive logical linking of the environment, configuration and log files in your SAP backup/archive system, we recommend using the TSM system client option `ERRORLOGName` rather than the environment variable `DSMI_LOG`. The main advantages are:

- As opposed to `DSMI_LOG`, `ERRORLOGName` allows including the SID in the file name. This can speed up problem determination by simplifying identification of the correct error log file and matching its name with the active user client options file name, which should also contain the SID and be stored in environment variable `DSMI_CONFIG`. This is especially useful on machines with several SIDs.

The following is the suggested setup for Data Protection for SAP for DB2 on AIX:

1. For each "SERVER *servername*" section in the profile `initSID.utl`, create a corresponding "S`Ervername servername`" stanza in the system client options file `/usr/tivoli/tsm/client/api/bin64/dsm.sys`, where `SID` designates the DB2 database name as returned by "echo `$DB2DBDFT`". One SID may use several "S`Ervername servername`" stanzas, but we do not recommend the use of a "S`Ervername <servername>`" stanza by several SIDs.
2. In all "S`Ervername servername`" stanzas belonging to the same SID, add option "ERRORLOGName `/writeable_path/dsierror_SID.log`". Write permission problems can usually be avoided by specifying a directory below `$HOME` of the DB2 instance owner as `writeable_path`.
3. Create one user options file for each DB2 SID with the filename `/usr/tivoli/tsm/client/api/bin64/dsm_SID.opt` containing option "S`Ervername servername`". `servername` should point to the stanza in `/usr/tivoli/tsm/client/api/bin64/dsm.sys` that is designated by the first "SERVER *servername*" section in `initSID.utl`. Add variable `DSMI_CONFIG=/usr/tivoli/tsm/client/api/bin64/dsm_SID.opt` to the environment of the user who is running the SAP backups, usually `db2SID` or `SIDadm`, or both in case of doubt.

With this recommended setup, you obtain the following logical interlinking:

- environment variable `DSMI_CONFIG` is exported from the login shell
- environment variable `DSMI_CONFIG` points to client user options file `/usr/tivoli/tsm/client/api/bin64/dsm_SID.opt`
- client user option "SERVER *servername*" in `dsm_SID.opt` points to the "SERVER *servername*" stanza in `/usr/tivoli/tsm/client/api/bin64/dsm.sys`
- the "SERVER *servername*" stanza contains the option "ERRORLOGName `/writeable_path/dsierror_SID.log`"

If the variable `DSMI_LOG` already exists in your environment from an earlier setup, its will be overridden by `dsm.sys` option `ERRORLOGName` as configured above. However, in order to avoid confusion, make sure the `DSMI_LOG` path is identical to the path in `ERRORLOGName`. Alternatively, you can remove `DSMI_LOG` completely from your environment.

Set Tivoli Storage Manager client options on Windows:

Tivoli Storage Manager clients on Windows are configured by setting options in the file *server_a.opt* (where *server_a* is the logical server name in the *initSID.utl* file). The *include/exclude* file is also used to define which files are included or excluded during backup, archive, or hierarchical storage processing. Examples of an *include/exclude* file for Windows can be found in “Include/Exclude List Sample (Windows)” on page 151. Perform these tasks to configure the Tivoli Storage Manager backup/archive clients to operate in an SAP® environment:

1. Install the Tivoli Storage Manager client software on the SAP database server machine.
2. For each logical Tivoli Storage Manager server, a corresponding client option file is required. In this example, the file name must be *server_a.opt* since *server_a* is the logical server name:

TCPPort	1500
TCPServeraddress	xxx.xxx.xxx.xxx
InclExcl	c:\tivoli\tsm\baclient\incl excl.list
Compression	OFF

In addition, the environment variable *DSMI_CONFIG* must specify the corresponding client options file (for example *c:\tivoli\tsm\api\server_a.opt*).

3. Specify *TCPServeraddress* 127.0.0.1 or loopback if the server and client are on the same machine. This improves TCP/IP communication speed.
4. Specify *InclExcl* if you want Tivoli Storage Manager to include or exclude the files listed in *incl excl.list*. This is optional. You may want to exclude all database files that are processed by the DB2 database backup.
5. Throughput improves when tape drives attached to the Tivoli Storage Manager server provide hardware compression. However, combining hardware compression and Tivoli Storage Manager client software compression (*Compression ON*) is not recommended.

A Tivoli Storage Manager error log (required for each client) can be specified for each process regardless of the number of Tivoli Storage Manager client option files *server.opt* involved. The Tivoli Storage Manager error log is determined by these rules:

1. The Tivoli Storage Manager Client log is written to the file specified by the *DSMI_LOG* environment variable.
2. If the *DSMI_LOG* environment variable is absent or is not writeable, the Tivoli Storage Manager client log is written to the file specified with keyword *ERRORlogname* in the client system options file *dsm.opt*.
3. If there is no *ERRORlogname* in *dsm.opt* or if it is not writeable, the Tivoli Storage Manager client log is written to file *dsierror.log* in the local path.

It is recommended to set up the Tivoli Storage Manager client so that different processes write to separate error logs. Therefore, the error log path should be defined in the *DSMI_LOG* environment variable if the client options files are shared among processes.

Tivoli Storage Manager server tasks

Data Protection for SAP for DB2 requires these tasks to be performed for the Tivoli Storage Manager server as part of the product configuration.

Configure the Tivoli Storage Manager server

Tasks required to set up the Tivoli Storage Manager server, general server configurations, and specific server configurations (such as setup of storage devices) are provided. Although the task examples use Tivoli Storage Manager commands, these tasks can also be performed using the Tivoli Storage Manager Web client GUI.

Consider these performance-related guidelines before installing the Tivoli Storage Manager server:

Tivoli Storage Manager server host machine

The Tivoli Storage Manager server should be installed on an exclusive machine. The tasks presented in this section avoid concurrent processes and disk I/O access with other applications. A single Tivoli Storage Manager server is sufficient for a single SAP® system landscape. If the Tivoli Storage Manager server will be used to back up and restore other clients, consider installing the server on a large machine or using several Tivoli Storage Manager servers.

Network topology

Network topologies such as Fast Ethernet and Gigabit Ethernet work well with the Tivoli Storage Manager server. Fast network topologies should be used to prevent bottlenecks during backup and restore operations. The Tivoli Storage Manager server supports multiple network adapters. This support increases server throughput by providing multiple connections to the same network or by providing several physically distinct networks for the same server.

In the AIX: LPAR environment

An LPAR node can be used for a Tivoli Storage Manager server. The use of a High Performance Switch network can improve backup and performance.

These steps are considered complete once the Tivoli Storage Manager server is successfully installed:

- Recovery log volume has been allocated and initialized.
- Recovery log mirror volume has been allocated and initialized.
- Database volume has been allocated and initialized.
- Database mirror volume has been allocated and initialized.
- Additional labeled volumes for the backup and archive storage pools have been allocated and initialized (disks, tapes or combinations).
- Licenses have been registered.
- The Tivoli Storage Manager server has been started.

The latest code fixes for Tivoli Storage Manager can be found at:
<ftp://index.storsys.ibm.com/tivoli-storage-management/maintenance>

1. Specify a Tivoli Storage Manager server:

Perform these tasks to add a Tivoli Storage Manager server:

1. Add a new server statement to the Data Protection for SAP for DB2 profile.
2. Adapt the Tivoli Storage Manager options files as described in “8. Verify the Tivoli Storage Manager server name” on page 64.
3. Set and save the Tivoli Storage Manager password for the new server as described in “Set the Tivoli Storage Manager password” on page 62.

2. Specify a storage device:

A storage device defines a device class which handles the type of media, such as tape libraries or jukeboxes. The default device class defined for disks is DISK and is considered sufficient. Verify that these items are established within the Tivoli Storage Manager server after installation:

- Query the defined library:

```
q library
```

- Query the defined drives:

```
q drive
```

- Query the defined device class:

```
q devclass
```

3. Define a storage pool:

A storage pool is a named collection of storage volumes that are associated with one device class. Each storage pool represents a collection of volumes that are the same media type. The storage pool setup defines the storage hierarchy for the appropriate environment. In an SAP® environment, these data types can be backed up:

- SAP system data
- SAP database data (containers, offline log files)

To separate this data within the Tivoli Storage Manager server, define appropriate storage pools for each of these data collections. Log on as the Tivoli Storage Manager Administrator using the *Admin Command Line* or the *Web Admin* and run these commands to define storage pools:

1. Define a storage pool for the SAP system data:

```
define stgpool sap_incr  
device_class_name maxscr=5
```
2. Define a storage pool for the containers:

```
define stgpool sap_db  
device_class_name maxscr=20
```
3. Define a storage pool for the offline log files:

```
define stgpool sap_log1  
device_class_name maxscr=3
```

When a library tape device is associated, the maximum number of *scratch volumes* (labeled volumes which are empty or contain no valid data) that this storage pool will be allowed to use (parameter `maxscr`) must be defined. The maximum number of scratch tapes depends on the size of the database, the capacity of the tapes, the number of scratch volumes available, and how many versions of the backup must be retained. Replace these values with appropriate estimates.

4. Define a policy:

Tivoli Storage Manager policies specify how files are backed up, archived, migrated from client node storage, and also how they are managed in server storage. A policy definition includes the definition of a *policy domain*, a *policy set*, *management classes*, and *copy groups*. After setting definitions, a default policy set must be assigned, validated, and activated. For the policy definition, log on as a Tivoli Storage Manager Administrator using the *Admin Command Line* or the *Web Admin* and run these commands:

1. Define a policy domain and policy set:

```
define domain sap_c21
define policyset sap_c21 p_c21
```

2. Define a management class for file system backups, data files, offline log files and copies of offline log files:

```
define mgmtclass sap_c21 p_c21 mdefault
define mgmtclass sap_c21 p_c21 mdb
define mgmtclass sap_c21 p_c21 mlog1
define mgmtclass sap_c21 p_c21 mlog2
```

If you are planning to use this Tivoli Storage Manager server with multiple SAP® systems, use a set of different management classes for each system.

3. Define a copy group:

```
define copygroup sap_c21 p_c21 mdefault type=backup destination=sap_incr
define copygroup sap_c21 p_c21 mdefault type=archive destination=archivepool
define copygroup sap_c21 p_c21 mdb type=archive destination=sap_db retver=nolimit
define copygroup sap_c21 p_c21 mlog1 type=archive destination=sap_log1 retver=nolimit
define copygroup sap_c21 p_c21 mlog2 type=archive destination=sap_log2 retver=nolimit
```

Data Protection for SAP for DB2 uses its own *version control* mechanism for managing SAP database backups by backing up all data to only those management classes for which an archive copy group has been defined (parameter type set to archive). In addition, to prevent backed up files within Tivoli Storage Manager server storage from being deleted because of their expiration date (Tivoli Storage Manager deletes expired files), the copy group parameter *retver* (specifies the number of days a file is to be kept) should be set to unlimited (9999 or *nolimit*).

4. Assign the default management class:

```
assign defmgmtclass sap_c21 p_c21 mdefault
```

5. Validate and activate the policy set:

```
validate policyset sap_c21 p_c21
activate policyset sap_c21 p_c21
```

5. Register a node:

The Tivoli Storage Manager server views its registered clients, application clients, host servers, and source servers as nodes. To register a node, log on as the Tivoli Storage Manager administrator using the *Admin Command Line* or the *Web Admin* and run this command:

```
register node C21 passwd domain=sap_c21 maxnummp=8
```

When using two or more tape drives, the `maxnummp` parameter settings can affect the nodes. It defines the maximum number of mount points that one node can use. The default value is `1`. If one node should use more than one mount point, the parameter must be set to the desired number of mount points. This parameter should not be set higher than the total number of drives available on the Tivoli Storage Manager server.

7. Determine the Tivoli Storage Manager password method:

There are three methods of password handling:

No password required

No authentication is performed on the Tivoli Storage Manager server. Each user connected to the backup server can access Tivoli Storage Manager data without a password. This method is only recommended if adequate security measures are established. For example, no password might be acceptable when the Tivoli Storage Manager server is only used for SAP®, no other clients are registered, and authentication and authorization is performed at the operating system level.

Manual password handling

A password is required for each connection to the Tivoli Storage Manager server. In this method, Data Protection for SAP for DB2 stores the encrypted password in its configuration files. As long as the password does not expire and is not changed on the Tivoli Storage Manager server, Data Protection for SAP automatically uses the stored password when connecting to Tivoli Storage Manager. This method provides password security and can be set up easily. Whenever the password expires or is changed, the new password must be set with this command:

```
backom -c password [-x]
```

(See “Password Command (Verify and Save Tivoli Storage Manager Password)” on page 129).

If setting the password is to be automated (such as in a script), enter this information on the command line:

```
backom -e full path/initSID.utl  
-c password serverA:nodeA:passwordA serverB:nodeB:passwordB [-x]
```

where *passwordA* is the password for Tivoli Storage Manager node *nodeA* on Tivoli Storage Manager server *serverA*.

Note:

1. The interactive password prompt is omitted only if the passwords for *all* server stanzas in the `.utl` file are specified.

2. There is a potential security risk involved in recording Tivoli Storage Manager passwords in a script.

Automatic password handling

A password is required for each connection to the Tivoli Storage Manager server. After the first connection, the password is managed by Tivoli Storage Manager. The Tivoli Storage Manager client stores the current password locally. When the password expires, the password is changed and stored automatically. If you schedule your backups or restores from a system user different from the database owner, you need to grant access permissions to your data files on disk for this user. You need to specify the Tivoli Storage Manager password currently in effect before you start using Data Protection for SAP in order to connect to the server for the first time and whenever the password is changed manually on the Tivoli Storage Manager server (command update node). You do this with the command

```
backom -c password [-x]
```

(See “Password Command (Verify and Save Tivoli Storage Manager Password)” on page 129). This method is recommended for an automated production environment.

Set the Tivoli Storage Manager password:

Data Protection for SAP for DB2 should be installed after the Tivoli Storage Manager installation has been completed. Tivoli Storage Manager provides different password methods to protect data. Data Protection for SAP must use the same method as specified within Tivoli Storage Manager. The default password method during Data Protection for SAP installation is PASSWORDACCESS prompt. The default parameters for Data Protection for SAP are set according to this default value. If a different password method is set within Tivoli Storage Manager, refer to “7. Determine the Tivoli Storage Manager password method” on page 61 in order to adjust the Data Protection for SAP parameters.

Provide Data Protection for SAP for DB2 with the password for the Tivoli Storage Manager node by entering this command:

```
backom -c password
```

For more information on the password, refer to “7. Determine the Tivoli Storage Manager password method” on page 61.

Password Configuration Matrix (UNIX or Linux):

Once the preferred method of password handling is determined, review these steps for direction as to how to set the keywords and parameters in the various profiles. Detailed information regarding password handling methods is available in “7. Determine the Tivoli Storage Manager password method” on page 61.

After you have selected the suitable password handling method, follow this configuration matrix to set the keywords and parameters accordingly. Proceed as indicated by the step number.

Table 10. Password Handling for UNIX or Linux

Step	Profile/Action	Parameter	Password		
			No	Manual	Set by Tivoli Storage Manager
1	Tivoli Storage Manager admin	AUTHENTICATION EXPIRATION PERIOD (see note 1)	OFF	ON <i>n days</i> (see note 2)	ON <i>n days</i>
2	dsm.sys	PASSWORDACCESS PASSWORDDIR (see note 5) NODENAME	Unavailable	PROMPT Unavailable Unavailable.	GENERATE <i>path</i> <i>nodename</i>
3	Tivoli Storage Manager admin	UPDATE NODE (see notes 1, 6)	Unavailable	<i>password</i>	<i>password</i>
4	Data Protection for SAP for DB2 profile (initSID.utl)	For each SERVER statement specify:PASSWORDREQUIRED ADSMNODE	NO <i>nodename</i>	YES <i>nodename</i>	NO (see note 4)
6	Command line	backom -c password	Unavailable	<i>password</i> (see notes 3,7)	<i>password</i> (see notes 3,7)

Note:

1. See appropriate Tivoli Storage Manager documentation.
2. If you are using manual password generation during testing, make sure that the expiration period is set to an appropriate period of time.
3. This password must be the one that is effective on the Tivoli Storage Manager server for the node.
4. ADSMNODE must not be set when PASSWORDACCESS generate is set.
5. The users *SIDadm* and *db2SID* must have read and write permission for the path specified.
6. This step is only necessary if the password is expired (manual handling only) or needs to be changed on the Tivoli Storage Manager server.
7. A password must be entered for each server statement in the Data Protection for SAP profile.
8. (No longer applicable.)
9. When using PASSWORDACCESS GENERATE, the operations must always be performed with the same user ID provided in step 5 (setting of passwords).

Password Configuration Matrix (Windows):

Once the preferred method of password handling is determined, review these steps for direction as to how to set the keywords and parameters in the various profiles. Detailed information regarding password handling methods is available in “7. Determine the Tivoli Storage Manager password method” on page 61.

After you have selected the suitable password handling method, follow this configuration matrix to set the keywords and parameters accordingly. Proceed as indicated by the step number.

Table 11. Password Handling for Windows

Step	Profile/Action	Parameter	Password		
			No	Manual	Set by Tivoli Storage Manager
1	Tivoli Storage Manager admin	AUTHENTICATION EXPIRATION PERIOD (see note 1)	OFF	ON <i>n days</i> (see note 2)	ON <i>n days</i>
2	<i>server.opt</i>	PASSWORDACCESS PASSWORDDIR (see note 5) NODENAME	Unavailable	PROMPT Unavailable Unavailable	GENERATE <i>path</i> <i>nodename</i>
3	Tivoli Storage Manager admin	UPDATE NODE (see notes 1,6)	Unavailable.	<i>password</i>	<i>password</i>
5	Data Protection for SAP profile <i>initSID.utl</i>	For each SERVER statement specify: PASSWORDREQUIRED ADSMNODE	NO <i>nodename</i>	YES <i>nodename</i>	NO (see note 4)
6	Command line	<code>backom -c password</code>	Unavailable	<i>password</i> (see notes 3,7)	<i>password</i> (see notes 3,7)

Note:

1. See Tivoli Storage Manager documentation.
2. If you are using manual password generation during testing, make sure that the expiration period is set to an appropriate period of time.
3. For an initial setup, this password must be the same password specified when the node was registered to Tivoli Storage Manager. The password must be changed first on the Tivoli Storage Manager server and then on Data Protection for SAP.
4. ADSMNODE must not be set when PASSWORDACCESS generate is set.
5. The users *SIDadm* and *sapserviceSID* must have read and write permission for the path specified .
6. This step is only necessary if the password is expired (manual handling only) or needs to be changed on the Tivoli Storage Manager server.
7. A password must be entered for each server statement in the Data Protection for SAP profile.

8. Verify the Tivoli Storage Manager server name:

Review the Tivoli Storage Manager client options files to make sure that the server name matches the name specified in the server statement of the *initSID.utl* file. review that other parameters are set correctly. These depend on the password method selected. (See "7. Determine the Tivoli Storage Manager password method" on page 61).

On UNIX or Linux, define the Tivoli Storage Manager server in the Tivoli Storage Manager client system options file (*dsm.sys*). The server stanza specified in *dsm.sys* must match the entry in *initSID.utl*.

On Windows, you must define a client options file *servername.opt*. This file must be in the directory that contains *dsm.opt*. The value of *servername* is the server

name specified in `initSID.utl`.

Chapter 6. Protecting SAP® data with Data Protection for SAP for DB2 V6.3

Information needed to back up and restore your SAP® data is provided.

Review the information carefully before performing a backup or restore operation.

Backing up SAP® data

Instructions about how to back up your SAP® data is provided.

Perform the tasks required for your operating system.

Implementing the Strategy by Scheduling Automated Backup Runs

Scheduling (or automating) backup and archive operations helps ensure that the data is backed up regularly at a specified time. These products provide scheduled operations:

SAP® scheduler

The SAP® Computer Center Management System (CCMS) provides a scheduler for database administration and backup planning on a single database server. The scheduler can be started from the SAP GUI command line (transaction code db13) or with the SAP GUI menu function Tools -> CCMS -> DB administration -> DBA scheduling.

Scheduler (Windows) or Crontab (UNIX or Linux)

Automating backups at the database server level is available using either the Schedule Services feature (on Windows) or the crontab command (for UNIX or Linux). See “UNIX or Linux Crontab Example” on page 132 for more information.

Tivoli Storage Manager scheduler

Tivoli Storage Manager also provides a scheduler function for all of its clients. As a result, automation can be performed for multiple database servers. The Tivoli Storage Manager administrative client GUI provides a user-friendly wizard for defining schedules. Information on how to define Tivoli Storage Manager schedules can be found in the *Tivoli Storage Manager Administrator's Reference* manual.

IBM Tivoli Workload Scheduler

The IBM Tivoli Workload Scheduler provides event-driven automation, monitoring, and job control for both local and remote systems. More information can be found at <http://www.ibm.com/software/tivoli/products/scheduler/>.

Sample Backup Strategy for Daily Backup Processing

This figure illustrates the sequence of backup operations to consider for a daily backup schedule.

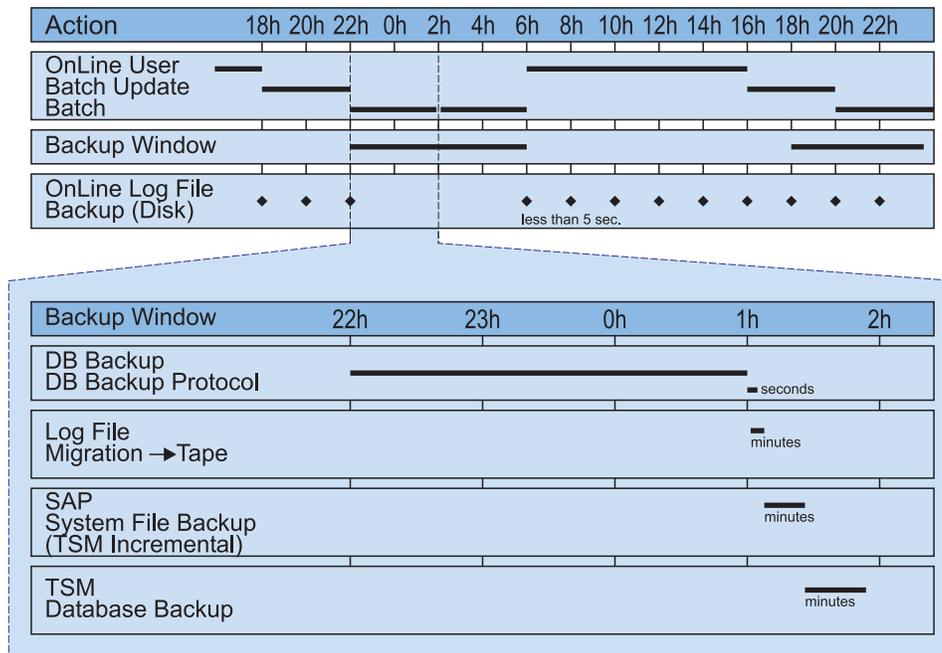


Figure 10. Production Backup Example

The automated backup example (shown in Figure 3) displays these common tasks:

- A full database backup (offline or without application load) performed each night.
- Offline log files are backed up to disk during online hours. This has the advantage of eliminating the need for extra tape mounts for relatively small files.
- The Tivoli Storage Manager server migrates archived log files from disk to tape after the full database backup.
- SAP system files are backed up incrementally with the Tivoli Storage Manager backup-archive client.
- The last backup in the daily cycle is the backup of the Tivoli Storage Manager database. This should always be performed.

Backups can be performed to disk storage as well as to tape media. The Tivoli Storage Manager server manages the data regardless of the storage media. However, backing up the SAP database directly to tape is the preferred media.

Windows Scheduling Example

On Windows systems, the schedule service must be running in order to start automated backup jobs. Issue this command to start the schedule service:

```
net start schedule
```

Use the `at` command to schedule jobs when the schedule service is running. This command launches the batch file `backup.cmd`. In this example, the command runs the schedule every Friday at 8:00 p.m.:

```
at 20:00 /every:f cmd /c c::\db2\C21\sapscripts\backup.cmd
```

Starting Backups in a Non-Partitioned Database Environment

The following examples show how you can start DB2 database/tablespace backups from the command line using DB2 CLP.

To start a DB2 backup or restore with Data Protection for SAP for DB2, log on as user `db2SID` or `SIDadm`. In these examples, the variable *shared library* represents the full path of the Data Protection for SAP shared library (UNIX and Linux) or DLL (Windows). DB2 database and tablespace backups are performed as follows:

- Full online backup (database parameter LOGRETAIN has to be activated):

```
db2 backup db dbname online load shared library
```

- Online tablespace backup (database parameter LOGRETAIN has to be activated):

```
db2 backup db dbname online tablespace (tablespace_name#1, ...)
      load shared library
```

Using DB2 Single System View for Backup

DB2 Version 9.5 (or later) provides the Single System View (SSV) function which allows backing up multiple database partitions at once. In earlier releases, partitioned databases needed to be backed up one partition at a time which can be time consuming and prone to errors. Backing up a partitioned database one partition at a time also failed to include the log files in the backup image. These log files are required to restore and recover the data. Restoring multiple partitions that were backed up individually is complicated as well, because the backup timestamp for each partition is slightly different. Because of this, identifying all database partitions belonging to the same backup is difficult, and determining the minimum recovery time for the backup that contains these partitions is difficult. Use of `db2_all` simplifies the backup of partitioned databases. However, backup and restore operations restrictions still exist that complicate these tasks. More information is available in “Backups and Restores in Partitioned Database Environments” on page 123.

With DB2 Version 9.5 (or later), when you perform a backup operation of a partitioned database, you can specify which partitions to include in the backup, or specify that all the database partitions should be included. The specified partitions are backed up simultaneously and the backup timestamp associated with all specified database partitions is the same. Also, by default, database logs are included in an SSV backup image. Finally, when you restore from an SSV backup image, you can specify to roll forward to end of logs, which is the minimum recovery time calculated by the database manager. See the *DB2 Command Reference* for additional information.

Multiple DB2 Log File Copies

Backing up multiple copies of a log file in a single archive operation helps protect against this data in the event of a storage hardware failure or disaster recovery situation. These copies can be located on different physical Tivoli Storage Manager volumes or on different Tivoli Storage Manager servers. When a log file copy is unavailable at restore time, Data Protection for SAP for DB2 automatically switches to another copy and continues restoring the log file from that copy. The description of the profile keyword `REDOLOG_COPIES` on page “Tivoli Storage Manager for ERP for DB2 profile parameter descriptions” on page 134 provides detailed

information about creating and using DB2 Log Copies.

Schedule Batch Sample

```
@echo off
rem -----
rem file name: schedule.sample
rem -----
rem Task:
rem Submits backup/archive commands at regularly scheduled intervals
rem using two simple batch files containing backup/archive commands.
rem -----
rem ***** NOTE ***** NOTE ***** NOTE *****
rem
rem This file is intended only as a model and should be
rem carefully tailored to the needs of the specific site.
rem
rem ***** NOTE ***** NOTE ***** NOTE *****
rem -----
rem For a full reference of the AT command please see the Windows NT
rem help.
rem -----
rem
rem For the following examples, the system ID of the DB2 database
rem is assumed to be "C21".
rem
rem -----
rem Full database backup, scheduled every Friday at 8:00 p.m.
rem
rem at 20:00 /every:f cmd /c c:\db2\C21\sqllib\scripts\backup.cmd
rem
rem -----
rem Save redo logs, scheduled twice a day at 11:30 a.m. and at 5:30 p.m.
rem Monday through Friday
rem
rem at 11:30 /every:m,t,w,th,f cmd /c c:\db2\C21\sqllib\scripts\archive.cmd
rem ----- end of schedule.sample -----
```

Full Offline Backup Batch File Sample

```
@echo off
rem Full Offline Backup batch file:
rem -----
rem file name: backup.cmd
rem -----
rem Sample DB2 backup batch file for 3264bit environments
rem -----
rem Task:
rem Invokes a DB2 backup in order to perform a full offline backup of
rem all DB2 tablespaces
rem -----
rem ***** NOTE ***** NOTE ***** NOTE *****
rem
rem This script is intended only as a model and should be
rem carefully tailored to the needs of the specific site.
rem
rem ***** NOTE ***** NOTE ***** NOTE *****
rem -----
rem
rem For the following examples, the system ID of the DB2 database
rem is assumed to be "C21".
rem
rem -----COMMAND-----
rem db2 backup db C21 load 'C:\Program Files\tivoli\tsm\tdp_r3\db264\tdpdb2.dll'
```

Full Offline Backup Shell Script Sample

```
#!/bin/ksh
# -----
# backup.ksh:
# Sample DB2 backup shell script for 3264bit environments
# -----
# Task:
# Invokes a DB2 backup in order to perform a full offline backup of
# all DB2 tablespaces
# -----
#          *****      NOTE          *****      NOTE          *****      NOTE          *****
#
#          This script is intended only as a model and should be
#          carefully tailored to the needs of the specific site.
#
#          *****      NOTE          *****      NOTE          *****      NOTE          *****
# -----
#
# For the following examples, the system id (alias) of the DB2 database is
# assumed to be 'C21'.
#
# -----COMMAND-----
su - db2c21 -c "db2 backup db C21
               load /usr/tivoli/tsm/tdp_r3/db264/libtdpdb264.a"
```

Segmenting large backup objects

To assist backing up and restoring of database objects that are larger than 1 TB, use the IBM Tivoli Storage Manager for Enterprise Resource Planning *SEGMENTSIZ*E keyword parameter for each DB2 backup session to be partitioned into multiple segments. These segments are stored on Tivoli Storage Manager as individual backup objects. The value of the *SEGMENTSIZ*E keyword parameter determines the maximum allowable size of a backup segment on Tivoli Storage Manager storage.

Each DB2 backup session is assigned its own backup segment group. A *backup segment group* is a collection of all segments of a backup session generated by Tivoli Storage Manager for ERP during a database backup operation. For example, two DB2 backup sessions (s1, s2) that contain two segments for each session (seg1, seg2), is assigned two backup segment groups (sg1, sg2). The first backup segment group (sg1) contains segments s1:seg1, s1:seg2. The second backup segment group (sg2) contains segments s2:seg1, s2:seg2.

When you specify segmentation, the session number substring of the backup image name is used to identify the backup object as part of a segmented data stream. The session number substring *segment number* is added to the backup image name, separated by a colon (:). For example:

DB2 instance.db alias.type.partition number.DB2 backup ID.session number:segment number

When Tivoli Storage Manager for ERP initiates a change of Tivoli Storage Manager objects, the segment number, for the new backup object segment, increases by one.

For integrity check processing of the backup segment group, an additional zero-byte backup object, the so-called commit object, is generated. This will be used by Tivoli Storage Manager for ERP to check the integrity of the related backup segment group. The naming convention of the commit object is as follows:

The character C following the colon (:) character identifies the backup image as a committed object. These committed objects are stored on Tivoli Storage Manager at the very end of each participating backup session. Also, the *last segment number* identifies the number of segments that must exist on Tivoli Storage Manager for all segments for that session to be restored. As a result, this update to the backup image name ensures that the correct object is assigned to the correct DB2 backup session. However, when one or more committed objects are missing, the integrity of the backup segment group is not guaranteed. For this reason, the database restore is not started by Tivoli Storage Manager for ERP

You can verify whether backup object segmentation was activated by using either of these methods:

Tivoli Storage Manager for ERP log entries

An information message that identifies that the maximum segment size is logged to this file. The session number substring *:segment number* is included in the backup image name, and in an information message indicating that a commit object (containing substring *Clast segment number*) was generated.

DB2 Backup Object Manager

The session number substring *:segment number* is visible in the backup image that is displayed by the `q_all -m detailed`, `q_db -m detailed` or `q_raw` command.

Administration Assistant function for Tivoli Storage Manager for ERP

The session number substring (*:segment number*) is visible in the backup image that is displayed in the operations monitoring or performance monitoring views.

Segmentation and Tivoli Storage Manager server

If segmentation is used for backup operations, the Tivoli Storage Manager server might issue the error message, "ANS0326E This node has exceeded its maximum number of mount points. This happens because there might be a short delay before the Tivoli Storage Manager actually closes client sessions. To overcome this problem the Tivoli Storage Manager server MAXNUMMPP parameter for the Tivoli Storage Manager node is set to twice the number of Tivoli Storage Manager sessions that are used for the backup. The number of active parallel sessions for Tivoli Storage Manager for ERP to the Tivoli Storage Manager server is limited by the Tivoli Storage Manager for ERP parameters SESSION, in the SERVER stanza, and MAX_SESSIONS.

For example, if two Tivoli Storage Manager sessions are needed for the database backup, the MAXNUMMPP parameter for the used node is set to four in the Tivoli Storage Manager server. In this example, Tivoli Storage Manager for ERP sends the data by using two Tivoli Storage Manager sessions only.

Segmentation and backup processing

Review the following backup characteristics before applying segmentation to your DB2 backup operations:

- The data stream sent from DB2 is segmented during a DB2 database backup.

- When implemented, segmentation is applied to every participating DB2 backup session.
- Back up and restore sessions are isolated from each other. As a result, segments that are generated by Tivoli Storage Manager for ERP are isolated on a per session basis. Therefore, segments cannot be mixed between different sessions. All segments backed up within the same session are restored in the same session.
- DB2 logs are not partitioned in to multiple segments.

Segmentation and restore processing

Review the following restore characteristics before applying segmentation to your DB2 restore operations:

- Metadata associated with the backup object indicates whether the object is part of a segmented data stream. If the backup object is part of a segmented data stream, Tivoli Storage Manager for ERP automatically joins the segments to the object DB2 expects to receive from Tivoli Storage Manager during the restore operation.
- Back up and restore sessions are isolated from each other. As a result, segments that are generated by Tivoli Storage Manager for ERP are isolated on a per session basis. Therefore, segments cannot be mixed between different sessions. All segments backed up within the same session are restored in the same session.
- Do not use segmentation into two or more segments for a backup that is to be restored to DB2 using the Backup Object Manager command **backom -c r_raw** . . . This backup can be restored to the destination directory, but not into DB2 if two or more segments were created. If the backup was created by using a single segment, it can be restored to DB2 from the destination directory after retrieval from Tivoli Storage Manager. There is no limitation concerning segmentation for other restore methods.

Restoring SAP® data

Instructions about how to restore your SAP® data is provided.

Perform the tasks required for your operating system.

Starting Restores in a Non-Partitioned Database Environment

The following examples show how you can start DB2 database/tablespace restores from the command line using DB2 CLP.

Every successful backup run generates a timestamp that is required for later restore operations. These timestamps are written to the DB2 Recovery History file (RHF), which can be queried with DB2 commands. The timestamps of backup images currently stored on the Tivoli Storage Manager server can be queried using the Backup Object Manager query commands. If no timestamp is specified in a restore command, the latest backup image found on Tivoli Storage Manager will be restored. See “Managing Backup Objects” on page 125 for details.

DB2 database and tablespace restores are performed as follows:

- Full restore to a certain point in time:

```
db2 restore db dbname load shared library taken at timestamp
```

OR

```
backom -c r_db -a dbname -t timestamp
```

- Online tablespace restore

```
db2 restore db dbname tablespace (tablespace_name#1, ...) online  
load shared library taken at timestamp
```

OR

```
backom -c r_ts -a dbname -t timestamp -0
```

- Recovery History File restore

```
db2 restore db dbname history file online load shared library
```

OR

```
backom -c r_hfile -a dbname
```

Data Protection for SAP for DB2 process results can be checked by analyzing the Data Protection for SAP log files. These log files may contain success, warning, and error messages.

Redirected Restore in Automatic Mode

Backup Object Manager provides an automatic cloning function which creates an exact copy of the original SAP® database in a different location. The physical database layout (tablespaces, tablespace number, and size of the tablespace containers) of the target database is identical to that of the source system. The path names of the new tablespace containers are constructed by replacing the original SID with the SID of the target system. In addition, modifications to the sizes of all or selected tablespace containers of the target database can be made to optimize the I/O performance. Backup Object Manager provides automated tablespace resizing and automated tablespace normalizing for these modifications as described in “Automated Tablespace Adaptations” on page 111. The Backup Object Manager automatic mode redirected restore function can be used to resize tablespace containers of the source database. This is accomplished by performing a redirected restore in automatic mode with the same SID set as both the original and the target SID and requesting scaling or normalizing (or both) during the operation.

Issue this command on the target system to perform a redirected restore in automatic mode:

```
backom -c rr_db_clone -a DB2 source alias,DB2 target alias -t timestamp
```

The complete command syntax is provided in “DB2 Backup Object Manager utility” on page 4.

Backup Object Manager performs these steps during a redirected restore in automatic mode:

1. Backup Object Manager retrieves the TDI for the requested backup from Tivoli Storage Manager into memory.
2. Backup Object Manager replaces the source database alias with the target database alias. If no target database alias is specified, Backup Object Manager uses the original database alias as the target database alias.
3. Using the modified TDI, Backup Object Manager performs basic plausibility checks as described in “Redirected Restore Plausibility Checks” on page 83.
4. Backup Object Manager uses the modified TDI to create the necessary tablespace containers on the target system. If the target database alias is the same as the original database alias, the database is restored to the original database alias and SID. When restoring to the original system, Backup Object Manager attempts to overwrite the original database. Overwriting the original database requires approval by the administrator.
5. Backup Object Manager calls the DB2 redirected restore function.

Tablespace Definition Information (TDI)

In order to automate a redirected restore as much as possible, Backup Object Manager requires information on the tablespaces and the tablespace containers used in the original database. This information is used to create the tablespace containers of the target database accordingly. This information is required for each tablespace:

- The ID and name of the tablespace.
- The type of the tablespace. For example, whether the tablespace is system (SMS) or database managed (DMS).
- The page size in bytes.
- The extent size in pages.
- The number of pages used. This number can help the administrator when resizing containers. Backup Object Manager also calculates the numbers of total pages and of usable pages from the data stored for each tablespace container. Is automatic storage used?
- Information on the tablespace containers used for the tablespace.

This information must be available for each tablespace container:

- The ID of the tablespace container.
- The name of the tablespace container. For example, whether the directory contains an SMS container or the file contains a DMS container, respectively.
- The type of the tablespace container. For example, whether a database managed container is stored in a file or on a raw device.
- For DMS tablespaces, the total number of pages stored in the container.

The TDI and the DB2 backup images are stored together on the Tivoli Storage Manager server. They are associated using the combination of the instance name of the database, the database alias, the database node number, and the timestamp of the backup. The name of the TDI is constructed in this format: *DB2 instance-<DB2 alias>-DB2 node numbertimestamp.tdi*. The TDI can be retrieved from Tivoli Storage Manager separately with the Backup Object Manager command 'r_tdi' and can be stored as an ASCII file in a specified file system. The availability of TDI in the filesystem is a prerequisite for the Backup Object Manager redirected restore in batch mode.

These changes can be done to the TDI file to prepare for a batch-mode redirected restore:

- add or remove of tablespace containers from dedicated tablespaces
- modify names (locations) of tablespace containers
- modify the size of a DMS tablespace container, whereby the sum of container sizes has to have at least the number of pages used plus $((\text{number of containers} + 1) * \text{extent})$, where extent is the extent size in pages.
- add an automatic storage path, if at least one automatic storage path is already present
- change the location of an existing automatic storage path
- remove one or more existing automatic storage paths, whereby in any case at least one automatic storage path needs to exist

tablespace

The following is a sample TDI file:

```

; IBM Tivoli Storage Manager for Enterprise Resource Planning
; Data Protection for SAP(R) for DB2
; - Tablespace Definition Information (TDI) -
;
; The following TDI sections can be modified manually:
; - Automatic_Storage_Path
; - Container
;
; An automatic storage path section consists of the following format:
;
; Automatic_Storage_Path = path#1
; ...
; Automatic_Storage_Path = path#n
;
; It is possible to add or remove an automatic storage path entry. For already existing
; automatic storage path entries the assigned path can be updated.
;
; A tablespace section consists of the following format:
;
; [Tablespace ID "tbsp. name" type page size extent size in pages
; used pages yes|no]
; Container[ID 1] = definition ;
; .
; .
; Container[ID n] = definition
;
; where the definition of a container statement is characterized by its tablespace:
; - SMS tablespace: "path"
; - DMS tablespace: file | "path/container name" | size in pages
;
; If the tablespace containers are modified manually (add or remove container,
; adjust container path or size) at least the following conditions have to be
; guaranteed for ensuring the TDI integrity:
; 1) Any new container specified requires empty brackets '[]'. The ID is calculated
; internally.
;
; 2) Each tablespace block has to have at least one container specification
;
; 3) The sum of container sizes of a DMS tablespace has to have at least the number
; of used pages plus ((number of containers + 1) * extent).
;
; !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
; ! DO NOT EDIT ANYTHING ELSE EXCEPT THE SECTIONS !
; ! - Automatic_Storage_Path (if present) !
; ! - Container !
; !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
[TDI]
Version = 1.1
Generator = Data Protection for SAP(R) 6100

[Backup]
Alias = T01
Instance = db2t01
Node = 0
Timestamp = 20081130094352
Database_Path = /db2/T01/sapdata1/db2t01/NODE0000/SQL00001/
Automatic_Storage_Path = /db2/T01/sapdata1
Automatic_Storage_Path = /db2/T01/sapdata2

[Tablespace 0 "SYSCATSPACE" dms 16384 4 9264 yes]
Container[0] = file | "/db2/T01/sapdata1/db2t01/NODE0000/T01/T0000000/C0000000.CAT" | 8192
Container[1] = file | "/db2/T01/sapdata2/db2t01/NODE0000/T01/T0000000/C0000001.CAT" | 8192

[Tablespace 1 "TEMPSPACE1" sms 16384 32 1 no]
Container[0] = "/db2/T01/saptemp1"

[Tablespace 10 "T01#USER1D" dms 16384 2 540 no]
Container[0] = file | "/db2/T01/sapdata1/NODE0000/T01#USER1D.container000" | 448
Container[1] = file | "/db2/T01/sapdata2/NODE0000/T01#USER1D.container001" | 448

[Tablespace 11 "T01#USER1I" dms 16384 2 540 no]
Container[0] = file | "/db2/T01/sapdata1/NODE0000/T01#USER1I.container000" | 448
Container[1] = file | "/db2/T01/sapdata2/NODE0000/T01#USER1I.container001" | 448

```

Be aware of the following considerations regarding the TDI file:

- The "[TDI]" header block is used to identify the data as TDI and holds some meta-information about it. The "Version" key holds the version of the TDI syntax. The "Generator" key denotes some product information.
- The [Backup] block holds various kinds of information about the database backup the TDI is associated with. This information must be kept within the TDI file so that it is available even when the file has been renamed.

[Backup] additionally includes the database path where database meta data is stored as well as all automatic storage paths the database provides for tablespaces supporting automatic storage. It is possible to add or remove an automatic storage path entry in that section. Optionally, for already existing automatic storage path entries the assigned path can be updated.

- The [Tablespace] block marks the start of the container definitions of a specific tablespace.
- The block header contains the following items in exactly this order: the ID of the tablespace, its name, its type, the page size in bytes, the extent size in pages and the number of used pages in the tablespace. Do not change any data within the tablespace block header.
- Each container statement defines one container of a tablespace according to the following rules:
 - The ID is denoted in square brackets if the line was written by the system. If a new container is to be added to a tablespace, the ID is not yet known. Therefore, the administrator specifies a new container without an ID, just entering consecutive brackets.
 - For an SMS tablespace, only the fully qualified path is specified.
 - For a DMS tablespace, the type, location and size of the container are specified, in this order, and separated by a vertical bar ("|"). The type is given by one of the strings "file" or "device". The size will be interpreted as a number of pages unless a unit is specified. In this case, the unit is used.
 - Names of tablespaces and paths must be quoted strings.

Redirected Restore in Batch Mode

Backup Object Manager provides a redirected restore batch mode function where the TDI for the target database is modified before starting the redirected restore. The TDI image to be used must be available as an ASCII file in the file system. For example, a TDI image created during an interactive redirected restore can be used as target TDI for a redirected restore in batch mode. Batch mode can also be used for multiple redirected restores to different locations with identical changes of the physical database structure. As with the interactive mode, the original TDI is used to test whether the changes of tablespace container sizes and locations made are valid. In addition, modifications to the sizes of all or selected tablespace containers of the target database can be made to optimize the I/O performance. Backup Object Manager provides automated tablespace resizing and automated tablespace normalizing for these modifications as described in "Automated Tablespace Adaptations" on page 111. A sample TDI file and modification guidelines are available in "Tablespace Definition Information (TDI)" on page 75.

Before beginning a redirected restore in batch mode, the TDI for the target database must be available. This is accomplished by providing the target TDI image of a previous interactive redirected restore as a file in the file system or by retrieving the original TDI from Tivoli Storage Manager. Issue the following command to retrieve a TDI image from Tivoli Storage Manager into the file system:

```
backom -c r_tdi -a DB2 source alias -t timestamp -d target directory of TDI
```

This original TDI image can be renamed and modified. The complete command syntax is provided in “DB2 Backup Object Manager utility” on page 4.

Issue the following command on the target system to perform a redirected restore in batch mode:

```
backom -c rr_db_batch -a DB2 source alias,DB2 target alias -t <timestamp,...  
...-f full qualified path and name of target TDI file
```

Backup Object Manager performs these steps during a redirected restore in batch mode:

1. Backup Object Manager replaces the alias specified in the target TDI file with the alias of the target database. If no target database alias is specified, Backup Object Manager uses the original database alias as the target database alias.
2. Backup Object Manager retrieves the original TDI from Tivoli Storage Manager and verifies whether the target TDI defines tablespace containers that are sufficient to replace the original tablespace containers.
3. Backup Object Manager uses the target TDI and the original TDI, to perform basic plausibility checks as described in “Redirected Restore Plausibility Checks” on page 83.
4. Backup Object Manager uses the target TDI to create the necessary tablespace containers on the target system. If the target database alias is the same as the original database alias, the database is restored to the original database alias and SID. When restoring to the original system, Backup Object Manager attempts to overwrite the original database. Overwriting the original database requires approval by the administrator.
5. Backup Object Manager calls the DB2 redirected restore function.

Redirected Restore in Interactive Mode

Backup Object Manager interactive mode is a menu-driven dialog where the table space container layout is redefined by adding, deleting, moving, or resizing items. Backup Object Manager compares the tablespace definitions entered in the menu dialog with the original database layout (as documented in the original TDI) and provides immediate feedback regarding potential configuration problems. In addition, modifications to the sizes of all or selected tablespace containers of the target database can be made to optimize the I/O performance. Backup Object Manager provides automated tablespace resizing and automated tablespace normalizing for these modifications as described in “Automated Tablespace Adaptations” on page 111.

Issue this command on the target system to perform a redirected restore in interactive mode:

```
backom -c rr_db_interactive -a DB2 source alias,DB2 target alias...  
...-t timestamp -f target TDI file
```

The complete command syntax is provided in “DB2 Backup Object Manager utility” on page 4.

Backup Object Manager performs these steps during a redirected restore in interactive mode:

1. Backup Object Manager retrieves the TDI for the requested backup from Tivoli Storage Manager into memory.
2. Backup Object Manager replaces the source database alias with the target database alias. If no target database alias is specified, Backup Object Manager uses the original database alias as the target database alias.
3. Backup Object Manager determines if specific containers need to be redefined.
4. Backup Object Manager displays the main menu which show a list of sorted tablespaces for the database to be restored. A '!' mark in front of a tablespace or tablespace container indicates a warning regarding a potential problem. Although the redirected restore can still begin, the problem should be resolved before proceeding. A '!!' character in front of a tablespace or tablespace container indicates an error was detected, such as a problem concerning their location or size. The redirected restore will not succeed until the error is first resolved.
5. The administrator can select tablespaces or tablespace containers to be changed by using their IDs. When all modifications of the physical database layout are completed and no more errors ('!!') are displayed, the redirected restore can be started by entering 'c' from the main menu. The administrator can also end the redirected restore from any menu dialog by entering 'a'.
6. When the -f option is specified during the redirected restore, the modified physical database layout of the target database is stored in an ASCII file in the file system. This file can be used afterwards as input for a redirected restore in batch mode at another location, where the same physical changes to the restored database must be applied.
7. Backup Object Manager uses the modified TDI and the original TDI to perform basic plausibility checks as described in “Redirected Restore Plausibility Checks” on page 83.
8. Backup Object Manager uses the modified TDI to create the necessary tablespace containers on the target system. If the target database alias is the same as the original database alias, the database is restored to the original database alias and SID. When restoring to the original system, Backup Object Manager attempts to overwrite the original database. Overwriting the original database requires approval by the administrator.
9. Backup Object Manager calls the DB2 redirected restore function.

Sample Work Flow for Redirected Restore

This is a sample work flow for a redirected restore with Data Protection for SAP for DB2 Backup Object Manager. In order to clone the SAP® production database (PRD) to a test system (TST) on a different machine, apply the following procedure:

1. Make sure that the administrator account to be used has the appropriate rights on the target system, such as permission to allocate files of a size greater than 2 GB.
2. Verify that the source database PRD meets the prerequisites for a redirected restore operation as identified in “Redirected Restore Prerequisites” on page 85.
3. Set up Data Protection for SAP for DB2 on the target machine. Verify that these environment variables specify these values:
 - XINT_PROFILE specifies the Data Protection for SAP profile.

- DB2_VENDOR_LIB specifies the Data Protection for SAP shared library.
 - TDP_DIR specifies the path for the Data Protection for SAP process log files.
4. For the restore process, customize the Data Protection for SAP profile (initTST.utl) on the test system with these settings:
 - Use BACKUPIDPREFIX as specified on the source system: PRD__.
 - Use the Tivoli Storage Manager server specified on the source system. This may include adding the appropriate Tivoli Storage Manager server stanza to the client system options file (dsm.sys) on the test system.
 - Use the ADMSNODE specified on the source system.
 - Use BRBACKUPMGTCLASS as specified on the source system.
 5. Issue the following command to record the password of the appropriate node on the Tivoli Storage Manager server:

```
backom -c password
```

This creates or updates the Data Protection for SAP configuration file `initTST.bki`.

6. Issue the following command to check the Data Protection for SAP database backup images on Tivoli Storage Manager:

```
backom -c q_db
```

Verify that the TDI flag is set to yes for the backup image to be restored. Information regarding how to create Tablespace Definition Information is provided in “Create Tablespace Definition Information” on page 85.

7. Issue the following command with the `-C` option to call the BackOM built-in check routine:

```
backom -c rr_db_clone -a PRD,TST -t timestamp -C
```

This checks for logical and physical integrity of the test system.

8. Issue the following command to start the redirected restore:

```
backom -c rr_db_clone -a PRD,TST -t timestamp
```

9. If the database is in rollforward pending mode and needs to be recovered, there are two possibilities for retrieving the required logs.

- automatically by the DB2 Log Manager during the recovery process, or
- manually with BackOM before the DB2 rollforward process is started.

The automatic log file retrieval requires some extra configuration parameters to enable Data Protection for SAP to find the logs on the TSM server, because the logs were archived under a different database name (the source database) but the rollforward process tries to find them based on the target database name. Therefore, two additional Data Protection for SAP configuration parameters help to find and retrieve the required logs.

The configuration parameters are:

- SRC_DB_INSTANCE
- SRC_DB_ALIAS

where SRC_DB_INSTANCE specifies the name of the DB2 instance of the source database and SRC_DB_ALIAS the name of the source database itself. These two parameters must be added to the DB2 vendor environment file, which will be used as the option (DB2 database configuration parameter LOGARCHOPT1 or LOGARCHOPT2) for the appropriate activated DB2 log archive method, for example:

```
XINT_PROFILE=/db2/TST/tdpr3/initTST.ut1
TDP_DIR=/db2/TST/tdpr3/tdplog
BACKOM_LOCATION=/usr/tivoli/tsm/tdp_r3/db264/backom
SRC_DB_INSTANCE=DB2PRD
SRC_DB_ALIAS=PRD
```

Activate the DB2 Log Manager on the test system (if not already done) in combination with Data Protection for SAP (here, log archive method 1 is used to service log requests):

```
db2 update db cfg for TST using LOGARCHMETH1 VENDOR:/<fully qualified name
of shared library>
```

Set LOGARCHOPT1 to the modified DB2 vendor environment file (vendor.env) created during the Data Protection for SAP installation:

```
db2 update db cfg for TST using LOGARCHOPT1 <fully qualified name of
DB2 vendor environment file>
```

The logs required for the database recovery can be either retrieved automatically, which required the Data Protection for SAP® parameters SRC_DB_INSTANCE and SRC_DB_ALIAS set in the DB2 vendor environment file or they can be retrieved manually with BackOM. In the latter case, the TSM server must first be checked for the logs already archived, where logs of a database will be lgrouped by their associated log chain number. Issue the following command:

```
backom -c q_log -a PRD
```

10. In order to retrieve the log files, issue:

```
backom -c r_log -a PRD -l log number range -k log chain number
-d destination directory
```

The database log directory or a different location may be specified for the destination directory.

11. Start the DB2 rollforward process. In case the log files were retrieved manually by BackOM to a location other than the database log directory, start the DB2 rollforward procedure and use the overflow log path option to specify the location of the retrieved log files.
12. After the redirected restore completes successfully and before backing up the new test system, modify the Data Protection for SAP profile `initTST.ut1` to match the values of the new test system. This modification might involve these keywords:
 - BACKUPIDPREFIX
 - SERVER
 - ADSMNODE

- BRBACKUPMGTCLASS
- BRARCHIVEMGTCLASS

If the DB2 vendor environment file was updated using the parameters SRC_DB_INSTANCE and SRC_DB_ALIAS for recovery purposes, remove those parameters from that file.

Attention: Be sure not to back up the test system with the BACKUPIDPREFIX of the production system.

13. Perform the following tasks to update the DB2 database configuration of the test system:

- Set VENDOROPT to the vendor environment file created during the Data Protection for SAP installation.

```
db2 update db cfg for TST using VENDOROPT
    fully qualified name of DB2 vendor environment file
```

- If the DB2 Log Manager is used in combination with Data Protection for SAP and is not yet configured, set the appropriate log archive method and its assigned option field in the database configuration as follows::

```
db2 update db cfg for TST using LOGARCHMETH1 VENDOR:
    fully qualified name of shared library
db2 update db cfg for TST using LOGARCHOPT1
    fully qualified name of DB2 vendor environment file
```

Redirected Restore Plausibility Checks

Regardless of the mode of the redirected restore operation (automatic, interactive, batch), Backup Object Manager performs these checks before the DB2 redirected restore operation begins:

- All paths of tablespace containers must be fully qualified.
- On Windows, all drives used for storing tablespace containers must be available.
- On UNIX or Linux, the volumes used for storing tablespace containers must be available.
- There must be sufficient space in the various locations of the tablespace containers in the target system for storing them.
- Backup Object Manager tests whether there exist other files or directories at the desired locations of the tablespace containers. A warning is issued when a directory for an SMS container already exists but is not attached to a different database. An error is issued when one of these situations is detected:
 - A directory for an SMS container already exists and is attached to a different database.
 - A file for a DMS container already exists in the target location.
- The tablespace containers must be provide sufficient storage space for the restored data.

For all modes of redirected restore, Backup Object Manager provides a test-only option that performs validation checking without actually starting a restore. This option is used to determine in advance whether a specific redirected restore will succeed. The Backup Object Manager test-only option is activated by adding the command option -C to a redirected restore command. For example, issue this command to test whether a redirected restore in batch mode will succeed with the provided target TDI file:

```
backom -c rr_db_batch -a DB2 source alias,DB2 target alias -t timestamp...
...-f full qualified path and name of target TDI file -C
```

If the test determines that the redirected restore will not succeed, check the Backup Object Manager log for error and warning messages.

DB2 Redirected Restore Using Backup Object Manager

The DB2 Backup Object Manager provides redirected restore functions such as these:

- Restore a DB2 database to a different location.
- Change the physical database layout of a restored database, including the location of tablespace containers, the number of tablespace containers, their names, and their sizes.
- Clone a database, changing both the name and the location of the database.

Backup Object Manager uses a simple set of commands to perform a redirected restore of a database and also performs some plausibility checks before actually starting the operation.

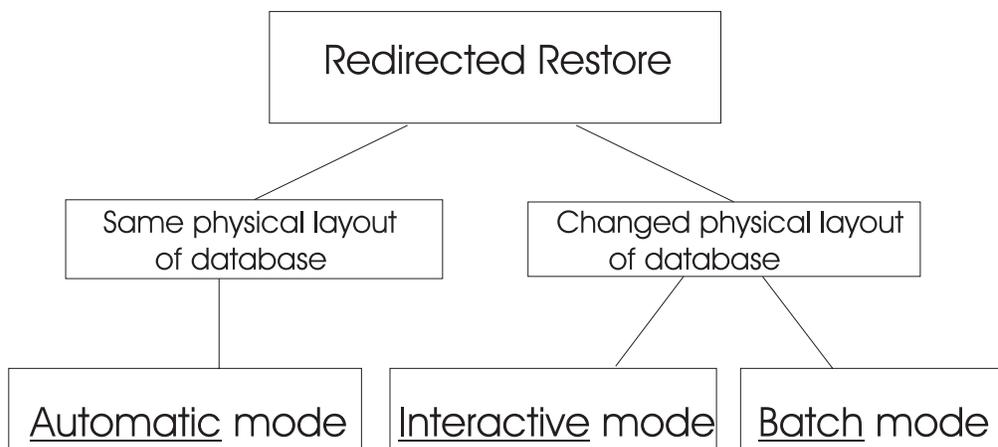


Figure 11. Redirected Restore Overview

Backup Object Manager provides these redirected restore modes:

Automatic

Restore a database to a different name and location while keeping the general database layout. However, scaling and normalizing of tablespace containers can be requested with an automatic redirected restore.

Batch Restore a database to a different location and database layout that is defined in a configuration file.

Interactive

Restore a database to a location and database layout specified by the administrator in a dialog.

In all modes, Backup Object Manager can also perform additional automated adaptations to tablespaces. For example, perform tablespace scaling to provide tablespaces with appropriate free space or perform tablespace normalizing to optimize the parallel I/O performance of the restored database.

Create Tablespace Definition Information

A TDI image is created after a full database backup completes successfully and is stored on the Tivoli Storage Manager server with the database backup image. Only database backups for which a corresponding TDI is available can be used for redirected restore with Backup Object Manager.

The Backup Object Manager must be used in order to create a TDI for an offline database backup image. For example, this command starts an offline database backup using two sessions:

```
backom -c b_db -a database alias -S 2
```

A TDI is not created when using the DB2 command line interface to perform an offline database backup.

There are two methods available to create a TDI for an online database backup image:

- One method is to use the Backup Object Manager backup function. For example, this command starts an online database backup using two sessions:

```
backom -c b_db -a database alias -S 2 -0
```

- Another method is to use the DB2 command line interface to start an online backup with the BACKOM_LOCATION parameter set in the vendor environment file. When the value of this parameter specifies the backom executable, the TDI is stored on the Tivoli Storage Manager server after the backup completes successfully. This statement must be included in the DB2 vendor environment file:

```
BACKOM_LOCATION=fully qualified path and name of the backom executable
```

The DB2 backup command can then be issued on the DB2 command line interface:

```
db2 backup db database alias online load shared library open 2 sessions
```

Use the Backup Object Manager query function to verify whether a TDI image is available for a Data Protection for SAP for DB2 backup image.

Redirected Restore Prerequisites

These requirements must be met in order for the Backup Object Manager to successfully perform a redirected restore:

- Only a backup of type FULL can be used for a redirected restore.
- A TDI image must be available for the backup to be restored.
- The database must not have a tablespace container that is a raw device.
- DMS tablespace containers of the original system are available in these locations:
 - UNIX or Linux: /db2/SAPSID/sapdatan
 - Windows: drive:\db2\SAPSID\sapdatan (*n* is an integer)
- SAPSID must be the database alias (SAPSID) and must consist of all upper case characters or digits
- SMS tablespace containers of the original system are available in these locations:

- UNIX or Linux: */db2/SAPSID/...*
- Windows: *drive:\db2\SAPSID\...*

Tablespace Definition Information

In order to perform a redirected restore, Backup Object Manager requires information about the physical layout of the original database, such as the tablespace containers used by the original database. In particular, the size of database managed containers (DMS) must be available in order to create new tablespace containers with sufficient space. Backup Object Manager keeps information on all tablespaces of a database backed up for every backup image on Tivoli Storage Manager. The following information is collected for each tablespace:

- The ID and name of the tablespace.
- Whether the tablespace type is system (SMS) or database managed (DMS).
- Whether the tablespace is managed by automatic storage.
- The page size in bytes.
- The extent size in pages.
- The number of pages used.
- The tablespace containers used for the tablespace:
 - The ID of the tablespace container.
 - The name of the tablespace container (the directory containing an SMS container or the file containing a DMS container, respectively).
 - Whether the tablespace container type is a database managed container stored in a file or on a raw device.
 - For DMS tablespaces, the total number of pages stored in the container.

This information about the physical database layout is referred to as the Tablespace Definition Information (TDI) and is stored along with the production data. The TDI is required for Backup Object Manager redirected restore operations. A TDI image is identified with its corresponding DB2 backup by the combination of DB2 instance name, database alias, database node number, and the timestamp of the backup as shown here:

```
DB2 instance-DB2 alias-DB2 node number-timestamp.tdi
```

The TDI is stored in ASCII format to allow for read and edit usability. For example, the number of used pages recorded in the TDI image can help identify the correct sizes to request when resizing containers as described in “Automated Tablespace Adaptations” on page 111. Backup Object Manager also calculates the number of total pages and used pages from the data stored for each tablespace container. Editing may be necessary when requesting a redirected restore in batch mode, as described in “Redirected Restore in Batch Mode” on page 78. A sample TDI file is provided in “Tablespace Definition Information (TDI)” on page 75.

Protecting SAP® data with the Administration Assistant function for Data Protection for SAP

Instructions about how to protect your SAP® data with the Administration Assistant function for Data Protection for SAP is provided.

Perform these tasks in order to protect your SAP® data with the Administration Assistant function for Data Protection for SAP.

Administering User IDs

The Administer users function allows accounts to be created or deleted and user permissions to be granted or revoked. Note that profiles for authorized users need to be created when the Administration Assistant is started for the first time. The online help provides details on creating profiles. For each SID in the system landscape, the following permissions can be granted:

- **Simulate backup/restores:** to initiate simulations
- **Configure groups:** to configure display groups to be used with function "Monitor backup states"
- **Problem support:** to send support request mail
- **Operations monitoring:** to view backup status information
- **User administration:** to manage user accounts
- **Performance monitoring:** to view performance data
- **Configuration:** to modify the configuration of Data Protection for SAP for DB2

Additionally, a user can be granted permission to configure parts of the internal logic of the Monitor backup states function.

Specifying a new Administration Assistant function for Tivoli Storage Manager for ERP

If the Administration Assistant function for Tivoli Storage Manager for ERP has not been installed, you can establish a connection when needed by following these instructions.

If you need to specify a new Administration Assistant function for Tivoli Storage Manager for ERP Server component, perform the following steps on the SAP® database server:

(UNIX or Linux)

1. On a Linux system, find the entry for daemon ProLE in `/etc/inittab`. Modify the entry to read as follows:

```
.../prole -p tpr3db264 Server component hostname port
```

where `<Server component hostname>` is the name or IP address of the host running the Administration Assistant Server component and `port` is the port the Server component is listening to for connects from Tivoli Storage Manager for ERP for DB2 (default 5126). If `upstart` is configured, add the `<Server component hostname>` and `port` to the init script `/etc/init/prole_db2.conf`.

2. Make sure that Tivoli Storage Manager for ERP is not running, and use the `kill` command to stop the ProLE daemon. The ProLE daemon is restarted automatically with the new parameters.

(Windows)

1. Log in as a user with administrator authority.
2. Enter this command from a command prompt:

```
prole -update -p tdpr3db264 Server component hostname port
```

where *<Server component hostname>* is the name or IP address of the host running the Administration Assistant Server component and *port* is the port that the Server component is listening to for connects from Data Protection for SAP (default 5126).

Generating Reports Using Report Templates

Once report templates are available, the Administration Assistant function for Data Protection for SAP reports can be started automatically at given points in time using a preferred scheduler. The scheduler must call the scheduler interface Sched_Main which can be started from a scheduling client as described in “4. Configuring a scheduling client to create reports” on page 49.

The scheduling interface is called by using this command syntax:

```
java -cp $CLASSPATH com.ibm.bkit.schedulerIF.Sched_Main Server component hostname...  
... RMI registry port template name userid password...  
... directory=local directory log=log path
```

- *Server component hostname*: The name or IP address of the host running the Administration Assistant Server component.
- *RMI registry port*: The number of the RMI registry port of the Administration Assistant Server component as defined in its configuration file (*assist.cfg*). The default value is 1099.
- *template name*: The name of the appropriate report template to be used. It must be available in the user template path in the Administration Assistant Server component.
- *userid*: The Administration Assistant account of the template owner.
- *password*: The password associated with *userid*.
- *local directory*: The local path in the system of the scheduling client where the requested reports are to be stored. If the local directory is not specified, the reports are not stored in the local file system. In order to access the report, the administrator needs file system access to the Administration Assistant server where the report is kept for 24 hours.
- *log path*: The local path in the system of the scheduling client where the scheduling client saves its own log files.

Consider creating a command file that sets the correct environment and schedules one (or more reports) on the scheduling client system as described in “4. Configuring a scheduling client to create reports” on page 49. If a large number of clients try to connect to the Administration Assistant server simultaneously, some of them may not immediately connect. In this case, the scheduling client waits for a random time between 15 and 45 seconds before another attempt is made. After the second unsuccessful attempt, the scheduling client creates an error log and exits.

Requesting a Report from the Administration Assistant function for Data Protection for SAP Client

A report is requested by selecting the Create Report button on the Monitor Backup States, Backup State - Detailed View, View Performance Data (History Mode), and Available Simulation Results panels of the Administration Assistant function for Data Protection for SAP graphical user interface. Reports requested from the Backup State - Detailed View, View Performance Data (History Mode), and Available Simulation Results panels always pertain to the single SID currently displayed on the panel. Reports requested from the Monitor Backup States panel contain information on all SIDs displayed on the panel. Selections made in the table of systems do not have an impact on the report created. However, active filters or the activation of a display group is reflected in the report. A time interval can be specified in the report. Backup operations are included in the report if they completed within the specified time interval. Also, some reports can include information about log files.

Starting and Stopping the Administration Assistant function for Data Protection for SAP Manually

You can manually start or stop the Administration Assistant function for Data Protection for SAP by using these command files (located in the installation directory):

- Issue this command to start or stop the Administration Assistant Server component:

(UNIX and Linux):

```
sadma.sh start|stop Server component configuration file
```

(Windows):

```
sadma.cmd start|stop Server component configuration file
```

- Issue this command to start or stop the Administration Assistant Database Agent:

(UNIX and Linux):

```
sdba.sh start|stop Database Agent configuration file
```

(Windows):

```
sdba.cmd start|stop Database Agent configuration file
```

- When using the bundled Apache Derby, issue this command to start or stop the Administration Assistant Database component:

(UNIX and Linux):

```
sdb.sh start|stop
```

(Windows):

```
sdb.cmd start|stop
```

- When using the IBM DB2 data server, use DB2 built-in utilities or commands to start or stop the database. Refer to your IBM DB2 data server documentation for complete instructions.

Important: When the Server or Database Agent components are started, a lock file (.lockAA and .lockDBA, respectively) is created. If either of these components are terminated or restarted using the delivered scripts, the respective lock file is also deleted. If for some reason the lock file still exists when the component is started, the request will fail with an error message. In this case, first verify that the process is not already active. If it is not active, the lock file must be deleted manually and the start request reissued.

Changing the Password for the Administration Assistant function for Data Protection for SAP Database User ID

The password for accessing the internal Administration Assistant function for Data Protection for SAP database can be changed using the changeSettings.jar program. This program was added to the installation directory in the `utils` subdirectory:

1. Change to the `utils` directory and issue the command

```
java -cp changeSettings.jar run
```

2. Select the type of database you are using with the Administration Assistant function for Data Protection for SAP (Apache Derby or IBM DB2).
3. Enter the directory containing the encrypted password file (pass.enc).
4. Enter the user ID and the existing password.
5. Enter the new password.
6. For Apache Derby only: To apply the new password to the database, check the box provided. Otherwise, the password file is updated but the database change must then be performed manually.
7. Click Next to complete the change.

Chapter 7. Performance tuning for Data Protection for SAP for DB2

Information needed to fine-tune Data Protection for SAP for DB2 performance is provided.

Overview of a balanced system

Descriptions on how to proceed when tuning your system according to your needs is discussed. This is done by employing a combination of functions provided in the Administration Assistant function for Data Protection for SAP.

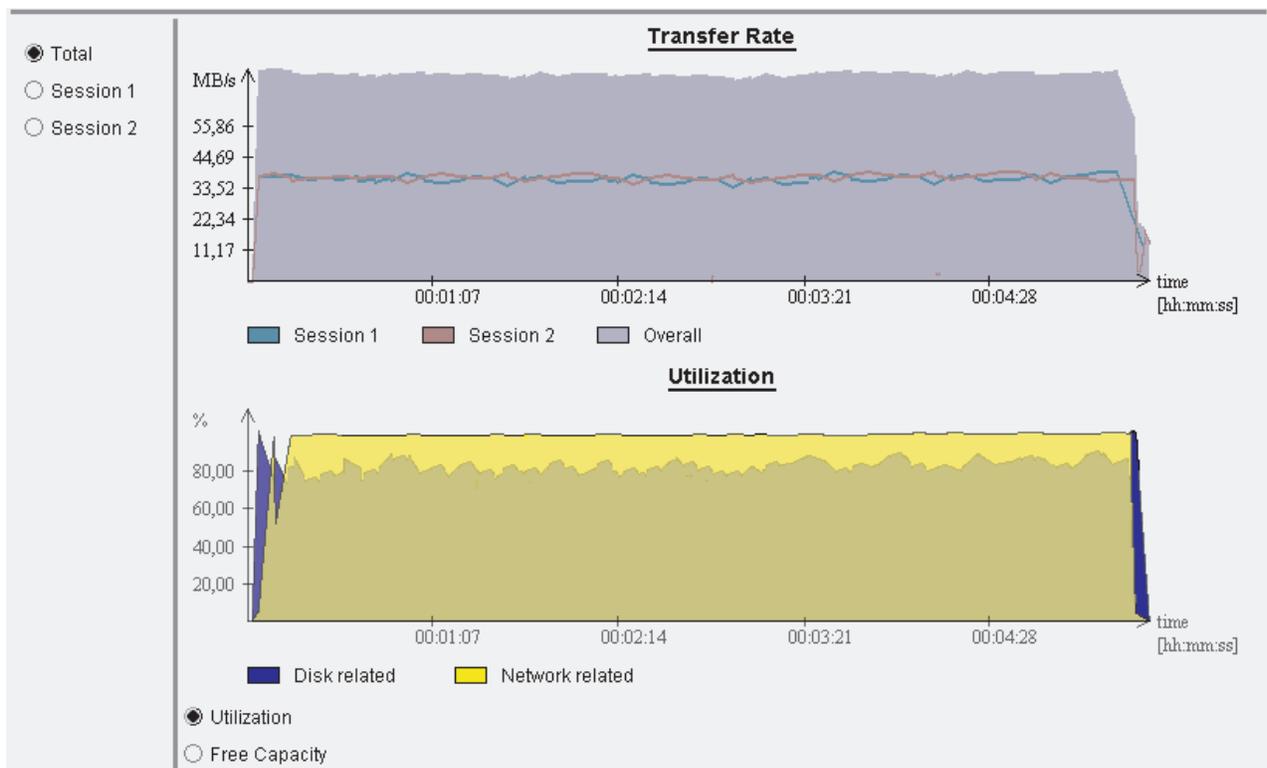


Figure 12. Indicating a Balanced Configuration

A system is considered balanced when the threads on both the disk and the network side are similarly busy throughout the backup and resource utilization is good. In an optimum setup, tapes are maintained in streaming mode. This means that the network is at least as fast as the tape and there is no idle time on the network side. Thus, a slight network bottleneck is desired. Under certain conditions, the degree of imbalance cannot be determined from the graphical presentation. Depending on your system characteristics (system buffering, buffer sizes, etc.), utilization may reduce to almost zero in the graphical presentation although the system is actually balanced. In this case, slight modifications can yield a change of bottleneck without significant throughput changes. However, whether the system is disk or network, tape constraints are always shown correctly. To improve overall throughput, consider adding more resources to create

a balanced system. A balanced system, however, does not necessarily mean that the data throughput cannot be improved further. Adding new resources can still improve the throughput rate.

Example of a disk bottleneck

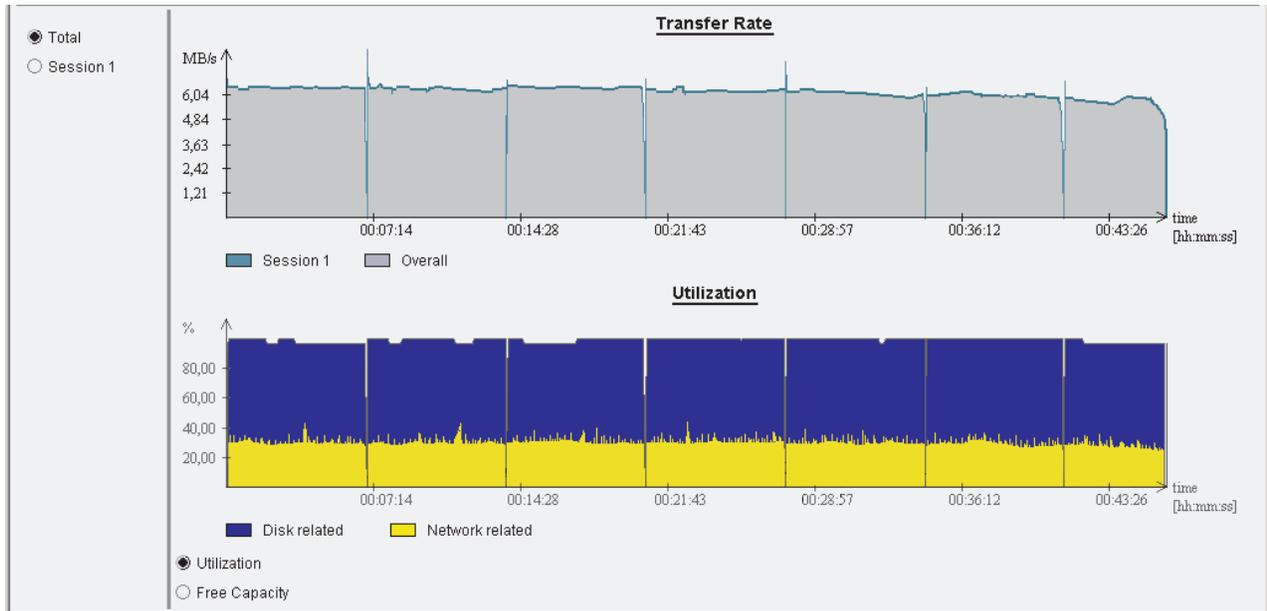


Figure 13. Indicating a Disk Bottleneck

A disk bottleneck occurs when data is processed by the network and Tivoli Storage Manager server faster than the data can be read from disk. As a result, overall throughput is limited by the disk I/O rate and the network thread is idle.

Although internal buffering causes network threads to return very quickly, the network utilization might be reduced to almost zero in this situation. Both the network and the storage media are not used to their capacity. When tapes are used, they are not kept in streaming mode when this type of bottleneck occurs. Overall throughput can be improved by increasing multiplexing (which accelerates disk reading) or making sure data compression is not used. By reducing the number of sessions to the Tivoli Storage Manager server and the number of tapes used for the backup while also increasing multiplexing at the same time, resources (such as tape drives) are used more efficiently.

Example of a network or Tivoli Storage Manager bottleneck

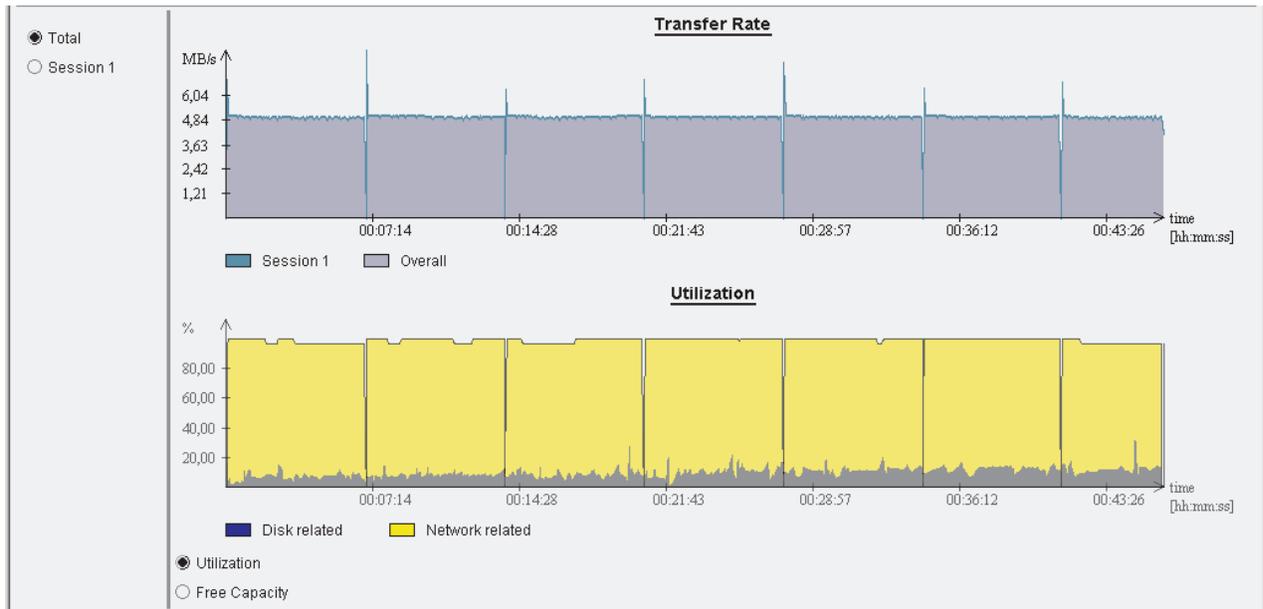


Figure 14. Indicating a Network or Tivoli Storage Manager Bottleneck

A network or Tivoli Storage Manager bottleneck occurs when data is read from the disk faster than the network or Tivoli Storage Manager can process the data. Consequently, throughput is limited either by the network capacity or by the disk or tape storage media rate. In depth analysis is usually required in order to identify the exact cause of the bottleneck. However, some insight is obtained from the Data Protection for SAP for DB2 performance analysis. Overall throughput might be improved by implementing any of these guidelines:

- If the tape is the bottleneck, increase the number of sessions to the Tivoli Storage Manager server.
- Use multiple paths to the Tivoli Storage Manager server or use multiple Tivoli Storage Manager servers.
- Use RL compression in order to reduce the amount of data to be sent to storage.

Also, to better exploit the resources, consider reducing multiplexing so that less data is read simultaneously from the disk. If the database is configured for file-online backup, reducing multiplexing will also reduce the number of redo logs created during the backup.

Viewing performance data



Figure 15. Showing Data Throughput and I/O Utilization

The Administration Assistant View Performance Data function provides a graphical representation of the data throughput rate at any point in time during the backup. Aligned with this representation, the utilization rates of the disk (presented in blue by the Administration Assistant) and network threads (presented in yellow by the Administration Assistant) are displayed. Optionally, the free capacity of these threads can also be displayed. These rates displayed can be displayed for all Tivoli Storage Manager sessions used in the backup or display rate on a per-session basis only. Time intervals that require further analysis are selected for viewing in replay mode as described in “Drilling Down on Special Situations.” Data Protection for SAP for DB2 performance sensor results are displayed using the Administration Assistant View Performance Data function. The Administration Assistant collects history data during each backup run for later analysis. In order to find the results, select View Performance Data, then select History Data. In the list of eligible backups, select the backup to be analyzed. Press the Review button to view the performance data summary panel.

Drilling Down on Special Situations

When looking at the diagrams in the View Performance Data function, you might find points in time when throughput or the utilization of a resource decreases significantly. To better understand what happened, you may drill down on these time intervals. In most cases you will find that a session is ending or a shorter file was multiplexed with longer files.

Using reports

After a backup completes, Data Protection for SAP for DB2 creates a report that contains statistical information such as the number of bytes transferred and the effective data throughput. The Administration Assistant program also provides detailed performance information that assists when optimizing your system. Reports can be provided in XML- or HTML-format for display and printing. Complete report information is available in “Reporting on Data Protection for SAP for DB2 Activities” on page 101.

Performance Analysis

The Administration Assistant provides performance data for all components involved in the data transfer. Graphical representations are provided that help identify problem areas and resource use. Performance optimization is discussed in detail in “Overview of a balanced system” on page 91.

Tracing

Trace information can be recorded in a file to help analyze problems that occur. However, contact your Data Protection for SAP support before attempting to use this function.

Monitoring the Backup Status

Backup status of multiple SAP® database servers is available by using the Administration Assistant. See “Reporting on Backup Status” on page 97 for complete details.

Reporting on the Performance of Backup Operations

The performance data of a single backup are included in the Performance Report. Although data is presented in the same manner as in the View Performance Data (History Mode) panel, the transfer rate and the utilization of adapters for each session are also displayed. The report is requested from the View Performance Data (History Mode) panel.

Performance-Report

TST (gladiator.boeblingen.de.ibm.com)

System Status: success

Type of run: full , data

Start Date	Start Time	Backup Type	Status	Throughput	End Date	End Time
18.11.2005	22:00:33	full	Success	113.11GB/h	18.11.2005	22:51:02

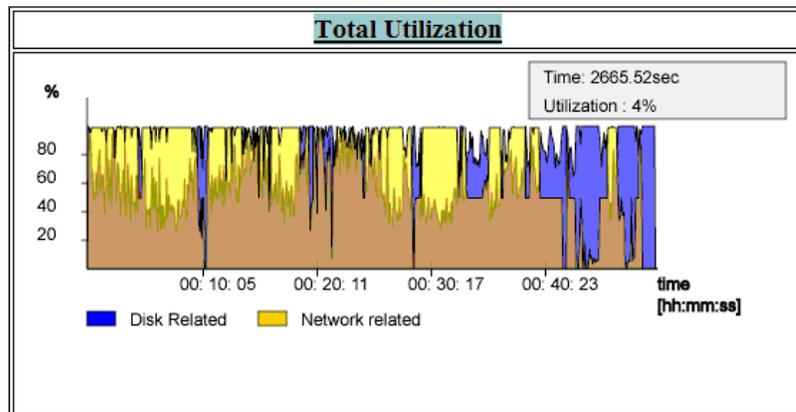
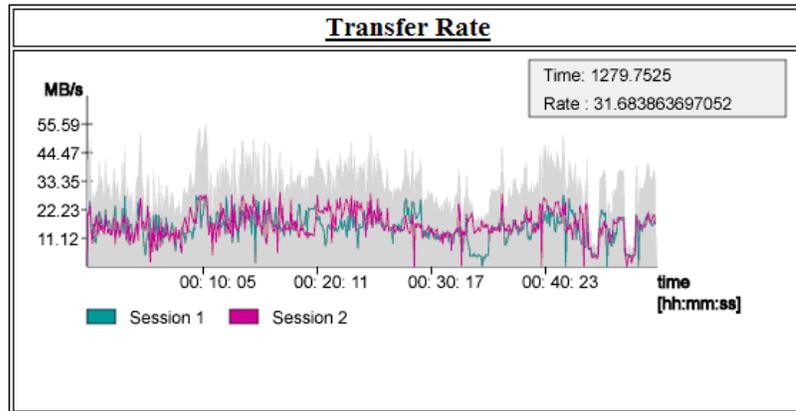


Figure 16. Performance Report - Graphical Presentation Section

Performance-Report

TST (gladiator.boeblingen.de.ibm.com)

System Status: success

Type of run: full , data

Start Date	Start Time	Backup Type	Status	Throughput	End Date	End Time
18.11.2005	22:00:33	full	Success	113.11GB/h	18.11.2005	22:51:02

[...]

Start Time	Filename	Session #	Orig. Filesize	Compr. Rate	Data Rate	Finished At
22:00:22	/oracle/TST/sapdata5/tst_5/tst.data5	2	10485768192 bytes	1.474	55.277GB/h	22:10:58
22:10:58	/oracle/TST/sapdata2/tst_2/tst.data2	2	10485768192 bytes	1.354	60.199GB/h	22:20:42
22:20:42	/oracle/TST/sapdata1/tst_1/tst.data1	2	10485768192 bytes	1.383	68.132GB/h	22:29:18

[...]

Messages

infos: 0 warnings: 0 errors: 0 undefined: 0

Type	Message
------	---------

Created: 24.11.2005 11:03:06

end of report

Figure 17. Performance Report - Tabular Presentation Section

Reporting on Backup Status

The Administration Assistant function for Data Protection for SAP provides information on the backup status of the monitored SAP® database servers. Administrators access this information by using the Monitor Operations, Monitor Backup Status function. Reports containing status information in tabular form are requested from this panel. The overview information provided in the Monitor Backup States panel is provided in the Status Report.

Status Report

System Status	System ID	Hostname	Conn.Status	DB Type	Date of Backup	Time of Backup	Backup Status	GMT Off.	Group
Success	LUS	lucius.boeblingen.de.ibm.com	offline	oracle	2005.11.23	14:18:16	Success	1	
Faehre	TST(0)	radon.boeblingen.de.ibm.com	offline	db2	2005.11.23	18:10:09	Success	1	

Created: 24.11.2005 11:40:10end of report

Figure 18. Status Report

Creating a Report

Consider this information when planning to create a report:

- Reports are requested from the Administration Assistant function for Data Protection for SAP client using the graphical user interface panels that contain the information to be included.
- Reports can also be generated from a scheduling client using a command line interface without any user interaction.
- Each report is produced as an XML file, an HTML file, with possibly one (or more) graphic files in SVG format. The HTML and the SVG files are displayed in the browser.
- All files created can be printed or saved to the local file system using the browser functionality. All reports are temporarily stored for 24 hours on the Administration Assistant server in these subdirectories:

```
Administration Assistant install dir/reports/  
report_type_time_stamp_userid/
```

File system access to the Administration Assistant server is required in order to access reports stored in the report cache.

Reporting on Failed Actions

Information on failed backup operations is provided in the Operations – Failure Report which is accessible from the Monitor Backup States panel. Administrators can choose to include information on failed backups of log files in this report.

Operations-Failure Report														
Reported failures between : 2005.11.22 11:41:10 and 2005.11.24 11:40:10														
System ID	Hostname	Conn. status	DB Type	Start Date	Start Time	BackupID	Size	Backup Type	Mode	End Date	End Time	Data RC	Control File RC	Catalog File RC
TCT	julius.boeblingen.de.ibm.com	offline	oracle	2005.11.24	06:37:43	A0EGEM5M5X	10987811	full	offline	2005.11.24	06:37:46	2	N/A	N/A
TST(0)	admiral.boeblingen.de.ibm.com	offline	db2	2005.11.23	18:10:09	A0EGDVHN65	110592	restore	restore	2005.11.23	18:10:28	2	N/A	N/A

Created: 24.11.2005 11:41:13

end of report

Figure 19. Operations - Failure Report

Modifying Report Output

Consider this information when modifying a report:

- All report requests result in the information being written to an XML file. Style sheets (which can be customized) reside with the Administration Assistant function for Data Protection for SAP Server component and are used to generate the information to different types of reports in HTML or SVG format. They determine the appearance and contents of a specific report.
- To generate a report, at least one report-specific style sheet is necessary for the transformation from XML to HTML. If a report contains graphics, each graphic is transformed to an SVG file which requires a separate style sheet. In this scenario, a single report needs a set of style sheets.

- The Administration Assistant provides two types of style sheet file sets. One set is contained in file *Admt.jar* and is used as the default. The second set of style sheets resides on the Administration Assistant server in the *Admin. Assistant install dir/styles/* directory.
- Style sheet names must be of the format *report_name_<file format>.xsl* where *file format* denotes the file type (HTML or SVG) and *report name* denotes the name of the file to be created. For example, *Picture1_svg.xsl* will generate a file named *Picture1.svg*. Note that the name of the HTML file must always be 'report'.
- The styles directory currently contains four subdirectories (Overview, Detailed, History, Simulation) that specify reports based on different XML data sources as provided in the corresponding Administration Assistant panels Monitor Backup States, Backup State – Detailed View, and View Performance Data (History Mode). The names of these folders are displayed in the list of selectable report types within the Create Report dialogs.
- For every report type, an additional file *config.xml* exists in the styles subdirectory. This file specifies default settings of the Create Reports dialogs. For example, the Operations – Daily Report has a reporting interval of 24 hours. Therefore the end of the time frame does not need to be specified, and the corresponding button will be hidden.
- All style sheets contained either in file *Admt.jar* or in the styles directory are displayed for selection in the Create Report dialogs of the Administration Assistant. Style sheets contained in *Admt.jar* are marked by the addition '(built-in)'.

Reporting on Operations Details

Detailed information on the latest backup operations for a single SID can be obtained with the Operations - Detailed Report requested from the Backup State - Detailed View panel of the Administration Assistant function for Data Protection for SAP. This panel is reached by selecting a single SID in the Monitor Backup States panel.

Operations - Detailed Report

All jobs between: 2005.11.23 12:00:00 and 2005.11.24 12:05:59

Expand All Collapse All

LUS (lucius.boeblingen.de.ibm.com)

System Status: success

Expand	Number	Start Date	Start Time	Backup ID	Size	Backup Type	Mode	Status	Throughput	End Date	End Time																																		
[-]	<u>1</u>	23.11.2005	14:06:04	LUS__A0EGDOVX4I	155.75 MB	full	online	Success	3.64GB/h	23.11.2005	14:08:41																																		
<p>Number: 1</p> <p>BackupID: LUS__A0EGDOVX4I</p> <p>Type: full</p> <table style="width: 100%; border: none;"> <tr> <td style="width: 50%; vertical-align: top;"> <p>Backup of Data Files</p> <table border="1" style="width: 100%;"> <tr><td>Run ID</td><td>LUS__A0EGDOVX4I</td></tr> <tr><td>Start of Data File Run</td><td>2005:11:23 14:06:04</td></tr> <tr><td>Duration</td><td>00:02:37</td></tr> <tr><td>Total Data</td><td>155.75 MB</td></tr> <tr><td>Throughput</td><td>3.64 GB/h</td></tr> <tr><td>Avg. Comp. Factor</td><td>1.000</td></tr> <tr><td>ReturnCode</td><td>0</td></tr> <tr><td>Sessions</td><td>1</td></tr> </table> </td> <td style="width: 50%; vertical-align: top;"> <p>Backup of Control Files</p> <table border="1" style="width: 100%;"> <tr><td>Run ID</td><td>LUS__A0EGDOZHB4</td></tr> <tr><td>Start of Data File Run</td><td>2005:11:23 14:08:48</td></tr> <tr><td>Duration</td><td>00:00:01</td></tr> <tr><td>Total Data</td><td>0.03 MB</td></tr> <tr><td>Throughput</td><td>0.09 GB/h</td></tr> <tr><td>Avg. Comp. Factor</td><td>0.965</td></tr> <tr><td>ReturnCode</td><td>0</td></tr> <tr><td>Sessions</td><td>1</td></tr> </table> </td> </tr> </table>												<p>Backup of Data Files</p> <table border="1" style="width: 100%;"> <tr><td>Run ID</td><td>LUS__A0EGDOVX4I</td></tr> <tr><td>Start of Data File Run</td><td>2005:11:23 14:06:04</td></tr> <tr><td>Duration</td><td>00:02:37</td></tr> <tr><td>Total Data</td><td>155.75 MB</td></tr> <tr><td>Throughput</td><td>3.64 GB/h</td></tr> <tr><td>Avg. Comp. Factor</td><td>1.000</td></tr> <tr><td>ReturnCode</td><td>0</td></tr> <tr><td>Sessions</td><td>1</td></tr> </table>	Run ID	LUS__A0EGDOVX4I	Start of Data File Run	2005:11:23 14:06:04	Duration	00:02:37	Total Data	155.75 MB	Throughput	3.64 GB/h	Avg. Comp. Factor	1.000	ReturnCode	0	Sessions	1	<p>Backup of Control Files</p> <table border="1" style="width: 100%;"> <tr><td>Run ID</td><td>LUS__A0EGDOZHB4</td></tr> <tr><td>Start of Data File Run</td><td>2005:11:23 14:08:48</td></tr> <tr><td>Duration</td><td>00:00:01</td></tr> <tr><td>Total Data</td><td>0.03 MB</td></tr> <tr><td>Throughput</td><td>0.09 GB/h</td></tr> <tr><td>Avg. Comp. Factor</td><td>0.965</td></tr> <tr><td>ReturnCode</td><td>0</td></tr> <tr><td>Sessions</td><td>1</td></tr> </table>	Run ID	LUS__A0EGDOZHB4	Start of Data File Run	2005:11:23 14:08:48	Duration	00:00:01	Total Data	0.03 MB	Throughput	0.09 GB/h	Avg. Comp. Factor	0.965	ReturnCode	0	Sessions	1
<p>Backup of Data Files</p> <table border="1" style="width: 100%;"> <tr><td>Run ID</td><td>LUS__A0EGDOVX4I</td></tr> <tr><td>Start of Data File Run</td><td>2005:11:23 14:06:04</td></tr> <tr><td>Duration</td><td>00:02:37</td></tr> <tr><td>Total Data</td><td>155.75 MB</td></tr> <tr><td>Throughput</td><td>3.64 GB/h</td></tr> <tr><td>Avg. Comp. Factor</td><td>1.000</td></tr> <tr><td>ReturnCode</td><td>0</td></tr> <tr><td>Sessions</td><td>1</td></tr> </table>	Run ID	LUS__A0EGDOVX4I	Start of Data File Run	2005:11:23 14:06:04	Duration	00:02:37	Total Data	155.75 MB	Throughput	3.64 GB/h	Avg. Comp. Factor	1.000	ReturnCode	0	Sessions	1	<p>Backup of Control Files</p> <table border="1" style="width: 100%;"> <tr><td>Run ID</td><td>LUS__A0EGDOZHB4</td></tr> <tr><td>Start of Data File Run</td><td>2005:11:23 14:08:48</td></tr> <tr><td>Duration</td><td>00:00:01</td></tr> <tr><td>Total Data</td><td>0.03 MB</td></tr> <tr><td>Throughput</td><td>0.09 GB/h</td></tr> <tr><td>Avg. Comp. Factor</td><td>0.965</td></tr> <tr><td>ReturnCode</td><td>0</td></tr> <tr><td>Sessions</td><td>1</td></tr> </table>	Run ID	LUS__A0EGDOZHB4	Start of Data File Run	2005:11:23 14:08:48	Duration	00:00:01	Total Data	0.03 MB	Throughput	0.09 GB/h	Avg. Comp. Factor	0.965	ReturnCode	0	Sessions	1												
Run ID	LUS__A0EGDOVX4I																																												
Start of Data File Run	2005:11:23 14:06:04																																												
Duration	00:02:37																																												
Total Data	155.75 MB																																												
Throughput	3.64 GB/h																																												
Avg. Comp. Factor	1.000																																												
ReturnCode	0																																												
Sessions	1																																												
Run ID	LUS__A0EGDOZHB4																																												
Start of Data File Run	2005:11:23 14:08:48																																												
Duration	00:00:01																																												
Total Data	0.03 MB																																												
Throughput	0.09 GB/h																																												
Avg. Comp. Factor	0.965																																												
ReturnCode	0																																												
Sessions	1																																												
[+]	<u>2</u>	23.11.2005	14:18:16	LUS__A0EGDPBKIP	299.8 MB	full	offline	Success	11.33GB/h	23.11.2005	14:19:49																																		

Created: 24.11.2005 end of report

Figure 20. Operations – Detailed Report

Reporting on Backup Operation Trends

This report type contains general information about the backups of a single SID. Data is represented in graphical and tabular form. A daily report produces a graphic that displays the amount of data saved for a single day. A monthly report produces a graphic that displays the backup duration, amount of data saved, throughput, and log file data for a specified time interval. These reports are requested from the Administration Assistant function for Data Protection for SAP Backup State – Detailed View panel which is accessible by selecting a single SID in the Monitor Backup States panel.

Operations - Daily Report

Reporting Period : 2005.11.23 12:00:00 and 2005.11.24 12:00:00

LUS (lucius.boeblingen.de.ibm.com)

[...]

Start Date	Start Time	Backup ID	Size	Backup Type	Mode	Status	Througput	End Date	End Time
23.11.2005	14:06:04	A0EGDOVX4I	163315712	full	online	Success	3.64GB/h	23.11.2005	14:08:41
23.11.2005	14:18:16	A0EGDPBKIP	314361856	full	offline	Success	11.33GB/h	23.11.2005	14:19:49

Totally saved data volume	455.55MB
Total Number of data backups	2
% Failed	00.0 %
Total Number of log backups	0
% Failed	00.0 %
Total Number of restores	0

Configuration History for Backups:

Date	Sessions	Compression	Mux	TSM Server	Mgmt Class
23.11.2005	1	On	1	MIRACULIX	MDBDISK1
23.11.2005	2	On	1	MIRACULIX	MDBDISK1

Created: 24.11.2005 11:42:53

end of report

Figure 21. Operations Daily Report

Reporting on Data Protection for SAP for DB2 Activities

The Administration Assistant function for Data Protection for SAP obtains, monitors, and administers backup configuration and performance information performed with Data Protection for SAP for DB2 and the corresponding backup status of SAP® database servers. The Administration Assistant Server and Database Agent components collect status, performance, and configuration data from several SAP database servers and retains it for a limited time. Reports can be created in XML or HTML format (or printed) by the Administration Assistant. This is useful when there is no access to the Server component.

Types of Reports

Administration Assistant reports contain the same information that is displayed by the Administration Assistant Monitor Operations and View Performance Data functions. All information is provided in XML format. In addition, the Administration Assistant provides style sheets used when generating these reports in HTML format:

- Status Report
- Operations - Detailed Report
- Operations - Daily Report
- Operations - Monthly Report
- Operations - Failure Report
- Performance Report

All built-in reports are created in English.

Working with Report Templates

A template must be created before a report can be generated without user interaction (for example, using a scheduled script). Templates are created in the same way as reports are requested from the Administration Assistant function for Data Protection for SAP panels. Whenever the Create Report button is used, you are prompted to create a report or use the corresponding template. Each template must be given a unique name which is used when referencing the template. The template is stored in a file with the given name in path <Administration Assistant install dir>/templates/userid/ where <Administration Assistant install dir> is the Administration Assistant server installation path. The file extension depends on the type of report requested. A single template can be used to generate reports on several SIDs. Note that a template is owned by the user account that creates it and cannot be accessed from a different account. In order to view, change, or delete owned templates, an administrator can use the Manage templates function in the Administration Assistant View pull-down menu.

Server-related tuning

You can manage the data stored on the Tivoli Storage Manager server for Tivoli Storage Manager for ERP. You can manage which servers are used to store data.

Managing Data on the Backup Server

The Data Protection for SAP for DB2 Backup Object Manager can search for backup objects on the Tivoli Storage Manager server in order to restore or delete them. Complete information is located in “DB2 Backup Object Manager utility” on page 4.

Alternate Network Paths and Servers

Multiple network paths and multiple backup servers can be used as an alternate instead of in parallel. When the number of available sessions to multiple servers exceeds the number of sessions allowed, Data Protection for SAP for DB2 uses the first sessions it can establish. It continues to use the number of sessions allowed as defined by the MAX_SESSIONS keyword (as described on page “Tivoli Storage Manager for ERP for DB2 profile parameter descriptions” on page 134). This allows data to be backed up even when a resource (such as a Tivoli Storage Manager server or its network interface) is unavailable. The servers used for the backup must be available in order to restore the data. Note that the days of the week that a server is used can also be specified as described for the USE_AT keyword on page “Tivoli Storage Manager for ERP for DB2 profile parameter descriptions” on page 134.

Options

You can use Tivoli Storage Manager for ERP options to tune performance.

Performance Options of Data Protection for SAP for DB2

These three components have the greatest impact on data transfer rates:

- the type of disks on which the database resides
- the network capabilities accessed by the database host and the Tivoli Storage Manager server
- the type of storage device that contains the backup

Data Protection for SAP for DB2 provides these options to help optimize the data transfer rate for these components.

Parallel (Multiple) Sessions

Data Protection for SAP can back up or restore data to multiple tape drives in parallel. Parallelism is achieved by using more than one session to send data to a backup server. Details are provided in “Multiple Sessions” on page 109.

Multiple (Parallel) Network Paths and Multiple (Parallel) Servers

Improve performance by configuring Data Protection for SAP to distribute a database backup across two or more Tivoli Storage Manager servers. In addition, you can balance network traffic by providing two (or more) separate network connections between the SAP® database host and the Tivoli Storage Manager server. Detail information regarding these features is available in “Multiple Network Paths” on page 110 and “Multiple Servers” on page 108.

Incremental and Delta Backup

Data Protection for SAP supports incremental and delta backups of DB2 databases. Depending on the system environment, incremental backups might decrease backup processing time.

RL_COMPRESSION

The RL_COMPRESSION profile keyword compresses a partially filled database. This can result in reduced network traffic and fewer tapes required for backup. See “Compression” on page 104 for complete details.

Adjustments to Data Protection for SAP for DB2 for Improving Performance of Data Transfer

Data Protection for SAP for DB2 is configured (by default) to send uncompressed data to the Tivoli Storage Manager server using a single session.

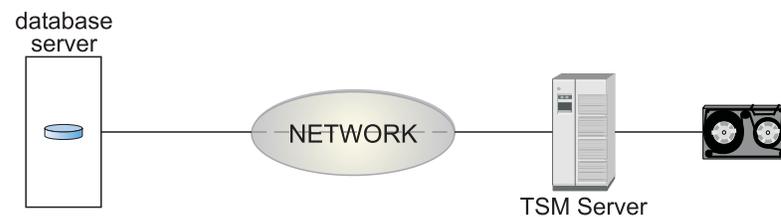


Figure 22. Data Transfer for a Backup / Restore

A single configuration that is best for all environments is not possible or realistic. However, the information provided in this section can help in determining which configuration is best for your environment. The Administration Assistant function for Data Protection for SAP provides the View Performance Data feature which provides information about performance characteristics and how they change with your configuration. Information about tuning a system with the Administration Assistant is available in “Overview of a balanced system” on page 91.

Buffer Copies

Data Protection for SAP for DB2 uses internal buffers to store and exchange data with Tivoli Storage Manager. When sending data from one component to another, data buffers are copied (by default). Data Protection for SAP can prevent copying the data buffers by sending the original data buffers. This reduces the CPU load of the database server. However, if client compression or client encryption are specified in the Tivoli Storage Manager options file (`dsm.sys` or `dsm.opt` on UNIX or Linux or `server.opt` on Windows), the original data buffers are sent. See the description of `BUFFCOPY` keyword on page “Tivoli Storage Manager for ERP for DB2 profile parameter descriptions” on page 134 for more information.

Buffer Size

Data Protection for SAP for DB2 allows the size of the internal data buffers to be adjusted. These buffers are used for both reading the disk and sending data to the Tivoli Storage Manager client API. The default values typically produce acceptable performance. It is recommended to optimize the buffer size for disk I/O. For disk subsystems, the best transfer rates have been achieved when the buffer size was set equal to the stripe size. Before increasing the size of internal buffers, however, make sure that sufficient storage is available for the number of buffers specified by Data Protection for SAP. This number correlates to the number of sessions requested. Be aware that number of buffers doubles when compression is specified. See the description of `BUFFSIZE` keyword on page “Tivoli Storage Manager for ERP for DB2 profile parameter descriptions” on page 134 for more information.

Compression

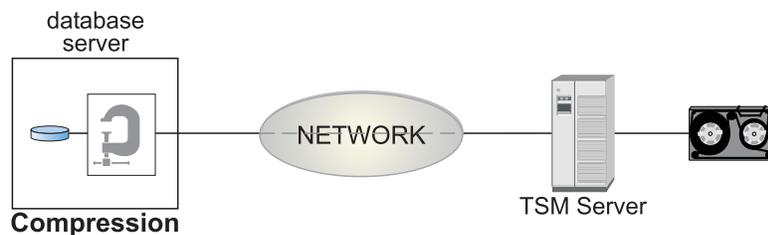


Figure 23. Null Block Compression

Data Protection for SAP for DB2 can decrease the amount of data sent to the Tivoli Storage Manager server by compressing zero-byte blocks. Although compression can increase the CPU load on the database server, it can improve performance in situations when the network is the point of constraint. Compression is most effective with database files that contain large portions of null blocks. See the description of the `RL_COMPRESSION` keyword on page “Tivoli Storage Manager for ERP for DB2 profile parameter descriptions” on page 134 for details on how to activate Data Protection for SAP compression.

Automation Options for Data Protection for SAP for DB2

Administrative productivity can be improved by using these Data Protection for SAP for DB2 automation options.

Selectable Management Classes

Specify different Tivoli Storage Manager management classes for back up data and archive data. It is recommended to configure Data Protection for SAP to back up directly to a tape storage pool and to archive DB2 log files to a disk storage pool. Multiple management classes can also be specified to use in conjunction with multiple DB2 log files. The profile keywords BRARCHIVEMGTCLASS and BRBACKUPMGTCLASS in “Tivoli Storage Manager for ERP for DB2 profile parameter descriptions” on page 134 provide information about specifying management classes.

Retain Backups by Version

Retaining backups by version limits the number of full backups retained on the Tivoli Storage Manager server. When the number of full backups on the Tivoli Storage Manager server exceeds the specified number, the oldest version is deleted. Retaining backups keeps track of all DB2 log files, and all incremental and delta backups, associated with a full backup. All these objects are removed together with the full backup.

Multiple Redo Log Copies

Backing up multiple copies of a log file in a single archive operation helps protect against this data in the event of tape defects or disaster recovery situation. These copies can be located on different physical Tivoli Storage Manager volumes or on different Tivoli Storage Manager servers. When a log file copy is unavailable at restore time, Data Protection for SAP automatically switches to another copy and continues restoring the log file from that copy. The description of the profile keyword REDOLOG_COPIES in “Tivoli Storage Manager for ERP for DB2 profile parameter descriptions” on page 134 provides detailed information about creating and using multiple Redo Log Copies.

Alternate Network Paths and Servers

The availability of backed up data can be improved by configuring Data Protection for SAP to use multiple Tivoli Storage Manager servers or multiple network connections to a single Tivoli Storage Manager server. In this configuration, Data Protection for SAP checks all servers and network connections for availability and then performs the backup even if some resources are unavailable. Policies can also be set that use different Tivoli Storage Manager servers for different days of the week.

Messaging

Policies can be created that enable Data Protection for SAP to send different classes of log messages to the Tivoli Storage Manager server.

Frontend/Backend Processing

Frontend and backend processing calls programs at specified times during backup processing. See the description of the profile keywords BACKEND and

Data transfer

Review information about ways to improve performance for data transfer.

Observations on the Data Protection for SAP for DB2 Data Throughput

Throughput rates differ widely among various environments due to different disk, network bandwidth, server platforms, number of tapes, and configuration settings. The information provided in this section concentrates on selected elements involved in the movement of data. This information should assist in determining how to use existing resources to their maximum efficiency and provide insight as to how throughput can be improved.

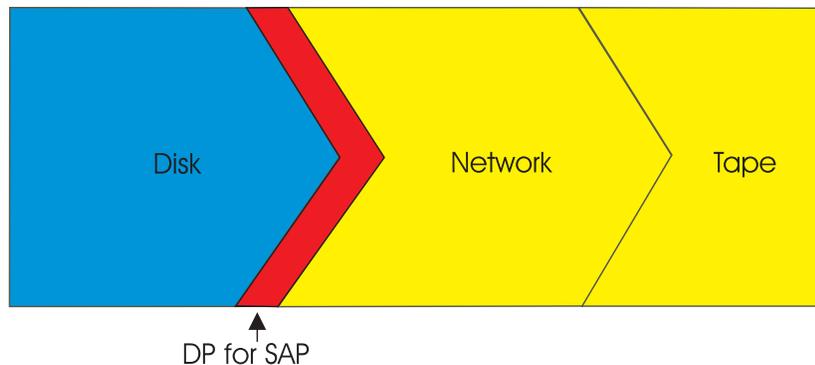


Figure 24. High-level View of the Data Flow During Backup

From a high-level view, the data packages need to send these elements when doing a backup with Data Protection for SAP for DB2: Data is read from disk, processed by Data Protection for SAP, and sent through the network to tape or disk storage media. If the system is not balanced, the disk I/O, network bandwidth, and storage media rates might create a bottleneck which can cause other resources to remain idle. Overall data throughput is typically measured per file or per entire backup operation. The results are documented as an average throughput rate in a log file. However, identifying bottlenecks based upon log file messages is difficult. To assist in this analysis effort, Data Protection for SAP provides performance sensors that indicate whether there is a bottleneck located either in the elements represented in blue (for disk) or in yellow (for network and tape respectively) in the this graphic. Data Protection for SAP configuration options that can be adjusted to improve performance is described in “Performance Options of Data Protection for SAP for DB2” on page 103. Additional performance issues are available in “General Performance Considerations” on page 107.

Data Protection for SAP for DB2 Performance Sensors

The method of transferring data packages depends on how Tivoli Storage Manager is configured. In a standard configuration, the data packages are sent from the Tivoli Storage Manager API Client through the network to the backup server. In an environment configured for LAN-free operations, the data packages are processed by the Tivoli Storage Manager API Client and the Tivoli Storage Manager Storage Agent.

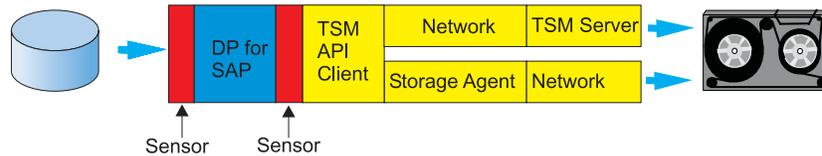


Figure 25. Performance Optimizing by Using Sensors

Data Protection for SAP for DB2 uses sensors that observe incoming and outgoing data streams. They measure throughput and the idle time of the I/O threads in comparison to the duration of the backup. This provides a way to determine whether the streams of incoming and outgoing Data Protection for SAP data are balanced. Be aware that once a backup operation begins, the buffers need to be filled before the effects of a bottleneck are viewable.

General Performance Considerations

Figure 26 on page 108 provides a high level overview of these three main components involved during a Data Protection for SAP for DB2 data transfer:

- The SAP® database server.
- The network.
- The Tivoli Storage Manager server which is also referred to as a backup server.

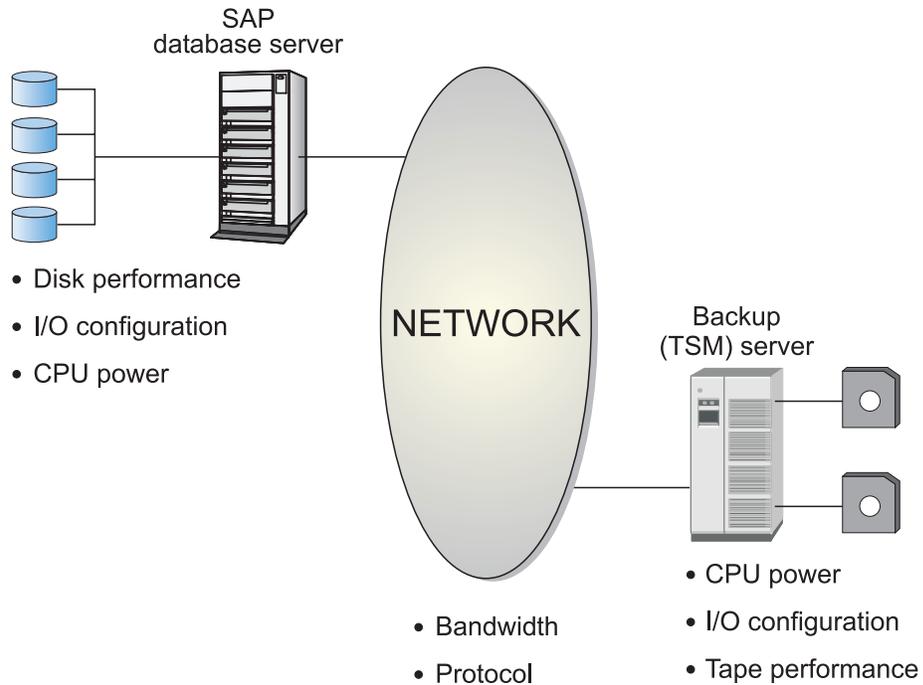


Figure 26. Data Protection for SAP Data Transfer

A continuous stream of data is generated among these components during a backup or restore operation. The weakest component in this stream decreases the overall data transfer rate. The guidelines provided are based on experience gathered from many installations and should be considered when designing a backup/restore infrastructure that will be efficient.

Multiple Servers

Data Protection for SAP for DB2 supports multiple servers which can distribute backup data among two (or more) backup servers. This feature helps eliminate constraints that are frequently encountered among backup servers.

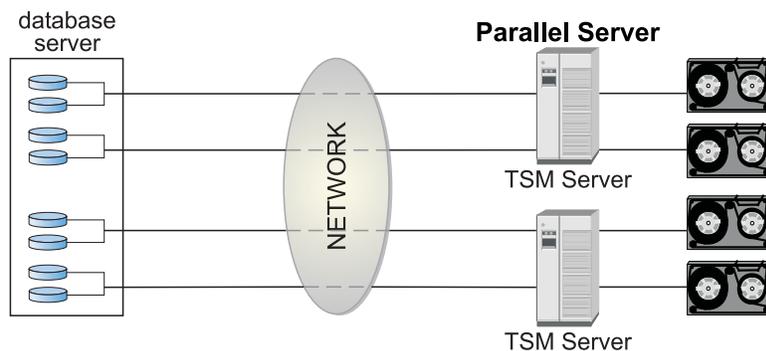


Figure 27. Multiple Servers

A server statement must be entered in the Data Protection for SAP profile for each adapter of the backup server as described for the SERVER keyword in “Tivoli Storage Manager for ERP for DB2 profile parameter descriptions” on page 134. The value of the MAX_SESSIONS keyword is not greater than the sum of all SESSION values specified for the SERVER statements used concurrently.

Multiple Sessions

Data Protection for SAP for DB2 allows use of multiple tape drives simultaneously in order to increase the transfer rate to or from the Tivoli Storage Manager server. The keyword `MAX_SESSIONS` is used for defining the number of parallel sessions

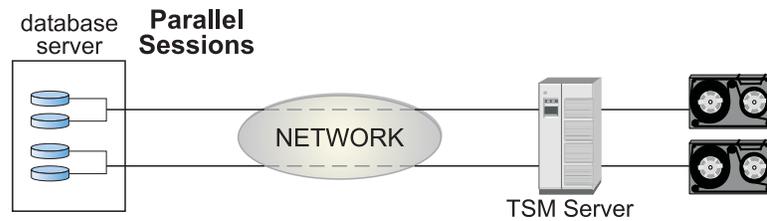


Figure 28. Parallel (Multiple) Sessions

to be established with the Tivoli Storage Manager server for database backup, archive (backup of log files) and restore. For a detailed description of how to use this keyword, see page "Tivoli Storage Manager for ERP for DB2 profile parameter descriptions" on page 134. When performing a database backup, the data is typically written directly to tape drives on the Tivoli Storage Manager server. The parameter specified in the `MAX_SESSIONS` keyword must match the number of tape drives used simultaneously. These must be available to the management class defined as `BRBACKUPMGTCLASS` in the Data Protection for SAP profile as described on page "Tivoli Storage Manager for ERP for DB2 profile parameter descriptions" on page 134.

When setting up the Tivoli Storage Manager server, make sure not to activate collocation in the (tape) storage pool defined for the management class chosen as `BRBACKUPMGTCLASS`. In addition, make sure as many tape drives for this management class are available as the number of sessions defined in `MAX_SESSIONS` as multiple access to the same tape might slow down data transfer.

When DB2 log backups are running, either disk or tape storage pools can be utilized. These must be available to the management class defined as `BRARCHIVEMGTCLASS` in the Data Protection for SAP profile. If you are using tape pools as primary pools for this management class, this consideration for database backups also applies to disk storage pools:

Several DB2 log archive sessions can simultaneously utilize one or two independent disk storage pool(s).

The number of storage pools that are required depends on the number of backup copies requested for a DB2 log file. See keyword `REDOLOG_COPIES` in "Tivoli Storage Manager for ERP for DB2 profile parameter descriptions" on page 134.

Multiplexing

Multiplexing is using parallel access to data on the database server. This is recommended when using a tape drive during database backup operations on the backup server.

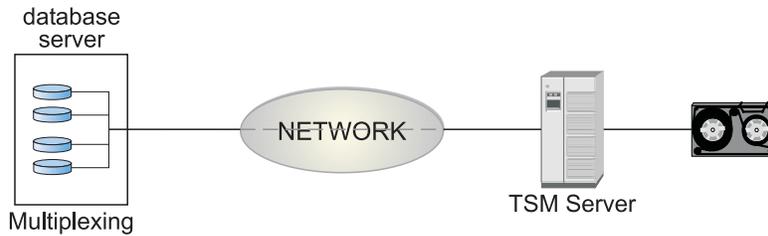


Figure 29. Multiplexing

This feature is provided by the PARALLELISM parameter available with the BACKUP DATABASE and RESTORE DATABASE commands. Refer to your *DB2 Command Reference* for details about these commands and the PARALLELISM parameter.

Multiple Network Paths

Data Protection for SAP for DB2 allows you to use multiple network connections (paths) for data transfer between the database server and the backup server.

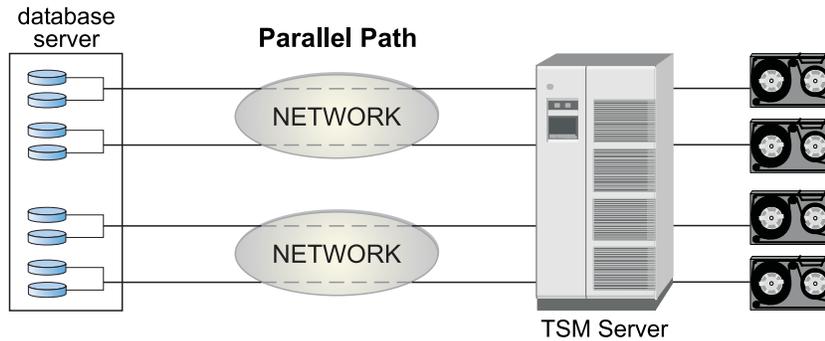


Figure 30. Parallel (Multiple) Paths

Parallel paths can be used to eliminate network points of constraint. For each additional path, additional network adapters are required on both the production and the backup server. A server statement must be entered in the Data Protection for SAP profile for each adapter of the backup server as described for the SERVER keyword on page “Tivoli Storage Manager for ERP for DB2 profile parameter descriptions” on page 134. The value of the MAX_SESSIONS keyword is not greater than the sum of all SESSION values specified for the SERVER statements used concurrently. Detailed information regarding setting up multiple parallel network paths is described in detail in “Alternate or parallel backup paths and backup servers” on page 18.

Storage space

You can manage aspects of storage space to tune the performance of Tivoli Storage Manager for ERP.

Automated Tablespace Adaptations

Backup Object Manager can adapt the sizes of tablespace containers when creating the containers of the target databases during a redirected restore operation. For example, tablespace container sizes might be increased in order to provide more space or decreased in order to use storage more efficiently. Tablespaces can also be allocated with similar sizes in order to make parallel I/O operations more efficient. These features are supported by Backup Object Manager resizing and normalizing functions.

Tablespace Normalizing

In order to achieve optimal parallel I/O operation performance for a database, all containers of a tablespace should be the same size. During tablespace maintenance, containers may be added or extended which creates different container sizes. As a result, data is unevenly distributed among the containers which can result in decreased parallel I/O operation performance during table scans (sequential prefetching). Backup Object Manager provides an automated tablespace normalizing function that allows the location and size of tablespace container to be redefined. This also helps prevent I/O-intensive tablespace rebalancing that can be detected by DB2.

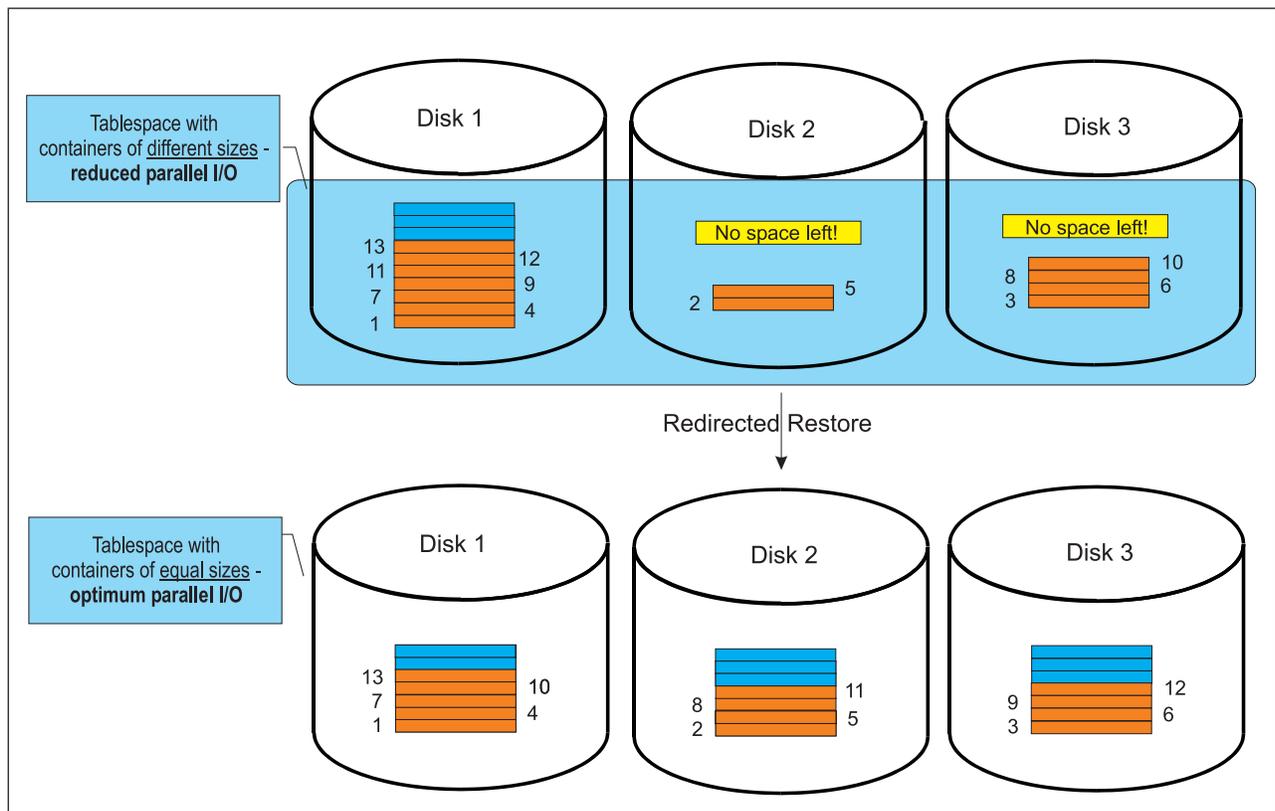


Figure 31. Tablespace Normalizing

This graphic shows that the original system tablespace consists of different sized tablespace containers. Even though sequential prefetching allows three processes to simultaneously read the data during table scans, the different container sizes and uneven data distribution prevent parallel I/O during part of the scan. The degree of parallelism decreases over time. To counteract this, adjust the containers for each tablespace so that they are the same size. This type of adjustment, in combination

with a redirected restore operation, requires no further resizing tasks for the tablespace containers after the restore completes.

Backup Object Manager simplifies the tablespace container resizing process by providing an automatic tablespace normalizing function. It can be used in combination with any mode of redirected restore facility by specifying the `-N` option. Issue this command to resize each container of a tablespace to the average size of all containers within the same tablespace during a redirected restore:

```
backom -c rr_db_type -aDB2 source alias,DB2 target alias -t timestamp -N
```

After the redirected restore completes successfully, all containers of a tablespace of the target database are the same size. As a result, the continuous parallel I/O performance of the physical layout of the restored database is optimized.

Tablespace Scaling

Sufficient free space must be available in a tablespace for the database to function properly. Backup Object Manager provides an automated tablespace scaling function that allows the location and size of tablespace container to be redefined. This also helps prevent I/O-intensive tablespace rebalancing that can be detected by DB2.

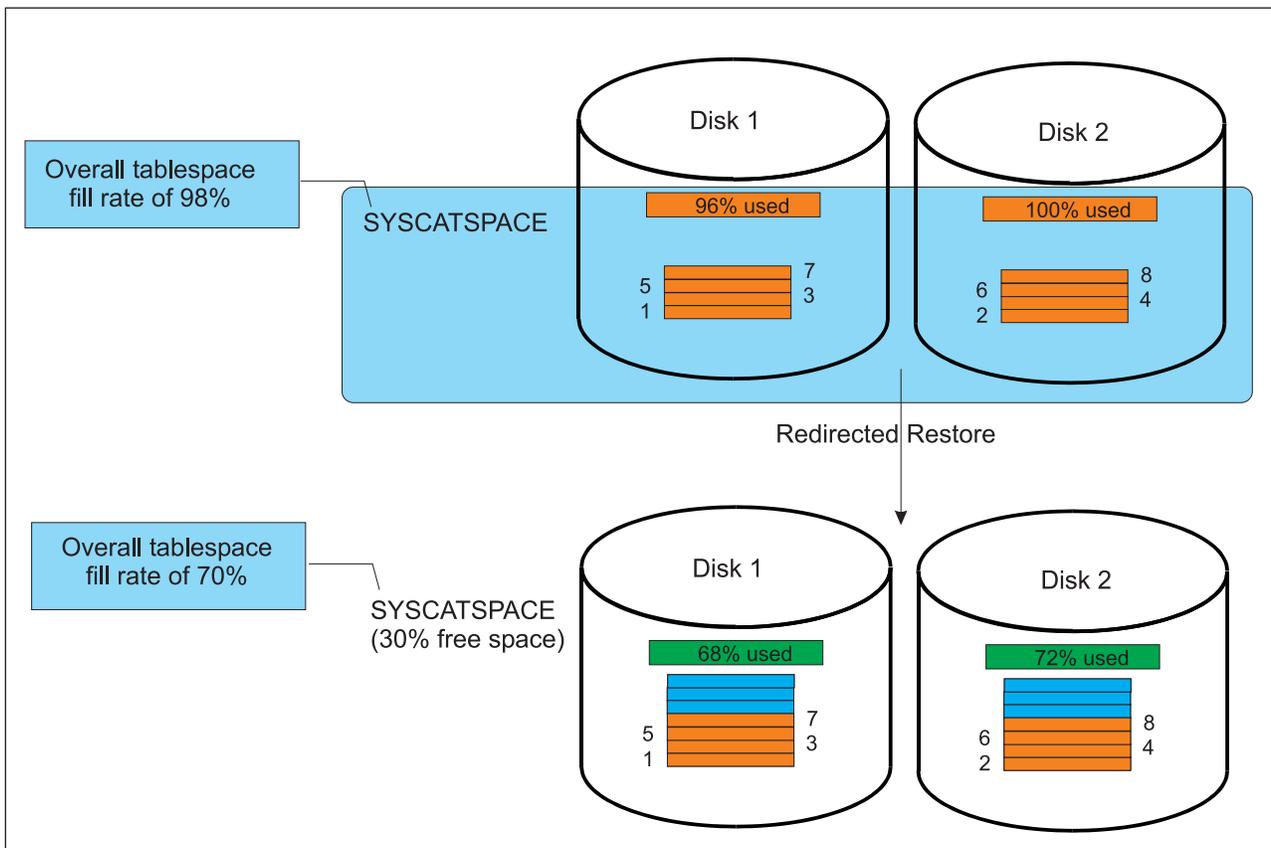


Figure 32. Tablespace Scaling

This graphic shows that 98% of the SYSCATSPACE tablespace of the original system is being used. Disk 1 has 4% free space while 100% of the second container is used. The free space available in a tablespace can be increased as part of the

redefinition feature during a redirected restore. The goal is to achieve an overall filling rate for the target side of 70%. This can be achieved by manually increasing the amount of free space the first container must provide at 20% and 40% for the second container. This type of adjustment, in combination with a redirected restore operation, requires no further resizing tasks for the tablespace containers after the restore completes.

Backup Object Manager simplifies the tablespace resizing process by providing an automatic tablespace scaling function. It can be used in combination with any mode of redirected restore by specifying the `-S` option with a floating point sizing factor. Consider these factors when resizing tablespaces:

- A value of '1' indicates that the target tablespace is 100% the size of the original (nothing is changed).
- A value greater than 1 increases the target tablespace. For example, a value of 1.1 increases the target tablespace by 10% to a target value of 110% of the original.
- A value less than 1 decreases the target tablespace. For example, a value of 0.9 decreases the target tablespace by 10% to a value of 90% of the original.

Therefore, manual adaptation of the tablespace containers described above can be replaced by the following procedure using Backup Object Manager redirected restore:

1. Issue this Backup Object Manager query to determine the original fill rate of the tablespace:

```
backom -c q_tdi -a DB2 source alias -t timestamp -m detailed
```

2. Calculate the tablespace scaling factor using this formula:

```
scaling factor = original fill rate / new fill rate
```

For example:

```
scaling factor = 0.98 / 0.7 = 1.4
```

3. Issue this command to begin the Backup Object Manager redirected restore:

```
backom -c rr_db_type -a DB2 source alias,DB2 target alias -t timestamp ...  
...-T SYSCATSPACE -s 1.4
```

After the redirected restore has completed successfully, the SYSCATSPACE tablespace on the target side is increased by 40% during tablespace container redefinition. The new overall fill rate of the SYSCATSPACE in the target database is now 70%.

Chapter 8. Troubleshooting IBM Tivoli Storage Manager for Enterprise Resource Planning

Information on how to resolve errors that might occur during IBM Tivoli Storage Manager for Enterprise Resource Planning operations is provided.

Troubleshooting IBM Tivoli Storage Manager for Enterprise Resource Planning common problems

Information on how to resolve errors that might occur during IBM Tivoli Storage Manager for Enterprise Resource Planning operations is provided.

Random problems

If a problem occurs inconsistently, try to determine what the difference is when the problem occurs, if any. Compare the log files of the application in question ((`tdpdb2.SID.nodename.log`, `db2diag.log`, Tivoli Storage Manager server activity log, etc.) to find out the differences between successful and unsuccessful operations. Look for one of these patterns when the problem occurs:

- The problem always occurs at the same time. If this is true, view the appropriate log files to determine review if there are any scheduled processes occurring simultaneously such as virus checker, automatic updates, or batch jobs.
- The problem always occurs after another operation is performed or the same operation is performed.
- The problem occurs when another application or process is performed in parallel.

Reproducible (repeatable) problems

When encountering a problem that occurs during an operation that has previously performed successfully, consider these possible causes:

- The Data Protection for SAP for DB2 setup changed.
- One (or more) of the DB2, SAP, Tivoli Storage Manager, operating system, network, or hardware components changed.
- Patches or updates to one (or more) of the components were applied.
- Changes originated by the system have occurred such as these:
 - Check if the disks are running full with the UNIX `df` command.
 - If network performance has decreased, check if additional hosts, additional applications, or defects in software or hardware occurred. Compare operation runs in the Administration Assistant Performance Monitor history view or compare the `tdpdb2.SID.nodename.log`.
 - If Tivoli Storage Manager server processing has decreased, check if additional clients or additional operations were added. Information is also available in the Tivoli Storage Manager server activity log.

When none of these possible causes has occurred, view the last modified time stamp of the configuration files (`vendor.env`, `initSID.utl`, `dsm.sys`, `dsm.opt`, `/etc/services`, `/etc/inittab`, ...). This UNIX command lists all files in the `/etc` directory which have been modified during the previous five days:

```
find /etc -type f -ctime 5 -print
```

If you are able to identify changes made to the system, roll them back one at a time and try to reproduce the problem. This method frequently reveals which change or set of changes caused the problem.

Internet Protocol version 6 (IPv6) support

Data Protection for SAP for DB2 supports both IPv4 and IPv6 for internal communication in that it will run in IPv4, IPv6, and mixed environments on AIX and Linux. However, these products do not exploit new IPv6 functionality. In a mixed environment, the communication depends on the adapter network settings. There is no option to enforce the use of a specific protocol other than by network configuration. Specifically, the ProLE or acsd service will listen for both IPv4 and IPv6 connection requests if the system is configured accordingly. Connection requests to ProLE are made for the addresses returned by the system for the respective port on the local host. Connection requests to other machines such as the Administration Assistant function for Data Protection for SAP are made for the addresses specified by the user. IPv6 addresses are supported when TCP/IP addresses are specified in a command line or in a profile parameter such as TCP_ADDRESS. However, when the IP address and port are specified in the *IPv4 address:service or port* format, then the format needs to be changed to *service or port@<IP address>* if the IP address is specified in the IPv6 notation. In the case of a dotted decimal IPv4 address, the traditional format can still be used.

The specification of IPv6 addresses assumes that Data Protection for SAP is used in an environment in which IPv6 is supported by all hardware and software components involved and has been adequately tested in this environment.

Understanding the Setup

Review these considerations to better understand the installation setup on UNIX or Linux systems:

- Make sure all files are installed as described in “Prerequisites” on page 24.
- Make sure an entry similar to this example is defined in the `/etc/inittab` file:

```
pd64:2:respawn:/usr/tivoli/tsm/tdp_r3/db264/prole -p tdp3db264
Server component hostname 5126
```

The purpose of this entry is to start a daemon process for ProLE. This process listens on the Data Protection for SAP for DB2 port `tdpr3db264` for connections with the shared library and sends performance-related information to the Administration Assistant Server component. The port can have a different name; however, the name must match the name in the `/etc/services` file as shown in this example:

```
tdpr3db264      57324/tcp
```

These lines are added to the `/etc/services` file by the installer.

See Figure 33 on page 117 for an overview of the configuration files on a UNIX or Linux system.

Review these considerations to better understand the installation setup on Windows systems:

- Make sure all files are installed as described in “Prerequisites” on page 24.
- Verify that service ProLE Service is running and set to automatic startup. If this service is not running, Data Protection for SAP does not function properly.
- The installer adds lines to the %SYSTEMROOT%\system32\drivers\etc\services file similar to these:

```
tdpr3db264      57324/tcp
```

- Make sure the Data Protection for SAP configuration file `initSID.utl` is located in the directory to pointed by the `TDP_DIR` environment variable.
- The vendor environment file `vendor.env` must contain the fully qualified path and file name of the `initSID.utl` file.
- The vendor environment file `vendor.env` should contain the path of the location where the Data Protection for SAP run logs are written. If this location is not specified, temporary directory of the machine is used.

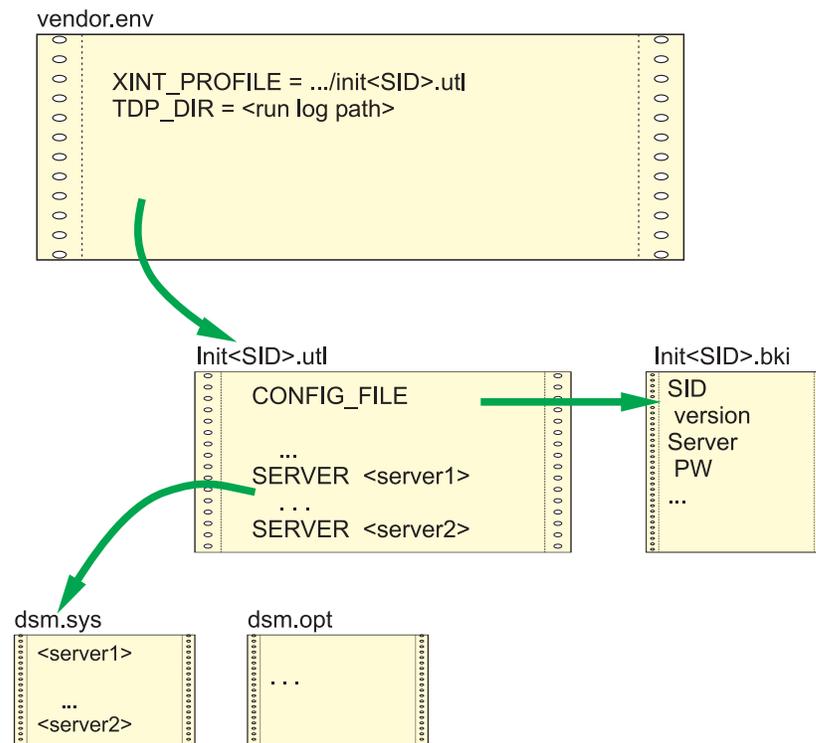


Figure 33. SAP® and Data Protection for SAP configuration files on UNIX or Linux

On UNIX or Linux systems, the names of the Tivoli Storage Manager servers specified in `initSID.utl` must match the names in the `dsm.sys` file. If the Tivoli Storage Manager API or Tivoli Storage Manager Backup Archive Client were installed into their default locations, then the `DSMI_*` variables do not need to be set. If the variables are set, however, make sure they specify the correct directories and files. The user ID that runs the backups must have the correct permissions to access all of files and directories specified by these variables. Also verify that write permissions exist for the `initSID.bki` file as this is the only file to which Data Protection for SAP writes persistent information.

On Windows systems, the `dsm.opt` file is used instead of the `dsm.sys` file. However, the content of this file is not relevant to Data Protection for SAP. The directory that contains the `dsm.opt` file must also contain a `server.opt` file for each server specified in the `initSID.utl` file. The environment variable `DSMI_CONFIG` must specify an option file within this directory. `DSMI_CONFIG` should specify the `dsm.opt` file in this directory. The `DSMI_DIR` environment variable must also specify the directory where the Tivoli Storage Manager API message text file resides. This is typically the `c:\Program Files\Tivoli\tsm\api64` directory.

Providing information to IBM or Tivoli support

Provide this information when contacting IBM or Tivoli support:

- The Data Protection for SAP for DB2 version.
- The operating system level and patches that were applied.
- The DB2 version
- The Tivoli Storage Manager server version.
- The Tivoli Storage Manager server operating system level.
- Data Protection for SAP configuration file (`vendor.env`, `initSID.utl`) including Tivoli Storage Manager client configuration files (`dsm.sys`, `dsm.opt`)
- Data Protection for SAP profile (`initSID.utl`)
- The change history of the system components (if the process worked previously).

Additional information might also be requested from the service representative.

Troubleshooting Data Protection for SAP for DB2 problems

Information on how to resolve errors that might occur during Data Protection for SAP for DB2 operations is provided.

General problem resolution

The following graphic (Figure 34 on page 119) will help you to isolate problems that occur when backing up or restoring of your DB2 database.

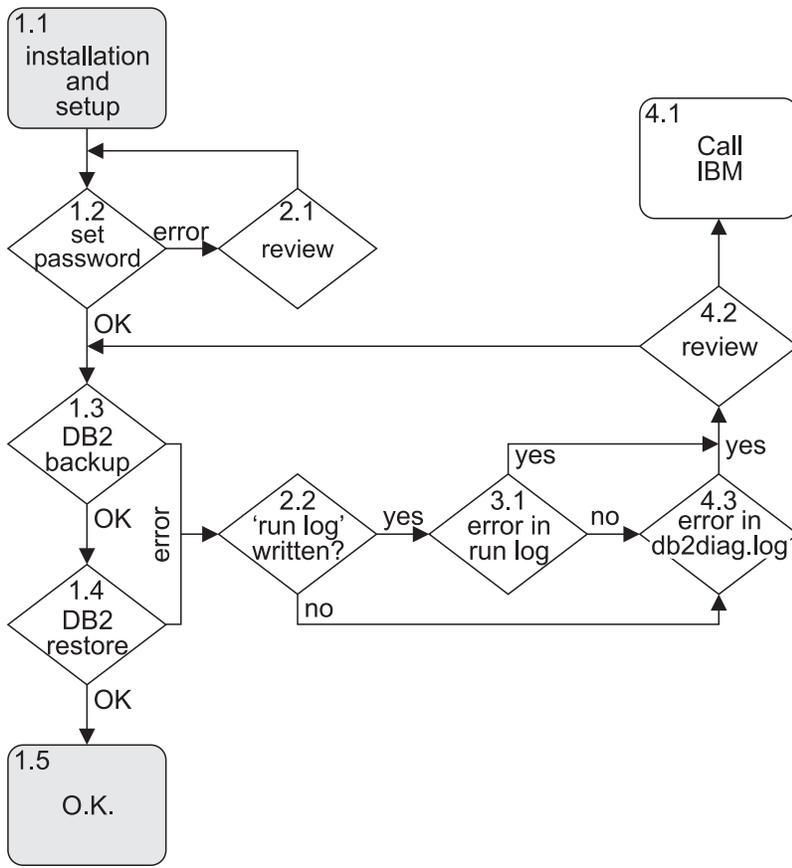


Figure 34. General Problem Isolation

After installation is completed (Step 1.1) and manual password handling is specified, set the password (Step 1.2) as described in “Set the Tivoli Storage Manager password” on page 62. When the operation completes successfully, the informational messages BKI0051I: Password successfully verified for node *NODENAME* on server *SERVERNAME* and BKI0024I: Return code is: 0. display for each server configured within the *initSID.utl* file. An error message displays when a problem occurred. The Administration Assistant can also be used. The Configurator feature loads the configuration of the node on which problems are encountered and allows the Administration Assistant to check the configuration.

These errors are frequently encountered at Step 1.2:

BKI2001E: Socket error while connecting to ProLE at IP-Address:PORT: Connection refused

On Windows, verify that the ProLE Service is running by viewing the Computer Management Services screen or issue this command:

```
net start
```

A list of all running services displays. On UNIX or Linux, verify that the background daemon is running by issuing this command:

```
ps -ef | grep prole
```

Check the entry in `/etc/services` (UNIX or Linux) and `%SYSTEMROOT%\system32\drivers\etc\services` (Windows). Compare the port number from the error message with the port number within `/etc/services`. Also check the entry in `/etc/inittab` (UNIX or Linux). If another port was set using the option `-pPORT`, check this as well. If all of this will not help, start the ProLE from another shell on UNIX or Linux with this command:

```
prole -p PORT
```

Issue this command on Windows:

```
prole -console -p PORT
```

Attempt to start backom again.

BKI5001E: Tivoli Storage Manager Error: Server not found in configuration file On UNIX or Linux, the Tivoli Storage Manager server defined in the `initSID.utl` file does not match the server specified in the `dsm.sys` file. On Windows, the `server.opt` file might be missing.

BKI5001E: Tivoli Storage Manager Error: ANS1353E (RC53) Session rejected: Unknown or incorrect ID entered

This message can display when the node in the server stanza of the UTL file is not valid on the server.

HANG If backom hangs after the password is entered, the server IP address specified in the UNIX or Linux `dsm.sys` file might be incorrect.

When Step 1.2 (setting the password) is successful, proceed to Step 1.3 and perform a backup using the DB2 backup command to verify the settings are correct as described in “Backing up SAP® data” on page 67. If the backup was successful you will see a message from DB2:

```
Backup successful. The timestamp for this backup image is: timestamp
```

If an error message displays, find the message and information about how to resolve the error in *Tivoli Storage Manager for ERP Messages*.

When an error occurs, always view the Data Protection for SAP run log `tdpdb2.SID.nodename.log` first. This log file is located in the directory specified by the `TDP_DIR` environment variable. If the variable is not specified, the log file resides in the system temporary directory. If the `tdpdb2.SID.nodename.log` file does not exist (Step 2.2), then either DB2 was unable to load the shared library that contains the DB2 connector for Data Protection for SAP or an error was encountered before calling the Data Protection for SAP library. In both cases, a DB2 error message should display on the command line that begins with the SQL prefix and is also written in the DB2 diagnostic log `db2diag.log` (Step 4.3). DB2 provides detailed error descriptions by issuing this command:

```
db2 ? SQLnumber
```

Replace *number* with the appropriate message number. Try to resolve this problem using the DB2 documentation.

If the `tdpdb2.SID.nodename.log` file exists, search for a message beginning with `BKIXXXXY` where `XXXX` is a four digit number and `Y` is the letter I, W, or E. When such a message occurs, the DB2 connector for Data Protection for SAP loaded correctly was called by DB2. In Step 3.1, the `tdpdb2.SID.nodename.log` file is created and an error message starting with BKI is recorded.

Location of log files

Text displayed on the screen during DB2 backup, DB2 restore, and BackOM operations are typically written to log files. DB2 also writes messages of internal operations, events, or status in the administration notification log file (`db2SID.nfy`) and diagnostic log file (`db2diag.log`). These log files reside in the directory specified with the DB2 database management configuration parameter `DIAGPATH`. Query the DB2 database management configuration with this command:

```
db2 get dbm cfg
```

Information about how to locate these log files is available in .

How to find files containing message output (log files)

Data Protection for SAP for DB2 process results are logged in files.

These files are located in the path indicated by the `TDP_DIR` environment variable. After the installation, `TDP_DIR` points to the subdirectory `tdplog` of the path for the Data Protection for SAP configuration files. If `TDP_DIR` is not set, or if a log file cannot be created in the path pointed to by `TDP_DIR`, the log files are created in path `/tmp` (UNIX or Linux) or in the path pointed to by environment variable `TEMP` (Windows). Information on how to set or change the `TDP_DIR` value is available in "Prerequisites" on page 24. The Data Protection for SAP shared library writes to the `tdpdb2.SID.<node name>.log` log file. The Backup Object Manager writes to the `backom.log` log file.

DB2 vendor reason codes

Data Protection for SAP for DB2 uses these reason codes which might also be displayed or logged by DB2 in the case of problems.

Table 12. DB2 Vendor Reason Codes

Reason Code	Explanation	User Response
1	The library specified could not be loaded.	Check the DB2 diagnostic log for further details.
2	Communication error between shared library and ProLE	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.
6	Object specified cannot be found on Tivoli Storage Manager.	There is no backup image on Tivoli Storage Manager matching the given search criteria.
10	Invalid options specified with the options parameter of the DB2 backup/restore command.	Check the options string specified and check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.

Table 12. DB2 Vendor Reason Codes (continued)

Reason Code	Explanation	User Response
11	Initialization procedure for shared library failed.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> and the DB2 diagnostic log file for further details.
17	During end processing of either backup/archive or restore/retrieve session(s), an error occurred.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.
18	An error occurred during reading or writing data from or to Tivoli Storage Manager.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.
26	An error occurred during deleting data from Tivoli Storage Manager.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.
29	An abort request from DB2 could not be handled correctly.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> and the DB2 diagnostic log file for further details.
30	A severe error occurred.	Check the DB2 diagnostic log file for further details.

Chapter 9. Data Protection for SAP for DB2 reference information

Data Protection for SAP for DB2 reference information is provided here.

Commands used with Data Protection for SAP for DB2

A list of various commands that are used with Data Protection for SAP for DB2 operations is provided.

Versioning

When Tivoli Storage Manager for ERP versioning is active (as defined by the MAX_VERSIONS keyword), versioning information is stored on the Tivoli Storage Manager Server. The version number is increased only after successful backups.

Backups and Restores in Partitioned Database Environments

DB2 Version 9.5 (or later) provides the Single System View (SSV) function, which allows the backup of all partitions to be triggered with a new dedicated DB2 backup command option ON ALL DBPARTITIONNUMS (see “Using DB2 Single System View for Backup” on page 69 for more information). Furthermore, partitioned database backups and restores can be made by using the already established DB2 framework for partitioned databases called db2_all. A requirement of DB2 is to back up and restore the catalog partition separately from all other DB2 partitions. Thus, backup and restore operations of a partitioned database are two-step scenarios, whereby the first step is to backup/restore the catalog partition and the second step to backup and restore all other partitions in parallel. Data Protection for SAP uses the DB2 command db2_all for separation and parallelization of the backup and restore commands. The db2_all command provides special characters for handling partitions and for running commands in parallel or sequentially. The DB2 *Administration Guide* contains further information.

The DB2 db2_all command supports the following special characters or character combinations:

- <<+N< Runs a command only on partition N
- <<-N< Runs a command on all partitions except on partition N
- " Substitute occurrences of () by the machine index and substitute occurrences of ## by the partition number
- ; Runs the commands in parallel in the background and terminates the command after all remote commands are completed.

By using these characters, each partition of the database can be backed up or restored with its special adapted environment.

This list presents possible partitioned database backups and restores using Data Protection for SAP for DB2:

Full offline backup

Full offline backup of all database partitions with two sessions, starting with the catalog partition (shown as partition 0) followed by all other partitions in parallel:

```
db2_all '<<+0< db2 backup db SID load shared library open 2 sessions'  
db2_all '<<-0<; db2 backup db SID load shared library open 2 sessions'
```

Full restore of latest backup

Full restore of latest backup starting with the catalog partition (shown as partition 0) followed by all other partitions (shown with three partitions) in parallel using different number of sessions for some of the partitions. Before starting the restore, a temporary environment script (this example uses `/db2/SID/EEEenv.sh`) has to be created. The `/db2/SID` directory must be an NFS share between all hosts where the partitions reside. For this example, it would contain the following:

```
export SESSION0=2  
export SESSION1=2  
export SESSION2=4  
export SESSION3=4
```

This means that partitions 0 and 1 will be restored with two sessions, and partitions 2 and 3 will be restored with four sessions. As a result, this is the restore command:

```
db2_all '<<+0<" ./db2/SID/EEEenv.sh; db2 restore db SID load shared library open  
$SESSION##'  
db2_all '<<-0<" ./db2/SID/EEEenv.sh; db2 restore db SID load shared library open  
$SESSION##'
```

The string `'$SESSION##'` in the `db2_all` command will be replaced while running it with the value provided by the environment variables `SESSION0` to `SESSION3`.

Full online backup

Full online backup of all database partitions, starting with the catalog partition (shown as partition 1) followed by all other partitions in parallel using different Data Protection for SAP profiles for each partition. To support this scenario, one Data Protection for SAP profile (`initSID.ut1`) must be created and maintained for each partition. Each profile can have different settings for the Tivoli Storage Manager node and management class. These are the profiles needed for this example:

- `initSID.ut1.1` (profile for partition 1)
- `initSID.ut1.2` (profile for partition 2)
- `initSID.ut1.3` (profile for partition 3)
- `initSID.ut1.4` (profile for partition 4)

In addition, a Data Protection for SAP vendor environment file (`vendor.env`) must be created and maintained for each partition. Each vendor environment file has an `XINT_PROFILE` entry which refers to the corresponding Data Protection for SAP profile. For example:

```
vendor.env.1 (environment for partition 1 ->  
XINT_PROFILE=/db2/SID/vendor.env.1)  
vendor.env.2 (environment for partition 2 ->  
XINT_PROFILE=/db2/SID/vendor.env.2)  
vendor.env.3 (environment for partition 3 ->  
XINT_PROFILE=/db2/SID/vendor.env.3)  
vendor.env.4 (environment for partition 4 ->  
XINT_PROFILE=/db2/SID/vendor.env.4)
```

As a result, this is the backup command:

```
db2_all '<<+1<' db2 backup db SID online load shared library options
/db2/SID/vendor.env.##'
db2_all '<<-1<'" db2 backup db SID online load shared library options
/db2/SID/vendor.env.##'
```

The string `vendor.env.##` in the `db2_all` command will be replaced while executing it with the values `vendor.env.1` - `vendor.env.4`, whereby the special characters `'##'` will be substituted by the corresponding partition number.

Managing Backup Objects

This is the Backup Object Manager syntax:

`backom [-?]` displays the syntax help.

Note: For the C shell, enclose the option string in quotes (`backom '-?'`).

`backom -h [password|query|backup|restore|delete]` displays the command online help.

```
backom -c command [ command option ...]
```

where 'command' is one of:

for Password: `password`

for Query: `q_all | q_db | q_ts | q_log | q_tdi | q_raw`

for Backup: `b_db`

for Restore: `r_db | r_ts | r_log | r_hfile | r_raw | r_tdi | rr_db_interactive | rr_db_batch | rr_db_clone`

for Delete: `d_db | d_ts | d_log | d_raw`

where 'command option' is one of:

```
-i instance
-a alias name
-n node number
-u userid
-p password -t timestamp | timerange
-l log number | log number range
-k log chain | log chain range
-f file name
-d destination directory
-e execution profile
-b buffer size
-s scaling factor
-N
-S sessions
-B number of buffers
-P parallelism
-D target database
-T tablespace
-R full | incremental | delta
-O
```

-L
-C
-x
-v
-m *output mode*

Backup Object Manager Commands

Details regarding these six types of Backup Object Manager commands are provided:

- Password command
- Query commands
- Backup command
- Restore commands
- Delete commands

Details regarding these commands and their syntax requirements are provided in this section. Optional command options are listed in brackets [], parameter descriptions, that must be replaced, are listed in angle brackets .

Backup Object Manager Command Options

The following options may be specified together with Backup Object Manager commands:

- *-a database alias* or *-a original database alias,target database alias*. Denotes the name of the database for which an operation is requested. In the case of a redirected restore to a different database or of database cloning, the database aliases of both the original and the target databases must be specified and separated by a comma. When a redirected restore is requested that specifies a single database alias, the database is restored to the original database.
- *-b buffer size* Denotes the size of DB2 backup or restore buffers, in 4 KB allocation units (pages). The minimum is eight units. The buffer size is limited by the memory available.
- *-B number of buffers* Denotes the number of DB2 buffers to be used for backup or restore. The minimum number is 2. The number of buffers is limited by the available memory.
- *-C* If specified for a redirected restore, this option indicates that the Backup Object Manager should only run a test of the setup but not start copying data.
- *-d destination directory* Denotes the destination path for restoring a file to the file system.
- *-D target database directory* For a redirected restore, this option denotes the fully qualified name of the target database directory. This command option is ignored when the database alias of the target database is the same as the database alias of the original database.
- *-e execution profile* Denotes the complete path of the Tivoli Storage Manager for ERP for DB2 profile to be used with the Backup Object Manager. This option overrides the profile name set in the XINT_PROFILE environment variable.
- *-f file name* Denotes the name of a file in the file system. Unless the file denotes a TDI image, the following wild card characters are accepted: ? denotes any single character * denotes any number of any characters.

- *-i DB2 instance* Used in query commands to limit the database or tablespace data to be displayed to a specific DB2 instance. With all other commands, this command option is used to override the DB2 instance name defined in the DB_INSTANCE environment variable.
- *-k log chain | log chain range* where *log chain range = chain1 - chain2*. Denotes the log chain number(s) of DB2 log file(s). DB2 log chains can be specified either in the format *Cnnnnnnnn*, where *nnnnnnnn* is a string of 7 decimal digits or in the format *mmmmmmmm*, where *mmmmmmmm* is a string of up to 7 decimal digits denoting the log chain number.
- *-l log number | log number range* where *log number range = log number 1 - log number 2* Denotes the log serial numbers of DB2 log files. DB2 log numbers can be specified either in the format *Snnnnnnn.log* (DB2 log file name), where *nnnnnnn* is a string of 7 decimal digits, or in the format *mmmmmmm*, where *mmmmmmm* is a string of up to 7 decimal digits denoting the log serial number.
- *-L* If specified for a database backup, the DB2 log files are saved to Tivoli Storage Manager with the database backup.
- *-m output mode* where *output mode = short | normal | detailed* Denotes the detail of information requested with a query command. The default is "short" for information related to DB2 log files, "normal" for all other kinds of information. If you need to override the default values generally, you may set environment variables FULL_OUTPUT (for information on database backups), TABLESPACE_OUTPUT (for information on tablespace backups) and LOG_OUTPUT (for information on DB2 log file backups) to the values desired.
- *-N* If specified, this command option causes all containers of a tablespace to be allocated with the same size during a redirected restore.
- *-n node number* Denotes the DB2 node number. For the password command in a DB2 partitioned environment: If for only one DB2 node/partition the password has to be set/changed, specify the command option *-n node number*. If the node/partition number is not specified, the new password is saved to all available node/partitioned based Data Protection for SAP configuration files. For all other commands: If the node number is not specified, node NODE0000 is assumed.
- *-O* If specified when requesting a backup or tablespace restore operation, an online backup or an online tablespace restore is performed.
- *-p password* The password of the user ID specified in option *-u*.
- *-P parallelism* Denotes the degree of parallelism within DB2, i.e. the number of buffer manipulator processes reading from or writing to tablespaces at the same time. The minimum parallelism is 1, the maximum is 1024.
- *-R backup type* where *backup type = full | incremental | delta* . Denotes the type of backup requested. If no backup type is specified a full backup is performed.
- *-S sessions* Denotes the number of I/O sessions that are to be started by DB2. The value of this command option must be less than or equal to the value of the keyword MAX_SESSIONS in the Data Protection for SAP profile.
- *-s scaling factor* Denotes the positive floating point factor to be used for resizing all containers of a tablespace during redirected restore. The default is 1, indicating that the new tablespace is exactly the size of the original.
- *-t timestamp | time range* where *time range = timestamp1-timestamp2* Denotes the time when a backup object was created. For database and tablespace backups, this timestamp matches the timestamp listed in the DB2 Recovery History File. It consists of 14 decimal digits and has the format:

yyyymmddhhmmss where yyyy is the year mm is the month of the year, 01 through 12 dd is the day of the month, 01 through 31 hh is the hour of the day, 00 through 23 mm is the minute of the hour, 00 through 59 ss is the second of the minute, 00 through 59. For restore commands, an **exact** timestamp must be given. For query and delete commands, a time range can be specified, or the timestamp might contain wild card characters. The following wild card characters are accepted:

- ? denotes any single digit
- * denotes any number of any digits.

If a timestamp is not specified for a query, the result will contain all eligible backup object available on the Tivoli Storage Manager server. If a timestamp is not specified for a restore, the newest object is retrieved from Tivoli Storage Manager.

- -T *tablespace list* where *tablespace list* = *tablespace[,tablespace list]*
For a backup request, denotes the names of the tablespace(s) to be backed up. Tablespace names are separated by commas. If there is no tablespace list specified, a full database backup is performed.
- -u *userid* Denotes the DB2 user ID used for backing up or restoring a DB2 database, tablespace, or recovery history file if it is different from the current login user ID.
- -v If set, all log messages will also be displayed on STDOUT.
- -x If specified, this option suppresses all confirmation requests. Otherwise, confirmation requests will be issued for restore commands that would overwrite existing data, and for delete requests.

In conjunction with the "-c password" option, -x causes the password to be changed on all database partitions.

Backup Command (Backup Database Data)

With the Backup Object Manager backup command, you can backup a complete database or selected tablespaces of a database. (For a detailed description of the command options, refer to "Backup Object Manager Command Options" on page 126.)

- Backup the database data denoted by the command options: `backom -c b_db -a database alias [-T tablespace list] [-R backup type] [-i instance] [-n node number] [-u userid] [-p password] [-b buffer size] [-B number of buffers] [-S sessions] [-P parallelism] [-e execution profile] [-O] [-L] [-v]`

Delete Commands (Remove Backup Objects from Tivoli Storage Manager)

The Backup Object Manager delete commands remove backup objects from Tivoli Storage Manager that were sent to Tivoli Storage Manager by Tivoli Storage Manager for ERP for DB2. (For a detailed description of the command options, see "Backup Object Manager Command Options" on page 126.)

- Delete the database backup(s) specified by the command options from Tivoli Storage Manager: `backom -c d_db -a database alias -t timestamp | time range [-i instance] [-n node number] [-e execution profile] [-x] [-v]`
- Delete the tablespace backup(s) specified by the command options from Tivoli Storage Manager: `backom -c d_ts -a database alias -t timestamp | time range [-i instance] [-n node number] [-e execution profile] [-x] [-v]`

- Delete the DB2 log file backup(s) specified by the command options from Tivoli Storage Manager: `backom -c d_log -a database alias -l log number|log number range[-n node number] [-t timestamp|time range] [-e execution profile] [-x] [-v]`
- Delete the file(s) specified by command option `-f` from Tivoli Storage Manager: `backom -c d_raw -f file name [-e execution profile] [-x] [-v]`

Password Command (Verify and Save Tivoli Storage Manager Password)

The password command connects to the backup server, prompts for a new password, and verifies the password entered with the backup server. If the verification is successful, the new password is encrypted and stored in the Data Protection for SAP for DB2 configuration file. Successful password verification requires that the password entered must be the active password for the corresponding node on the Tivoli Storage Manager server. Issue this command to verify and save a Tivoli Storage Manager password:

```
backom -c password [-x] [-a DB2 alias name] [-n DB2 node number]
[-e execution profile]
```

Information regarding when to use the password command is provided in “7. Determine the Tivoli Storage Manager password method” on page 61.

Query Commands (List Backup Objects)

The query commands list backup objects that were sent to Tivoli Storage Manager by Tivoli Storage Manager for ERP for DB2. The objects to be displayed can be filtered by using appropriate command options (see also “Backup Object Manager Command Options” on page 126).

- List all backup objects related to DB2 (database or tablespace backups and DB2 log file backups): `backom -c q_all [-i instance] [-a database alias] [-n node number] [-t timestamp | time range] [-l log number | log number range] [-e execution profile] [-m output mode] [-v]`
- List database backups: `backom -c q_db [-i instance] [-a database alias] [-n node number] [-t timestamp | time range] [-e execution profile] [-m output mode] [-v]`
- List tablespace backups: `backom -c q_ts [-i instance] [-a database alias] [-n node number] [-t timestamp | time range] [-e execution profile] [-m output mode] [-v]`
- List tablespace definition information (TDI) images related to a full DB2 database backup: `backom -c q_tdi -a database alias -t timestamp [-i instance] [-n node number] [-e execution profile] [-m output mode] [-v]`
- List DB2 log file backups: `backom -c q_log [-a database alias] [-n node number] [-t timestamp | time range] [-l log number | log number range] [-e execution profile] [-m output mode] [-v]`
- List backup objects available on Tivoli Storage Manager (database or tablespace backups, DB2 log file backups, and file backups): `backom -c q_raw [-f file name] [-e execution profile] [-m output mode] [-v]`

Restore Commands (Restore Backup Objects)

With the Backup Object Manager restore commands, you can restore any backup object that was created by Tivoli Storage Manager for ERP for DB2. (For a detailed description of the command options, refer to “Backup Object Manager Command Options” on page 126.)

- Restore the database denoted by the command options: `backom -c r_db -a database alias [-n node number] [-u userid] [-p password] [-t timestamp] [-b buffer size] [-B number of buffers] [-S sessions] [-P parallelism] [-R restore type] [-0] [-e execution profile] [-x] [-v]`
- Restore the database denoted by the command options to a different location (redirected restore) in automatic mode: `backom -c rr_db_clone -a original database alias,target database alias [-i instance] [-n node number] [-u userid] [-p password] [-t timestamp] [-b buffer size] [-B number of buffers] [-S sessions] [-P parallelism] [-D target database directory] [-e execution profile] [-s scaling factor] [-N] [-C] [-v]` For a detailed discussion of the redirected restore function read “DB2 Redirected Restore Using Backup Object Manager” on page 84.
- Restore the database denoted by the command options to a different location (redirected restore) in batch mode: `backom -c rr_db_batch -a original database alias,target database alias -f TDI image [-i instance] [-n node number] [-u userid] [-p password] [-t timestamp] [-b buffer size] [-B number of buffers] [-S sessions] [-P parallelism] [-D target database directory] [-e execution profile] [-s scaling factor] [-N] [-C] [-v]` For a detailed discussion of the redirected restore function read “DB2 Redirected Restore Using Backup Object Manager” on page 84.
- Restore the database denoted by the command options to a different location (redirected restore) in interactive mode: `backom -c rr_db_interactive -a original database alias,target database alias [-i instance] [-n node number] [-u userid] [-p password] [-t timestamp] [-f modified TDI image] [-b buffer size] [-B number of buffers] [-S sessions] [-P parallelism] [-D target database] [-e execution profile] [-s scaling factor] [-N] [-C] [-v]` For a detailed discussion of the redirected restore function read “DB2 Redirected Restore Using Backup Object Manager” on page 84.
- Restore the tablespaces denoted by the command options: `backom -c r_ts -a database alias [-n node number] [-u userid] [-p password] [-t timestamp] [-b buffer size] [-B number of buffers] [-S sessions] [-P parallelism] [-R restore type] [-0] [-e execution profile] [-x] [-v]`
- Restore the DB2 Recovery History File of the database denoted by the command options: `backom -c r_hfile -a database alias [-n node number] [-t timestamp] [-u userid] [-p password] [-b buffer size] [-B number of buffers] [-S sessions] [-P parallelism] [-e execution profile] [-x] [-v]`
- Restore the tablespace definition information (TDI) denoted by the command options: `backom -c r_tdi -t timestamp -a database alias [-d destination directory] [-e execution profile] [-x] [-v]`
- Retrieve the DB2 log files denoted by the command options: `backom -c r_log -a database alias -l log number|log number range -d destination directory [-n node number] [-t timestamp|time range] [-e execution profile] [-x] [-v]`
- Retrieve the file(s) specified by command option `-f` to the path specified by command option `-d`:
`backom -c r_raw -f file name -d destination directory [-e execution profile] [-x] [-v]`

This command can restore data to the destination directory. If a single segment was used during the backup, the data can be restored to DB2 from the destination directory after retrieval. If two or more segments were used during the backup, the data can be restored to the destination directory but cannot be restored to DB2.

BACKOM command examples

These examples show which commands can be used to perform certain tasks:

- Issue the following command to verify and save a Tivoli Storage Manager password:

```
backom -c password
```

- Issue the following command to create a list of all available backup objects sent to Tivoli Storage Manager by Tivoli Storage Manager for ERP for DB2:

```
backom -c q_all
```

- Issue the following command to create a list of all DB2 log files for database SAMPLE, with a log number greater than 123, and created in November 2002, with normal output detail level:

```
backom -c q_log -a SAMPLE -l 124-9999999 -t 200211* -m normal  
backom -c q_log -a SAMPLE -l S0000124.log-S9999999.log -t 200211* -m normal
```

- Issue the following command to create a list of DB2 log files for log chains 5 to 15 for database SAMPLE with log numbers from 98 to 180, archived between 4 p.m. and 8.30 p.m.:

```
backom -c q_log -a SAMPLE -k C0000005-C0000015 -l 98-180 -t ????????16*-????????2030*
```

- Issue the following command to create a list of all tablespace backups for partition NODE0001, of database SAMPLE, that were created in November 2002 between 4 p.m. and 5 p.m.:

```
backom -c q_ts -a SAMPLE -n NODE0001 -t 200211??16*
```

- Issue the following command to backup online database SAMPLE using two I/O sessions and four backup buffers:

```
backom -c b_db -a SAMPLE -S 2 -B 4 -0
```

- Issue the following command to backup the tablespaces SYSCATSPACE and USERSPACE1 of database SAMPLE, using the execution profile 'initSAMPLE.utl' located at /db2/SAMPLE/config:

```
backom -c b_db -a SAMPLE -T SYSCATSPACE,USERSPACE1 -e /db2/SAMPLE/config/initSAMPLE.utl
```

- Issue the following command to restore a tablespace of database SAMPLE with the tablespace backup created on November 27, 2002, at 6:32:15 p.m.:

```
backom -c r_ts -a SAMPLE -t 20021127183215
```

- Issue the following command to restore database SAMPLE with the latest backup:

```
backom -c r_db -a SAMPLE
```

- Issue the following command to delete all DB2 log files for database SAMPLE that were created before June 2002:

```
backom -c d_log -a SAMPLE -t 1900*-20020601000000
```

- Issue the following command to delete all versions of files containing "tmp" in their path or file names that were sent to Tivoli Storage Manager by Tivoli Storage Manager for ERP:

```
backom -c d_raw -f *tmp*
```

UNIX or Linux Crontab Example

UNIX or Linux cron jobs can be scheduled with the crontab command. This command launches an editing session that allows you to create a crontab file. The cron jobs and the appropriate times are defined within the crontab. The crontab can be customized with this command:

```
crontab -e
```

In this example, a cron job starts the shell script backup.ksh at 11:30 p.m. Monday through Friday and uses DB2 backup to back up the SAP® database. This is the entry in the crontab that starts the script for this scenario:

```
30 23 * * 1,2,3,4,5 /usr/bin/su - db2c21 -c "/db2/C21/sapscripsts/backup.ksh"
```

The content of backup.ksh is available in “Full Offline Backup Shell Script Sample” on page 71.

Crontab File Sample

```
# -----  
# crontab.sample:  
# Sample crontab file to be included in the root crontab jobs.  
# -----  
# Task:  
# Submits backup/archive commands at regularly scheduled intervals  
# using two simple shell scripts containing backup/archive commands  
# and TSM commands.  
# -----  
#          *****      NOTE          *****      NOTE          *****      NOTE          *****  
#  
#          This file is intended only as a model and should be  
#          carefully tailored to the needs of the specific site.  
#  
#          *****      NOTE          *****      NOTE          *****      NOTE          *****  
# -----  
#  
# Remarks on the crontab file format:  
#  
# Each crontab file entry consists of a line with six fields, separated  
# by spaces and tabs, that contain, respectively:  
#   o The minute (0 through 59)  
#   o The hour (0 through 23)  
#   o The day of the month (1 through 31)  
#   o The month of the year (1 through 12)
```

```

# o The day of the week (0 through 6 for Sunday through Saturday)
# o The shell command
# Each of these fields can contain the following:
# o A number in the specified range
# o Two numbers separated by a dash to indicate an inclusive range
# o A list of numbers separated by commas
# o An * (asterisk); meaning all allowed values
#
# -----
#
# For the following examples, the system id (alias) of the DB2 database
# is assumed to be 'C21' and the username 'db2c21'.
#
# -----
# Full database backup, scheduled every Friday at 8:00 p.m.
#
0 20 * * 5
  /usr/bin/su - db2c21 -c "/db2/C21/sql1lib/scripts/backup.ksh"
#
# -----
# Save redo logs, scheduled twice a day at 11:30 a.m. and at 5:30 p.m.
# Monday through Friday
#
30 11,17 * * 1,2,3,4,5
  /usr/bin/su - db2c21 -c "/db2/C21/sql1lib/scripts/archive.ksh"

```

The Data Protection for SAP for DB2 Profile

The Data Protection for SAP for DB2 profile provides keyword parameters that customize how Data Protection for SAP operates. A sample profile `initSID.utl` is provided on the product media. During installation on Windows systems, the sample profile (along with all other files) is placed in the `C:\Program Files\Tivoli\TDP4SAP` directory.

The profile is copied to the profile path (during installation) if no other profile exists there. Data Protection for SAP reads the profile pointed to by environment variable `XINT_PROFILE` (shared library, BackOM) or sent as a parameter (BackOM) immediately prior to a backup or restore operation.

These rules apply to the keyword syntax:

- Each line is analyzed separately.
- Keywords can start in any column of the line.
- Keywords must not be preceded by any string, except blanks.
- If a keyword is encountered several times, the last one is used.
- File processing ends when the `END` keyword is encountered or the end of file is reached.
- The comment symbol is the pound sign (`#`). Scanning of the current line stops when the comment symbol is encountered. No comment is allowed between the keyword and the value(s). For example:

```

#BRARCHIVEMGTCLASS  MLOG1                <-- correct
BRARCHIVEMGTCLASS  MLOG1 #                <-- correct
BRARCHIVEMGTCLASS # MLOG1                <-- incorrect

```

- Although some keywords are required, most are optional. Each of the optional keywords has a preset default value.
- Additional profile information is provided in “Enable ProLE on Windows to access configuration files on a remote share” on page 29.

Tivoli Storage Manager for ERP for DB2 profile parameter descriptions

The default value is underlined in these descriptions and applies if the parameter is not specified.

ADSMNODE *node_name*

Specifies a *node_name* that is registered to the Tivoli Storage Manager server as a Tivoli Storage Manager node. This parameter must be defined in conjunction with the respective SERVER statement, as shown in the sample profile. You can assign a different node name to your database system with this option. It is used if you have several SAP® for DB2 database systems in your network with the same name, for example, *SID*, and they all use the same Tivoli Storage Manager server. This keyword must not be set when automated password handling is selected. It should be set for manual password handling as described in “7. Determine the Tivoli Storage Manager password method” on page 61.

BACKEND *pgmname* [*parameterlist*]

Specifies a program *pgmname* that is called by IBM Tivoli Storage Manager for Enterprise Resource Planning for DB2 after the backup function completed and before program control is returned to DB2. If *pgmname* is not a fully qualified path, the default search path is used to locate the program. If not specified, no backend processing is done.

Example for UNIX or Linux:

```
BACKEND write operator@remotesite Backup of SAP database object completed.
```

This sends a message to a remote user when the backup has finished.

BACKUPIDPREFIX *6-charstring* | SAP___

Specifies a six-character prefix that is used to create a backup identifier for each archived object. If not specified, the default value is SAP___. All partitions of a partitioned DB2 database have the same BACKUPIDPREFIX.

BRARCHIVEMGTCLASS *management_class* [*management_class...*]

Specifies the Tivoli Storage Manager management class(es) that Tivoli Storage Manager for ERP uses to back up offline DB2 log files. Each parameter string can consist of up to 30 characters. Specify a separate BRARCHIVEMGTCLASS for each log file copy requested. As a result, make sure the number of different BRARCHIVE management classes specified must be greater than or equal to the number of log file copies (keyword REDOLOG_COPIES on page “Tivoli Storage Manager for ERP for DB2 profile parameter descriptions.” This parameter must be defined with the respective SERVER statement, as shown in the sample profile.

To use different Tivoli Storage Manager servers for backup and archive data, the value ':SKIP:' can be used to define a server stanza with no archive management classes. This value is allowed for the parameter BRARCHIVEMGTCLASS only.

BRBACKUPMGTCLASS *management_class* [*management_class...*]

Specifies the Tivoli Storage Manager management class(es) Tivoli Storage Manager for ERP uses to back up the DB2 database. The parameter string can consist of up to 30 characters. This parameter must be defined with the respective SERVER statement, as shown in the sample profile.

BUFFCOPY SIMPLE | PREVENT | AUTO

This optional parameter controls how Tivoli Storage Manager for ERP uses

the internal buffers for transferring data during a backup. If set to SIMPLE, data buffers are copied when they are sent between Tivoli Storage Manager components. This is the default. If set to PREVENT, the original data buffers are sent between Tivoli Storage Manager components. For this mode, BUFFSIZE is restricted to a maximum of 896 KB. Furthermore, it cannot be selected when the Tivoli Storage Manager client encryption or client compression features are activated. If set to AUTO, Tivoli Storage Manager for ERP runs in PREVENT mode whenever the configuration supports it. Otherwise, SIMPLE mode is automatically selected. This parameter has no effect on restore operations.

BUFFSIZE *n* | **131072**

This parameter specifies the block size (in bytes) for the buffers used when communicating with DB2. The size of the buffers sent to the Tivoli Storage Manager API is the value of BUFFSIZE increased by approximately 20 bytes. The valid range is from 4096 (4 KB) to 32 MB. Inappropriate values are adjusted automatically. If BUFFCOPY is set to PREVENT, the value of BUFFSIZE must not exceed 896 KB. If not specified, the default value is 131072 (128 KB) for UNIX or Linux systems and 32768 (32 KB) for Windows systems. In most cases, these values are appropriate. If you plan to increase the size of internal buffers make sure that sufficient storage is available. The number of buffers acquired by Tivoli Storage Manager for ERP correlates to the number of sessions (keyword SESSIONS). By activating RL_COMPRESSION, the number of buffers is doubled.

CONFIG_FILE *path/SID.bki*

Specifies the configuration file *initSID.bki* for Tivoli Storage Manager for ERP to store all variable parameters such as passwords and the date of the last password change. During processing, the string %DB2NODE is replaced automatically by the current DB2 node of a partitioned database or by 'NODE0000' otherwise. This parameter is required.

END Specifies the end of the parameter definitions. Tivoli Storage Manager for ERP stops searching the file for keywords when END is encountered.

FRONTEND *pgmname* [*parameterlist*]

Specifies a program *pgmname* that is called by Tivoli Storage Manager for ERP in a backup run before the connection to the Tivoli Storage Manager server is established. If *pgmname* is not a fully qualified path, the default search path is used to find the program. If not specified, no frontend processing is not performed.

Example for UNIX or Linux:

```
FRONTEND write operator@remotesite Backup of SAP database  
object is starting.
```

This sends a message to a remote user before backup begins.

LOG_SERVER *servername* [*verbosity*]

The *servername* value specifies the name of the Tivoli Storage Manager server to which log messages are sent. The *servername* must match one of the servers listed in a SERVER statement in order for Tivoli Storage Manager for ERP messages to be logged in the Tivoli Storage Manager server activity log. The *verbosity* value can be one of these specifications: ERROR, WARNING, or DETAIL. This value determines which messages are sent. The default value is WARNING, which means that error and warning messages are sent. ERROR sends only error messages. DETAIL sends all message types (errors, warnings, and informational messages). If

there is no LOG_SERVER statement in the profile, log messages are not sent to any of the Tivoli Storage Manager servers.

MAX_SESSIONS *n*

Specifies the maximum number of parallel Tivoli Storage Manager client sessions that Tivoli Storage Manager for ERP establishes. For a direct backup or restore on tape drives, the number of sessions must be less than or equal to the number of tape drives available for the backup. Make sure that the `mountlimit (mountl)` parameter in the device class is set to the number of available tape drives. Make sure that the `maxumps` parameter of the node is set to the number of available tape drives. The value of keyword MAX_SESSIONS must be less than or equal to the sum of the SESSIONS values specified in the SERVER statements of the currently available servers.

MAX_VERSIONS *n* | 0

The *n* value defines the maximum number of database backup versions to be kept in backup storage. The default setting for this value is 0, meaning that backup version control is disabled. Every time a full backup completes successfully, the version count is increased by an increment of 1 and stored in the Tivoli Storage Manager for ERP configuration file. This value is also assigned to the tablespace files and to all subsequent DB2 log file backups. If the number of versions kept in backup storage is larger than the specified maximum number of backup versions (stored by the parameter MAX_VERSIONS), the oldest version is deleted, together with the corresponding tablespace, incremental and log file backups until only the specified maximum number of most recent versions remain. For partitioned DB2 databases, backup version control is done on a partition basis. Therefore, full backups should always be initiated for all partitions at the same time, for example by the DB2 script `db2_a11`. For details on the `db2_a11` script, see your DB2 documentation. Also, consider these characteristics:

- When Tivoli Storage Manager for ERP deletes an old full backup, all partial backups older than this full backup are also deleted.
- If the backups are distributed over multiple Tivoli Storage Manager servers and one of the servers is temporarily unavailable at the time of a new full backup, it will not be possible to find all the backup versions. This may result in retaining a backup that would otherwise have been deleted.
- Every database partition needs its own configuration file. Partitions of a partitioned database should have the same BACKUPIDPREFIX.

Tivoli Storage Manager uses the value of the RETVER parameter (specified when defining a copy group) to give files an expiration date. Use only one of these methods to control how long you keep backups:

- If you use Tivoli Storage Manager for ERP backup version control, you need to bypass this expiration function. Set the Tivoli Storage Manager parameter `RETVER=9999` so that the files are not considered expired and are not deleted by Tivoli Storage Manager.
- If you use the Tivoli Storage Manager expiration function, you need to turn off the maximum number of full database backup backup version control. Deactivate Tivoli Storage Manager for ERP backup version control by setting `MAX_VERSIONS=0`.

Information about defining a copy group is available in “4. Define a policy” on page 60.

PASSWORDREQUIRED NO | YES

Specifies whether Tivoli Storage Manager requires a password to be supplied by the Tivoli Storage Manager client. This depends on the Tivoli Storage Manager installation. If not specified, the default is **PASSWORDREQUIRED YES** which implements manual password handling. This parameter must be defined in conjunction with the respective **SERVER** statement, as shown in the sample profile. Further details are described in “7. Determine the Tivoli Storage Manager password method” on page 61.

REDOLOG_COPIES *n* | 1

Specifies the number of copies Tivoli Storage Manager for ERP stores for each processed DB2 log file. The valid range is from 1 to 9. If not specified, Tivoli Storage Manager for ERP stores one copy of each log file. The number of different management classes for archived logs (keyword **BRARCHIVEMGTCLASS** specified must be greater than or equal to the number of log file copies specified. The number of different management classes specified must be greater than or equal to the number of log file copies specified.

RL_COMPRESSION NO | YES

If set to **YES**, Tivoli Storage Manager for ERP performs a null block compression of the data before they are sent over the network. Although RL compression introduces additional CPU load, throughput can be improved when the network is the bottleneck. It is not recommended to use RL compression together with the Tivoli Storage Manager API compression. If not specified, the default value is **NO** meaning null block compression is not performed. **RL_COMPRESSION** is only performed if a full database backup was started. The offline log files are not compressed.

SEGMENTSIZe *size*[GB | TB]

This keyword specifies the maximum size of the segments that are split from large backup objects. For details on the impact of using this parameter please refer to Segmenting Large Backup Objects. The required *size* value must be a positive integer equal to or greater than 1. Consider these characteristics when specifying this parameter:

- The scale units (GB or TB) are not required. GB is the default value.
- When specifying the scale units (GB or TB), you can use lower case letters, upper case letters, or a combination of both cases. However, you cannot specify single-character abbreviations (G or T).
- When this parameter is not specified, one backup object per DB2 backup session is transferred to Tivoli Storage Manager.
- If the specified segment size is less than the DB2 block size specified during the backup, the specified segment size is ignored at runtime. The specified DB2 backup block size is used instead.

The following example sets the maximum size of the backup objects segments on Tivoli Storage Manager to 100 GB:

```
SEGMENTSIZ  100 GB
```

You can also specify the command in the following way:

```
SEGMENTSIZ  100
```

SERVER *servername*

This keyword specifies the name of the Tivoli Storage Manager server to which Tivoli Storage Manager for ERP backups are to be stored. This statement begins a server section in the Tivoli Storage Manager for ERP profile. At least one server section is required. Server sections are located at

the end of the profile. A server section ends before a following SERVER keyword, before the END keyword, or at the end of the profile. These dependent keywords are applicable in a server section:

- AD SMNODE
- BRARCHIVEMGTCLASS
- BRBACKUPMGTCLASS
- PASSWORDREQUIRED
- SESSIONS
- TCP_ADDRESS
- USE_AT

The server name must be defined in the Tivoli Storage Manager profiles `dsm.sys` (UNIX and Linux) or `servername.opt` (for Windows). In order to set up alternate or parallel paths, each path is denoted by its own logical server name and corresponding server section, although these logical names refer to the same server. In this case, the Tivoli Storage Manager profiles specify the same TCP/IP address for these server names. In order to set up alternate or parallel servers, each server is represented by one or more server statements and the corresponding server sections (depending on the number of paths to the server). In this case, the Tivoli Storage Manager profiles specify different TCP/IP addresses for the different servers. Different server names result in different server entries in the Administration Assistant View Tivoli Storage Manager Server Utilization function while identical server names are considered to point to the same Tivoli Storage Manager server even if they are specified in different Tivoli Storage Manager for ERP profiles throughout the system landscape. Do NOT use any profile keywords, AD SM, or TSM as the servername.

SESSIONS *n* | 1

The *n* value specifies the number of parallel sessions Tivoli Storage Manager for ERP uses for the server. This keyword is required in every server section. This parameter must be defined in conjunction with the respective SERVER statement, as shown in the sample profile.

TCP_ADDRESS | IP *address of server*

Specifies the IP address of the Tivoli Storage Manager server in dotted decimal notation. This parameter overrides the value for the parameter TCPSERVERADDRESS in the Tivoli Storage Manager client system options file (`dsm.sys`) on UNIX or Linux or in the client options file (`servername.opt`) on Windows. This parameter must be defined in conjunction with the respective SERVER statement, as shown in the sample profile.

TRACE FILEIO_MIN | FILEIO_MAX | COMPR_MIN | COMPR_MAX | MUX_MIN | MUX_MAX | TSM_MIN | TSM_MAX | ASYNC_MIN | ASYNC_MAX | APPLICATION_MIN | APPLICATION_MAX | SYSCALL_MIN | SYSCALL_MAX | COMM_MIN | COMM_MAX | DEADLOCK_MIN | DEADLOCK_MAX | PROLE_MIN | PROLE_MAX | BLAPI_MIN | BLAPI_MAX | SOCKET_DATA | ALL | OFF

This parameter writes trace information to the file specified with the TRACEFILE parameter. Arguments to TRACE can be any combination of the possible components and levels separated by spaces. A trace will only be written if both TRACE and TRACEFILE are specified. Do not use this parameter unless instructed to use it by Tivoli Storage Manager for ERP support. Using it can significantly deteriorate the performance of Tivoli Storage Manager for ERP.

TRACEFILE *path*

Specifies the name and location of the trace file for Tivoli Storage Manager for ERP to store all trace information. When TRACE is used, *path* specifies the full path and the name of file. If the value of TRACEFILE contains the string %BID, this string is replaced by the backup ID to get the path and name of the trace file actually used. For example, specifying /tmp/%BID.trace will yield a trace file /tmp/myBackup.trace for backup ID myBackup. A trace will only be written if both TRACE and TRACEFILE are specified.

TRACEMAX *n*

Specifies the maximum size of the trace file in KB. The valid range is 4096 (4MB) to unlimited. If not specified, the trace file size is unlimited.

USE_AT *days*

Specifies the days that the Tivoli Storage Manager server (specified with the corresponding SERVER keyword) is used. The *days* value can be numbers from 0 (Sunday) to 6 (Saturday). Multiple numbers can be used when separated by spaces. If not specified, the default is to use the Tivoli Storage Manager server on all days. This parameter must be defined with the respective SERVER statement, as shown in "Example of SERVER statement with alternate servers" on page 42. The parameter USE_AT has no effect on actions other than backup.

Sample Tivoli Storage Manager for ERP for DB2 Profile for UNIX or Linux

The sample profile (initSID.utl) is included in the Tivoli Storage Manager for ERP for DB2 installation package. Although the UNIX, Linux, and Windows versions are similar, all example versions are provided.

```
#-----
#
# Data Protection for SAP (R) interface for DB2 UDB
#
# Sample profile for Data Protection for SAP (R) Version 6.2
#
#-----
#
# See the 'Data Protection for SAP (R) Installation &
# User's Guide' for a full description.
#
# For a comment symbol the character '#' can be used.
# Everything following this character will be interpreted as comment.
#
# Data Protection for SAP (R) accesses its profile
# in "read only" mode. All variable parameters like passwords, date of
# last password change, current version number will be written into the file
# specified with the CONFIG_FILE parameter. The passwords will be encrypted.

#-----
# Prefix of the 'Backup ID' which is stored in the description field of
# the Tivoli Storage Manager archive function.
# Maximum 6 characters.
# Default: none.
#-----
BACKUPIDPREFIX SID___

#-----
# Number of parallel sessions to be established.
```

```

# Note: This number must not exceed the number of tape drives simultaneously
# available to the node on the Tivoli Storage Manager servers to be accessed.
# The valid range of MAX_SESSIONS is from 1 and 32.
# Default: none.
#-----
MAX_SESSIONS 1 # Tivoli Storage Manager client sessions

#-----
# Number of backup copies of the DB2 log files.
# The valid range of REDOLOG_COPIES is from 1 to 9.
# Default: 1.
#-----
#REDOLOG_COPIES 2

#-----
# Specifies the block size for disk I/O (in bytes).
# The default values have been chosen from our performance experiments in
# standard hardware environments.
# The valid range of BUFFSIZE is from 4KB to 32MB.
# Default: 131072 (128 KB) on UNIX, 32768 (32 KB) on Windows.
#-----
BUFFSIZE 131072 # block size in bytes

#-----
# This optional parameter controls how Data Protection for SAP(R) uses
# the internal buffers for transferring data during a backup.
# Valid values: SIMPLE | PREVENT | AUTO
# Default: SIMPLE
#-----
#BUFFCOPY AUTO

#-----
# Name of a program to be called before the backup task is started.
# Default: none.
#-----
#FRONTEND pgmname parameterlist

#-----
# Name of a program to be called after the backup task is completed.
# Default: none.
#-----
#BACKEND pgmname parameterlist

#-----
# Maximum number of data base backup versions to be kept.
# Note: Version control by Data Protection for SAP (R) is only activated
# only activated if the parameter MAX_VERSION is not 0.
# The valid range of MAX_VERSIONS is from 0 to 9999.
# A value of 0 means no versioning.
# Default: 0, no versioning.
#-----
#MAX_VERSIONS 4

#-----
# Specifies whether a null block compression of the data is to be performed
# before transmission to Tivoli Storage Manager.
# Although RL compression introduces additional CPU load, throughput can be
# improved when the network is the bottleneck. RL compression in Data
# Protection for SAP(R) should not be used together with
# Tivoli Storage Manager API compression.

```

```

# Default: NO
#-----
#RL_COMPRESSION YES # NO is default

#-----
# Controls generation of a trace file.
# Note: We recommend using the trace function only in cooperation with
# Data Protection for SAP (R) support.
# Default: OFF.
#-----
#TRACE OFF
#TRACEFILE /db2/C21/sqllib/log/tdpr3.trace

#-----
# Denotes the maximum size of the trace file in KB.
# If not specified, the trace file size is unlimited.
#-----
#TRACEMAX max. size # trace file size in KB

#-----
# Specify the full path of the configuration file.
# Default: none.
#-----
CONFIG_FILE /db2/C21/sqllib/%DB2NODE/initSID.bki

#-----
# Denotes if Data Protection for SAP (R) shall send error/status
# information to a Tivoli Storage Manager server.
# The servername must match one of the servers listed in a SERVER statement.
# Valid values for verbosity are ERROR | WARNING | DETAIL.
# Default: none.
#-----
#LOG_SERVER servername [verbosity]
#LOG_SERVER server_a ERROR

*****
# Statement for servers and paths.
# Multiple servers may be defined.
*****

SERVER server_a # Servername, as defined in dsm.sys
SESSIONS 2 # Maximum number of sessions
# to server_a
PASSWORDREQUIRED YES # Use a password
ADSMNODE NODE # Tivoli Storage Manager Nodename
BRBACKUPMGTCCLASS MDB # Mgmt-Classes for database backup
BRARCHIVEMGTCCLASS MLOG1 MLOG2 # Mgmt-Classes for redo log backup
# TCP_ADDRESS 192.168.1.1 # IP address of network interface
# on server_a
# USE_AT 0 1 2 3 4 5 6 # Overrides IP address of dsm.sys
# Days when server_a is used for
# backup
*****
# USE_AT : 0=Su 1=Mo 2=Tu 3=We 4=Th 5=Fr 6=Sa
# The valid range of USE_AT is from 0 to 6.
# Default: all days
*****

#SERVER server_b # Servername, as defined in dsm.sys
# SESSIONS 2 # Maximum number of sessions
# to server_b
# PASSWORDREQUIRED YES # Use a password

```

```

# ADSMNODE          NODE          # Tivoli Storage Manager Nodename
# BRBACKUPMGTCCLASS  MDB          # Mgmt-Classes for database backup
# BRARCHIVEMGTCLASS  MLOG1 MLOG2  # Mgmt-Classes for redo log backup
# TCP_ADDRESS        192.168.1.1  # IP address of network interface
#                   # on server_b
#                   # Overrides IP address of dsm.sys
# USE_AT             0 1 2 3 4 5 6  # Days when server_b is used for
#                   # backup
*****
# USE_AT : 0=Su 1=Mo 2=Tu 3=We 4=Th 5=Fr 6=Sa
# Default: all days
*****

#-----
# End of profile

END

```

Sample Data Protection for SAP for DB2 Profile for Windows

```

#-----
#
# Data Protection for SAP (R) interface for DB2 UDB
#
# Sample profile for Data Protection for SAP (R)
# Version 6.2 for Windows 2000/2003
#
#-----
#
# See the 'Data Protection for SAP (R) Installation & User's Guide' for
# a full description.
#
# For a comment symbol the character '#' can be used.
# Everything following this character will be interpreted as comment.
#
# Data Protection for SAP (R) accesses its profile in "read only" mode.
# All variable parameters like passwords, date of last password
# change, current version number will be written into the file specified
# with the CONFIG_FILE parameter. The passwords will be encrypted.

#-----
# Prefix of the 'Backup ID' which is used for communication with the
# SAP® BR*Tools and stored in the description field of the
# Tivoli Storage Manager archive function.
# Must be 6 characters.
# Default: none.
#-----
BACKUPIDPREFIX  SID___

#-----
# Number of parallel sessions to be established.
# Note: This number must not exceed the number of tape drives simultaneously
# available to the node on the Tivoli Storage Manager servers to be accessed.
# The valid range of MAX_SESSIONS is from 1 and 32.
# Default: none.
#-----
MAX_SESSIONS  1 # Tivoli Storage Manager client sessions

#-----
# Number of backup copies of the DB2 log files.
# The valid range of REDOLOG_COPIES is from 1 to 9.

```

```

# Default: 1.
#-----
#REDOLOG_COPIES 2

#-----
# Specifies the block size for disk I/O (in bytes).
# The default values have been chosen from our performance experiments in
# standard hardware environments.
# The valid range of BUFFSIZE is from 4KB to 32MB.
# Default: 131072 (128 KB) on UNIX, 32768 (32 KB) on Windows.
#-----
BUFFSIZE 32768          # block size in bytes

#-----
# This optional parameter controls how Data Protection for SAP(R) uses
# the internal buffers for transferring data during a backup.
# Valid values:  SIMPLE | PREVENT | AUTO
# Default: SIMPLE
#-----
#BUFFCOPY              AUTO

#-----
# Name of a program to be called before the backup task is started.
# Default: none.
#-----
#FRONTEND              pgmname parameterlist

#-----
# Name of a program to be called after the backup task is completed.
# Default: none.
#-----
#BACKEND              pgmname parameterlist

#-----
# Maximum number of data base backup versions to be kept.
# Note: Version control by Data Protection for SAP (R) is
# only activated if the parameter MAX_VERSION is not 0.
# The valid range of MAX_VERSIONS is from 0 to 9999.
# Default: 0
#-----
#MAX_VERSIONS 4

#-----
# Specifies whether a null block compression of the data is to be performed
# before transmission to Tivoli Storage Manager.
# Although RL compression introduces additional CPU load, throughput can be
# improved when the network is the bottleneck. RL compression in Data
# Protection for SAP(R) should not be used together with
# Tivoli Storage Manager API compression.
# Default: NO
#-----
#RL_COMPRESSION YES

#-----
# Controls generation of a trace file.
# Note: We recommend using the trace function only in cooperation with
# Data Protection for SAP (R) support.
# Default: OFF
#-----
#TRACE OFF

```

```

#TRACEFILE                c:\sqllib\tdp_r3\log\tdpr3.trace

#-----
# Denotes the maximum size of the trace file in KB.
# If not specified, the trace file size is unlimited.
#-----
#TRACEMAX                max. size                # trace file size in KB

#-----
# Specify the full path of the configuration file.
# Default: none.
#-----
CONFIG_FILE                c:\sqllib\tdp_r3\%DB2NODE\initSID.bki

#-----
# Denotes if Data Protection for SAP (R) shall send
# error/status information to a Tivoli Storage Manager server.
# The servername must match one of the servers listed in a SERVER statement.
# Valid values for verbosity are ERROR | WARNING | DETAIL.
# Default: none.
#-----
#LOG_SERVER                servername                [verbosity]
#LOG_SERVER                server_a                ERROR

#*****
# Statement for servers and paths.
# Multiple servers may be defined.
#*****

SERVER                server_a                # Servername, as defined in dsm.sys
SESSIONS                2                # Maximum number of sessions
                        # to server_a
PASSWORDREQUIRED        YES                # Use a password
ADSMNODE                NODE                # Tivoli Storage Manager Nodename
BRBACKUPMGTCCLASS        MDB                # Mgmt-Classes for database backup
BRARCHIVEMGTCCLASS        MLOG1 MLOG2        # Mgmt-Classes for redo log backup
# TCP_ADDRESS                192.168.1.1                # IP address of network interface
                        # on server_a
                        # Overrides IP address of dsm.sys
# USE_AT                0 1 2 3 4 5 6                # Days when server_a is used for
                        # backup

#*****
# USE_AT : 0=Su 1=Mo 2=Tu 3=We 4=Th 5=Fr 6=Sa
# The valid range of USE_AT is from 0 to 6.
# Default: all days
#*****

#SERVER                server_b                # Servername, as defined in dsm.sys
# SESSIONS                2                # Maximum number of sessions
                        # to server_b
# PASSWORDREQUIRED        YES                # Use a password
# ADSMNODE                NODE                # Tivoli Storage Manager Nodename
# BRBACKUPMGTCCLASS        MDB                # Mgmt-Classes for database backup
# BRARCHIVEMGTCCLASS        MLOG1 MLOG2        # Mgmt-Classes for redo log backup
# TCP_ADDRESS                192.168.1.1                # IP address of network interface
                        # on server_b
                        # Overrides IP address of dsm.sys
# USE_AT                0 1 2 3 4 5 6                # Days when server_b is used for
                        # backup

#*****
# USE_AT : 0=Su 1=Mo 2=Tu 3=We 4=Th 5=Fr 6=Sa
# Default: all days
#*****

```

```
#-----
# End of profile

END
```

Defining the Custom SQL file

Note: The custom SQL file is intended to be implemented or modified only by IBM support personnel with a detailed knowledge of the process involved and the internal Administration Assistant function for Data Protection for SAP database. This section does not discuss this process in detail.

The custom SQL file must be named customSQLFile.txt and placed in the installation directory (or folder) of the Administration Assistant. For example:
C:\Program Files\tdpr3assi\customSQLFile.txt

The custom SQL file contains this structure:

```
# CUSTOM SQL FILE Comment

sqlSQL statement/sqldescription ... /param
sqlSQL statement/sqldescription ... /param
...
```

As an aid to explaining the entry structure, it is shown in the following with each tag set in a separate line:

```
sqlSQL statement/sql
descriptionDescription of the SQL statement/description
programid0/programid
actionid0/actionid
displaygroup1,3/displaygroup
backuptype2/backuptype
executionmode0/executionmode
paramparameter-value1/param
paramparameter-value2/param
...
paramparameter-valuen/param
```

Each entry must be coded in a single line.

The tag definitions are as follows:

Table 13. Contents of the Custom SQL File

Tag	Definition
#	Comment line
sql	An SQL statement that defines which data is to be sent. Note: 1. Only SELECT statements will be executed. 2. A semicolon at the end of the line is not permitted. 3. The maximum line length is 400 characters.
description	Description of the SQL statement (maximum length: 300 characters)

Table 13. Contents of the Custom SQL File (continued)

Tag	Definition
<i>programid</i>	Specifies the program that handles the result of the SQL statement. <ul style="list-style-type: none"> • <i>programid</i> 0: Administration Assistant
<i>actionid</i>	Defines the way the result will be handled, depending on the <i>programid</i> (currently, the only value for <i>actionid</i> is 0): <ul style="list-style-type: none"> • (<i>programid</i> 0: Administration Assistant): Send e-mail when threshold exceeded (SQL statement returns data)
<i>displaygroup</i>	List of display group IDs separated by commas, or "ALL" for all display groups.
<i>system</i>	List of system IDs separated by commas, or "ALL" for all systems.
<i>backuptype</i>	List of backup types separated by commas, or "ALL" for all backup types. <ul style="list-style-type: none"> • 0: Archive • 1: Partial backup • 2: Incremental backup • 3: Full backup
<i>executionmode</i>	<i>executionmode</i> sets the time the entry will be performed (i.e., the SQL statement issued): <ul style="list-style-type: none"> • 0: Entry will be performed after each backup run • 1: Entry will be performed periodically
<i>param</i>	Parameters needed by the programs. The number of parameters depends on the selected program and action. Multiple parameters are coded using repeating <i>param/param</i> tag pairs. <ul style="list-style-type: none"> • (<i>programid</i> 0: Administration Assistant): <ul style="list-style-type: none"> – One parameter, consisting of the e-mail address list (separated by semicolons)

Consider these facts about the custom SQL file:

- Each entry in the file must be on a single line.
- If *executionmode* is 1, the *system*, *displaygroup*, and *backuptype* tags are ignored, and the SQL statement will be executed periodically.
- If *executionmode* is 0, the SQL statement will be executed after the backup completes, but only if the *system* tag matches the system on which the backup was performed, or the *displaygroup* tag matches the displaygroup the system belongs to. Furthermore, the *backuptype* tag must match the backup type of the backup performed.
- The *system* and *displaygroup* tags are mutually exclusive.
- The custom SQL file will be reloaded periodically by the Administration Assistant Server component. The server does not need to be restarted.

Defining Thresholds Using the Custom SQL File

A custom threshold can be defined in the custom SQL file. The corresponding entry has the following values for the indicated tags:

Table 14. Tags for Defining Thresholds in the Custom SQL File

Tag	Value
<i>sql</i>	An SQL statement that will return data when the threshold is exceeded.
<i>programid</i>	0 (Administration Assistant)
<i>actionid</i>	0 (send e-mail when threshold exceeded)
<i>executionmode</i>	1 (run periodically)
<i>param</i>	(Optional) One or more e-mail addresses, separated by semicolons. If no e-mail address is given, only a panel indication is given that the threshold has been exceeded. Note: Multiple e-mail addresses are given in a single <i>param/param</i> tag pair, not in multiple pairs.

Sample Custom SQL File

This is a sample of a custom SQL file.

```
# CUSTOM SQL FILE FOR THE ADMINISTRATION ASSISTANT
#
# This file should only be changed by an IBM Employee
# After the changes you have to check this file using CustomSQLFilecheck
#
# NOTE: Each entry must be coded in one line. The multi-line format
# shown below is for illustration purposes only.
#
# Sample threshold definition: backup size > 500 GB, display group 1, backup type 2
#
sqlselect * from AdminAssistant.tsmrun where amount > 500000000000/sql
descriptionAmount over 500 GB/description
programid0/programid
actionid0/actionid
displaygroup1/displaygroup
backuptype2/backuptype
executionmode0/executionmode
paramemailAdress@email.com/param
#
```

Data Protection for SAP for DB2 files and samples

Use these file samples to assist with Data Protection for SAP for DB2 operations.

Sample Shell Script for Scheduling a Report from a UNIX Scheduling Client

The scheduledReport.sh file is provided in the Data Protection for SAP for DB2 package and is copied to the Administration Assistant function for Data Protection for SAP installation path.

```
#-----  
#  
# Tivoli Storage Manager for ERP. Data Protection for SAP for DB2  
#  
# Sample command file for the Administration Assistant scheduling client  
#  
#-----  
#      *****   NOTE      *****   NOTE      *****   NOTE      *****  
#  
#          This script is provided as a model and should be  
#          carefully tailored to the needs of the specific site.  
#  
#      *****   NOTE      *****   NOTE      *****   NOTE      *****  
#-----  
export CLASSPATH=/reporting/Admt.jar:$CLASSPATH  
export PATH=/usr/bin:$PATH  
java -classpath $CLASSPATH com.ibm.bkit.schedulerIF.Sched_Main xxx.xxx.xxx.xxx...  
... 1099 myReport ADMIN admin directory=/myreports log=/tmp/reportlogs
```

Sample Command File for Scheduling a Report from a Windows Scheduling Client

The scheduledReport.cmd file is provided in the Data Protection for SAP for DB2 package and is copied to the Administration Assistant function for Data Protection for SAP Server component installation path.

```
#-----  
#  
# Tivoli Storage Manager for ERP. Data Protection for SAP for DB2  
#  
# Sample command file for the Administration Assistant scheduling client  
#  
#-----  
#      *****   NOTE      *****   NOTE      *****   NOTE      *****  
#  
#          This script is provided as a model and should be  
#          carefully tailored to the needs of the specific site.  
#  
#      *****   NOTE      *****   NOTE      *****   NOTE      *****  
#-----  
set CLASSPATH=C:\ProgramFiles\reporting\Admt.jar  
set PATH=C:\Program Files\IBM\Java142\jre\bin;%PATH%  
java -cp %CLASSPATH% com.ibm.bkit.schedulerIF.Sched_Main xxx.xxx.xxx.xxx ...  
... 1099 myReport ADMIN admin directory=C:\reports log=C:\reportlogs
```

Client User Options File Sample (dsm.opt) UNIX and Linux

```
*****
* Tivoli Storage Manager                                     *
*                                                         *
* Sample Client User Options file for Unix platforms      *
*****

SErvername      server_a
Replac          On
Tapeprompt     No
DOM            /usr/sap /sapmnt/C21 /usr/sap/trans /db2/C21
```

Client User Options File Sample (dsm.opt) Windows

Data Protection for SAP for DB2 requires a client options file dsm.opt to be present in the location indicated by environment variable DSMI_CONFIG. The specific options used by Data Protection for SAP for each server however are taken from files *server.opt* residing in the same path.

```
*****
*
* DSM.OPT (for Data Protection for SAP (R) )
*
* This file is intentionally left empty. It must be present in the location
* indicated by environment variable DSMI_CONFIG. The specific options used
* by Data Protection for SAP for each server however are taken from files
* server.opt residing in the same path.
*
* Please note: This client options file is not meant to be used by other
*               TSM clients.
*
*****
```

Client System Options File Sample (dsm.sys)

```
*****
* IBM Tivoli Storage Manager                               *
*                                                         *
* Sample Client System Options file for Unix platforms    *
*****

SErvername      server_a
COMMmethod      TCPip
TCPport        1500
TCPserveraddress your_ITSM_server_1
TCPbuffsize    32
TCPwindowsize  24
Compression    Off
InclExcl       /usr/lpp/adsm/bin/inclExcl.list

SErvername      server_b
COMMmethod      TCPip
TCPport        1500
TCPserveraddress your_ITSM_server_2
TCPbuffsize    32
TCPwindowsize  24
Compression    Off
InclExcl       /usr/lpp/adsm/bin/inclExcl.list
```

Include/Exclude List Sample (UNIX and Linux)

```
* -----
* incl excl.list:
* Sample include/exclude list
* -----
* Task:
* Include/Exclude list of files and directories for TSM incremental backups
* -----
*      *****      NOTE      *****      NOTE      *****      NOTE      *****
*
*      This file is intended only as a model and should be
*      carefully tailored to the needs of the specific site.
*
*      *****      NOTE      *****      NOTE      *****      NOTE      *****
* -----
*
* For all UNIX systems
*
* exclude /unix
* exclude /.../core
* exclude /u/.../*.sh_history
* exclude /home/.../*.sh_history
*
* Note: It is recommended to perform system backups on a regular
*       basis. Consequently, you can exclude at least the following
*       directories:
*
* exclude /usr/games/.../*
* exclude /usr/bin/.../*
* exclude /usr/sbin/.../*
* exclude /usr/mibin/.../*
* exclude /usr/sbin/.../*
* -----
*
* For those using AFS, exclude the cache filesystem or file
*
* exclude /usr/vice/cache/*
* exclude /var/vice/cache/*
* or
* exclude /afscfs
* -----
*
* This stuff is either not worthwhile to be included or should be backed up
* using DB2 backup techniques and the SAP utility brarchive.
*
* exclude /db2/C21/log_archive/C21/*
* exclude /db2/C21/sapreorg/.../* (There may be important scripts
*                               located, check it out and decide.)
* exclude /db2/C21/sapdata*/.../*
* exclude /db2/C21/sapraw*/.../*
* -----
*
* With the above include/exclude list we implicitly include everything not
* excluded above. Especially for DP for SAP, this means including:
*
* /sapmnt/C21      > 300 MB
* /usr/sap         > 50 MB
* /db2/C21        > 200 MB
* and UNIX related > 350 MB
* -----
```

Include/Exclude List Sample (Windows)

This sample include/exclude list is intended for the standard client user option file. The purpose is to exclude files that are easy to restore or that are already saved by Data Protection for SAP for DB2 from routine Tivoli Storage Manager incremental backups. Typically, such files are Windows system files and DB2 database files.

```
*****
* This Include-Exclude list is used for incremental backups of file
* systems by the Tivoli Storage Manager command-line backup client.
* Therefore the name of this file has to be set under the keyword InclExcl
* in the standard Tivoli Storage Manager client user option file "dsm.opt".
*
* Since the backup of the DB2 database is done by
* Data Protection for SAP(R) and not by Tivoli Storage
* Manager command-line backup client, the DB2 database should be excluded
* from backups by the Tivoli Storage Manager command-line backup client.
*
* Note 1:
* The environment variable DSM_CONFIG contains the full file name of
* the Tivoli Storage Manager client user option file "dsm.opt".
* Note 2:
* This Include-Exclude is not used by Data Protection for SAP(R)
*
*****
Exclude *:...\*.swp
Exclude *:...\*.obj
Exclude *:...\*.csm
Exclude *:...\*.dsk
Exclude *:...\*.bak
Exclude *:...\win386.swp
Exclude *:...\386spart.par
Exclude *:...\pagefile.sys
Exclude *:...\*.par
Exclude *:...\SYSTEM32\CONFIG\*.
Exclude *:...\SYSTEM32\CONFIG\...\*
Exclude *: \IBMBIO.COM
Exclude *: \IBMDOS.COM
*
*Exclude the following DB2 database files:
*
Exclude *: \db2\C21\log_archive\C21\...\*
Exclude *: \db2\C21\sapreorg\...\*
Exclude *: \db2\C21\sapdata*\...\*
```

Client Options Files Sample (*server.opt*)

Data Protection for SAP for DB2 requires a corresponding client option file *server.opt* for each Tivoli Storage Manager server. These files must reside in the same directory. This directory must also contain the client options file *dsm.opt*, which is specified in the environment variable *DSMI_CONFIG*. The contents of this (second) *dsm.opt* file is ignored by Data Protection for SAP.

```
*****
*
* SERVER.OPT
*
* Data Protection for SAP (R) obtains the necessary information about
* a Tivoli Storage Manager server 'server' from a client option file
* called 'server.opt'. For each Tivoli Storage Manager server a
* corresponding client option file is required.
*
* Note: This file contains the client options for the Tivoli Storage Manager
* server called 'server_a'.
```

```

*
* Please see the Tivoli Storage Manager documentation for details.
*
*****
COMMethod      TCP/IP
COMpression    OFF
*NODENAME      C21
TCPport        1500
TCPserveraddress xxx.xxx.xxx.xxx
PASSWORDACCESS PROMPT
TCPBUFFSIZE    31
TCPWINDOWSIZE  32

```

Sample DB2 Vendor Environment File

A DB2 vendor environment file (`vendor.env`) is created from the information entered in the installation dialog panels during installation. A sample DB2 vendor environment file is included in the Data Protection for SAP for DB2 installation package.

Note

Ensure that there are no blanks within the paths specified for the vendor-specific environment variables of the vendor environment file. DB2 is currently unable to handle embedded blanks. Note that in the case of a standard Windows installation, the Data Protection for SAP profile is located at

```
c:\Program Files\Tivoli\tsm\tdp_r3\db264\initSID.utl
```

Sample DB2 vendor environment file for UNIX or Linux:

```
XINT_PROFILE=/db2/C21/tdpr3/initC21.utl
TDP_DIR=/db2/C21/tdpr3/tdplog
BACKOM_LOCATION=/usr/tivoli/tsm/tdp_r3/db264/backom
```

Sample DB2 vendor environment file for Windows:

```
XINT_PROFILE=c:\db2\C21\tdpr3\initC21.utl
TDP_DIR=c:\db2\C21\tdpr3\tdplog
BACKOM_LOCATION=c:\tivoli\tsm\tdp_r3\db264\backom.exe
```

Data Protection for SAP for DB2 planning sheets

Uses these planning sheets to assist with installing and configuring Data Protection for SAP for DB2.

Data Protection for SAP for DB2 (base product) planning sheet

Collect the information in this planning sheet before installing Data Protection for SAP for DB2. This table is also provided in file form as `planning_sheet_db2` for UNIX and Linux and `planning_sheet_db2.txt` for Windows.

Table 15. Installation Parameters for Data Protection for SAP

UNIX or Linux	Windows	Installation Parameter
X	X	DB2 database SID:
X	X	Tivoli Storage Manager server name or IP address:
X	X	Tivoli Storage Manager node name: Tivoli Storage Manager node configured on the Tivoli Storage Manager server named for the backup of the SID denoted above. For details, refer to “5. Register a node” on page 61.
X	X	Tivoli Storage Manager management classes for database and log file backups. Management classes configured for the database backup and for the backup of log files. For details, refer to “4. Define a policy” on page 60. Default: MDB for database backups, MLOG1 and MLOG2 for log file backups.
	X	Path where the Tivoli Storage Manager API resides (contents of environment variable DSMI_DIR): Default: C:\Program Files\Common Files\tivoli\TSM\api64
	X	Path to client option file of Tivoli Storage Manager (contents of environment variable DSMI_CONFIG). For details refer to the Tivoli Storage Manager documentation.
	X	Path to Tivoli Storage Manager log files (contents of environment variable DSMI_LOG): The Tivoli Storage Manager API will create the file dserror.log< in this path. For details, refer to the Tivoli Storage Manager documentation. Default: C:\temp
	X	Installation path for Data Protection for SAP executable files: C:\Program Files\Tivoli\TSM\tdp_r3\db264
X	X	Options: <ul style="list-style-type: none"> • Use of the Administration Assistant (see “Administration Assistant function for Data Protection for SAP” on page 9 and Table 16 on page 154). The Administration Assistant should be installed prior to Data Protection for SAP so that the interface between the two can be automatically established. • Use of DB2 Log Manager for log archiving.

Administration Assistant function for Data Protection for SAP planning sheet

Collect the information in this planning sheet before installing the Administration Assistant function for Data Protection for SAP. This table is also provided in file form as `planning_sheet_aa` for UNIX and Linux, and `planning_sheet_aa.txt` for Windows.

Table 16. Installation Parameters for the Administration Assistant function for Data Protection for SAP

Installation Option	Installation Parameter
Installation type.	Decision as to whether the Administration Assistant is to be installed on a single host (typical installation) or distributed across multiple hosts (custom installation). Default: Single-host
Server/client communication mode	Decision as to whether the Administration Assistant Server component and clients communicate in nonsecure mode via HTTP or secure mode via HTTPS. Default: Nonsecure
Database type	Decision as to which DBMS the Administration Assistant should use. Select either the installation of the bundled Apache Derby package or the use of an existing Apache Derby or IBM DB2 installation. Default: Installation of Apache Derby as bundled with product.
Data migration	If you want to migrate data from an existing Administration Assistant environment, enter the directory containing the *.aa files. Default: No migration.
Software language	Decision as to whether to install only the English version of the program or all national language versions. Default: English-only

Table 16. Installation Parameters for the Administration Assistant function for Data Protection for SAP (continued)

Installation Option	Installation Parameter
<p>Parameters applying to the Administration Assistant Server component</p>	<p>Hostname or IP address: Default: Hostname of current system</p> <p>Port number for Data Protection for SAP for DB2 (ProLE) connect. This port number must be made known to all instances of Data Protection for SAP that are to be managed and monitored by this Server component instance. Default: 5126</p> <p>Port number for client connect in non-secure mode (HTTP). Default: 80</p> <p>Port number for client connect in secure mode (HTTPS). Default: 443</p> <p>RMI registry port number Default: 1099</p> <p>Port number for performance data from Database Agent Default: 5129</p> <p>Port number for communication with Database Agent Default: 5128</p>
<p>Parameters applying to the Administration Assistant Database Agent component</p>	<p>Hostname or IP address: Default: Hostname of current system</p> <p>Port number for Data Protection for SAP (ProLE) connection Default: 5125</p> <p>Port number for communication with Administration Assistant Server component Default: 5127</p>
<p>Parameters applying to the Administration Assistant Database component (Apache Derby)</p>	<p>Hostname or IP address: Default: Hostname of current system</p> <p>Port number for database connect Default: 1527</p> <p>User ID and password to access internal database.</p>

Table 16. Installation Parameters for the Administration Assistant function for Data Protection for SAP (continued)

Installation Option	Installation Parameter
<p>Parameters applying to the Administration Assistant Database component (IBM DB2)</p>	<p>Hostname or IP address: Default: Hostname of current system</p> <p>Port number for database connect Default: 50000</p> <p>User ID and password of the system user for which the DB2 instance should be installed that the internal database should access.</p>
<p>Installation directory</p>	<p>Installation directory (on each host) Default: /opt/tivoli/tsm/tdp_r3_assist on UNIX and Linux, or C:\Program Files\tdpr3assi on Windows.</p>
<p>Product Support</p>	<p>Location of mail.jar (Java Mail)</p>
<p>Product Support</p>	<p>Location of activation.jar (Java Beans Activation Framework):</p>
<p>History file</p>	<p>History file directory (on Server component host) Default: history (in installation directory)</p> <p>History file retention time (days). Can be changed via the Administration Assistant client. Default: 14</p>

Table 16. Installation Parameters for the Administration Assistant function for Data Protection for SAP (continued)

Installation Option	Installation Parameter
Secure Communication	<p>Information on the public key infrastructure (PKI):</p> <ul style="list-style-type: none"> • <i>Keystore name</i>. Keystore containing the private and public keys of the Administration Assistant Server component when running in secure mode. If you do not yet have a public key infrastructure, the keystore can be created during the installation process. • <i>Keystore password</i>. Password ensuring the consistency of the keystore. The server's key pair must be protected by the same password. • <i>Truststore name</i>. Truststore containing a set of trusted certificates. When running in secure mode, the Administration Assistant's server certificate must be verified against this truststore when the server is started. • <i>Truststore password</i>. Password ensuring the consistency of the truststore. This is only required if a trusted certificate needs to be imported into the truststore during the installation process. • <i>Certificate file</i>. Path of the certificate file in case you already have a server certificate issued by a certificate authority. • <i>Certificate creation information</i>. Information on the X.500 distinguished name (common name, organizational unit, organization name, locality name, state name, and country code) and on the validity period required in case a new self-signed certificate is to be created during the installation process. For details on this information, refer to the X.500 and X.509 standards. • <i>New certificate file name</i>. If the public key of a newly created server key pair needs to be distributed to client machines it will be exported to this file. • <i>CSR file name</i>. If the newly created server key pair will be used to request a certificate signed by a Certificate Authority, the Certificate Signing Request will be written to this file.
Internal database managed by DB2	<p>DB2 JDBC Universal Driver. The corresponding packages are bundled with your IBM DB2 installation.</p> <ul style="list-style-type: none"> • db2jcc.jar location:Default: None • db2jcc_license_cu.jar location:Default: None <p>The Administration Assistant database is enabled for automatic storage and has a set of one or more associated storage paths. Enter at least one disk or path that DB2 is allowed to assign and allocate for its table space containers.</p> <p>Default: None</p> <p>The name of the internal database is predefined and cannot be changed.</p> <p>Default: AADB</p>

Table 16. Installation Parameters for the Administration Assistant function for Data Protection for SAP (continued)

Installation Option	Installation Parameter
Internal database managed by Apache Derby	Database directory: Default: aaDBSupport (in installation directory)
	Name of the internal database Default: 'adminAssistant'
	Retention time for data in database (days). (To save this data, the backup facilities offered by Derby can be used.) Default: 175
Documentation	Option: English-only or all languages
	Default: English-only

Tips for network settings

Helpful information to assist with adjusting your network is provided.

Network Settings of the Tivoli Storage Manager

The performance adjustments for Tivoli Storage Manager are performed by editing these configuration files:

- Tivoli Storage Manager server option file `dsmserv.opt`
- Tivoli Storage Manager backup-archive client option file `dsm.sys` (UNIX and Linux systems) or `server.opt` (Windows systems).

This table shows the corresponding Tivoli Storage Manager configuration file attributes with the recommended values.

Table 17. Tuning Tivoli Storage Manager Configuration File Attributes

Attributes	Value	Description
TCPBuffsize	32	Specifies the size, in kilobytes, of the buffer used for TCP/IP send requests. This option affects whether or not Tivoli Storage Manager sends the data directly from the session buffer or copies the data to the TCP buffer. A 32K buffer size forces Tivoli Storage Manager to copy data to its communication buffer and flush the buffer when it fills.
TCPNODElay	YES	Specifies whether the server should send small amounts of data or allow TCP/IP to buffer the data. Disallowing buffering may improve throughput but more packets will be sent over the network.
TCPWindowSize	640 (AIX) 63 (others)	Specifies the size, in kilobytes, which will be used for the TCP/IP sliding window for the client node. This is the size of the buffer used when sending or receiving data. The range of values is 0 to 2048.

Additional information can be found at: <http://www-306.ibm.com/software/tivoli/products/storage-mgr-erp/>.

Networks with Large Bandwidth-Delay Product

For networks with a large bandwidth-delay product, it is recommended to activate the TCP enhancements as specified in RFC1323. For example, the network on an AIX machine can be configured with the `no` command. This command sets or displays current network attributes in the kernel. Details about the `no` command are available in the `man` page of `no` of your operating system.

This table shows the network attributes with their recommended values:

Table 18. Tuning of Network Settings

Attributes	Value	Description
<code>rfc1323</code>	1	Enables TCP enhancements as specified by RFC 1323, TCP Extensions for High Performance. The default is 0. A value of 1 specifies that all TCP connections will attempt to negotiate the RFC enhancements.
<code>sb_max</code>	131072	Specifies the maximum buffer size allowed for a socket. The default is 65536 bytes. From the point of view of performance recommendations, the <code>sb_max</code> value should be twice the <code>TCPWindow size</code> set within the Tivoli Storage Manager configuration file <code>dsm.sys</code> .

Set these values issuing these commands by the root user on the appropriate machine:

```
no -o rfc1323=1
no -o sb_max=131072
```

The `no` command does not perform range checking. It therefore accepts all values. If used incorrectly, the command might cause the system to become inoperable. These changes will be lost at system reboot. To make changes permanent, edit the `/etc/rc.net` file.

SP Switch (RISC 6000)

If an SP switch (RISC 6000) is used, the following two values should be set as shown in this table:

Table 19. Tuning of SP Switch Buffer Pools

Attributes	Value	Description
<code>rpoolsz</code>	1048576	The receive pool is a buffer pool for incoming data. The size for values is in bytes.
<code>spoolsz</code>	1048576	The send pool is a buffer for outgoing data. The size for values is in bytes.

The buffer pool settings can be changed using the `chgcsc` command. After the changes, it is necessary to reboot the node.

Appendix. Accessibility features for the Tivoli Storage Manager product family

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

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

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

The Tivoli Storage Manager Information Center, and its related publications, are accessibility-enabled. The accessibility features of the information center are described at http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/topic/com.ibm.help.ic.doc/iehs36_accessibility.html.

Keyboard navigation

On Windows, the Tivoli Storage Manager product family follows Microsoft conventions for all keyboard navigation and access. Drag and Drop support is managed 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).

On other operating systems, these products follow the operating-system conventions for keyboard navigation and access.

Vendor software

The Tivoli Storage Manager product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who want to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
U.S.A.*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample

programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

SAP and SAP NetWeaver are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

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

This glossary includes terms and definitions for IBM Tivoli Storage Manager and IBM Tivoli Storage FlashCopy Manager products.

To view glossaries for other IBM products, go to <http://www.ibm.com/software/globalization/terminology/>.

The following cross-references are used in this glossary:

- *See* refers the reader from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- *See also* refers the reader to a related or contrasting term.

A

absolute mode

In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup even if the file has not changed since the last backup. See also *modified mode*.

access control list (ACL)

In computer security, a list associated with an object that identifies all the subjects that can access the object and their access rights. For example, an access control list is associated with a file that identifies the users who can access that file and their access rights.

access mode

An attribute of a storage pool or a storage volume that specifies whether the server can write to or read from the storage pool or storage volume. The access mode can be read/write, read-only, or unavailable. Volumes in primary storage pools can also have an access mode of destroyed. Volumes in copy storage pools can also have an access mode of offsite.

acknowledgment

The transmission of acknowledgment characters as a positive response to a data transmission.

ACL See *access control list*.

activate

To validate the contents of a policy set and then make it the active policy set.

active-data pool

A named set of storage pool volumes that contain only active versions of client backup data.

active file system

A file system to which space management has been added. With space management, tasks for an active file system include automatic migration, reconciliation, selective migration, and recall. Contrast with *inactive file system*.

active policy set

The activated policy set that contains the policy rules in use by all client nodes that are assigned to the policy domain. See also *policy domain* and *policy set*.

active version

The most recent backup copy of a file stored. The active version of a file cannot be deleted until a backup process detects that the user has either replaced the file with a newer version or has deleted the file from the file server or workstation. Contrast with *inactive version*.

activity log

A log that records normal activity messages that are generated by the server. These messages include information about server and client operations, such as the start time of sessions or device I/O errors.

adaptive subfile backup

A type of backup that sends only changed portions of a file to the server, instead of sending the entire file. Adaptive subfile backup reduces network traffic and increases the speed of the backup.

administrative client

A program that runs on a file server, workstation, or mainframe that administrators use to control and monitor the Tivoli Storage Manager server. Contrast with *backup-archive client*.

administrative command schedule

A database record that describes the

planned processing of an administrative command during a specific time period. See also *client schedule*.

administrative privilege class

See *privilege class*.

administrative session

A period of time during which an administrator user ID communicates with a server to perform administrative tasks. Contrast with *client node session*.

administrator

A user who is registered to the server as an administrator, and who is authorized to perform tasks and issue commands through the assignment of an administrative privilege class.

Advanced Program-to-Program Communication (APPC)

An implementation of the SNA LU 6.2 protocol that allows interconnected systems to communicate and share the processing of programs.

agent node

A client node that has been granted proxy authority to perform operations on behalf of another client node, which is the target node.

aggregate

An object, stored in one or more storage pools, consisting of a group of logical files that are packaged together. See also *logical file* and *physical file*.

aggregate data transfer rate

A performance statistic that indicates the average number of bytes that were transferred per second while processing a given operation.

APPC See *Advanced Program-to-Program Communication*.

application client

A program that is installed on a system to protect an application. The Tivoli Storage Manager server provides backup services to an application client.

archive

To copy programs, data, or files to other storage media, usually for long-term storage or security. Contrast with *retrieve*.

archive copy

A file or group of files that was archived to server storage.

archive copy group

A policy object containing attributes that control the generation, destination, and expiration of archived files.

archive-retention grace period

The number of days that the storage manager retains an archived file when the server is unable to rebind the file to an appropriate management class. See also *bind*.

association

(1) The defined relationship between a client node and a client schedule. An association identifies the name of a schedule, the name of the policy domain to which the schedule belongs, and the name of a client node that performs scheduled operations.

(2) On a configuration manager, the defined relationship between a profile and an object such as a policy domain. Profile associations define the configuration information that is distributed to a managed server when it subscribes to the profile.

audit

To check for logical inconsistencies between information that the server has and the actual condition of the system. The storage manager can audit information about items such as volumes, libraries, and licenses. For example, when a storage manager audits a volume, the server checks for inconsistencies between information about backed-up or archived files that are stored in the database and the actual data that are associated with each backup version or archive copy in server storage.

authentication

The process of checking a user's password before permitting user access to the Tivoli Storage Manager server. Authentication can be turned on or off by an administrator with system privilege.

authentication rule

A specification that another user can use to either restore or retrieve files from storage.

authority

The right to access objects, resources, or functions. See also *privilege class*.

authorization rule

A specification that permits another user to either restore or retrieve a user's files from storage.

authorized user

A user who has administrative authority for the Tivoli Storage Manager client on a workstation. This user changes passwords, performs open registrations, and deletes file spaces.

AutoFS

See *automounted file system*.

automatic detection

A feature that detects, reports, and updates the serial number of a drive or library in the database when the path from the local server is defined.

automatic migration

The process that is used to automatically move files from a local file system to storage, based on options and settings that are chosen by a root user on a workstation. See also *threshold migration* and *demand migration*.

automatic reconciliation

The process that is used to reconcile file systems at regular intervals. The intervals are set by a user with root user authority. See also *reconciliation*.

automounted file system (AutoFS)

A file system that is managed by an automounter daemon. The automounter daemon monitors a specified directory path, and automatically mounts the file system to access data.

B**backup-archive client**

A program that runs on a workstation or file server and provides a means for users to back up, archive, restore, and retrieve files. Contrast with *administrative client*.

backup copy group

A policy object containing attributes that control the generation, destination, and expiration of backup versions of files. A backup copy group belongs to a management class.

backup-retention grace period

The number of days the storage manager retains a backup version after the server is unable to rebind the file to an appropriate management class.

backup set

A portable, consolidated group of active versions of backup files that are generated for a backup-archive client.

backup set collection

A group of backup sets that are created at the same time and which have the same backup set name, volume names, description, and device classes. The server identifies each backup set in the collection by its node name, backup set name, and file type.

backup version

A file or directory that a client node backed up to server storage. More than one backup version can exist in server storage, but only one backup version is the active version. See also *active version* and *inactive version*.

bind To associate all versions of a file with a management class name. See *rebind*.

bindery

A database that consists of three system files for a NetWare server. The files contain user IDs and user restrictions.

C

cache To place a duplicate copy of a file on random access media when the server migrates a file to another storage pool in the hierarchy.

cache file

A snapshot of a logical volume created by Logical Volume Snapshot Agent. Blocks are saved immediately before they are modified during the image backup and their logical extents are saved in the cache files.

CAD See *client acceptor*.

central scheduler

A function that permits an administrator to schedule client operations and administrative commands. The operations can be scheduled to occur periodically or on a specific date. See *client schedule* and *administrative command schedule*.

client A software program or computer that requests services from a server.

client acceptor

An HTTP service that serves the applet for the web client to web browsers. On Windows systems, the client acceptor is installed and run as a service. On AIX, UNIX, and Linux systems, the client acceptor is run as a daemon, and is also called the *client acceptor daemon (CAD)*.

client acceptor daemon (CAD)

See *client acceptor*.

client domain

The set of drives, file systems, or volumes that the user selects to back up or archive data, using the backup-archive client.

client node

A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

client node session

A session in which a client node communicates with a server to perform backup, restore, archive, retrieve, migrate, or recall requests. Contrast with *administrative session*.

client options file

An editable file that identifies the server and communication method, and provides the configuration for backup, archive, hierarchical storage management, and scheduling.

client option set

A group of options that are defined on the server and used on client nodes in conjunction with client options files.

client-polling scheduling mode

A method of operation in which the client queries the server for work. Contrast with *server-prompted scheduling mode*.

client schedule

A database record that describes the planned processing of a client operation during a specific time period. The client operation can be a backup, archive, restore, or retrieve operation, a client operating system command, or a macro. See also *administrative command schedule*.

client/server

Pertaining to the model of interaction in

distributed data processing in which a program on one computer sends a request to a program on another computer and awaits a response. The requesting program is called a client; the answering program is called a server.

client system-options file

A file, used on AIX, UNIX, or Linux system clients, containing a set of processing options that identify the servers to be contacted for services. This file also specifies communication methods and options for backup, archive, hierarchical storage management, and scheduling. This file is also called the *dsm.sys* file. See also *client user-options file*.

client user-options file

A file that contains the set of processing options that the clients on the system use. The set can include options that determine the server that the client contacts, and options that affect backup operations, archive operations, hierarchical storage management operations, and scheduled operations. This file is also called the *dsm.opt* file. For AIX, UNIX, or Linux systems, see also *client system-options file*.

closed registration

A registration process in which only an administrator can register workstations as client nodes with the server. Contrast with *open registration*.

collocation

The process of keeping all data belonging to a single-client file space, a single client node, or a group of client nodes on a minimal number of sequential-access volumes within a storage pool. Collocation can reduce the number of volumes that must be accessed when a large amount of data must be restored.

collocation group

A user-defined group of client nodes whose data is stored on a minimal number of volumes through the process of collocation.

commit point

A point in time when data is considered consistent.

Common Programming Interface for Communications (CPI-C)

A call-level interface that provides a consistent application programming interface (API) for applications that use program-to-program communications. CPI-C uses LU 6.2 architecture to create a set of interprogram services that can establish and end a conversation, send and receive data, exchange control information, and notify a partner program of errors.

communication method

The method by which a client and server exchange information. See also *Transmission Control Protocol/Internet Protocol*.

communication protocol

A set of defined interfaces that permit computers to communicate with each other.

compression

A function that removes repetitive characters, spaces, or strings of characters from the data being processed and replaces the repetitive characters with control characters. Compression reduces the amount of storage space that is required for the data.

configuration manager

A server that distributes configuration information, such as policies and schedules, to managed servers according to their profiles. Configuration information can include policy and schedules. See also *managed server* and *profile*.

conversation

A connection between two programs over a session that allows them to communicate with each other while processing a transaction.

copy backup

A full backup in which the transaction log files are not deleted so that backup procedures that use incremental or differential backups are not disrupted.

copy group

A policy object containing attributes that control how backup versions or archive copies are generated, where backup versions or archive copies are initially

located, and when backup versions or archive copies expire. A copy group belongs to a management class. See also *archive copy group*, *backup copy group*, *backup version*, and *management class*.

copy storage pool

A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class (for space-managed files). See also *primary storage pool* and *destination*.

CPI-C See *Common Programming Interface for Communications*.

D**daemon**

A program that runs unattended to perform continuous or periodic functions, such as network control.

damaged file

A physical file in which Tivoli Storage Manager has detected read errors.

data access control mode

A mode that controls whether a command can access a migrated file, see a migrated file as zero-length, or receive an input/output error if it attempts to access a migrated file. See also *execution mode*.

database backup series

One full backup of the database, plus up to 32 incremental backups made since that full backup. Each full backup that is run starts a new database backup series. A number identifies each backup series.

database snapshot

A complete backup of the entire database to media that can be taken off-site. When a database snapshot is created, the current database backup series is not interrupted. A database snapshot cannot have incremental database backups associated with it. See also *database backup series*. Contrast with *full backup*.

data deduplication

A method of reducing storage needs by eliminating redundant data. Only one instance of the data is retained on storage

media. Other instances of the same data are replaced with a pointer to the retained instance.

data manager server

A server that collects metadata information for client inventory and manages transactions for the storage agent over the local area network. The data manager server informs the storage agent with applicable library attributes and the target volume identifier.

data mover

A device that moves data on behalf of the server. A network-attached storage (NAS) file server is a data mover.

data storage-management application-programming interface (DSMAPI)

A set of functions and semantics that can monitor events on files, and manage and maintain the data in a file. In an HSM environment, a DSMAPI uses events to notify data management applications about operations on files, stores arbitrary attribute information with a file, supports managed regions in a file, and uses DSMAPI access rights to control access to a file object.

default management class

A management class that is assigned to a policy set. This class is used to govern backed up or archived files when a file is not explicitly associated with a specific management class through the include-exclude list.

deduplication

See *data deduplication*.

demand migration

The process that is used to respond to an out-of-space condition on a file system for which hierarchical storage management (HSM) is active. Files are migrated to server storage until space usage drops to the low threshold that was set for the file system. If the high threshold and low threshold are the same, one file is migrated.

desktop client

The group of backup-archive clients that includes clients on Microsoft Windows, Apple, and Novell NetWare operating systems.

destination

A copy group or management class attribute that specifies the primary storage pool to which a client file will be backed up, archived, or migrated.

device class

A named set of characteristics that are applied to a group of storage devices. Each device class has a unique name and represents a device type of disk, file, optical disk, or tape.

device configuration file

(1) For a server, a file that contains information about defined device classes, and, on some servers, defined libraries and drives. The information is a copy of the device configuration information in the database.

(2) For a storage agent, a file that contains the name and password of the storage agent, and information about the server that is managing the SAN-attached libraries and drives that the storage agent uses.

device driver

A program that provides an interface between a specific device and the application program that uses the device.

disaster recovery manager (DRM)

A function that assists in preparing and using a disaster recovery plan file for the server.

disaster recovery plan

A file that is created by the disaster recovery manager (DRM) that contains information about how to recover computer systems if a disaster occurs and scripts that can be run to perform some recovery tasks. The file includes information about the software and hardware that is used by the server, and the location of recovery media.

domain

A grouping of client nodes with one or more policy sets, which manage data or storage resources for the client nodes. See *policy domain* or *client domain*.

DRM See *disaster recovery manager*.

DSMAPI

See *data storage-management application-programming interface*.

dynamic serialization

A type of copy serialization in which a file or folder is backed up or archived on the first attempt regardless of whether it changes during a backup or archive.

E

EA See *extended attribute*.

EB See *exabyte*.

EFS See *Encrypted File System*.

Encrypted File System (EFS)

A file system that uses file system-level encryption.

enterprise configuration

A method of setting up servers so that the administrator can distribute the configuration of one of the servers to the other servers, using server-to-server communication. See also *configuration manager*, *managed server*, *profile*, and *subscription*.

enterprise logging

The process of sending events from a Tivoli Storage Manager server to a designated event server. The event server routes the events to designated receivers, such as to a user exit. See also *event*.

error log

A data set or file that is used to record error information about a product or system.

estimated capacity

The available space, in megabytes, of a storage pool.

event (1) An administrative command or a client operation that is scheduled to be run using Tivoli Storage Manager scheduling.

(2) A message that an Tivoli Storage Manager server or client issues. Messages can be logged using Tivoli Storage Manager event logging.

event record

A database record that describes actual status and results for events.

event server

A server to which other servers can send events for logging. The event server routes the events to any receivers that are enabled for the sending server's events.

exabyte (EB)

For processor storage, real and virtual storage, and channel volume, 1 152 921 504 606 846 976 bytes. For disk storage capacity and communications volume, 1 000 000 000 000 000 000 bytes.

exclude

The process of identifying files in an include-exclude list. This process prevents the files from being backed up or migrated whenever a user or schedule enters an incremental or selective backup operation. A file can be excluded from backup and space management, backup only, or space management only.

exclude-include list

See *include-exclude list*.

execution mode

A mode that controls the space-management related behavior of commands that run under the **dsmmode** command.

expiration

The process by which files, data sets, or objects are identified for deletion because their expiration date or retention period has passed.

expiring file

A migrated or premigrated file that has been marked for expiration and removal from storage. If a stub file or an original copy of a premigrated file is deleted from a local file system, or if the original copy of a premigrated file is updated, the corresponding migrated or premigrated file is marked for expiration the next time reconciliation is run.

extend

To increase the portion of available space that can be used to store database or recovery log information.

extended attribute (EA)

Names or value pairs that are associated with files or directories. There are three classes of extended attributes: user attributes, system attributes, and trusted attributes.

extent The part of a file that is created during the data-deduplication process. Extents are compared with other file extents to identify duplicates.

external library

A type of library that is provided by Tivoli Storage Manager that permits LAN-free data movement for StorageTek libraries that are managed by Automated Cartridge System Library Software (ACSL). To activate this function, the Tivoli Storage Manager library type must be EXTERNAL.

F**file access time**

On AIX, UNIX, or Linux systems, the time when the file was last accessed.

file age

For migration prioritization purposes, the number of days since a file was last accessed.

file device type

A device type that specifies the use of sequential access files on disk storage as volumes.

file server

A dedicated computer and its peripheral storage devices that are connected to a local area network that stores programs and files that are shared by users on the network.

file space

A logical space in server storage that contains a group of files that have been backed up or archived by a client node, from a single logical partition, file system, or virtual mount point. Client nodes can restore, retrieve, or delete their file spaces from server storage. In server storage, files belonging to a single file space are not necessarily stored together.

file space ID (FSID)

A unique numeric identifier that the server assigns to a file space when it is stored in server storage.

file state

The space management mode of a file that resides in a file system to which space management has been added. A file can be in one of three states: resident, premigrated, or migrated. See also *resident file*, *premigrated file*, and *migrated file*.

file system migrator (FSM)

A kernel extension that intercepts all file system operations and provides any space

management support that is required. If no space management support is required, the operation is passed to the operating system, which performs its normal functions. The file system migrator is mounted over a file system when space management is added to the file system.

file system state

The storage management mode of a file system that resides on a workstation on which the hierarchical storage management (HSM) client is installed. A file system can be in one of these states: native, active, inactive, or global inactive.

frequency

A copy group attribute that specifies the minimum interval, in days, between incremental backups.

FSID See *file space ID*.

FSM See *file system migrator*.

full backup

The process of backing up the entire server database. A full backup begins a new database backup series. See also *database backup series* and *incremental backup*. Contrast with *database snapshot*.

fuzzy backup

A backup version of a file that might not accurately reflect what is currently in the file because the file was backed up at the same time as it was being modified.

fuzzy copy

A backup version or archive copy of a file that might not accurately reflect the original contents of the file because it was backed up or archived the file while the file was being modified. See also *backup version* and *archive copy*.

G**General Parallel File System**

A high-performance shared-disk file system that can provide data access from nodes in a cluster environment.

gigabyte (GB)

In decimal notation, 1 073 741 824 when referring to memory capacity; in all other cases, it is defined as 1 000 000 000.

global inactive state

The state of all file systems to which

space management has been added when space management is globally deactivated for a client node. When space management is globally deactivated, hierarchical storage management (HSM) cannot perform migration, recall, or reconciliation. However, a root user can update space management settings and add space management to additional file systems. Users can access resident and premigrated files.

Globally Unique Identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system.

GPFS™

See *General Parallel File System*.

GPFS node set

A mounted, defined group of GPFS file systems.

group backup

The backup of a group containing a list of files from one or more file space origins.

GUID See *Globally Unique Identifier*.

H

hierarchical storage management (HSM)

A function that automatically distributes and manages data on disk, tape, or both by regarding devices of these types and potentially others as levels in a storage hierarchy that range from fast, expensive devices to slower, cheaper, and possibly removable devices. The objectives are to minimize access time to data and maximize available media capacity.

hierarchical storage management (HSM) client

A client program that works with the Tivoli Storage Manager server to provide hierarchical storage management (HSM) for a system. See also *hierarchical storage management* and *space manager client*.

HSM See *hierarchical storage management*.

HSM client

See *hierarchical storage management client*.

I

ILM See *information lifecycle management*.

image A file system or raw logical volume that is backed up as a single object.

image backup

A backup of a full file system or raw logical volume as a single object.

inactive file system

A file system for which space management has been deactivated. Contrast with *active file system*.

inactive version

A backup version of a file that is either not the most recent backup version, or that is a backup version of a file that no longer exists on the client system. Inactive backup versions are eligible for expiration processing according to the management class assigned to the file. Contrast with *active version*.

include-exclude file

A file containing statements to determine the files to back up and the associated management classes to use for backup or archive. See also *include-exclude list*.

include-exclude list

A list of options that include or exclude selected files for backup. An exclude option identifies files that should not be backed up. An include option identifies files that are exempt from the exclusion rules or assigns a management class to a file or a group of files for backup or archive services.

incremental backup

(1) A copy of all database data that has changed since the most recent successful full backup operation. An incremental backup is also known as a *cumulative backup image* because each incremental backup includes the contents of the previous incremental backup.

(2) The process of backing up information in the database that is new or changed since the last full backup. Contrast with *full backup*. See also *database backup series*.

(3) For Data Protection for Microsoft Exchange Server, a backup in which the transaction logs are backed up and then cleared.

individual mailbox restore

See *mailbox restore*.

information lifecycle management (ILM)

GPFS policy-based file management for storage pools and file sets.

inode The internal structure that describes the individual files on AIX, UNIX, or Linux systems. An inode contains the node, type, owner, and location of a file.

inode number
A number specifying a particular inode file in the file system.

IP address
A unique address for a device or logical unit on a network that uses the IP standard.

J

job file
A generated file that contains configuration information for a migration job. The file is XML format and can be created and edited in the hierarchical storage management (HSM) client for Windows client graphical user interface.

journal-based backup
A method for backing up Windows clients and AIX clients that exploits the change notification mechanism in a file to improve incremental backup performance by reducing the need to fully scan the file system.

journal daemon
On AIX, UNIX, or Linux systems, a program that tracks change activity for files residing in file systems.

journal service
In Microsoft Windows, a program that tracks change activity for files residing in file systems.

K

kilobyte (KB)
For processor storage, real and virtual storage, and channel volume, 210 or 1 024 bytes. For disk storage capacity and communications volume, 1 000 bytes.

L

LAN See *local area network*.

LAN-free data movement
The movement of client data between a client system and a storage device on a storage area network (SAN), bypassing the local area network. This process is also referred to as *LAN-free data transfer*.

LAN-free data transfer
See *LAN-free data movement*.

leader data
Bytes of data, from the beginning of a migrated file, that are stored in the file's corresponding stub file on the local file system. The amount of leader data that is stored in a stub file depends on the stub size that is specified.

library
(1) A repository for demountable recorded media, such as magnetic disks and magnetic tapes.
(2) A collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

library client
A server that uses server-to-server communication to access a library that is managed by another storage management server. See also *library manager*.

library manager
A server that controls device operations when multiple storage management servers share a storage device. See also *library client*.

local
(1) Pertaining to a device, file, or system that is accessed directly from a user system, without the use of a communication line.
(2) For HSM products, pertaining to the destination of migrated files that are being moved.

local area network (LAN)
A network that connects several devices in a limited area (such as a single building or campus) and that can be connected to a larger network.

local shadow volumes
Data that is stored on shadow volumes localized to a disk storage subsystem.

LOFS See *loopback virtual file system*.

logical file
A file that is stored in one or more server storage pools, either by itself or as part of an aggregate. See also *aggregate* and *physical file*.

logical occupancy
The space that is used by logical files in a

storage pool. This space does not include the unused space created when logical files are deleted from aggregate files, so it might be less than the physical occupancy.

logical unit (LU)

An access point through which a user or application program accesses the Systems Network Architecture (SNA) network to communicate with another user or application program.

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier that is used to differentiate devices, each of which is a logical unit (LU).

logical volume

A portion of a physical volume that contains a file system.

logical volume backup

A backup of a file system or logical volume as a single object.

Logical Volume Snapshot Agent (LVSA)

Software that can act as the snapshot provider for creating a snapshot of a logical volume during an online image backup.

loopback virtual file system (LOFS)

A file system that is created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

LU See *logical unit*.

LUN See *logical unit number*.

LVSA See *Logical Volume Snapshot Agent*.

M

macro file

A file that contains one or more storage manager administrative commands, which can be run only from an administrative client using the MACRO command. Contrast with *Tivoli Storage Manager command script*.

mailbox restore

A function that restores Microsoft Exchange Server data (from IBM Data Protection for Microsoft Exchange backups) at the mailbox level or mailbox-item level.

managed object

In Tivoli Storage Manager, a definition in the database of a managed server that was distributed to the managed server by a configuration manager. When a managed server subscribes to a profile, all objects that are associated with that profile become managed objects in the database of the managed server. In general, a managed object cannot be modified locally on the managed server. Objects can include policy, schedules, client option sets, server scripts, administrator registrations, server definitions, and server group definitions.

managed server

A Tivoli Storage Manager server that receives configuration information from a configuration manager using a subscription to one or more profiles. Configuration information can include definitions of objects such as policy and schedules. See also *configuration manager*, *subscription*, and *profile*.

management class

A policy object that users can bind to each file to specify how the server manages the file. The management class can contain a backup copy group, an archive copy group, and space management attributes. See also *copy group*, *space manager client*, *bind*, and *rebind*.

maximum transmission unit

The largest possible unit of data that can be sent on a given physical medium in a single frame. For example, the maximum transmission unit for Ethernet is 1500 bytes.

MB See *megabyte*.

media server

In a z/OS® environment, a program that provides access to z/OS disk and tape storage for Tivoli Storage Manager servers that run on operating systems other than z/OS.

megabyte (MB)

(1) 1 048 576 bytes (2 to the 20th power) when used in this publication.

(2) For processor storage, real and virtual storage, and channel volume, 2 to the power of 20 or 1 048 576 bits. For disk

storage capacity and communications volume, 1 000 000 bits.

metadata

Data that describes the characteristics of data; descriptive data.

migrate

To move data from one storage location to another. In Tivoli Storage Manager products, migrating can mean moving data from a client node to server storage, or moving data from one storage pool to the next storage pool defined in the server storage hierarchy. In both cases the movement is controlled by policy, such as thresholds that are set. See also *migration threshold*.

migrated file

A file that has been copied from a local file system to Tivoli Storage Manager storage. For HSM clients on UNIX or Linux systems, the file is replaced with a stub file on the local file system. On Windows systems, creation of the stub file is optional. See also *stub file* and *resident file*. For HSM clients on UNIX or Linux systems, contrast with *premigrated file*.

migrate-on-close recall mode

A mode that causes a migrated file to be recalled back to its originating file system temporarily. Contrast with *normal recall mode* and *read-without-recall recall mode*.

migration job

A specification of files to migrate, and actions to perform on the original files after migration. See also *job file*.

migration threshold

High and low capacities for storage pools or file systems, expressed as percentages, at which migration is set to start and stop.

mirroring

The process of writing the same data to multiple locations at the same time. Mirroring data protects against data loss within the recovery log.

mode

A copy group attribute that specifies whether to back up a file that has not been modified since the last time the file was backed up. See *modified mode* and *absolute mode*.

modified mode

In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup only if it has changed since the last backup. A file is considered a changed file if the date, size, owner, or permissions of the file have changed. See also *absolute mode*.

mount limit

The maximum number of volumes that can be simultaneously accessed from the same device class. The mount limit determines the maximum number of mount points. See also *mount point*.

mount point

On the Tivoli Storage Manager server, a logical drive through which volumes in a sequential access device class are accessed. For removable-media device types, such as tape, a mount point is a logical drive that is associated with a physical drive. For the file device type, a mount point is a logical drive that is associated with an I/O stream. The number of mount points for a device class is defined by the value of the mount limit attribute for that device class. See also *mount limit*.

mount retention period

The maximum number of minutes that the server retains a mounted sequential-access media volume that is not being used before it dismounts the sequential-access media volume.

mount wait period

The maximum number of minutes that the server waits for a sequential-access volume mount request to be satisfied before canceling the request.

MTU See *maximum transmission unit*.

N**Nagle algorithm**

An algorithm that reduces congestion of TCP/IP networks by combining smaller packets and sending them together.

named pipe

A type of interprocess communication that permits message data streams to pass between peer processes, such as between a client and a server.

NAS See *network-attached storage*.

NAS node

A client node that is a network-attached storage (NAS) file server. Data for the NAS node is transferred by a NAS file server that is controlled by the network data management protocol (NDMP). A NAS node is also called a NAS file server node.

native file system

A file system that is locally added to the file server and is not added for space management. The hierarchical storage manager (HSM) client does not provide space management services to the file system.

native format

A format of data that is written to a storage pool directly by the Tivoli Storage Manager server. Contrast with *non-native data format*.

NDMP

See *Network Data Management Protocol*.

NetBIOS

See *Network Basic Input/Output System*.

network-attached storage (NAS) file server

A dedicated storage device with an operating system that is optimized for file-serving functions. A NAS file server can have the characteristics of both a node and a data mover.

Network Basic Input/Output System (NetBIOS)

A standard interface to networks and personal computers that is used on local area networks to provide message, print-server, and file-server functions. Application programs that use NetBIOS do not have to handle the details of LAN data link control (DLC) protocols.

Network Data Management Protocol (NDMP)

A protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server, without installing vendor-acquired software on that file server.

network data-transfer rate

A rate that is calculated by dividing the total number of bytes that are transferred by the data transfer time. For example, this rate can be the time that is spent transferring data over a network.

node A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

node name

A unique name that is used to identify a workstation, file server, or PC to the server.

node privilege class

A privilege class that gives an administrator the authority to remotely access backup-archive clients for a specific client node or for all clients in a policy domain. See also *privilege class*.

non-native data format

A format of data that is written to a storage pool that differs from the format that the server uses for operations.

normal recall mode

A mode that causes a migrated file to be copied back to its originating file system when it is accessed.

O**offline volume backup**

A backup in which the volume is locked so that no other system applications can access it during the backup operation.

online volume backup

A backup in which the volume is available to other system applications during the backup operation.

open registration

A registration process in which users can register their workstations as client nodes with the server. Contrast with *closed registration*.

operator privilege class

A privilege class that gives an administrator the authority to disable or halt the server, enable the server, cancel server processes, and manage removable media. See also *privilege class*.

options file

A file that contains processing options. On Windows and NetWare systems, the file is called *dsm.opt*. On AIX, UNIX, Linux, and Mac OS X systems, the file is called *dsm.sys*.

originating file system

The file system from which a file was

migrated. When a file is recalled using normal or migrate-on-close recall mode, it is always returned to its originating file system.

orphaned stub file

A file for which no migrated file can be found on the Tivoli Storage Manager server that the client node is contacting for space management services. For example, a stub file can be orphaned when the client system-options file is modified to contact a server that is different than the one to which the file was migrated.

out-of-space protection mode

A mode that controls whether the program intercepts out-of-space conditions. See also *execution mode*.

P

pacing

In SNA, a technique by which the receiving system controls the rate of transmission of the sending system to prevent overrun.

packet In data communication, a sequence of binary digits, including data and control signals, that is transmitted and switched as a composite whole.

page A defined unit of space on a storage medium or within a database volume.

partial-file recall mode

A recall mode that causes the hierarchical storage management (HSM) function to read just a portion of a migrated file from storage, as requested by the application accessing the file.

password generation

A process that creates and stores a new password in an encrypted password file when the old password expires. Automatic generation of a password prevents password prompting. Password generation can be set in the options file (passwordaccess option). See also *options file*.

path An object that defines a one-to-one relationship between a source and a destination. Using the path, the source accesses the destination. Data can flow from the source to the destination, and back. An example of a source is a data

mover (such as a network-attached storage [NAS] file server), and an example of a destination is a tape drive.

pattern-matching character

See *wildcard character*.

physical file

A file that is stored in one or more storage pools, consisting of either a single logical file, or a group of logical files that are packaged together as an aggregate. See also *aggregate* and *logical file*.

physical occupancy

The amount of space that is used by physical files in a storage pool. This space includes the unused space that is created when logical files are deleted from aggregates. See also *physical file*, *logical file*, and *logical occupancy*.

plug-in

A self-contained software component that modifies (adds, or changes) the function in a particular system. When a plug-in is added to a system, the foundation of the original system remains intact.

policy domain

A grouping of policy users with one or more policy sets, which manage data or storage resources for the users. The users are client nodes that are associated with the policy domain.

policy privilege class

A privilege class that gives an administrator the authority to manage policy objects, register client nodes, and schedule client operations for client nodes. Authority can be restricted to certain policy domains. See also *privilege class*.

policy set

A group of rules in a policy domain. The rules specify how data or storage resources are automatically managed for client nodes in the policy domain. Rules can be contained in management classes. See also *active policy set* and *management class*.

premigrated file

A file that has been copied to Tivoli Storage Manager storage, but has not been replaced with a stub file on the local file system. An identical copy of the file resides both on the local file system and

in Tivoli Storage Manager storage. Premigrated files occur on UNIX and Linux file systems to which space management has been added. Contrast with *migrated file* and *resident file*.

premigrated files database

A database that contains information about each file that has been premigrated to Tivoli Storage Manager storage. The database is stored in a hidden directory named `.SpaceMan` in each file system to which space management has been added.

premigration

The process of copying files that are eligible for migration to Tivoli Storage Manager storage, but leaving the original file intact on the local file system.

premigration percentage

A space management setting that controls whether the next eligible candidates in a file system are premigrated following threshold or demand migration.

primary storage pool

A named set of volumes that the server uses to store backup versions of files, archive copies of files, and files migrated from client nodes. See also *destination* and *copy storage pool*.

privilege class

A level of authority that is granted to an administrator. The privilege class determines which administrative tasks the administrator can perform. See also *node privilege class*, *operator privilege class*, *policy privilege class*, *storage privilege class*, and *system privilege class*.

profile

A named group of configuration information that can be distributed from a configuration manager when a managed server subscribes. Configuration information can include registered administrator IDs, policies, client schedules, client option sets, administrative schedules, storage manager command scripts, server definitions, and server group definitions. See also *configuration manager* and *managed server*.

Q

quota (1) For HSM on AIX, UNIX, or Linux systems, the limit (in megabytes) on the

amount of data that can be migrated and premigrated from a file system to server storage.

(2) For HSM on Windows systems, a user-defined limit to the space that is occupied by recalled files.

R

randomization

The process of distributing schedule start times for different clients within a specified percentage of the schedule's startup window.

raw logical volume

A portion of a physical volume that is comprised of unallocated blocks and has no journaled file system (JFS) definition. A logical volume is read/write accessible only through low-level I/O functions.

read-without-recall recall mode

A mode that causes hierarchical storage management (HSM) to read a migrated file from storage without storing it back on the local file system. The last piece of information read from the file is stored in a buffer in memory on the local file system. Contrast with *normal recall mode* and *migrate-on-close recall mode*.

rebind

To associate all backed-up versions of a file with a new management class name. For example, a file that has an active backup version is rebound when a later version of the file is backed up with a different management class association. See also *bind*.

recall In Tivoli Storage Manager, to copy a migrated file from server storage back to its originating file system using the space management client. See also *transparent recall*, *selective recall*, and *recall mode*.

recall mode

A mode that is assigned to a migrated file with the `dsmatrr` command that determines how the file is processed when it is recalled. It determines whether the file is stored on the local file system, is migrated back to Tivoli Storage Manager storage when it is closed, or is read from Tivoli Storage Manager storage without storing it on the local file system.

receiver

A server repository that contains a log of server and client messages as events. For example, a receiver can be a file exit, a user exit, or the Tivoli Storage Manager server console and activity log. See also *event*.

reclamation

The process of consolidating the remaining data from many sequential-access volumes onto fewer, new sequential-access volumes.

reclamation threshold

The percentage of space that a sequential-access media volume must have before the server can reclaim the volume. Space becomes reclaimable when files are expired or are deleted.

reconciliation

The process of synchronizing a file system with the Tivoli Storage Manager server, and then removing old and obsolete objects from the Tivoli Storage Manager server.

recovery log

A log of updates that are about to be written to the database. The log can be used to recover from system and media failures. The recovery log consists of the active log (including the log mirror) and archive logs.

register

To define a client node or administrator ID that can access the server.

registry

A repository that contains access and configuration information for users, systems, and software.

remote

(1) Pertaining to a system, program, or device that is accessed through a communication line.

(2) For HSM products, pertaining to the origin of migrated files that are being moved.

resident file

On a Windows system, a complete file on a local file system that might also be a migrated file because a migrated copy can exist in Tivoli Storage Manager storage. On a UNIX or Linux system, a complete

file on a local file system that has not been migrated or premigrated, or that has been recalled from Tivoli Storage Manager storage and modified. Contrast with *stub file* and *premigrated file*. See *migrated file*.

restore

To copy information from its backup location to the active storage location for use. For example, to copy information from server storage to a client workstation.

retention

The amount of time, in days, that inactive backed-up or archived files are kept in the storage pool before they are deleted. Copy group attributes and default retention grace periods for the domain define retention.

retrieve

To copy archived information from the storage pool to the workstation for use. The retrieve operation does not affect the archive version in the storage pool.

roll back

To remove changes that were made to database files since the last commit point.

root user

A system user who operates without restrictions. A root user has the special rights and privileges needed to perform administrative tasks.

S

SAN See *storage area network*.

schedule

A database record that describes client operations or administrative commands to be processed. See *administrative command schedule* and *client schedule*.

scheduling mode

The type of scheduling operation for the server and client node that supports two scheduling modes: client-polling and server-prompted.

scratch volume

A labeled volume that is either blank or contains no valid data, that is not defined, and that is available for use.

script

A series of commands, combined in a file, that carry out a particular function when the file is run. Scripts are interpreted as

they are run. Contrast with *Tivoli Storage Manager command script*.

Secure Sockets Layer (SSL)

A security protocol that provides communication privacy. With SSL, client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery.

selective backup

The process of backing up certain files or directories from a client domain. The files that are backed up are those that are not excluded in the include-exclude list. The files must meet the requirement for serialization in the backup copy group of the management class that is assigned to each file. Contrast with *incremental backup*.

selective migration

The process of copying user-selected files from a local file system to Tivoli Storage Manager storage and replacing the files with stub files on the local file system. Contrast with *threshold migration* and *demand migration*.

selective recall

The process of copying user-selected files from Tivoli Storage Manager storage to a local file system. Contrast with *transparent recall*.

serialization

The process of handling files that are modified during backup or archive processing. See *dynamic serialization*, *static serialization*, *shared static serialization*, and *shared dynamic serialization*.

server A software program or a computer that provides services to other software programs or other computers.

server options file

A file that contains settings that control various server operations. These settings affect such things as communications, devices, and performance.

server-prompted scheduling mode

A client/server communication technique where the server contacts the client node when tasks must be done. Contrast with *client-polling scheduling mode*.

server storage

The primary, copy, and active-data storage

pools that are used by the server to store user files such as backup versions, archive copies, and files migrated from space manager client nodes (space-managed files). See also *active-data pool*, *primary storage pool*, *copy storage pool*, *storage pool volume*, and *volume*.

session

A logical or virtual connection between two stations, software programs, or devices on a network that allows the two elements to communicate and exchange data.

session resource usage

The amount of wait time, processor time, and space that is used or retrieved during a client session.

shared dynamic serialization

A value for serialization that specifies that a file must not be backed up or archived if it is being modified during the operation. Tivoli Storage Manager retries the backup or archive operation a number of times; if the file is being modified during each attempt, Tivoli Storage Manager will back up or archive the file on its last try. See also *serialization*. Contrast with *dynamic serialization*, *shared static serialization*, and *static serialization*.

shared library

A library device that is used by multiple storage manager servers.

shared static serialization

A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. Tivoli Storage Manager attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *static serialization*.

snapshot

An image backup type that consists of a point-in-time view of a volume.

space-managed file

A file that is migrated from a client node by the space manager client. The space manager client recalls the file to the client node on demand.

space management

The process of keeping sufficient free storage space available on a local file system for new data by migrating files to server storage. Synonymous with *hierarchical storage management*.

space manager client

A program that runs on a UNIX or Linux system to manage free space on the local file system by migrating files to server storage. The program can recall the files either automatically or selectively. Also called *hierarchical storage management (HSM) client*.

space monitor daemon

A daemon that checks space usage on all file systems for which space management is active, and automatically starts threshold migration when space usage on a file system equals or exceeds its high threshold.

sparse file

A file that is created with a length greater than the data it contains, leaving empty spaces for the future addition of data.

special file

On AIX, UNIX, or Linux systems, a file that defines devices for the system, or temporary files that are created by processes. There are three basic types of special files: first-in, first-out (FIFO); block; and character.

SSL See *Secure Sockets Layer*.

stabilized file space

A file space that exists on the server but not on the client.

stanza A group of lines in a file that together have a common function or define a part of the system. Each stanza is identified by a name that occurs in the first line of the stanza. Depending on the type of file, a stanza is ended by the next occurrence of a stanza name in the file, or by an explicit end-of-stanza marker. A stanza can also be ended by the end of the file.

startup window

A time period during which a schedule must be initiated.

static serialization

A copy-group serialization value that specifies that a file must not be modified

during a backup or archive operation. If the file is in use during the first attempt, the storage manager cannot back up or archive the file. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *shared static serialization*.

storage agent

A program that enables the backup and restoration of client data directly to and from storage attached to a storage area network (SAN).

storage area network (SAN)

A dedicated storage network that is tailored to a specific environment, combining servers, systems, storage products, networking products, software, and services.

storage hierarchy

(1) A logical order of primary storage pools, as defined by an administrator. The order is typically based on the speed and capacity of the devices that the storage pools use. The storage hierarchy is defined by identifying the next storage pool in a storage pool definition. See also *storage pool*.

(2) An arrangement of storage devices with different speeds and capacities. The levels of the storage hierarchy include: main storage, such as memory and direct-access storage device (DASD) cache; primary storage (DASD containing user-accessible data); migration level 1 (DASD containing data in a space-saving format); and migration level 2 (tape cartridges containing data in a space-saving format).

storage pool

A named set of storage volumes that are the destination that is used to store client data. A storage pool contains backup versions, archive copies, and files that are migrated from space manager client nodes. A primary storage pool is backed up to a copy storage pool. See also *primary storage pool*, *copy storage pool*, and *active-data pool*.

storage pool volume

A volume that has been assigned to a storage pool. See also *volume*, *active-data pool*, *copy storage pool*, and *primary storage pool*.

storage privilege class

A privilege class that gives an administrator the authority to control how storage resources for the server are allocated and used, such as monitoring the database, the recovery log, and server storage. See also *privilege class*.

stub A shortcut on the Windows file system that is generated by the hierarchical storage management (HSM) client for a migrated file that allows transparent user access. A stub is the sparse file representation of a migrated file, with a reparse point attached.

stub file

A file that replaces the original file on a local file system when the file is migrated to storage. A stub file contains the information that is necessary to recall a migrated file from Tivoli Storage Manager storage. It also contains additional information that can be used to eliminate the need to recall a migrated file.

stub file size

The size of a file that replaces the original file on a local file system when the file is migrated to Tivoli Storage Manager storage. The size that is specified for stub files determines how much leader data can be stored in the stub file. The default for stub file size is the block size defined for a file system minus 1 byte.

subscription

In a Tivoli environment, the process of identifying the subscribers that the profiles are distributed to. For Tivoli Storage Manager, a subscription is the process by which a managed server receives configuration information associated with a particular profile on a configuration manager. See also *managed server*, *configuration manager*, and *profile*.

system privilege class

A privilege class that gives an administrator the authority to issue all server commands. See also *privilege class*.

Systems Network Architecture (SNA)

The description of the logical structure, formats, protocols, and operational sequences for transmitting information through and controlling the configuration and operation of networks.

T**tape library**

A set of equipment and facilities that support an installation's tape environment. The tape library can include tape storage racks, mechanisms for automatic tape mounting, a set of tape drives, and a set of related tape volumes mounted on those drives.

tape volume prefix

The high-level-qualifier of the file name or the data set name in the standard tape label.

target node

A client node for which other client nodes (called agent nodes) have been granted proxy authority. The proxy authority allows the agent nodes to perform operations such as backup and restore on behalf of the target node, which owns the data.

TCA See *trusted communications agent*.

TCP/IP

See *Transmission Control Protocol/Internet Protocol*.

threshold migration

The process of moving files from a local file system to Tivoli Storage Manager storage based on the high and low thresholds that are defined for the file system. Contrast with *demand migration*, *selective migration*, and *migration job*.

throughput

In storage management, the total bytes in the workload, excluding overhead, that are backed up or restored, divided by elapsed time.

timeout

A time interval that is allotted for an event to occur or complete before operation is interrupted.

timestamp control mode

A mode that determines whether commands preserve the access time for a file or set it to the current time.

Tivoli Storage Manager command script

A sequence of Tivoli Storage Manager administrative commands that are stored in the database of the Tivoli Storage Manager server. The script can run from any interface to the server. The script can

include substitution for command parameters and conditional logic.

tombstone object

A small subset of attributes of a deleted object. The tombstone object is retained for a specified period, and at the end of the specified period, the tombstone object is permanently deleted.

Transmission Control Protocol/Internet Protocol (TCP/IP)

An industry-standard, nonproprietary set of communication protocols that provides reliable end-to-end connections between applications over interconnected networks of different types.

transparent recall

The process that is used to automatically recall a file to a workstation or file server when the file is accessed. See also *recall mode*. Contrast with *selective recall*.

trusted communications agent (TCA)

A program that handles the sign-on password protocol when clients use password generation.

U

UCS-2 A 2-byte (16-bit) encoding scheme based on ISO/IEC specification 10646-1. UCS-2 defines three levels of implementation: Level 1-No combining of encoded elements allowed; Level 2-Combining of encoded elements is allowed only for Thai, Indic, Hebrew, and Arabic; Level 3-Any combination of encoded elements are allowed.

UNC See *Universal Naming Convention name*.

Unicode

A character encoding standard that supports the interchange, processing, and display of text that is written in the common languages around the world, plus some classical and historical texts. The Unicode standard has a 16-bit character set defined by ISO 10646.

Unicode-enabled file space

Unicode file space names provide support for multilingual workstations without regard for the current locale.

Unicode transformation format 8

Unicode Transformation Format (UTF), 8-bit encoding form, which is designed for ease of use with existing ASCII-based

systems. The CCSID value for data in UTF-8 format is 1208.

Universal Naming Convention (UNC) name

A name that is used to access a drive or directory containing files shared across a network. The UNC name includes the system name and a SharePoint name that represents the shared drive or directory.

Universally Unique Identifier (UUID)

The 128-bit numeric identifier that is used to ensure that two components do not have the same identifier.

UTF-8 See *Unicode transformation format 8*.

UUID See *Universally Unique Identifier*.

V

validate

To check a policy set for conditions that can cause problems if that policy set becomes the active policy set. For example, the validation process checks whether the policy set contains a default management class.

version

A backup copy of a file stored in server storage. The most recent backup copy of a file is the active version. Earlier copies of the same file are inactive versions. The number of versions retained by the server is determined by the copy group attributes in the management class.

virtual file space

A representation of a directory on a network-attached storage (NAS) file system as a path to that directory.

virtual volume

An archive file on a target server that represents a sequential media volume to a source server.

volume

A discrete unit of storage on disk, tape or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control. See also *scratch volume*, and *storage pool volume*.

volume history file

A file that contains information about volumes that have been used by the server for database backups and for export of administrator, node, policy, or

server data. The file also has information about sequential-access storage pool volumes that have been added, reused, or deleted. The information is a copy of volume information that is recorded in the server database.

Volume Shadow Copy Service

A set of Microsoft application-programming interfaces (APIs) that you can use to create shadow copy backups of volumes, exact copies of files, including all open files, and so on.

VSS See *Volume Shadow Copy Service*.

VSS Backup

A backup operation that uses Microsoft Volume Shadow Copy Service (VSS) technology. The backup operation produces an online snapshot (point-in-time consistent copy) of Microsoft Exchange data. This copy can be stored on local shadow volumes or on Tivoli Storage Manager server storage.

VSS Fast Restore

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on local shadow volumes.

VSS Instant Restore

A volume-level hardware-assisted Microsoft Volume Shadow Copy Service (VSS) function where target volumes that contain the snapshot are copied back to the original source volumes.

VSS offloaded backup

A backup operation that uses a Microsoft Volume Shadow Copy Service (VSS) hardware provider (installed on an alternate system) to move IBM Data Protection for Microsoft Exchange data to the Tivoli Storage Manager server. This type of backup operation shifts the backup load from the production system to another system.

VSS Restore

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on

Tivoli Storage Manager server storage to their original location.

W**wildcard character**

A special character such as an asterisk (*) or a question mark (?) that can be used to represent one or more characters. Any character or set of characters can replace the wildcard character.

workstation

A terminal or personal computer at which a user can run applications and that is usually connected to a mainframe or a network.

worldwide name

A 64-bit, unsigned name identifier that is unique.

workload partition (WPAR)

A partition within a single operating system instance.

Index

A

- accessibility features 161
- activate policy set 60
- adaptation
 - of tablespace 111
- Adaptive Computing Environment 51
- Administration Assistant 95
 - administer users 12
 - administering userids 87
 - authorizing users 87
 - concepts 9
 - configuration tool 12
 - configure systems 12
 - configuring 87
 - customizing 12
 - Database Agent 9
 - Database component 9
 - installation planning sheets 154
 - Java 31
 - manage report templates 12
 - migration 36
 - monitor operations 12
 - overview 12
 - problem support 12
 - Server component 9
 - starting and stopping 89
 - threshold definitions 50
 - upgrading 36
 - view performance data 12
 - Web browser, remote method invocation 31
- Administration Assistant client
 - Java 31
 - prerequisites 31
 - setting up 48
 - verifying the installation 48
- Administration Assistant database
 - changing password for 90
- Administration Assistant scheduling client
 - command line interface 11
- Administration Assistant scheduling client
 - prerequisites 31
 - setting up 49
- Administration Assistant Scheduling Client
 - creating reports from 88
- Administration Assistant Server
 - specifying 87
- Administration Assistant Server component
 - secure connection to clients 47
- Administration Assistant Server component configuration file 11, 47, 48
- Administration Assistant server-level components
 - initial installation 32
 - installation planning 31
 - installing 32
 - prerequisites 31
 - uninstalling 32

- Administration Assistant server-level components (*continued*)
 - upgrade installation 32
 - verifying the installation 48
- administrative client 67
- ADSMNODE profile keyword 134
- Advanced Copy Services (ACS) xi
- alternate / parallel backup paths 103
- alternate network paths and servers 102
- alternate path 105
- alternate/parallel backup paths
 - example for availability 41
 - example for performance 41
 - reasons to use 18
- alternate/parallel backup servers 105
 - example 3 for disaster recovery 42
- archiving
 - inactive data 18
- authorizing Administration Assistant users 87
- automating backup and archive operations 67
- automation options
 - alternate network paths and servers 105
 - backup version control 105
 - frontend/backend processing 105
 - messaging 105
 - multiple DB2 log file copies 69
 - multiple redo log copies 105
 - selectable management classes 105
- availability improvements
 - with alternate/parallel paths and servers 18, 105
 - with multiple DB2 Log Files 69
 - with multiple redo log copies 105

B

- backend processing 105
- BACKEND profile keyword 134
- backup
 - automated 67, 68
 - database command 3
 - full offline 123
 - full online 124
 - incremental 8, 68, 103
 - objects 125
 - offline 85
 - online 85
 - operation reporting 99
 - parallel 103
 - path 103
 - performance optimization
 - multiplexing 109
 - simulating 139
 - start in non-partitioned environment 69
 - start in partitioned environment 123
 - start in partitioned environment via SSV 123

- backup (*continued*)
 - status reporting 97, 100
 - tape usage 17
 - terminating 18
 - version control 105
 - windows 8
- backup object
 - segmentation 46, 71
- Backup Object Manager 4
 - commands 125
- backup paths
 - alternate/parallel 103
 - profile examples 18
- backup scheduler
 - IBM Tivoli Workload scheduler 67
 - SAP scheduler 67
 - Tivoli Storage Manager scheduler 67
 - UNIX or Linux crontab / Windows scheduler 67
 - Windows sample 68
- backup server
 - alternate/parallel 105
 - managing data 102
 - performance optimization 16
 - profile examples 18
- backup status
 - monitoring 95
- backup strategy
 - Sample 67
- backup-archive client 1, 8, 68
- BACKUPIDPREFIX profile keyword 134
- books
 - See publications*
- BRARCHIVE 123
- BRARCHIVEMGTCLASS profile
 - keyword 105, 134
- BRBACKUP 123
- BRBACKUPMGTCLASS profile
 - keyword 105, 134
- BUFFCOPY profile keyword 134
- buffer copies 104
- buffer size 104
- BUFFSIZE profile keyword 135

C

- client options file 57, 64, 151
- client system options file
 - UNIX and Linux example 149
 - Windows example 151
- client user options file
 - UNIX and Linux example 149
- client/server connection paths 18
- cloning 80, 130
 - of a database 74
- collocation 17, 109
- Command Line Processor (CLP) 3
- command syntax
 - Backup Object Manager 125
- CommonStore 8, 18
- compression 16, 103

- compression (*continued*)
 - null block 104
 - Tivoli Storage Manager 104
 - when restoring 19
- CONFIG_FILE profile keyword 135
- configuration files
 - creation of 46
- configuration matrix for Tivoli Storage Manager password settings 62, 63
- copies of log files 69
- copy group 60
- crontab
 - file example 132
 - scheduling backups with 67
- custom SQL file 145
- custom SQL file (sample) 147
- customer support
 - contact xvi
- customization
 - Administration Assistant 12

D

- data block 3, 4
- data compression
 - and overall throughput 92
 - hardware vs. software 55, 57
 - null block 104
 - Tivoli Storage Manager 104
- Data Protection for SAP
 - architecture and properties 1
 - components 1
 - database utilities 1
 - installation planning sheets 152
 - installing 23
 - integration with Administration Assistant 2
 - integration with SAP 2
 - introduction 1
 - migration 35
 - overview 1
 - profile
 - keyword descriptions 133
 - keyword rules 133
 - migration 36
 - Windows sample 142
 - reporting 101
 - shared library 2, 4
 - upgrading 23, 35
- Data Protection for SAP installation package 24, 152
- Data Protection for SAP installation packages 24
- Data Protection for SAP profile 133
- data throughput 93, 106
- Database Server
 - performance optimization 15
- DB2
 - Single System View (SSV) backup 69
- DB2 CommonStore for SAP 18
- DB2 Failure Types 8
- DB2 log files
 - creating multiple copies 69
- DB2 Log Manager 6
- DB2 vendor environment file
 - sample 152

- DB2 Version 8.2
 - rules when using 44
- definition information
 - sample 76
- definitions 167
- device class 16, 59, 136
- disability 161
- disaster recovery
 - example 3 for disaster recovery 42
 - with alternate/parallel paths and servers 18
- disk
 - layout 15
- distributed file system 50
- DMS 86
- documentation
 - See* publications
- drill-down 94
- DSMI_CONFIG 28
- DSMI_DIR 28
- DSMI_LOG 28

E

- education
 - see* Tivoli technical training xiv
- END profile keyword 135
- Enterprise Storage Server 12
- environment variable
 - DB2_INSTANCE 127
 - DSMI_CONFIG 28
 - DSMI_DIR 28
 - DSMI_LOG 28
 - TDP_DIR 27, 28, 121
 - XINT_PROFILE 26, 28
- examples
 - alternate/parallel paths for availability 41
 - alternate/parallel paths for disaster recovery 42
 - alternate/parallel paths for increased performance 41
 - client system options file (UNIX and Linux) 149
 - client system options file (Windows) 151
 - client user options file (UNIX) 149
 - client user options file (Windows) 149
 - crontab file 132
 - full offline batch backup (Windows) 70
 - include/exclude list (UNIX and Linux) 150
 - include/exclude list (Windows) 151
 - offline backup shell script 71
 - scheduled batch backup (Windows) 70
 - Tivoli Storage Manager profiles for UNIX and Linux 149
 - Tivoli Storage Manager profiles for Windows 149

F

- fixes, obtaining xvi

- FlashCopy and snapshots
 - backup 12
 - devices 12
 - restore 12
- FlashCopy Manager xi
- frontend processing 105
- FRONTEND profile keyword 135
- full offline backup 123
- full offline batch backup
 - Windows example 70
- full online backup 124
- full restore 124

G

- glossary 167

H

- HACMP 19
 - sample stop script 53
- hardware compression 55, 57
 - compared with Tivoli Storage Manager client software compression 55, 57
- high availability 19

I

- IBM Publications Center xi, xiv
- inactive data
 - archiving of 8, 18
- include/exclude list
 - UNIX and Linux example 150
 - Windows example 151
- incremental backup 8, 103
- installation package 24, 152
- installing
 - Administration Assistant client, prerequisites 31
 - Administration Assistant server-level components 32
 - Administration Assistant server, prerequisites 31
 - Administration Assistant, installation verification 48
- installing Data Protection for SAP 23, 24
 - on Windows 28
 - password handling 62
 - planning for installation 24
 - prerequisites 24
 - TSM Option file 64
 - verifying the installation 30
- installing Tivoli Storage Manager for ERP
 - initial installation 26
 - migration 26
 - on AIX, Linux, Solaris 26
 - upgrade installation (migration) 26
- integration of Data Protection for SAP
 - with Administration Assistant 2
 - with SAP 2
- Internet, searching for problem resolution xv, xvi
- IPv6 116

J

Java

- prerequisite for Administration Assistant client 31
- prerequisite for Administration Assistant server 31

K

- keyboard 161
- keywords
 - ADSMNODE 134
 - BACKEND 134
 - BACKUPIDPREFIX 134
 - BRARCHIVEMGTCLASS 69, 105, 134
 - BRBACKUPMGTCLASS 105, 134
 - BUFCOPY 134
 - BUFSIZE 104, 135
 - CONFIG_FILE 135
 - END 135
 - FRONTEND 135
 - LOG_SERVER 135
 - MAX_SESSIONS 136
 - MAX_VERSIONS 136
 - PASSWORDREQUIRED 137
 - REDOLOG_COPIES 69, 137
 - RL_COMPRESSION 104, 137
 - SEGMENTSIZ 137
 - SERVER 137
 - SESSIONS 138
 - syntax for all keywords 133
 - TCP_ADDRESS 138
 - TRACE 138
 - TRACEFILE 139
 - TRACEMAX 139
 - USE_AT 139
- knowledge bases, searching xv

L

- LAN-free backup 15
 - performance optimization 16
- log files
 - creating multiple copies 69
 - how to find 121
 - location 121
- Log Manager
 - built in 44
- LOG_SERVER profile keyword 135
- LOGARCHMETHn configuration
 - parameters 27, 28
- logging
 - messages 105

M

- management classes 105
 - Tivoli Storage Manager server configuration 60
- manuals
 - See publications
- MAX_SESSIONS profile keyword 136
- MAX_VERSIONS 123
- MAX_VERSIONS profile keyword 136

- message file
 - how to find 121
- messages
 - logging of 105
- migration
 - Data Protection for SAP 35
 - of Administration Assistant 36
 - profile 36
- monitoring Data Protection for SAP 12
- mount points
 - maximum number per node 61
 - node parameter maxnummp 61
- multiple network paths 110
- multiple paths 159
- multiple redo log copies 105
- multiple servers 108
 - when restoring 19
- multiple sessions 103, 109
- multiplexing 16, 109

N

- network
 - performance optimization 16
- node
 - maxnummp 61
 - number of mount points 61
 - Tivoli Storage Manager server 61
- non-partitioned environment
 - starting backup 69
 - starting restore 73
- normalizing of tablespaces 111
- null block data compression 104

O

- offline backup shell script
 - example 71
- offline log file 8, 59, 68, 137
- optimizing
 - backup 106
 - restore 19

P

- parallel backup and restore
 - number of parallel sessions to specify 136
- parallel backup paths
 - sample 2 for increased performance 41
- parallel backup servers
 - alternate 105
- parallel path 103
 - example for increased performance 41
- parallel sessions 103, 159
- partition 46, 123, 124, 127, 131, 134, 135, 136
 - partitioned environment
 - starting backup / restore 123
- Passport Advantage xvii
- password handling
 - automatic generation 62
 - configuration matrix (UNIX or Linux) 62

- password handling (*continued*)
 - configuration matrix (Windows) 63
 - manual generation 61
 - no password usage 61
 - set the password 62
 - PASSWORDREQUIRED profile keyword 137
 - path
 - alternate 105
 - backup 103
 - performance analysis 95
 - performance monitoring 106
 - Administration Assistant 12
 - using sensors 107
 - performance optimization 93
 - Backup Server 16
 - by changing buffer size 104
 - by compressing data 104
 - by multiplexing 92, 109
 - by setting up proper environment 103
 - by using multiple network paths 103, 110
 - example 2 41
 - by using multiple servers 103, 108
 - by using multiple sessions 103, 109
 - CPU power 15, 16
 - data transfer 103
 - dedicated backbone network 16
 - disk layout 15
 - general considerations 107
 - I/O paths 15
 - LAN-free backup 16
 - network bandwidth 16
 - options for 103
 - settings for the Tivoli Storage Manager 158
 - size of database 15
 - tuning 103
 - volume manager 15
- permissions
 - granting for Administration Assistant client 31
- policy
 - definition 60
 - domain 60
 - set 60
 - activate 60
- policy file
 - granting permissions for Administration Assistant client 31
- printing reports 95, 101
- problem determination
 - describing problem for IBM Software Support xvii
 - determining business impact for IBM Software Support xvii
 - submitting a problem to IBM Software xviii
- production system 80
- productivity options
 - backup status monitoring 95
 - managing data on the backup server 102
 - performance analysis 95
 - reporting 95
 - tracing 95

- profile keywords
 - syntax 133
- profiles
 - example of Tivoli Storage Manager for UNIX and Linux 149
 - Tivoli Storage Manager example for Windows 149
- ProLE 26, 28, 87, 88
- publications
 - download xi
 - order xi
 - search xi
 - Tivoli Storage FlashCopy Manager xiv
 - Tivoli Storage Manager xii

R

- recovery
 - time for 8
- Recovery History File 4, 73, 74, 127, 128, 130
- redirected restore 85
 - automatic mode 74
 - batch mode 78
 - interactive mode 79
 - plausibility checks 83
 - prerequisite 85
 - using Backup Object Manager 84
- redo logs
 - multiple copies 105
- REDOLOG_COPIES 69
 - profile keyword 137
- remote shares (Windows)
 - implementing configuration files on 29
- reports 95
 - creating 98
 - modifying the output 98
 - on backup performance 95
 - on failed actions 98
 - types 101
 - UNIX scheduling sample 148
 - using templates 88, 102
 - using the Administration Assistant Client 89, 101
 - Windows scheduling sample 148
- restore
 - database command 4
 - full restore 124
 - performance optimization
 - multiplexing 109
 - redirected 78, 79, 85, 130
 - start in non-partitioned environment 73
 - start in partitioned environment 123
 - tablespaces 130
- retention 17, 18, 19
- RISC 6000
 - buffer pool settings 159
- RL_COMPRESSION 103
- RL_COMPRESSION profile keyword 104, 137

S

- sample Data Protection for SAP profile
 - UNIX or Linux 139
 - Windows 142
- SAP (backup) scheduler 67
- SAP-DB2 Administration Tools 2
- Scaling of tablespaces 112
- scheduled batch backup
 - Windows example 70
- scheduling
 - automatic backups 67
 - function of Tivoli Storage Manager 67
 - products 67
- scheduling client
 - command line interface 11
 - creating reports 88
 - prerequisites 31
 - setting up 49
- scope
 - of Data Protection for SAP 8
- security settings
 - authorizing Administration Assistant users 87
- SEGMENTSIZ profile keyword 137
- server
 - installing Administration Assistant components 32
- SERVER profile keyword 137
- servers
 - alternate 105
 - alternate/parallel 103
- sessions
 - multiple (parallel) 16, 17, 18, 41, 42, 92, 93, 103, 109, 123, 124, 136, 137, 138
 - single 17, 94
- SESSIONS profile keyword 138
- setting Tivoli Storage Manager passwords manually 61
- shared library 4
- SID 18, 26, 28, 40, 69, 115, 120, 121, 124, 133, 134, 153
- simulating
 - using USE_AT 139
- Single System View (SSV) backup 69
- snapshot
 - devices 12
- software compression vs. hardware compression
 - UNIX or Linux 55
 - Windows 57
- software support
 - describing problem for IBM Software Support xvii
 - determining business impact for IBM Software Support xvii
 - submitting a problem xviii
- Software Support
 - contact xvi
- SP Switch
 - buffer pool settings 159
- specifying management classes 105
- storage device setup
 - Tivoli Storage Manager server 59
- storage pool 58, 59, 105, 109
 - definition 59

- support contract xvii
- support information xiv
- support subscription xvii
- SYSCATSPACE 113

T

- tablespace 4, 15, 69, 73, 74, 127, 128, 129, 130, 131, 136
 - automated adaptations 111
 - container 75
 - definition information 75, 86, 129
 - creation of 85
 - restore of 130
 - normalizing 78, 111
 - scaling 112
- tape drives
 - using hardware compression with 55, 57
- tape usage
 - for backups 17
- TCP_ADDRESS profile keyword 138
- TDP_DIR 27, 28
- template
 - for creating reports 88
- terminating the backup job 18
- test system 80
- threshold definitions
 - via custom SQL file 145
- Tivoli Storage Manager
 - backup
 - version control 136
 - backup scheduler 67
 - client software compression 57
 - configuration file customization 158
 - copy group 136
 - data compression 104
 - expiration function 136
 - management classes 105
 - network settings 158
 - options files 64
 - passwords 61
 - performance optimization 158
 - profile example for UNIX and Linux 149
 - profile example for Windows 149
 - scheduling function 67
- Tivoli Storage Manager API
 - client configuration 24
- Tivoli Storage Manager client configuration
 - configuration on UNIX or Linux 55
 - configuration on Windows 57
 - software compression 55
- Tivoli Storage Manager for ERP 12
- Tivoli Storage Manager passwords
 - authentication off 61
 - automatic generation 62
 - configuration matrix to set keywords 62, 63
 - manual generation 61
- Tivoli Storage Manager server
 - adding 59
 - configuration 58
 - configuration, prerequisites 58
 - management classes 60
 - node definition 61

Tivoli Storage Manager server (*continued*)
 performance considerations 58
 policy definition 60
 storage device setup 59
 storage pool definition 59
 storing data on 17
Tivoli technical training xiv
TRACE profile keyword 138
TRACEFILE profile keyword 139
TRACEMAX profile keyword 139
tracing 95
training, Tivoli technical xiv
transaction rate
 of Data Protection for SAP 8
trimming the database 8
troubleshooting
 IBM support 118
 random problems 115
 reproducible problems 115
TSM for Advanced Copy Services 12
tuning 103

X
XINT_PROFILE 26, 28, 126

U

uninstalling Data Protection for SAP
 from UNIX 30
 from Windows 30
UNIX or Linux crontab, backup
 scheduler 67
upgrading
 Data Protection for SAP 23, 35
 Tivoli Storage Manager for ERP 26
USE_AT profile keyword 139
user authorization for Administration
 Assistant 87
user exit 6
userids
 administering 87

V

validate
 policy set 60
vendor API 3
vendor environment file 43, 44, 117, 124
VENDOROPT configuration
 parameter 27, 28
verifying
 installation of the Administration
 Assistant 48
 the Data Protection for SAP
 installation 30
versioning 123
volume 15, 58, 59, 69, 105
 manager settings 15

W

Web browser
 prerequisite for Administration
 Assistant 31
Windows, backup scheduler 67
with alternate/parallel paths and
 servers 18



Product Number: 5608-E05

Printed in USA

SC33-6341-12

