IBM Tivoli Storage Manager for Enterprise Resource Planning
Version 6.3

Data Protection for SAP®
Installation and User's Guide for Oracle

SC33-6340-12

IBM Tivoli Storage Manager for Enterprise Resource Planning
Version 6.3

Data Protection for SAP®
Installation and User's Guide for Oracle

SC33-6340-12

Note:
Before using this information and the product it supports, read the information in "Notices" on page 155.

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 SC33-6340-11.

Contents

Figures

Tables

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

Table 2. Tivoli Storage Manager storage agent publications

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

Table 3. Tivoli Storage Manager client publications

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

Table 3. Tivoli Storage Manager client publications (continued)

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

Table 4. Tivoli Storage Manager data protection publications

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

Table 5. IBM Tivoli Storage Manager troubleshooting and tuning publications

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

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). 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.

The following feature is new for Tivoli Storage Manager for ERP for Oracle Version 6.3:

The ability to create Oracle RMAN incremental backups from a FlashCopy image is provided for SAP environments by using Tivoli Storage Manager for ERP for Oracle Version 6.3 in combination with IBM Tivoli Storage FlashCopy Manager Version 3.1.

Chapter 1. Protection for SAP database servers

Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP for Oracle protects SAP® system data and is integrated with the database-specific utilities of Oracle and SAP BR* Tools, which are a set of database administration functions incorporated into SAP for Oracle databases. 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 Oracle overview

Data Protection for SAP for Oracle 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 Oracle-based SAP® systems. It is integrated with SAP backup and recovery utilities BRBACKUP, BRARCHIVE, BRRESTORE, and BRRECOVER, 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.

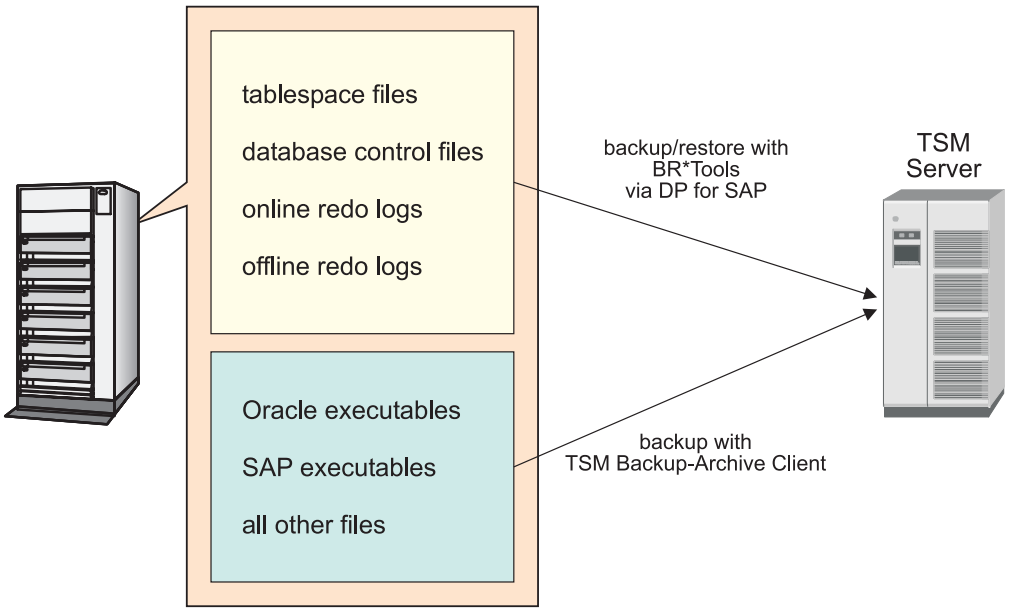

Figure 1. Scope of Data Protection for SAP for Oracle

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 data files, control files, and online or offline redo logs.

Other files (such as SAP and Oracle 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 Oracle 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 an transparent link between Oracle and BR*Tools and Tivoli Storage Manager.

Data Protection for SAP provides two adapters:

backint
This executable file is called directly by SAP and is used to perform full database backups (online and offline) and back ups of control and redo log files.

orasbt.dll
This shared media management library is dynamically linked by Oracle RMAN. When a backup is performed using this shared library, SAP communicates through Oracle RMAN instead of Data Protection for SAP. Incremental backups are also available when using RMAN with this shared library.

Both adapters share the `init`*`SID`*`.utl` profile file. This file contains information that describes how to perform backups and restores and can be customized for the Data Protection for SAP environment. Both adapters also communicate with the Tivoli Storage Manager server through an API that is shared with other IBM Data Protection products. These adapters require that the Data Protection for SAP ProLE background process is running.

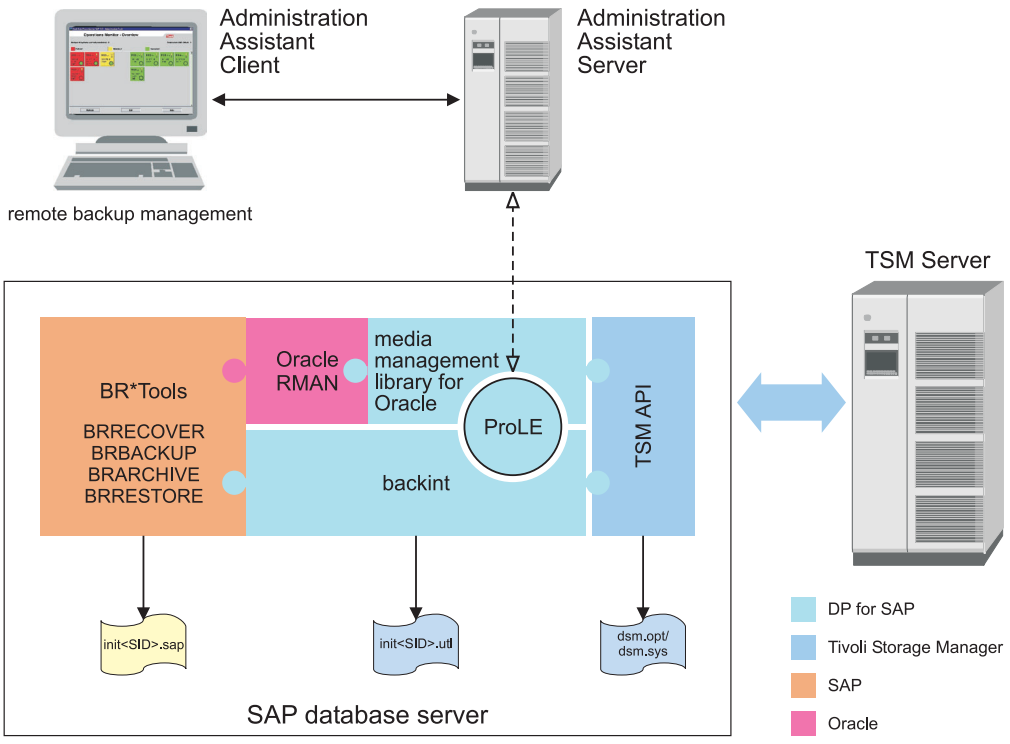

Figure 2. Integration of Data Protection for SAP with SAP

Data Protection for SAP also provides the Administration Assistant which is used to increase administrator productivity. The Administration Assistant 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 the SAP system, 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 73.

BACKINT interface

Data Protection for SAP for Oracle provides the BACKINT interface to perform full online and offline backups of Oracle databases, control files, and redo log files. The BACKINT interface communicates directly with SAP®.Figure 3 shows the interaction between BR*Tools, Data Protection for SAP, and the BACKINT interface when performing a backup or restore.

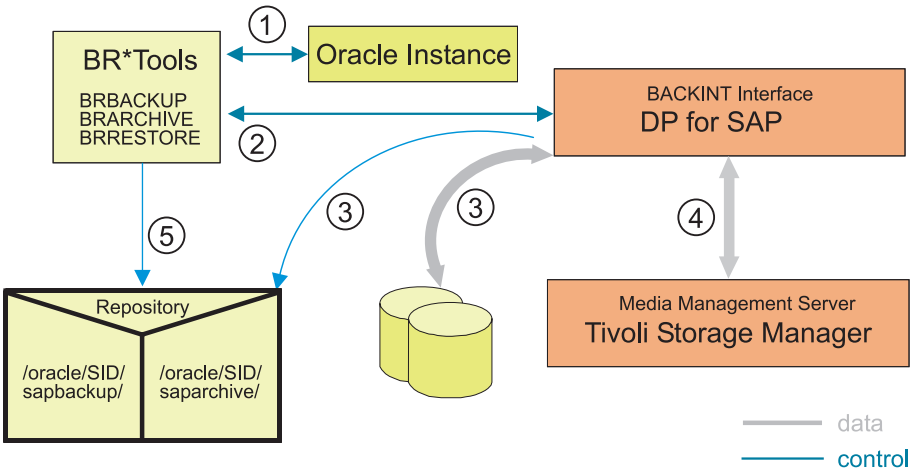

*Figure 3. Data Protection for SAP with BR*Tools using the BACKINT Interface*

The BR*Tools record the status of the Oracle data file backups and logfile backups by using tables contained within the Oracle database and system data. This information enables SAP to automatically restore the correct data files and their specific database transaction log files (redo log files), if necessary. The data files reside in the Oracle database (Oracle Instance). Data Protection for SAP runs as a separate process, independently from the database. It receives the data through the BACKINT interface and saves the data to the Tivoli Storage Manager server.

A backup operation proceeds in the following order (see circled numbers):

1. The BR*Tools utility BRBACKUP informs Oracle which data is to be backed up. It then places the database in the proper 'online or offline backup state.
2. BRBACKUP calls Data Protection for SAP through the BACKINT interface with a list of all files to be backed up.
3. Data Protection for SAP reads all requested files from the database and reports back to BRBACKUP. BRBACKUP adds these files to the repository that contains all processed backups.
4. BACKINT saves the data to the Tivoli Storage Manager server.
5. The BR*Tools update the file repository with status information about the files.

Oracle Backup/Restore and Data Protection for SAP

The SAP® database administration provides four tools (referred to as the BR*Tools) for Oracle databases:

- BRBACKUP: Provides online or offline partial or full backups of tablespaces.
- BRARCHIVE: Provides back ups of archived redo log files.
- BRRESTORE: Provides system-guided restore of Oracle backups.
- BRRECOVER: Provides recover capabilities.

These SAP database administration tools offer all the functions necessary to administer a database. Oracle also provides a Recovery MANager administration utility (referred to as RMAN) which is required to perform an incremental backup. Data Protection for SAP integrates with SAP BR*Tools and Oracle RMAN to provide unattended, 24-hour, 7-days-per-week production backup and restore tasks.

Oracle Recovery Manager (RMAN)

Oracle RMAN is used to perform a backup, restore, and recover operation of an Oracle database. RMAN is also required when performing an incremental backup.

Make sure to review SAP® support information regarding how to configure SAP on your operating system to perform a backup by using RMAN with your database version. SAP information is available at SAP Service Marketplace (*http://service.sap.com/*).

When operating RMAN, Tivoli Storage Manager for ERP is loaded by one (or more) Oracle processes as a shared library. These Oracle processes decide on how many parallel sessions are opened, when a session is opened and closed, and which data object (table space) is included in the session. Some of the above mentioned parameters must be configured for RMAN. Depending on how RMAN is used, these parameters can be configured either within the RMAN script or within the BR*Tools configuration file (init*SID*.sap).

If you want to use parallel sessions with RMAN, make sure that you configure at least the same number of sessions within the Tivoli Storage Manager for ERP configuration file as you configure for RMAN (see also “Multiple Sessions” on page 101 and “Multiple Servers” on page 100).

Figure 4 on page 5 shows the interaction between BR*Tools, Tivoli Storage Manager for ERP, Oracle RMAN, and Data Protection for SAP when performing a backup or restore.

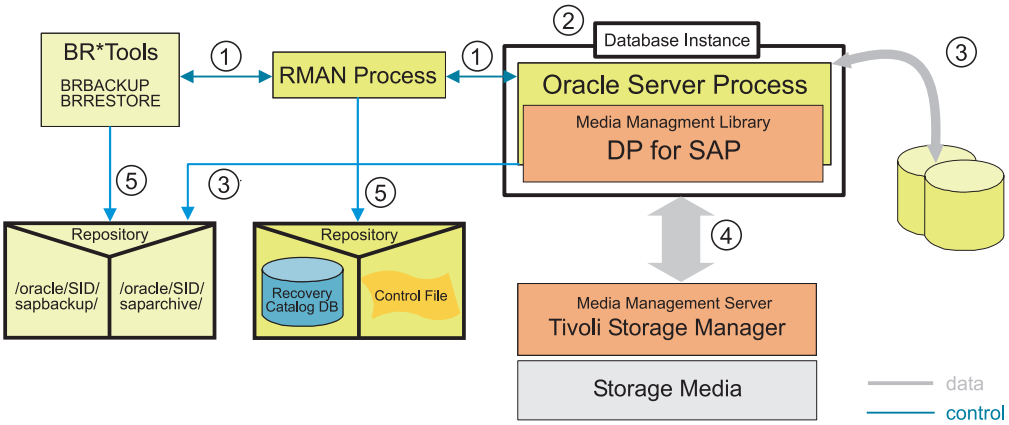

*Figure 4. Tivoli Storage Manager for ERP with BR*Tools using the RMAN Interface*

The BR*Tools use tables contained within the Oracle database and system data to record status information for the database and redo log backups.This information allows SAP to restore the correct data and their corresponding redo logs. The data files reside in the Oracle Instance of the Oracle database.

Tivoli Storage Manager for ERP runs as a linked library controlled by the Oracle Server Process.

A backup operation proceeds in the following order (see circled numbers):

1. The BR*Tools utility BRBACKUP informs Oracle RMAN which data is to be backed up. It then places the database in the proper online or offline backup state.
2. The Oracle server process loads Tivoli Storage Manager for ERP and communicates with it through the Oracle Media Management API.
3. Tivoli Storage Manager for ERP reads the requested data from the database and reports back to BRBACKUP. BRBACKUP adds this data to the repository that contains all processed backups.
4. Tivoli Storage Manager for ERP saves the data to the Tivoli Storage Manager server.
5. The BR*Tools update the file repository with status information about the data. RMAN uses a control file to maintain its own repository for a separate recovery catalog database.

A special configuration is required when Tivoli Storage Manager for ERP is used with Oracle RMAN for off-loaded backups to Tivoli Storage Manager from FlashCopy backups created with IBM Tivoli Storage FlashCopy Manager. An incremental backup is enabled by using profile parameters in the Tivoli Storage Manager for ERP configuration file, the `.utl` file. While all parameters required for this feature are described in "Tivoli Storage Manager for ERP for Oracle profile parameter descriptions" on page 120 in this publication, the configuration is detailed in the *IBM Tivoli Storage FlashCopy Manager for UNIX and Linux Installation and User's Guide*

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 Oracle 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. The use of DB2 is in no way dependent on use of the DB2 version of Data Protection for SAP. A DB2 Administration Assistant database can also be used in an Oracle environment. 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).

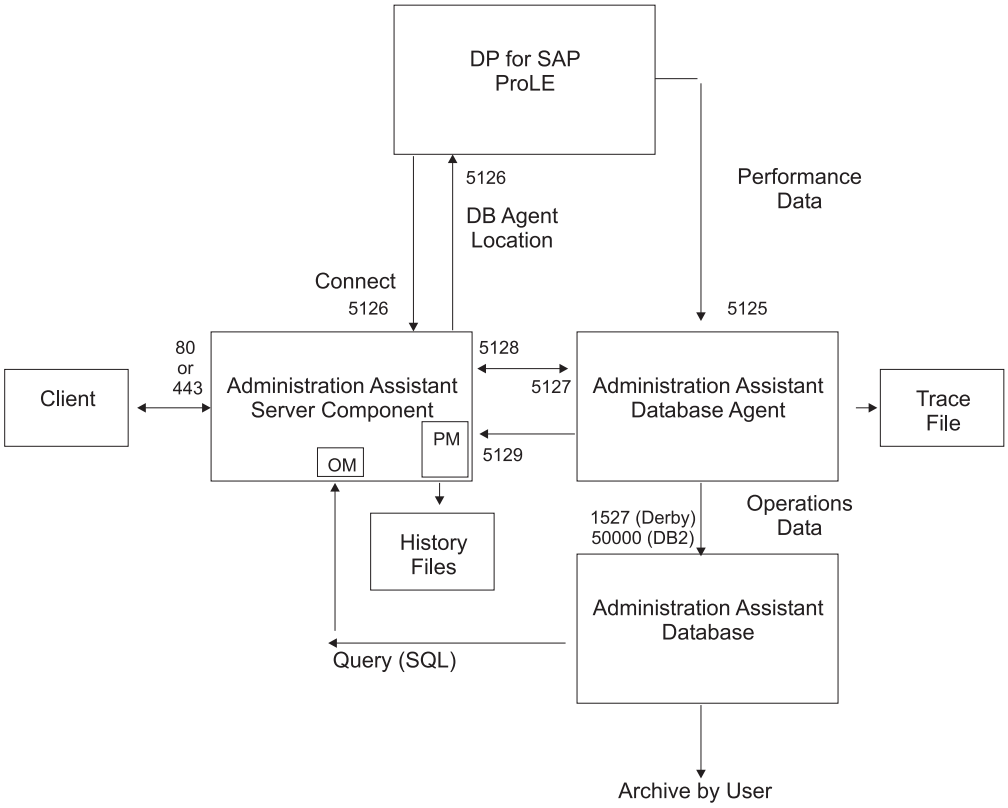

Figure 5. Administration Assistant function for Data Protection for SAP Components (with Default Port Numbers)

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 6 on page 8.

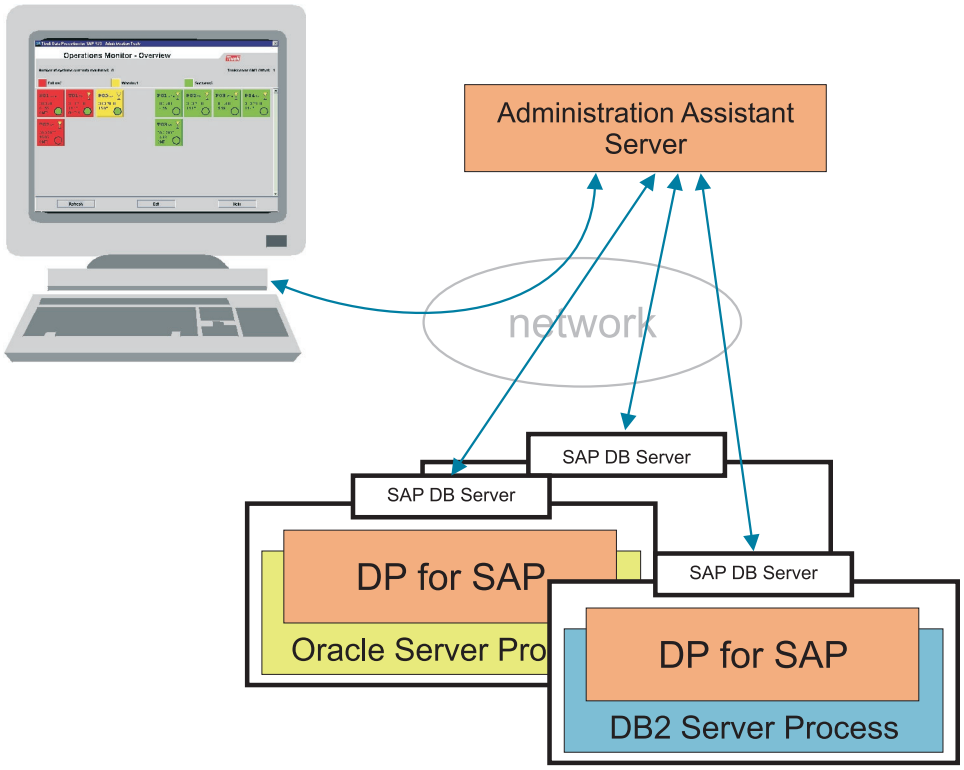

Figure 6. Example of an SAP Landscape

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 feature 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 Oracle backup or restore operations is displayed. The Administration Assistant also saves this performance data and provides a graphical presentation for later analysis.

Simulate backup / restore
Configuration changes or production restores can be tested without changing the production environment or compromising production data. This function is provided for Oracle databases in combination with the BACKINT interface.

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*, *View performance data*, and *Simulate backup / restore*.

Minimizing backup and restore processing with Tivoli Storage Manager for ERP

Although Tivoli Storage Manager for ERP for Oracle 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.

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 Oracle operations

Planning information regarding how to define an appropriate backup strategy for your SAP® system is provided.

The strategy you choose is dependent on your specific requirements. Consider these questions when reviewing this information:

- What type of events do you wish to protect your SAP® system against?
- How large is your database?
- What is the transaction rate of your database?
- How fast do you need to recover from a failure?
- What backup windows are available?

Planning a Backup Strategy for Your Oracle Database

To help prevent the loss of data associated with the Oracle database, back up these Oracle files and logs on a regular basis:

Data files
Files belonging to a specific tablespace (data files) are backed up by BRBACKUP. This is done at the file level, where offline or online backups are possible.

Control files
The control file is backed up by BRBACKUP whenever a tablespace backup occurs. Oracle provides mirroring of the control file to protect the running database system against corruption of this active file. The AIX LVM facilities can also be used to mirror these files.

Online redo logs
Online redo logs are backed up by BRBACKUP whenever a full offline database backup occurs. Oracle provides mirroring of the redo log files to protect the running database system against corruption of these active files. The AIX LVM facilities can also be used to mirror these files.

Offline redo logs
Offline redo logs are backed up by BRARCHIVE. You can specify that the redo logs be deleted from their original location when BRARCHIVE completes successfully. Additional information is available in "Cooperation of Data Protection for SAP for Oracle with BRARCHIVE" on page 117.

SAP® system data and Oracle system data should also be backed up on a regular basis using the Tivoli Storage Manager backup-archive client incremental backup feature.

Figure 7 on page 12 shows the various backup scenarios within an SAP database server machine.

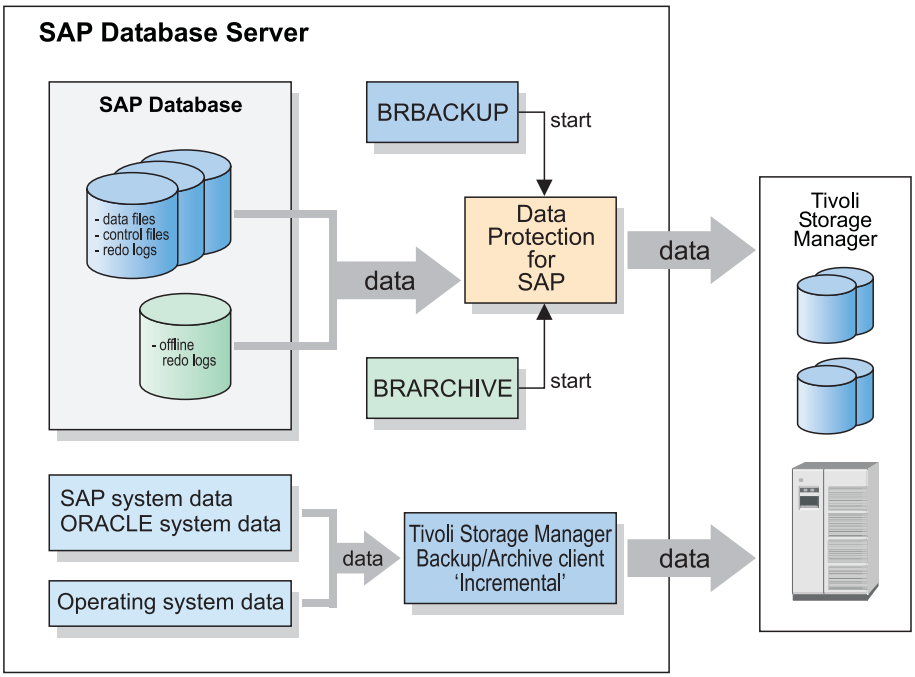

Figure 7. Backup Scenarios Within an SAP® Oracle Environment

Although the database is backed up with Data Protection for SAP, note that SAP, Oracle and operating system protocols are backed up directly by the Tivoli Storage Manager backup-archive client.

Planning a Backup Strategy for the Operating System

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 user dependent data.

Planning a Backup Strategy for Backup Protocols and Profiles

Every BRBACKUP and BRARCHIVE operation performs two actions. The first action backs up these type of objects:

- SAP® data files
- database control files
- online redo log files
- offline redo log files

Note that the type of object backed up depends on the action started (see previous section).

The second action backs up these profile and protocol files:

- BR*Tools profile (`init`*`SID`*`.dba`)
- BR*Tools Initialization profile (`init`*`SID`*`.sap`)
- Data Protection for SAP profile (`init`*`SID`*`.utl`)

- Oracle profile (init*SID*.ora)
- BRBACKUP summary log (back*SID*.log)
- BRARCHIVE summary log (arch*SID*.log)
- BRBACKUP/BRARCHIVE detailed log
- BR*Tools main log (reorg*SID*.log)
- Structure log (struct*SID*.log)

These profiles and protocols are required by the BR*Tools whenever a backup, restore, or recovery is performed. Unless these profile and protocol files have been backed up using a file system backup, they can only be restored using Data Protection for SAP for Oracle. Data Protection for SAP File Manager can be used to restore one of these files (when lost) before running a BRRESTORE or BRRECOVER command. See "Data Protection for SAP for Oracle File Manager" on page 69 for details. The directories and file systems that containing these profile and protocol files should be backed up separately. This helps prevent excessive recovery processing times when using Data Protection for SAP or the Data Protection for SAP File Manager in a disaster recovery situation.

Profile and protocol files are located in these directories:

UNIX or Linux

Directory	File
/oracle/*SID*/sapreorg	BR*Tools and structure logs
/oracle/*SID*/sapbackup	BRBACKUP logs
/oracle/*SID*/saparch	(BRARCHIVE logs
/oracle/*SID*/dbs	Profiles

Windows

Directory	File
drive:\oracle*SID*\sapreorg	BR*Tools and structure logs
drive:\oracle*SID*\sapbackup	BRBACKUP logs
drive:\oracle*SID*\saparch	BRARCHIVE logs
drive:\orant\database	Profiles

Planning a Backup Strategy for SAP® System Data

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 all relevant database data that has been backed up or archived by Data Protection for SAP for Oracle, such as data files, the control file, and online or offline redo logs. See "Include/Exclude List Sample (UNIX and Linux)" on page 144 or "Include/Exclude List Sample (Windows)" on page 145 for example include-exclude lists. Example include-exclude list files are also provided in the Data Protection for SAP installation directory.

The information saved with the BRBACKUP and BRARCHIVE utilities is usually located in the following file systems or directories. Add these directories to the list

of paths to be excluded in the Tivoli Storage Manager backup-archive client include-exclude list. These entries will prevent the database from being backed up twice.

UNIX or Linux

```
/oracle/SID/saparch/
/oracle/SID/sapdata1/
/oracle/SID/sapdata2/
/oracle/SID/sapdata3/
/oracle/SID/sapdata4/
/oracle/SID/sapdata5/
/oracle/SID/sapdata6/
     .
     .
     .
/oracle/SID/origlogA/
/oracle/SID/origlogB/
```

Windows

```
drive:\oracle\SID\saparch\
drive:\oracle\SID\sapdata1\
drive:\oracle\SID\sapdata2\
drive:\oracle\SID\sapdata3\
drive:\oracle\SID\sapdata4\
drive:\oracle\SID\sapdata5\
drive:\oracle\SID\sapdata6\
     .
     .
     .
drive:\oracle\SID\origlogA\
drive:\oracle\SID\origlogB\
```

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 Oracle" on page 94. 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 96.

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 Data Protection for SAP allows parallel access to database files 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*. This *CommonStore* product is database independent and therefore can also be used with Oracle. See "Archiving Inactive Data" on page 18 for additional information.

Size of the database files
When similar files are the same size, multiplexing can be used to improve data transfer rates.

Backup types.
Online backups save database files, control files, and redo logs non-disruptively. On the other hand, more data is saved to redo log files during an online backup. The amount of data saved to redo logs during an online backup may be decreased when using the file-online mode provided by SAP, while such a backup will take longer. Incremental backups will reduce the backup time and the amount of data to be sent to the backup server while restore time may be increased. For incremental backups, Oracle RMAN must be employed. For details on backup options, refer to your Oracle and SAP documentation.

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 Oracle 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 102.

Backup Server Considerations

Consider these items when setting up the Tivoli Storage Manager server. Note that Data Protection for SAP for Oracle 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 17 and “Configure the Tivoli Storage Manager server” on page 56.

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 data files, 'archive' (BRARCHIVE) means the backup of archived redo log files. Data Protection for SAP for Oracle 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. In addition, the backup file sequence is maintained for restore which improves restore processing time. A disk-tape storage hierarchy is recommended for backing up log files, each 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 20 MB 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 "Automation Options for Data Protection for SAP for Oracle" on page 97.

SAP administration tools can generate information about backups that reside on the Tivoli Storage Manager server. This is performed by viewing the local (detailed) backup log or using the Data Protection for SAP File Manager (`backfm`). In addition to viewing backups, File Manager also allows the administrator to bypass SAP tools in order to query, delete, or restore backups and files.

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 redo log backups, are automatically deleted.

Note: Be aware that the SAP backup log may still list deleted (expired) backups since this log cannot be updated by Data Protection for SAP.

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 Oracle 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 `init*SID*.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). See also "Sample Tivoli Storage Manager for ERP for Oracle Profile for UNIX or Linux" on page 128 or "Sample Data Protection for SAP for Oracle Profile for Windows" on page 133). 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 `TCPSERVERADDRESS` and can be located on one system (multiple paths) or several systems (multiple servers). `SESSIONS` denotes the number of parallel session 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 Oracle 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. Be aware that the DB2 CommonStore for SAP product is database independent and can be used with Oracle. 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.

Multiplexing
The same level of multiplexing as was used during backup is automatically applied during restore.

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

Creating multiple redo log copies

Data Protection for SAP for Oracle can save a number of copies of each redo log by using different Tivoli Storage Manager server management classes. By creating multiple redo log copies on separate physical media, the administrator can restore and recover a database even if a backup tape becomes corrupt. This list identifies Data Protection for SAP profile keywords that are relevant in this multiple redo log context:

- Keyword BRARCHIVEMGTCLASS denotes the Tivoli Storage Manager server management classes to be used when saving redo logs. With the use of different management classes, the backup media targeted for redo logs is separated from the backup media targeted for the database objects. Different redo log copies can also be saved to different backup media.
- Keyword REDOLOG_COPIES allows the administrator to initiate the creation of multiple backup copies of each redo log. By creating multiple copies on separate physical media, the database administrator is able to restore and recover an Oracle database in an SAP® environment even if a backup tape becomes corrupt or lost.
- Keyword MAX_SESSIONS specifies the maximum number of sessions that a single Data Protection for SAP instance is allowed to access to the Tivoli Storage Manager server.

These rules describe how Data Protection for SAP satisfies a request to back up redo log files:

- Data Protection for SAP creates as many backup copies of each redo log as are specified by the REDOLOG_COPIES keyword.
- Data Protection for SAP requires that there are as many archive management classes (as defined by the BRARCHIVEMGTCLASS keyword) as there are redo log copies requested. To best protect against the loss of data, it is important that

the different management classes are linked to different storage pools within Tivoli Storage Manager storage so that the various redo log copies reside different backup media.

- When RMAN is used, Data Protection for SAP requires that the maximum number of sessions (as defined by the MAX_SESSIONS keyword) is greater than or equal to the number of redo log copies requested. A setup with a smaller number of sessions is not recommended with the BACKINT interface.
- Data Protection for SAP cannot control the order in which Tivoli Storage Manager processes the requests. Therefore, an administrator cannot rely on sessions to be processed in the order they were started by Data Protection for SAP.

Planning for using IBM HACMP for AIX

This section provides information about Data Protection for SAP for Oracle 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:

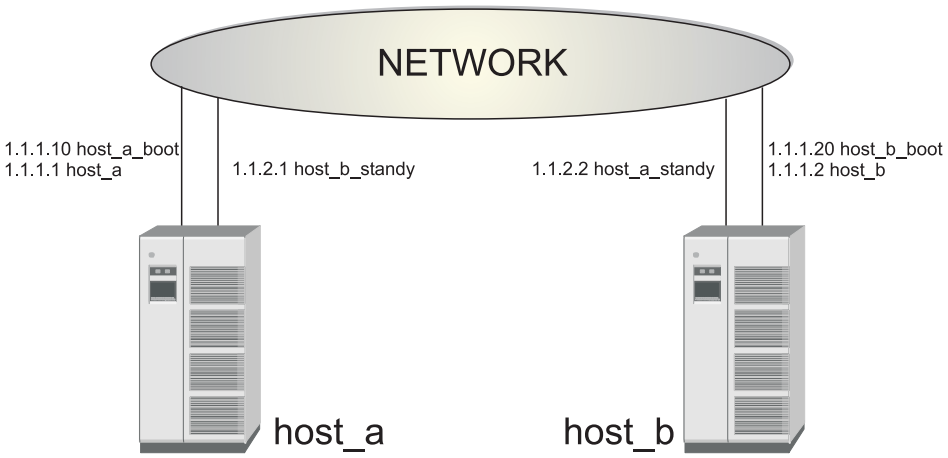

Figure 8. Sample Environment for HACMP Takeover

HACMP impact on Data Protection for SAP for Oracle

A list of Data Protection for SAP for Oracle 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`.
- Disk sorting 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 (`init`*`SID`*`.utl`, `init`*`SID`*`.bki`) in `$ORACLE_HOME/dbs`.

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

`Backint` connects to `ProLE` using the following procedure:

- Retrieves the IP address for localhost (should be 127.0.0.1 for IPv4).
- Retrieves the `backint` service (port 57323).
- Connects to 127.0.0.1:*backint service*.

Digital Signing of Executable Files for Windows Systems

Data Protection for SAP for Oracle 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-TSMERPORA-WinIA64.exe
  - *version*-TIV-TSMERPORA-WinX64.exe
- The Data Protection for SAP application executable files
  - backfm.exe
  - backint.exe
  - createinfo.exe
  - prole.exe
  - orasbt.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 Oracle for V6.3

Information needed to install the various IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP for Oracle 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 Oracle 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 Oracle 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 Oracle 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 Oracle (base product) planning sheet" on page 146.
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 9.
5. Install Data Protection for SAP as described in "Installing Tivoli Storage Manager for ERP for Oracle on UNIX or Linux" on page 26 or "Installing Tivoli Storage Manager for ERP for Oracle on Windows" on page 28. See "Upgrade the Data Protection for SAP for Oracle 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 53 and "Configure the Tivoli Storage Manager server" on page 56.
7. Verify the installation completed successfully as described in "Verify the installation" on page 39.

Note these additional requirements:

- Data Protection for SAP can be installed and operated for SAP® systems with Oracle databases employing a standard file system or raw logical volumes.

- Be aware of the minor differences between UNIX or Linux and Windows versions of Data Protection for SAP. For example, UNIX or Linux uses the path separator "/" and Windows uses the path separator "\" with a drive letter. Insignificant differences are documented where applicable.

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

Information needed to install the Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP for Oracle 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 Oracle 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 Oracle V6.3 base product" on page 35.

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

- Oracle Database
- SAP® R/3 or SAP e-business Solution, based on Oracle

  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 53. 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 Oracle in a Distributed File System" on page 48.

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 Oracle (base product) planning sheet" on page 146. 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 Oracle 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-TSMERPORA-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-TSMERPORA-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=/usr/tivoli/tsm/tdp_r3/ora64
SID=SID
SAP_CFG_FILE=/oracle/SID/dbs
SAP_BR_TOOL=/usr/sap/SID/SYS/exe/run
TSM_CFG_FILE=
TSMUTL_YES=1
TSMUTL_NO=0
TSMUTL_SERVERADRESSE=TSMServer
TSMUTL_NODE=R3NODE
TSMUTL_BACKUPMGM=MDB
TSMUTL_ARCHIVEMGM=MLOG1 MLOG2
TSMAPI_YES=
TSMAPI_NO=
TSMAPI_DSMI_DIR=
TSMAPI_DSMI_CONFIG=
TSMAPI_DSMI_LOG=
RMANYES=1
RMANNO=0
NAMEPORTAA_ADRESSE=
NAMEPORTAA_PORT=5126
```

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 Oracle on UNIX or Linux

IBM Tivoli Storage Manager for Enterprise Resource Planning for Oracle 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-TSMERPORA-*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

In the following description, you must replace the directory name ora*bit* in the installation path. Depending on the version of Tivoli Storage Manager for ERP you have installed, you must replace it with

Directory name	Bit-width version of Tivoli Storage Manager for ERP
ora	32
ora64	64

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. If the Oracle RMAN interface will be used, configure the Tivoli Storage Manager backup-archive client on your SAP database server as described in "Configure the Tivoli Storage Manager client options" on page 53.
3. Verify that the *DISPLAY* variable is set to view the installation prompts through a graphical X-Window.
4. Invoke the appropriate Tivoli Storage Manager for ERP installation file for your operating system and your Oracle database.
5. Perform these tasks if the Oracle RMAN interface was selected during the installation process:

a. Set the Tivoli Storage Manager for ERP password for Tivoli Storage Manager as described in "7. Determine the Tivoli Storage Manager password method" on page 59.
b. Make sure `/usr/lib` is specified in the library path environment of your system.
c. Customize the SAP backup profile `init`*`SID`*`.sap` to use RMAN by adding this text:

```
backup_dev_type=rman_util
rman_parms="ENV=(XINT_PROFILE=path/initSID.utl,
PROLE_PORT=portnumber,&BR_INFO)"
```

Locate the appropriate *ProLE* port number in the `/etc/services` file. Look for port name tdpr3ora64.

If Tivoli Storage Manager for ERP is not installed in the default path and backups are done by using Oracle RMAN, then the environment variable `XINT_NLS_CATALOG_PATH` must be added to the parameter `rman_pars` in the `initSID.sap` file. The value of `XINT_NLS_CATALOG_PATH` must be set to the new customized install path, otherwise the message catalog will not be found.

6. If the Oracle RMAN interface was not selected during the installation process, create these links:

```
cd $ORACLE_HOME/rdbms/lib
ln -s /usr/tivoli/tsm/tdp_r3/ora64/libtdp_r3.ext /usr/lib/libobk.ext
ln -s /usr/lib/libobk.ext $ORACLE_HOME/lib/libobk.ext
```

7. 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 internal communication.

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

```
backint
prole
createinfo
backfm
initSID.bki
libtdp_r3.ext
archive.ksh
backup.ksh
crontab.sample
dsm.opt
dsm.sys
gensortfile.sh
SanFSsetupFS.sh (AIX only)
inclexcl.list
README
README_TSMERPversionlanguage.html
```

```
TIPHINTS
agent.lic (Only after installation from disc or disc image. This file is not
present in the packages available on the FTP server.)
```

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 Tivoli Storage Manager for ERP configuration files are installed in the SAP directory (typically, /oracle/SID/dbs):

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

Installing Tivoli Storage Manager for ERP for Oracle 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-TSMERPORA-*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 the RMAN interface is to be installed:
   a. Stop the OracleService*SID* service.
   b. Configure the Tivoli Storage Manager backup-archive client on your SAP database server as described in "Configure the Tivoli Storage Manager client options" on page 53).
3. In Windows Explorer, go to the directory where the installation package is located.
4. Start the Tivoli Storage Manager for ERP executable file, and follow the instructions of the installation dialog.
5. Perform these tasks if the Oracle RMAN interface was selected during the installation process:
   a. Set the Tivoli Storage Manager for ERP password for Tivoli Storage Manager as described in "7. Determine the Tivoli Storage Manager password method" on page 59.
   b. Customize the SAP backup profile init*SID*.sap to use RMAN by adding this text:

```
backup_dev_type=rman_util
rman_parms="ENV=(XINT_PROFILE=path/initSID.utl,
PROLE_PORT=portnumber,&BR_INFO)"
```

Locate the appropriate *ProLE* port number in the *drive*:\WINNT\system32\drivers\etc\services file. Look for port name tdpr3ora64.

If Tivoli Storage Manager for ERP is not installed in the default path and backups are done by using Oracle RMAN, then the environment variable XINT_NLS_CATALOG_PATH must be added to the parameter rman_pars in the initSID.sap file. The value of XINT_NLS_CATALOG_PATH must be set to the new customized install path, otherwise the message catalog will not be found.

c. Restart the Oracle service: OracleService*SID*.

6. 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 background process is created.
- An entry required for internal communication is created in %WINDIR%\system32\drivers\etc\services.

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

```
backint.exe
prole.exe
createinfo.exe
orasbt.dll
backfm.exe
backup.cmd
server_a.opt
server_b.opt
inclexcl.list
schedule.sample
dsm.opt
README.txt
README_TSMERPversionlanguage.html
TIPHINTS
agent.lic (Only after installation from disc or disc image. This file is not
present in the packages available on the FTP server.)
```

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

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

```
initSID.bki
initSID.utl
agent.lic (Only after installation from disc or disc image. Not present in the
Web package.)
```

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 Oracle 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 Oracle under UNIX or Linux

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

1. Log in on the SAP® database server machine 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 Oracle is not running.
4. Launch the uninstall executable file and follow the instructions of the uninstall procedure. The uninstall executable file is located in one of the following directories:
   - AIX 64-bit:

     for version prior to 6.1:

     ```
     /usr/tivoli/tsm/tdp_r3/ora64/_uninst/uninstaller.bin  [-silent | -console]
     ```

     for Version 6.1 or later:

     ```
     /usr/tivoli/tsm/tdp_r3/ora64/Uninstall_TIV-TSMERPORA/
     Uninstall_TIV-TSMERPORA  [-i silent | -i console]
     ```

   - Other UNIX 64-bit or Linux 64-bit:

     for version prior to 6.1:

     ```
     /opt/tivoli/tsm/tdp_r3/ora64/_unist/uninstaller.bin [-silent | -console]
     ```

     for Version 6.1 or later:

     ```
     /opt/tivoli/tsm/tdp_r3/ora64/Uninstall_TIV-TSMERPORA/
     Uninstall_TIV-TSMERPORA [-i silent | -i console]
     ```

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

Perform these tasks to uninstall a previous version of Data Protection for SAP for Oracle 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.

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 Oracle 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 17 on page 147 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 Oracle 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

- 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 Adminstration 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 /etc/init.d
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 Oracle for V6.3

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

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

Upgrade the Data Protection for SAP for Oracle 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 Oracle 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. Make sure that the requirements for the new version of Data Protection for SAP are met as described in “Prerequisites” on page 24.
3. Make sure planning information is available as described in “Prerequisites” on page 24.
4. A full backup of the SAP® database should be performed before upgrading to the new version.
5. Uninstall the old version as described in “Uninstalling the Old Version of Tivoli Storage Manager for ERP for Oracle under UNIX or Linux” on page 30 or “Uninstalling the Old Version of Data Protection for SAP for Oracle under Windows” on page 30.
6. Install the new version of Data Protection for SAP as described in “Prerequisites” on page 24.
7. Verify the installation as described in “Verify the installation” on page 39.
8. A full backup should be performed after upgrading to the new version.
9. Following an upgrade and subsequent RMAN setup on Windows, start (or restart) service `OracleService`*`SID`* in order to activate the new Data Protection for SAP environment.

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 Oracle 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."
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 45.
8. Set up the Administration Assistant client as described in "2. Configuring the Administration Assistant function for Data Protection for SAP Client" on page 46.
9. Verify the installation as described in "3. Verifying the Administration Assistant function for Data Protection for SAP installation" on page 47.

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 Oracle 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 Oracle installation disc (or disc image) in the migration directory. This directory contains:
      - aaDerbyAdaption.jar
      - prepareExport.sql c.
      - 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 Oracle

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

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

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

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

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

Data Protection for SAP for Oracle requires that you complete certain configuration tasks before it performs a backup operation.

Verification tasks

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

Verify the installation

Preparing for the Verification for Initial and Upgrade Installations: Make sure the following considerations are met before verifying the Data Protection for SAP for Oracle installation:

- The SAP® Backup profile is configured properly. This profile can be found on UNIX or Linux systems in the path `$ORACLE_HOME/dbs` and on Windows systems in the path `%ORACLE_HOME%\database`. This configuration refers to the following keywords within that profile:

  backup_type
  Identifies the default type of the database backup. This parameter is only used by BRBACKUP (default is `offline`).

  backup_dev_type
  Determines the backup medium that will be used (default is tape). In order to use the backint interface, this parameter must be set either to 'util_file' or 'util_file_online'. For RMAN, this parameter is set to 'rman_util'

  util_par_file
  This parameter specifies the location of the parameter file. This file is required in order to perform a backup operation with an external backup program.

rman_parms
When backup_dev_type is set to "rman_util", this parameter defines various parameters required for RMAN operations.

Available values for the backup_dev_type and backup_type keywords.

Table 8. SAP Backup Profile Parameter Combinations

Operation	backup_dev_type	backup_type
Offline backup	util_file	offline
Online backup	util_file	online
Online backup with individual tablespace locking	util_file_online	online
Online backup via RMAN	rman_util	online

The SAP Backup profile parameter must be set or changed as follows in order to perform online backups with individual tablespace that lock with Data Protection for SAP:

```
backup_type      = online
backup_dev_type  = util_file_online
util_par_file    = ORACLE_HOME/dbs/initSID.utl
```

Performing the Verification for Initial and Upgrade Installations: Perform a tablespace backup using BR*Tools and then start a full online or offline backup using BRBACKUP:

```
brbackup -c -t online
brbackup -c -t offline
```

A complete restore or recovery of the entire SAP database is also recommended (using BR*Tools). However, a complete offline backup (using BRBACKUP) should be performed first. Step by step scenarios for backup and restore/recovery procedures of an SAP Oracle database using Data Protection for SAP can be found in the IBM Redbooks publication *R/3 Data Management Techniques Using Tivoli Storage Manager, SG24–5743*. IBM Redbooks can be found at http://www.redbooks.ibm.com. For backup tests, the BR*Tools utilities BRBACKUP and BRARCHIVE should be used. For restore or recovery test, only BRRECOVER should be used.

Verify the RMAN Setup on UNIX and Linux

In the following description you have to replace the directory name ora*bit* in the installation path. Depending on the version of Tivoli Storage Manager for ERP for Oracle you have installed, you must replace it with ora64 for the 64–bit version of Tivoli Storage Manager for ERP.

Perform these tasks verify that RMAN is set up correctly on the UNIX or Linux system:

1. Make sure that Oracle is linked to the correct library: /usr/lib/libobk.*ext* /usr/tivoli/tsm/tdp_r3/ora*bit*/libtdp_r3.*ext*. This link is not required in a distributed file system. See "Setting up Data Protection for SAP for Oracle with RMAN in a Distributed File System in an Adaptive Computing Environment" on page 50.
2. Remove the library specified in /$ORACLE_HOME/rdbms/lib/libobk.*ext*.

3. Make sure the installed Oracle Server is a 64-bit version.
4. Examine the `sbtio.log` located in the directory specified in the `user_dump_dest` keyword within the Oracle profile `init*SID*.ora`. This file is usually located at `oracle/SID/saptrace/usertrace/sbtio.log`.
5. Check the log file `sbtio.log` for lines starting with `BKI`. The first message for each RMAN session is:`BKI7060I: Data Protection for SAP<version and build number> session: process ID` If you cannot find any such message in the file, the library is not correctly linked with Oracle.
6. Examine the `dsierror.log` located in the directory specified with the environment variable DSMI_LOG or in the file denoted by keyword ERRORlogname in the first stanza of file `dsm.sys`.
7. To get a Tivoli Storage Manager API trace file, set the following entries in the client system options file `dsm.sys`: `tracefile */path/trace file* traceflags api api_detail config policy` The additional soft link might help: `ln -s /usr/tivoli/tsm/tdp_r3/ora*bit*/libtdp_r3.*ext* /usr/lib/libtdp_r3.*ext*.1`

See "RMAN problem resolution" on page 111 if problems or errors were encountered.

Verify the RMAN Setup on Windows

Perform these tasks to verify that the RMAN interface is setup correctly:

1. After an operation using RMAN, examine the `sbtio.log` located in the directory specified in the `user_dump_dest` keyword within the Oracle profile `init*SID*.ora`.. If the `sbtio.log` file does not exist or there is no line that begins with the letters 'BKI' within an existing `sbtio.log`, perform these tasks:
   a. Check if the shared library file `orasbt.dll` was found and loaded by Oracle.
   b. Put the shared library file `orasbt.dll` into the directory `%ORACLE_HOME%\bin`. This is the directory where `oracle.exe` resides.
   c. Stop the service OracleService*SID* and restart it.
2. Examine the `dsierror.log` located in the directory specified with the environment variable DSMI_LOG.
3. To create a Tivoli Storage Manager API trace file, set the following entries in the client options file: `tracefile *drive*:\*path*\<trace file> traceflags api`

See "RMAN problem resolution" on page 111 for additional assistance.

Profile tasks

Data Protection for SAP for Oracle 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 Oracle profile

The SERVER statement is specified in the Data Protection for SAP for Oracle 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.

Table 9. SERVER Statement and Appropriate Profile and Option File Settings.

Configuration possibilities	**Data Protection for SAP profile init*SID*.utl**	**Tivoli Storage Manager client option file dsm.sys or *server*.opt** [2]
single path; no password or manual password	SERVER *server* ADSMNODE *node*[1]	SERVERNAME *server* TCPSERVERADDRESS *address* NODENAME must not be specified
single path; automatic password by Tivoli Storage Manager	SERVER *server* ADSMNODE must not be specified	SERVERNAME *server* NODENAME *node* TCPSERVERADDRESS *address*
several paths/servers; no password or manual password	SERVER *server 1* ADSMNODE *node 1* • • • SERVER *server n* ADSMNODE *node n*	SERVERNAME *server 1* NODENAME must not be specified TCPSERVERADDRESS *address 1* • • • SERVERNAME *server n* NODENAME must not be specified TCPSERVERADDRESS *address n*
several paths/servers; automatic password by Tivoli Storage Manager[3]	SERVER *server 1* ADSMNODE must not be specified • • • SERVER *server n* ADSMNODE must not be specified	SERVERNAME *server 1* NODENAME *node 1* TCPSERVERADDRESS *address 1* • • • SERVERNAME *server n* NODENAME *node n* TCPSERVERADDRESS *address n*
several paths/servers; automatic password by Tivoli Storage Manager with Tivoli Storage Manager API 5.2 (or later) [4]	SERVER *server* ADSMNODE must not be specified TCP_ADDRESS *address 1* • • • SERVER *server n* ADSMNODE must not be specified TCP_ADDRESS *address n*	SERVERNAME <server NODENAME *node* TCPSERVERADDRESS *address*

Notes:

[1] If ADSMNODE is not specified, the host name is used.

[2] On UNIX and Linux, `dsm.sys` is the single client option file for all Tivoli Storage Manager servers. On Windows, there is a separate client option file *`server`*`.opt` for each Tivoli Storage Manager server.

[3] If two different physical machines have the same Tivoli Storage Manager node name or if multiple paths are defined on one node using several server stanzas, `passwordaccess generate` may only work for the first stanza that is used after password expiration. During the first client-server contact, the user is prompted for the same password for each server stanza separately, and a copy of the password is stored for each stanza. When the password expires, a new password is generated for the stanza that connects the first client-server contact. All subsequent attempts to connect through other server stanzas fail because there is no logical link between their copies of the old password and the updated copy generated by the first stanza used after password expiration. To avoid this situation, update

the passwords before they expire. When the passwords have already expired, perform these tasks to update the password:

1. Run dsmadmc and update the password on the 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 Oracle 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
  BRBACKUPMGTCLASS  mdb
  BRARCHIVEMGTCLASS mlog1 mlog2
# USE_AT            0 1 2 3 4 5 6

SERVER     server_b      # via network path 2
  ADSMNODE          C21
  SESSIONS          2
  PASSWORDREQUIRED  YES
  BRBACKUPMGTCLASS  mdb
  BRARCHIVEMGTCLASS 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 Oracle 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
  BRBACKUPMGTCLASS  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
  BRBACKUPMGTCLASS  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 Oracle 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
  BRBACKUPMGTCLASS  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
  BRBACKUPMGTCLASS  MDB
  BRARCHIVEMGTCLASS MLOG1 MLOG2 MLOG3 MLOG4
  USE_AT            5    # for Disaster Recovery
```

Administration Assistant function for Data Protection for SAP tasks

Data Protection for SAP for Oracle 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 47.

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 Oracle 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 73.

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 142 or "Sample Command File for Scheduling a Report from a Windows Scheduling Client" on page 142for 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 Oracle requires these tasks to be performed to configure Data Protection for SAP in a distributed file system.

Configuring Tivoli Storage Manager for ERP for Oracle 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 ora*SID* user ID instead of root. Perform these tasks to set up the ProLE service to run with the ora*SID* 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:

   ```
   po64:345:respawn:/usr/tivoli/tsm/tdp_r3/ora64/prole -p profile
   ```

   with this entry:

   ```
   po64:345:respawn:su - oraSID -c /usr/tivoli/tsm/tdp_r3/ora64/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 Oracle in a Distributed File System in an Adaptive Computing Environment

Certain setup tasks must be performed when Data Protection for SAP for Oracle 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 and links to the Data Protection for SAP executables 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/ora64 /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/ora64/SanFSsetupSID.sh
   ```

   The following information must be provided to the script:

   a. The SID for the SAP system to be backed up.
   b. If the Data Protection for SAP executable files reside in a location other than the default directory, specify that path when running the script.

c. The path for the Data Protection for SAP profile and configuration file (init*SID*.utl, init*SID*.bki).

d. 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.

Setting up Data Protection for SAP for Oracle with RMAN in a Distributed File System in an Adaptive Computing Environment

The SanFSsetupSID.sh script does not create the link /usr/lib/libobk.a to the Data Protection for SAP for Oracle shared library. Therefore, when configuring Oracle to use the Data Protection for SAP shared library (as described in "Verify the RMAN Setup on UNIX and Linux" on page 40), specify the full path and name of the library located in the directory residing in the distributed file system. Add this directory to the library path environment of the database instance owner. Do not link Oracle with the library in /usr/lib. This prevents the database from failing to start if the instance is moved to a different host.

HACMP tasks

Data Protection for SAP for Oracle 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 Oracle 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 Oracle must be defined as an application to HACMP. Although the *HACMP for AIX Installation Guide*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.

Start Script
Enter the full pathname of the script that starts the server (for example, /usr/sbin/cluster/events/utils/start_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 start the server.

Stop Script
Enter the full pathname of the script that stops the server (for example, /usr/sbin/cluster/events/utils/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 Oracle to an HACMP Resource Group:

A final step in enabling Data Protection for SAP for Oracle for HACMP failover is to define it to a cluster resource group. Although the *HACMP for AIX Installation Guide*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                                                              #
#                                                                                     #
#    Function: A sample shell script to update the disk information      #
after the SAP instance is unmounted.                                                  #
#                                                                                     #
# # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # # #

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


# Function to update all disk information for Data Protection for SAP
STOP_PROCESSING()

{
# 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 Protecion 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_PROCESSING


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 Oracle operation are provided.

Data Protection for SAP for Oracle 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 Oracle 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 Oracle.

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 144. 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
TCPPort           1500
TCPServeraddress  xxx.xxx.xxx.xxx or servername
InclExcl          /usr/tivoli/tsm/client/ba/bin/inclexcl.list
Compression       OFF
```

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 `inclexcl.list`. This is optional. You may want to exclude all database files that are processed by the BR*Tools.
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. It might be necessary to experiment with Tivoli Storage Manager client software compression settings to determine its impact in your environment. Tivoli Storage Manager client software compression generally improves performance only when network throughput is low.
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.ext` that points to the `libApiDS.ext` file in the Tivoli Storage Manager API installation directory (`/usr/tivoli/tsm/client/api/bin64`). See "Required installation tasks" on page 23 for a detailed description of the meaning of the *ext*.

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 suggested configuration prepares the system for TSM API Client tracing for both backint and RMAN operation. For more information, see "Verify the RMAN Setup on UNIX and Linux" on page 40.

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 145. 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\inclexcl.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 inclexcl.list. This is optional. You may want to exclude all database files that are processed by the BR*Tools.
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. It might be necessary to experiment with Tivoli Storage Manager client software compression settings to determine its impact in your environment. Tivoli Storage Manager client software compression generally improves performance only when network throughput is low.

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 Oracle 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 Oracle profile.
2. Adapt the Tivoli Storage Manager options files as described in "8. Verify the Tivoli Storage Manager server name" on page 63.
3. Set and save the Tivoli Storage Manager password for the new server as described in "Set the Tivoli Storage Manager password" on page 61.

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 (data files, online and offline redo log, control 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 data files : `define stgpool sap_db` *`device_class_name`* `maxscr=20`
3. Define a storage pool for the first copy of offline redo log files : `define stgpool sap_log1` *`device_class_name`* `maxscr=3`
4. It is strongly recommended that you back up the offline redo log files twice on two different Tivoli Storage Manager volumes. For this purpose, you have to

define an additional storage pool for the second copy of offline redo log files:
`define stgpool sap_log2 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 redo logs and copies of offlineredo logs:

```
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 Oracle 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.

6. Set the IdleTimeOut parameter:

For simulations of network transfer and media rates, the Tivoli Storage Manager server must be configured so that sessions do not time out during simulation. This is achieved by setting the parameter IdleTimeOut to a value higher than the time required for sending the largest table space file to Tivoli Storage Manager. For example:

```
setopt IdleTimeOut 60
```

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 Oracle 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:

  (UNIX or Linux):

```
backint -p full path to UTL file/initSID.utl -f password
```

(Windows):

```
backint -p full path to UTL file\initSID.utl -f password
```

On Windows, the path can also be specified in UNC notation (for example: `-p \\SERVER_A\dpsap\initSID.utl`. However, the password updates need to be synchronized on the Tivoli Storage Manager server with the `update node` command. These steps must also be repeated whenever the Tivoli Storage Manager password expires. Therefore, this method is only recommended during installation or testing, and a long password expiration period should be specified. manual password handling is not recommended for production operations.

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

```
backom -e 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 are planning to use Oracle RMAN and 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:

(UNIX or Linux):

```
backint -p full path to UTL file/initSID.utl -f password
```

(Windows):

```
backint -p full path to UTL file\initSID.utl -f password
```

On Windows, the path can also be specified in UNC notation (for example: `-p \\SERVER_A\dpsap\initSID.utl`This method is recommended for an automated production environment.

Set the Tivoli Storage Manager password:

Data Protection for SAP for Oracle 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 59 in order to adjust the Data Protection for SAP parameters.

Provide Data Protection for SAP for Oracle with the password for the Tivoli Storage Manager node by performing these steps in the shell:

1. Log in as the Oracle user.
2. Enter the following command for Windows:

```
backint -p full path to UTL file\initSID.utl -f password
```

   On Windows, the path can also be specified in UNC notation (for example: `-p \\SERVER_A\dpsap\initSID.utl`
3. Enter the following command for UNIX or Linux:

```
backint -p full path to UTL file/initSID.utl -f password
```

4. Enter the password when prompted. On HP-UX, the password is limited to 8 characters. Make sure that the Tivoli Storage Manager password for HP-UX clients does not exceed this limit.

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 59.

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 *n days* (see note 2)	ON *n days*

Table 10. Password Handling for UNIX or Linux (continued)

Step	Profile/Action	Parameter	Password		
			No	**Manual**	**Set by Tivoli Storage Manager**
2	dsm.sys	PASSWORDACCESS PASSWORDDIR (see note 5) NODENAME	Unavailable	PROMPT Unavailable Unavailable.	GENERATE *path* *nodename*
3	Tivoli Storage Manager admin	UPDATE NODE (see notes 1, 6)	Unavailable	*password*	*password*
4	Data Protection for SAP for Oracle profile (init*SID*.utl)	For each SERVER statement specify:PASSWORDREQUIRED ADSMNODE	NO *nodename*	YES *nodename*	NO (see note 4)
5	Data Protection for SAP command line	Specify in each SERVER statement: `backint -p init`*`SID`*`.utl -f password`	Unavailable	*password* (see notes 3,7,9)	*password* (see notes 3,7,9)

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 *SID*`adm` and `ora`*SID* must have read and write permission for the `path` specified in the `PASSWORDDIR` option in the Tivoli Storage Manager client options file.
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 Oracle RMAN with PASSWORDACCESS GENERATE, backups must always be started 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 59.

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 *n days* (see note 2)	ON *n days*
2	*server*.opt	PASSWORDACCESS PASSWORDDIR (see note 5) NODENAME	Unavailable	PROMPT Unavailable Unavailable	GENERATE *path* *nodename*
3	Tivoli Storage Manager admin	UPDATE NODE (see notes 1,6)	Unavailable.	*password*	*password*
4	Data Protection for SAP for Oracle profile `init`*`SID`*`.utl`	For each SERVER statement specify: PASSWORDREQUIRED ADSMNODE	NO *nodename*	YES *nodename*	NO (see note 4)
5	Data Protection for SAP command line	Specify in each SERVER statement: `backint -p init`*`SID`*`.utl -f password`	Unavailable	*password* (see note 1)	*password*

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 `SID`adm and `sapservice`*`SID`* must have read and write permission for the `path` specified in the `PASSWORDDIR` option in the Tivoli Storage Manager client options file.
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 `init`*`SID`*`.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 59).

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 `init`*`SID`*`.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 init*SID*.utl.

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

Information needed to back up, restore, and clone 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 118 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.

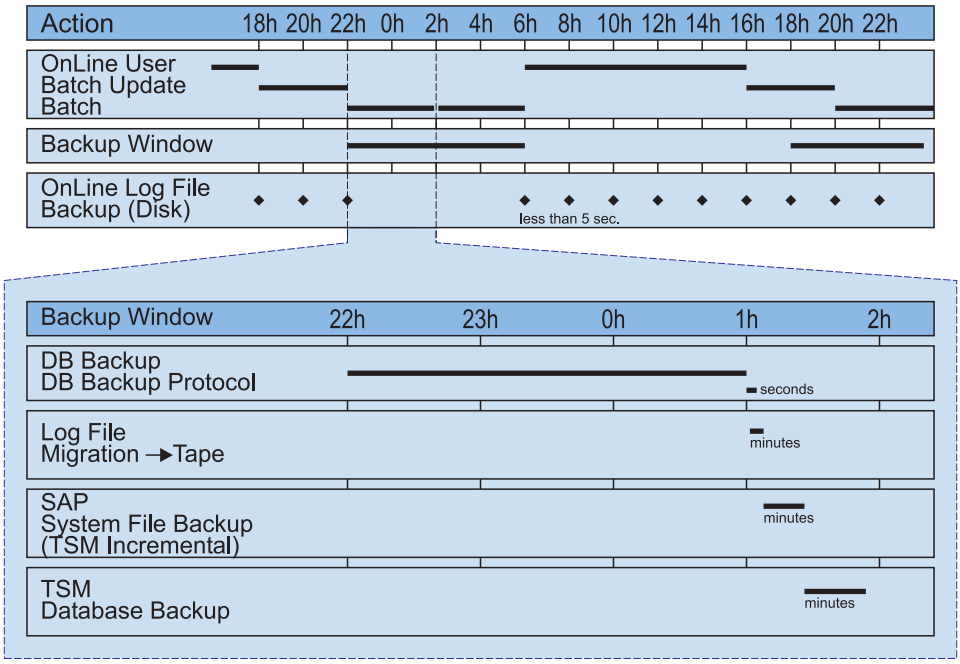

Figure 9. Production Backup Example

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 redo logs 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 drive:\oracle\SID\sapscripts\backup.cmd
```

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 SAP 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 ORACLE database
rem is assumed to be "C21".
rem
rem ------------------------------------------------------------------------
rem Full database backup, scheduled every Friday at 8:00 p.m.
rem
at 20:00 /every:f cmd /c c:\oracle\C21\sapscripts\backup\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
at 11:30 /every:m,t,w,th,f cmd /c c:\oracle\C21\sapscripts\backup\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 BRBACKUP batch file
rem ------------------------------------------------------------------------
rem Task:
rem Invokes the SAP utility BRBACKUP in order to perform a full offline
rem backup of all tablespaces using Data Protection for SAP (R)
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 ORACLE database
rem is assumed to be "C21".
rem
rem ------------------------------------------------------------------------
rem
rem First, let's do a full offline backup of the ORACLE database. This
rem includes at least files located in the following file systems:
rem c:\oracle\C21\sapdata0
rem c:\oracle\C21\sapdata1
rem c:\oracle\C21\sapdata2
```

```
rem c:\oracle\C21\sapdata3
rem c:\oracle\C21\sapdata4
rem
rem Remarks on the parameters of BRBACKUP:
rem
rem -u system/manager ORACLE username/password
rem -c run BRBACKUP in quiet mode
rem -m all backup all tablespaces
rem -t offline perform backup offline
rem
rem The following should be configured within the SAP profile
rem initC21.sap:
rem
rem backup_dev_type = util_file
rem causes BRBACKUP to use the external program
rem Data Protection for SAP (R)
rem util_par_file = %ORACLE_HOME%\database\initC21.utl
rem Data Protection for SAP (R) profile
rem ------------------------------COMMAND----------------------------------
brbackup -u system/manager -c -m all -t offline
```

Full Offline Backup Shell Script Sample

```
#!/bin/ksh
# ------------------------------------------------------------------------
# backup.ksh:
# Sample BRBACKUP shell script
# ------------------------------------------------------------------------
# Task:
# Invokes the SAP utility brbackup in order to perform a full offline
# backup of all tablespaces using Data Protection for SAP (R) technology.
# ------------------------------------------------------------------------
#      *****     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 of the ORACLE database
# is assumed to be 'C11'.
#
# ------------------------------------------------------------------------
#
# First, lets do a full offline backup of the ORACLE database. This includes
# at least files located in the following filesystems:
#      /oracle/C11/sapdata0
#      /oracle/C11/sapdata1
#      /oracle/C11/sapdata2
#      /oracle/C11/sapdata3
#      /oracle/C11/sapdata4
#
# Remarks on the parameters:
#
# -u system/manager     Oracle username/password
# -c                    run brbackup in quiet mode
# -m all                backup all tablespaces
# -t offline            perform backup offline
#
# The following should be configured within the SAP profile initC11.sap:
#
# backup_dev_type = util_file
#  causes brbackup to use the external program backint
# util_par_file =  initC11.utl
#  Data Protection for SAP profile
```

```
#
# --------------------------------COMMAND----------------------------------
brbackup -u system/manager -c -m all -t offline
```

Restoring SAP® data

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

Perform the tasks required for your operating system.

Data Protection for SAP for Oracle File Manager

The Data Protection for SAP for Oracle File Manager is a supplementary tool that simplifies the Data Protection for SAP inquire, restore, and delete operations. However, users with Oracle database restore and recovery experience and knowledge should use this tool. BR*Tools is the standard tool for restore operations. Consider these important characteristics before using the File Manager:

- The File Manager perform all operations by using the standard functions provided by Data Protection for SAP.
- The interface consists of a split window that is character based. In the left window, all backup IDs found on all Tivoli Storage Manager servers that match the backup ID prefix configured in the Data Protection for SAP profile are displayed. In the right window, all the files belonging to the selected backup ID are displayed. Individual backup IDs or multiple files are available for selection as shown in Figure 11 on page 70).

1. Start the File Manager with the path and name of the Data Protection for SAP profile. The user must be a member of the `dba` group: (UNIX or Linux):

   ```
   backfm -p /oracle/SID/dbs/initSID.utl [-o log file directory]
   ```

   (Windows):

   ```
   backfm -p drive:\orant\database\initSID.utl [-o log file directory]
   ```

   If the `-o` parameter is specified at startup, the default directory for log files will be changed.
2. The File Manager calls the `backint` executable file to connect to the Tivoli Storage Manager server configured in the Data Protection for SAP profile. If this call fails, the File Manager shows an error message but does not analyze the cause of the error. Use the backint inquire function as described in (“Inquire function” on page 115) to analyze the error.
3. An automatic inquire operation for all backup IDs is performed by the File Manager. Figure 10 on page 70 displays a set of backup IDs located by an inquiry procedure. If you mark the backup ID you are interested in and then press the **Tab** key to move the cursor to the right-hand panel, all file names belonging to the marked backup ID will be displayed as shown in Figure 11 on page 70.

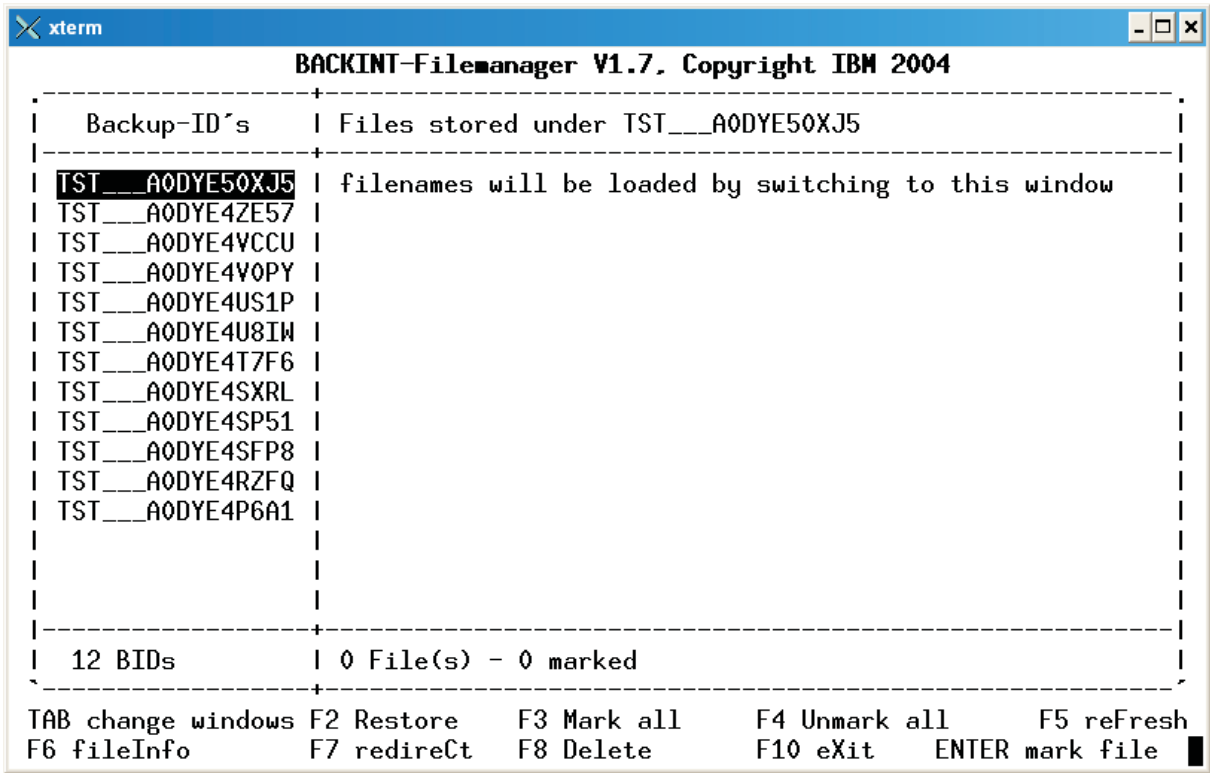

Figure 10. File Manager — Result of an Inquiry Procedure

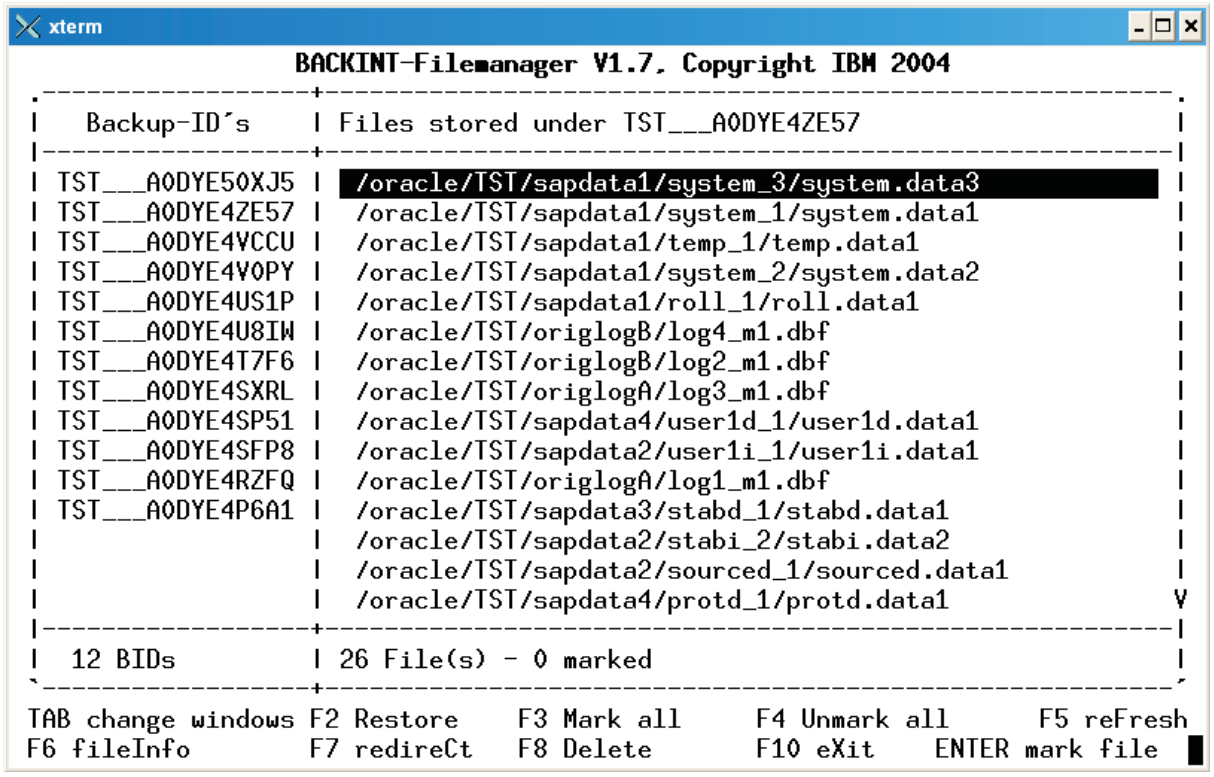

Figure 11. File Manager — Result of an Inquiry Procedure Showing File Names

The following function keys are defined for performing restore or delete operations:

Up, Down, Left, Right - Move cursor
Move the highlighted cursor in the direction indicated on the key.

Tab - Switch window side
Move the cursor between the left and right sides of the window.

F2 - Restore
Restore all marked files. Before the restore actually begins, you can specify a common destination path and you will be asked to confirm the restore process. If you specify a destination path, all marked files will be restored to that directory. Otherwise the files will be restored to the directory from which they were backed up.

For restore operations, the desired files first have to be marked. This can be done either with the **F3** function key to mark all the files that were found or with the **ENTER** key to mark only one desired file. Marked files can be identified by the symbol " * " in front of the filename. Only the marked files will be restored. For every restore operation, a log file will be created in the following location:

- (UNIX or Linux): $SAPDATA_HOME/sapbackup/backfm_*timestamp*.log
- (Windows): timestamp>.log

F3 - Mark all
All files belonging to the current backup ID will be marked.

F4 - Unmark all
Unmark all files belonging to the current backup ID.

F5 - Refresh
Refresh the list of backup IDs and file names.

F6 - Fileinfo
Opens a separate window to display file information.

For backup IDs, the sequence number is displayed (backup version count, for more information see on page "Tivoli Storage Manager for ERP for Oracle profile parameter descriptions" on page 120). For files, the Tivoli Storage Manager expiration date and time is displayed.

F7 - Redirected Restore
Restores the selected files to a new location. A new directory structure is created. The new path names are derived from the original paths by replacing the original SID with the target SID. Filenames are not modified. Redirected restore makes cloning of SAP® systems easier. See also "Cloning the SAP® System" on page 79. To clone a database you need to restore the database files to a different directory structure. In the path names of the new directory structure, the Oracle SID is replaced by the new SID. Please note that the file names are left untouched by this function. You first have to mark the files for restore. This can be done either with the **F3** function key to mark all files of a backup ID or with the **ENTER** key to mark only the highlighted file. Marked files can be identified by the symbol " * " in front of the filename. Press **F7** to start the redirected restore.

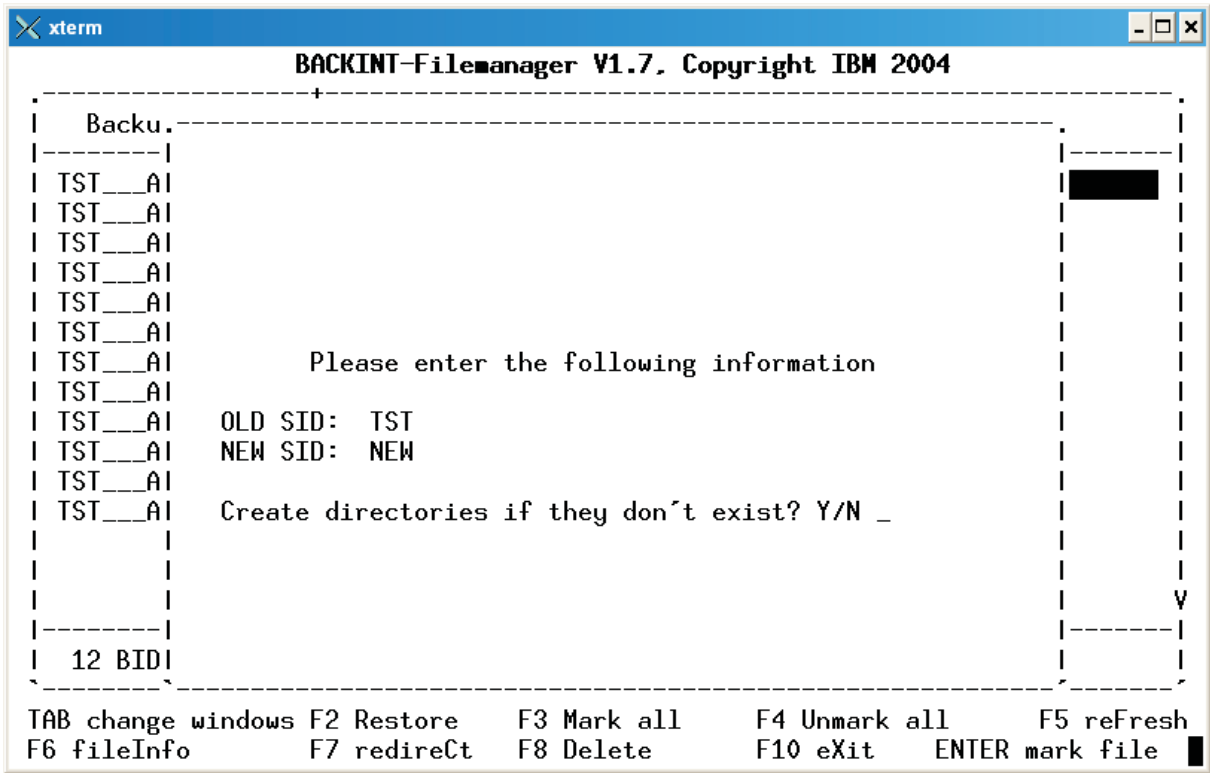

Figure 12. File Manager — Result of a Redirected Restore Procedure

F8 - Delete
Delete the selected backup ID and all corresponding files. The File Manager can delete backup IDs with all included files. It is not possible to delete single files within a backup ID. To delete a backup ID it must be highlighted. After pressing **F8** you have to confirm the deletion operation. The backup ID and all included files are deleted from the Tivoli Storage Manager server.

F10 - Exit
Exit from Data Protection for SAP File Manager

ENTER - Mark/unmark file
Mark or unmark the file below the cursor.

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 Oracle

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 tdpr3ora64 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 Oracle (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 tdpr3ora64 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 47.

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 47. 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.

Cyclic Procedure for Optimizing your Configuration

The Administration Assistant function for Data Protection for SAP Database User ID provides the ability to analyze performance, modify the configuration, and test the effects of configuration changes without having to modify the production environment.

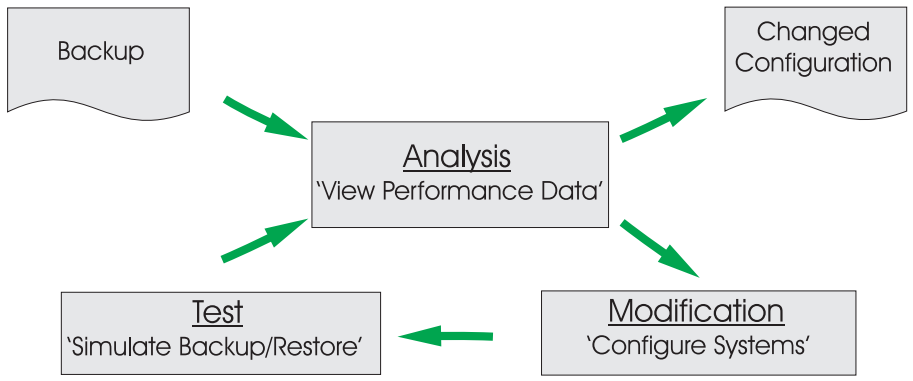

Figure 13. Optimizing your Configuration with the Administration Assistant function for Data Protection for SAP

The optimization cycle starts with a full backup of the database using the file interface (BRBACKUP). The performance data is analyzed using the View Performance Data function. This function provides insight as to possible Data Protection for SAP for Oracle configuration or infrastructure changes. These changes are temporarily implemented in a test profile with the Configure Systems function. Using the Simulate Backup/Restore function, another backup or restore is simulated to test the configuration changes. The View Performance Data function can then be used to verify whether the changes produced the desired results. This cycle can be implemented multiple times until the desired results are met. Once the configuration changes are confirmed, they can be propagated to the production system. Note that all configuration changes and simulation runs within the cycle are maintained separately from the production system.

Determining Throughput Rates

Table 12. Summary: How to Determine Throughput Rates

	Disk I/O Rate	**Storage Media Rate**	**Network Throughput Rate**
Simulation Type	No data moved to/from Tivoli Storage Manager	No data moved to/from disk	No data moved to/from disk
Disk Transfer Rate	-	infinite	infinite
Network Transfer Rate	infinite	-	-
Tape Rate	infinite	-	-
RL Compression	off	off	off
Multiplexing	1	1	1
Number of Sessions	1	1	maximum possible number

Determining the Actual Disk I/O Rate

Run a simulated backup of type No data moved to Tivoli Storage Manager in order to determine the actual disk reading rate. Both the Tape Transfer Rate and the Network Transfer Rate must be set to infinite in order to make sure there will not be a network bottleneck. Compression must be turned off and the View Performance Data function should show 100 % disk utilization. The overall throughput you get with this configuration is the rate at which data is read from disk. In order to determine the actual disk writing rate, run a simulated restore of type No data moved from Tivoli Storage Manager. Both the Tape Transfer Rate and the Network Transfer Rate must be set to infinite in order to make sure the system will not create a network bottleneck. Compression must be turned off and the View Performance Data function should show 100 % disk utilization. The overall throughput you get with this configuration is the rate at which data is written to disk. Be aware that an increased disk I/O rate is shown while data is written to the file system cache.

Determining the Actual Network Throughput Rate

Run a simulated backup of type No data moved from disk in order to determine the actual network throughput rate. The Disk Transfer Rate must be set to infinite in order to make sure there will not be a disk bottleneck. Increase the number of sessions to the maximum number possible (for example, the number of available tape drives). To be sure that the limiting factor is not the tape transfer rate, the throughput rate must be less than the media rate as provided in the 'Determining the Actual Throughput Rate of Storage Media' section, multiplied by the number of sessions. The View Performance Data function should show 100% network utilization. The overall throughput provided with this configuration is the network throughput rate.

Reporting on Simulations

An overview of simulation parameters and results for a single SID is contained in the Simulation Report. It is requested from the Available Simulation Results panel and displays a screen capture similar to this panel.

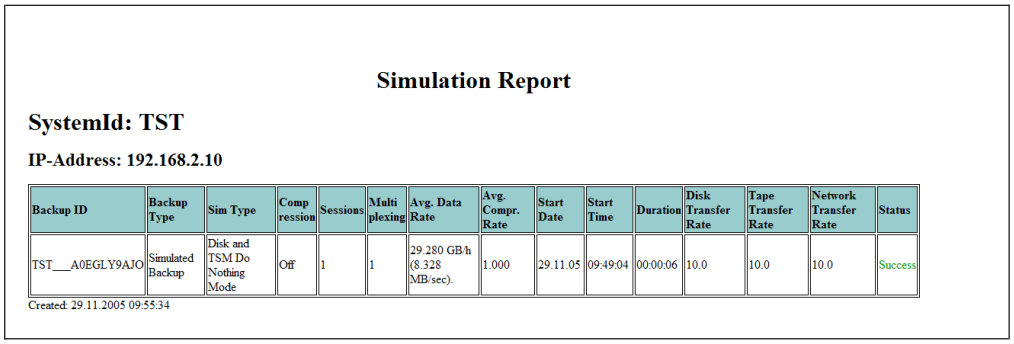

Simulation Report

SystemId: TST

IP-Address: 192.168.2.10

Backup ID	Backup Type	Sim Type	Comp ression	Sessions	Multi plexing	Avg. Data Rate	Avg. Compr. Rate	Start Date	Start Time	Duration	Disk Transfer Rate	Tape Transfer Rate	Network Transfer Rate	Status
TST___A0EGLY9AJO	Simulated Backup	Disk and TSM Do Nothing Mode	Off	1	1	29.280 GB/h (8.328 MB/sec).	1.000	29.11.05	09:49:04	00:00:06	10.0	10.0	10.0	Success

Created: 29.11.2005 09:55:34

Figure 14. Simulation Report

Simulating Backup and Restore

The Administration Assistant function for Data Protection for SAP Simulate Backup/Restore function requires a full backup of the database using the file interface (BRBACKUP). Backups done using the RMAN interface cannot serve as a basis for simulation. Both backups and restores can be simulated. A restore simulation might provide information regarding the duration of the restore operation but it will not affect your production system. When there are two (or more) eligible backups available, the latest one is used as the basis for simulation. Compression should be enabled for the base backup and the COMPR_INFO parameter should specify a valid file in the Data Protection for SAP for Oracle profile. These environment components are available for simulation:

Disk I/O

No data is read from the disk when simulating disk I/O for a backup. Data is generated in memory instead. When simulating disk I/O for a restore, data is consumed and is not written to disk. The disk I/O rate to be used for the simulation is set by the administrator as described in “Determining the Actual Disk I/O Rate” on page 77.

Network transfer and media rates

No data is sent through the network when simulating network transfer and media rates for a backup. The data is consumed instead. When simulating network transfer and media rates for a restore, no data is expected from the network and the data is generated in memory instead. However, a connection to the Tivoli Storage Manager server needs to be maintained during the simulation. Therefore, configure the Tivoli Storage

Manager server so that the sessions do not time out as described in "6. Set the IdleTimeOut parameter" on page 59. The network throughput rate and the media rate used for the simulation can be set by the administrator. "Determining the Actual Network Throughput Rate" on page 78 and "Determining the Actual Throughput Rate of Storage Media" provide information about determining these rates.

Configuration changes

When simulating configuration changes, performance parameters (in the Data Protection for SAP profile) can be modified to test for the optimum configuration within a given infrastructure. During a backup configuration simulation, data is read from disk and written to a special file space in Tivoli Storage Manager and does not affect production backups. During a restore configuration simulation, data is retrieved from the Tivoli Storage Manager server and written to disk before they are deleted. See "Performance Options of Data Protection for SAP for Oracle" on page 94 for information about Data Protection for SAP profile parameters that affect data throughput.

Determining the Actual Throughput Rate of Storage Media

Run a simulated backup of type No data moved from disk in order to determine the actual writing rate of a tape. The Disk Transfer Rate must be set to infinite in order to make sure there will not be a disk bottleneck. The number of sessions must be set to one, compression must be turned off, and the View Performance Data function should show 100 % disk utilization. The overall throughput provided with this configuration is the rate at which data is written to the storage media unless the network rate is lower than the media rate. In order to determine the actual reading rate of a tape, run a simulated restore of type No data moved to disk. The Disk Transfer Rate must be set to infinite in order to make sure there will not be a disk bottleneck. The number of sessions must be set to one. In order to exclude a CPU bottleneck, make sure that the View Performance Data function shows 100 % network utilization. The overall throughput provided with this configuration is the rate at which data is read from the storage media unless the network rate is lower than the media rate. Note that the throughput rate might not increase when the number of sessions is increased. In this situation, the network throughput rate is lower than the media rate, and the media rate cannot be determined with the Administration Assistant.

Cloning the SAP® System

This information regarding how to clone an SAP® system should be used to complement the primary SAP documentation *R/3 Homogeneous System Copy* and *R/3 Installation on UNIX / Windows - Oracle Database*. Make sure the SAP documentation is correct for the environment. SAP documentation is available at http://sapnet.sap.com and on the SAP Documentation Guides CD. Additional information about SAP system cloning can be found in the IBM Redbooks publication *SAP R/3 Data Management Techniques Using Tivoli Storage Manager*. The book can be downloaded at http://www.redbooks.ibm.com.

What is Cloning?

SAP system cloning refers to an operation where an exact copy of one source SAP system (original system) is copied to a target SAP system (destination system). The copy is considered an homogeneous system copy when the original system and destination system contain the same SAP release level, operating system, and database version. The copy is considered an heterogeneous system copy when the

SAP release level, operating system, and database version are not the same. Detailed information about these two system copy scenarios can be found in SAP Notes 86859 and 86860.

SAP system cloning considered appropriate in these situations:

- Setting up an SAP system landscape (development, quality assurance, and production system).
- After a hardware upgrade is completed.
- Creating multiple SAP test or demonstration systems.

Performing SAP® System Cloning when automatic password handling is used

Although this procedure is provided as a reference, SAP® documentation should be used as the primary instructions when cloning SAP systems. For SAP-specific changes, see also SAP Note 71254. This procedure assumes this environment:

- Two SAP R/3 systems are installed and operating on two different machines.
- Data Protection for SAP for Oracle is installed and operating on both SAP R/3 systems.

This procedure describes the tasks necessary to restore an Oracle SID to a different machine with a different SID. Use the procedure that reflects the password handling method for the environment.

Perform these tasks when automatic password handling (passwordaccess=generate) is used:

1. Make sure that the same nodename and password that are specified in the Tivoli Storage Manager client options file on the source system are specified on the target system.

   Note: Make sure the client uses the password that is stored on the Tivoli Storage Manager server. Although passwords are stored in different locations, the only original password is the one that resides on the Tivoli Storage Manager server.
2. Make a backup copy of the client option file on the target system.
3. Copy the client option file from the source system to the target system.
4. Edit the client option file and add `NODENAME source system` to the server stanza.
5. Reset the Tivoli Storage Manager password for the target system node on the server.
6. As root (UNIX or Linux) or administrator (Windows), set the new password on the client.
7. Make a backup copy of the `init`*`SID`*`.utl` file on the target system.
8. Copy the `init`*`SID`*`.utl` file from the source system to the target system. Rename the file from `init`*`SID`*`.utl` to `init`*`target_SID`*`.utl`.
9. Edit the `init`*`SID`*`.utl` file on the target system to reflect all the correct file and path names, especially for CONFIGFILE and TRACEFILE.
10. Restore the database under the SAP considerations.
11. After the restore, reset the client option file and `init`*`SID`*`.utl` file to their original and set the passwords on the target system.
12. Reset the passwords on the source system.

Detailed information regarding automatic password handling is available on page "7. Determine the Tivoli Storage Manager password method" on page 59.

Performing SAP® System Cloning when manual password handling is used

Although this procedure is provided as a reference, SAP® documentation should be used as the primary instructions when cloning SAP systems. For SAP-specific changes, see also SAP Note 71254. This procedure assumes this environment:

- Two SAP R/3 systems are installed and operating on two different machines.
- Data Protection for SAP for Oracle is installed and operating on both SAP R/3 systems.

This procedure describes the tasks necessary to restore an Oracle SID to a different machine with a different SID. Use the procedure that reflects the password handling method for the environment.

Perform these tasks when manual password handling (passwordaccess=prompt) is used:If you are using `passwordaccess=prompt`, you only need to set the nodename/password in the `init`*`SID`*`.utl` file:

1. Create a backup copy of the `init`*`SID`*`.utl` file on the target system.
2. Copy the `init`*`SID`*`.utl` file from the source system to the target system. Rename the file from `init`*`SID`*`.utl` to `init`*`target_SID`*`.utl`.
3. Edit the `init`*`SID`*`.utl` file on the target system to reflect all the correct file and path names, especially for CONFIGFILE and TRACEFILE.
4. As *`SID`*`adm` user, set the Data Protection for SAP for Oracle password on the target system: (UNIX or Linux):

   ```
   backint -p /oracle/SID/dbs/initSID.utl -f password
   ```

   (Windows):

   ```
   backint -p drive:\orant\database\initSID.utl -f password
   ```

   Issue the password when prompted. On Windows, the profile path can also be specified in UNC notation (for example: `-p \\SERVER_A\ora nt\database\init`*`SID`*`.utl`
5. Restore the database according to the SAP recommendation.
6. Reset the `init`*`SID`*`.utl` file and the password on the target system.

Detailed information regarding manual password handling is available on page "7. Determine the Tivoli Storage Manager password method" on page 59.

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

Information needed to fine-tune Data Protection for SAP for Oracle 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.

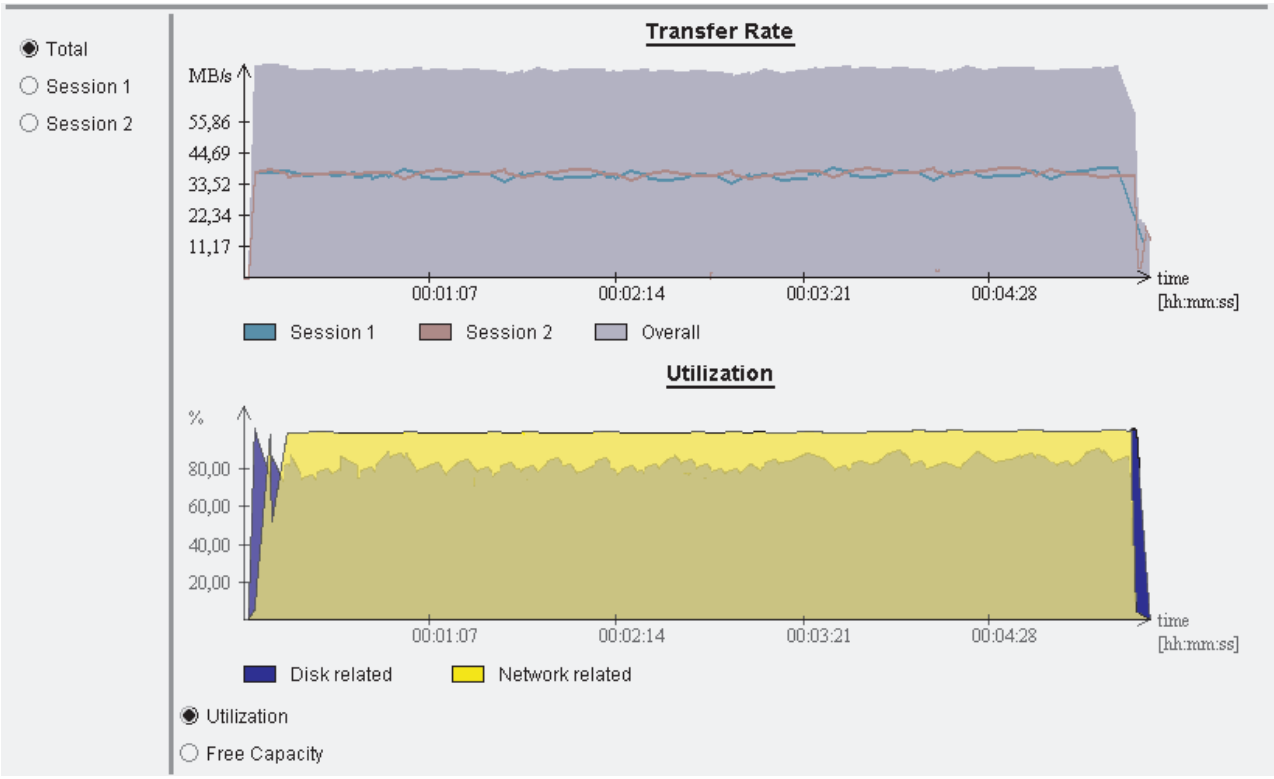

Figure 15. Indicating a Balanced Configuration

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

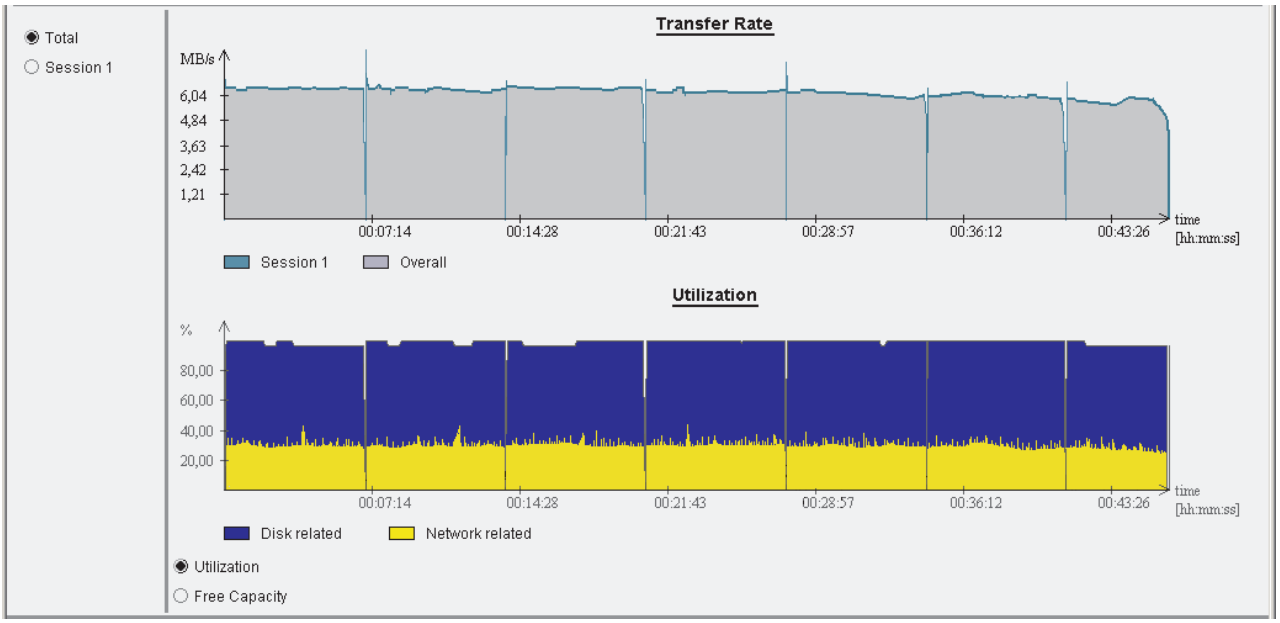

Figure 16. Indicating a Disk Bottleneck

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. See "Cyclic Procedure for Optimizing your Configuration" on page 76 for more information about optimizing your system.

Example of a network or Tivoli Storage Manager bottleneck

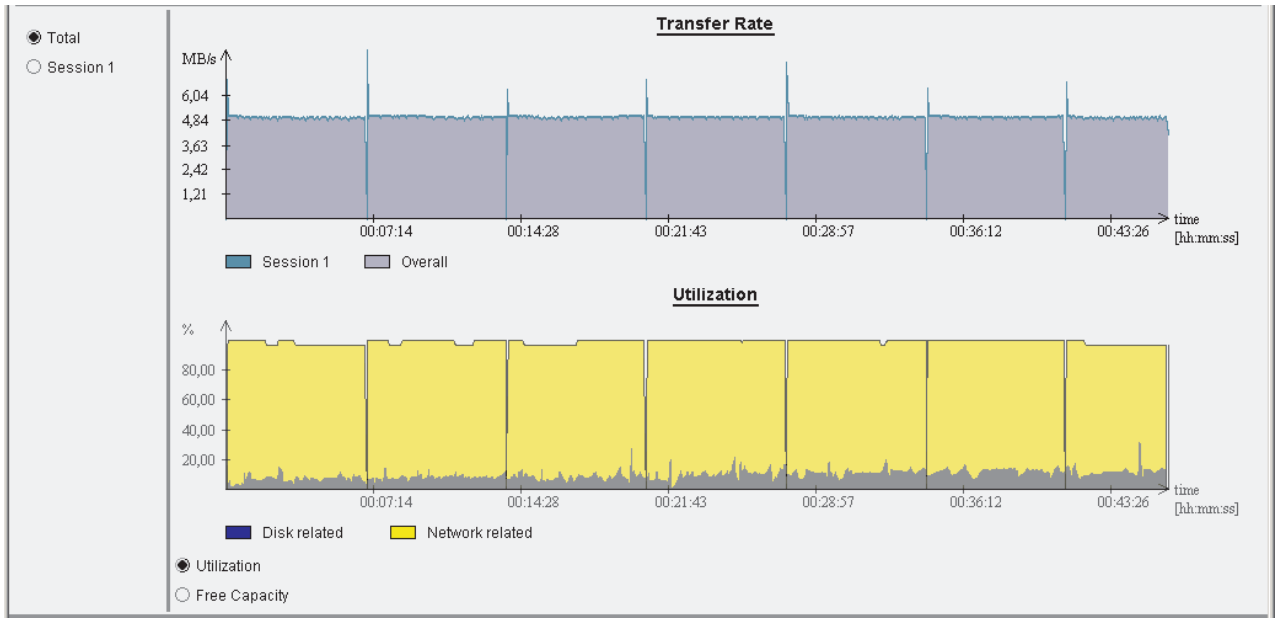

Figure 17. Indicating a Network or Tivoli Storage Manager Bottleneck

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 Oracle performance analysis, as described in "Simulating Backup and Restore" on page 78. 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. See "Cyclic Procedure for Optimizing your Configuration" on page 76 for more information about optimizing your system.

Viewing performance data

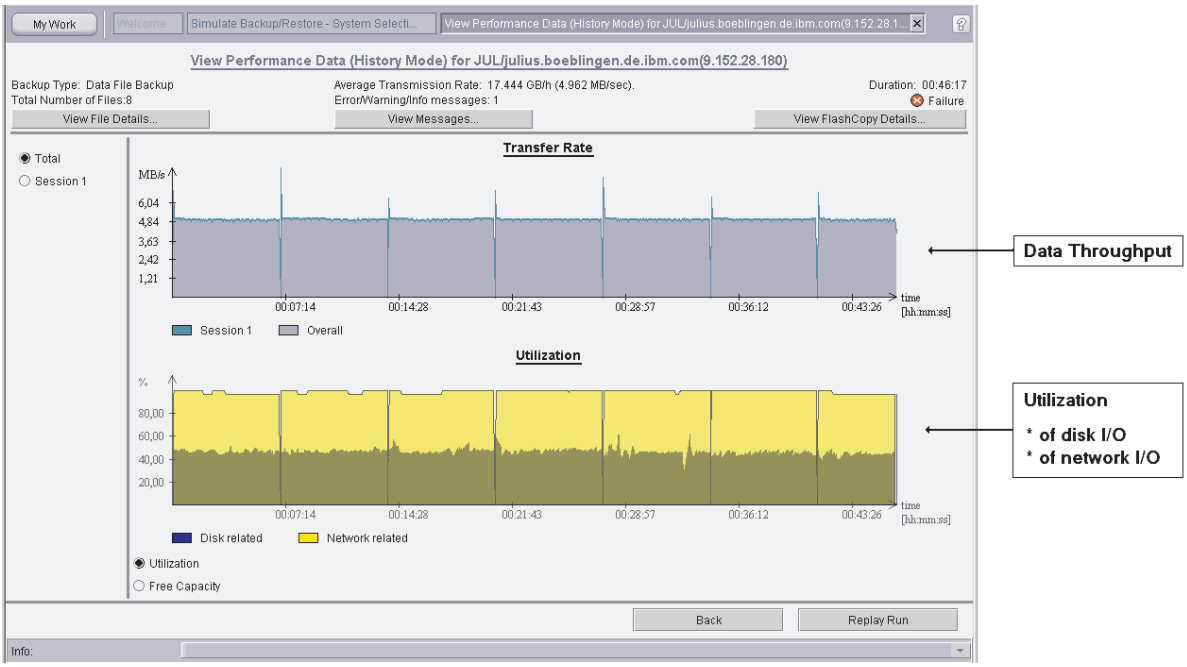

Figure 18. Showing Data Throughput and I/O Utilization

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 Oracle 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 Oracle creates a report that contains statistical information such as the number of bytes transferred and the effective data throughput. The profile keyword REPORT (on page "Tivoli Storage Manager for ERP for Oracle profile parameter descriptions" on page 120) can provide additional information on report levels. 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 Oracle Activities" on page 93.

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. Backup and restore simulations are also available. These simulations can test configuration changes or planned restore operations without compromising the production system. Performance optimization is discussed in detail in "Overview of a balanced system" on page 83.

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 89 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

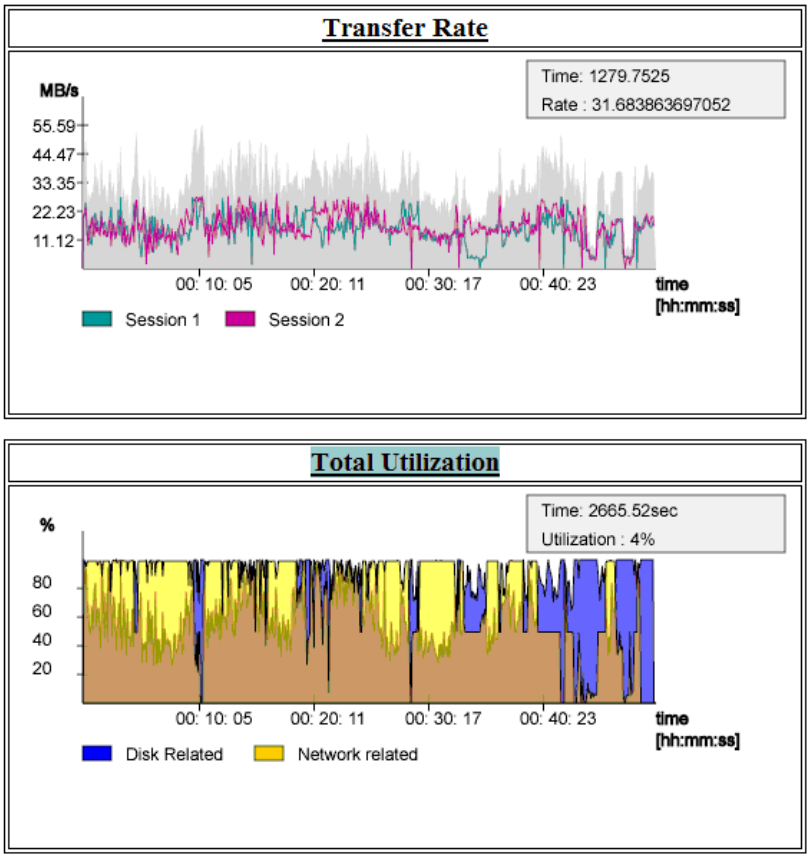

Figure 19. Performance Report - Graphical Presentation Section

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 20. 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	
Failure	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 21. 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.

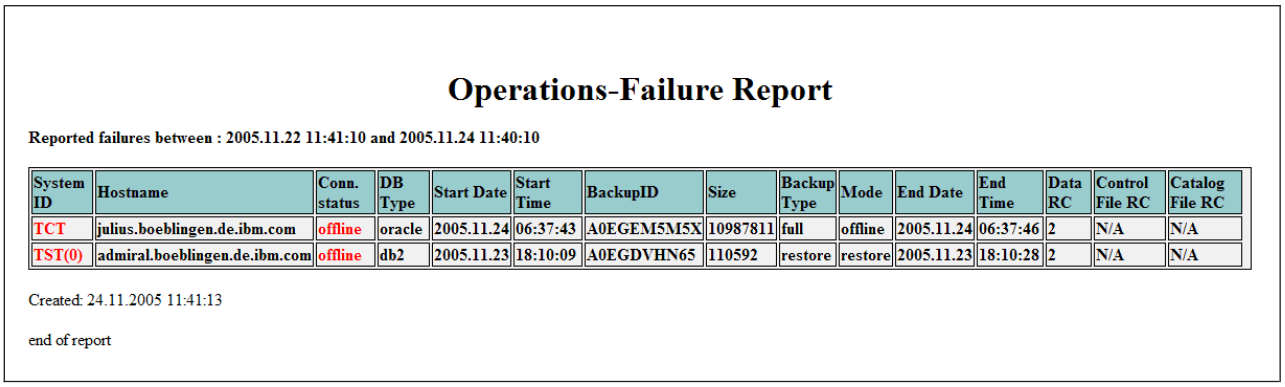

Operations-Failure Report

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 22. 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, View Performance Data (History Mode), and Available Simulation Results. 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.

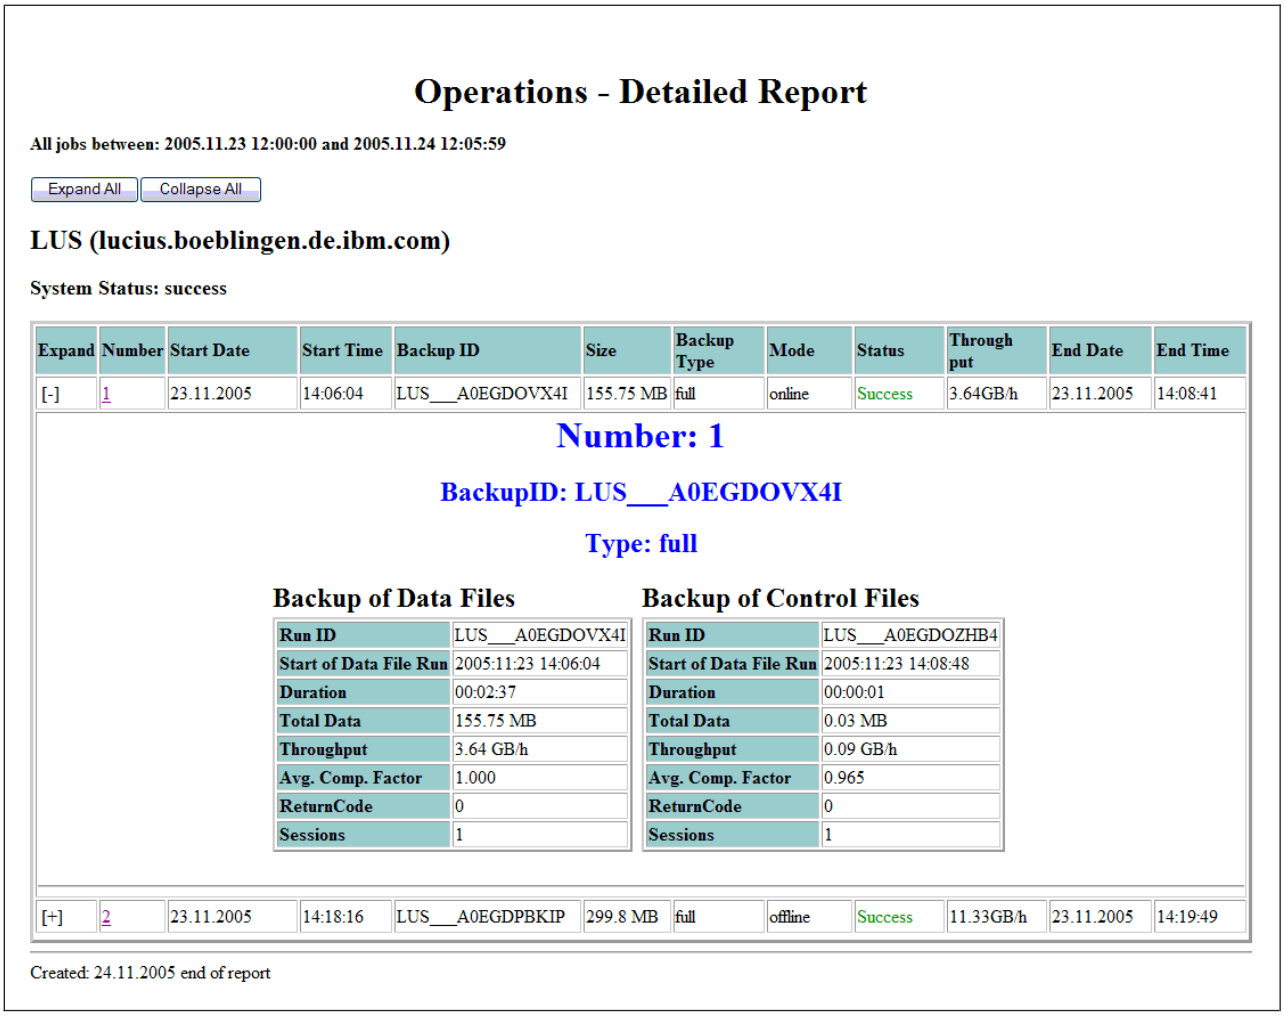

Operations - Detailed Report

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	Through put	End Date	End Time
[-]	1	23.11.2005	14:06:04	LUS___A0EGDOVX4I	155.75 MB	full	online	Success	3.64GB/h	23.11.2005	14:08:41

Number: 1

BackupID: LUS___A0EGDOVX4I

Type: full

Backup of Data Files

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

Backup of Control Files

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

[+]	2	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 23. 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.

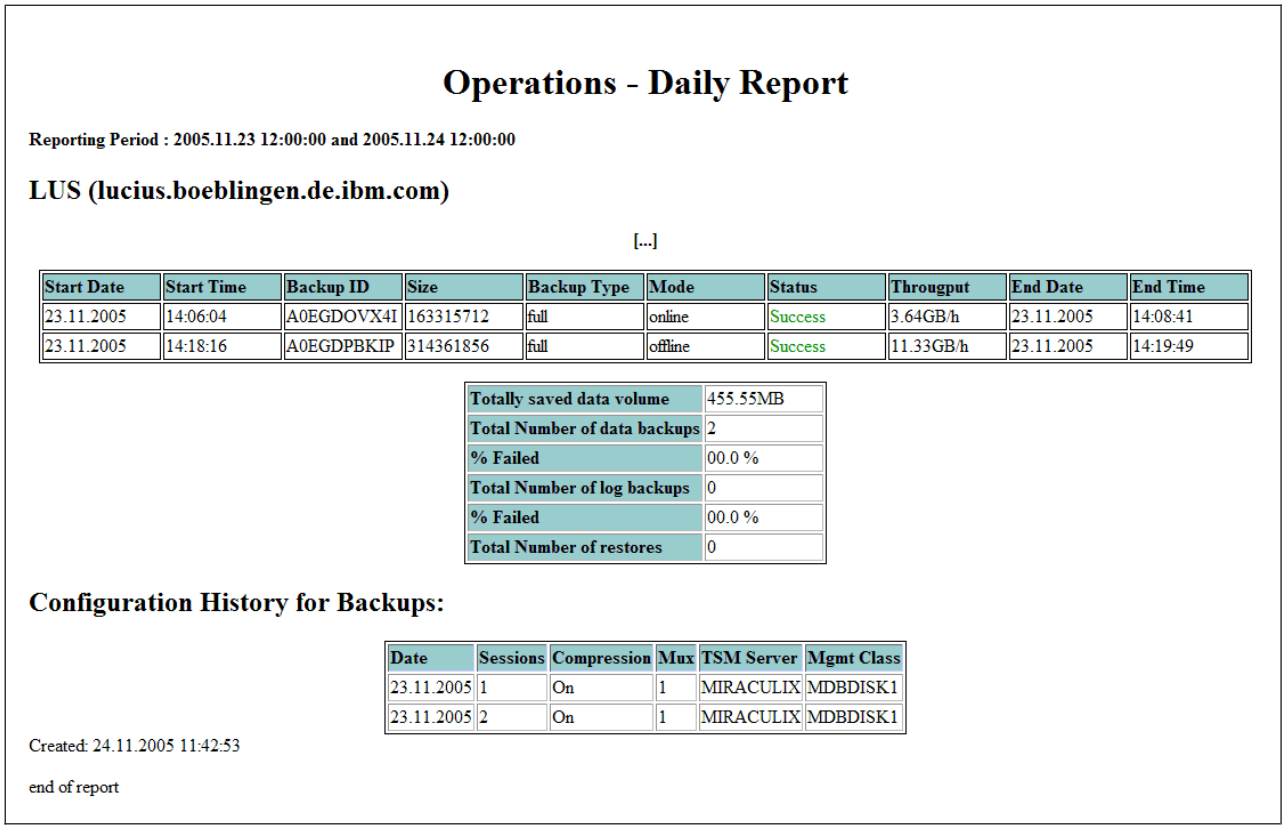

Operations - Daily Report

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 24. Operations Daily Report

Reporting on Data Protection for SAP for Oracle 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 Oracle 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, View Performance Data, and Simulate Backup/Restore 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
- Simulation 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.

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 Oracle 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 Oracle profile parameter descriptions" on page 120). 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 Oracle profile parameter descriptions" on page 120.

Options

You can use Tivoli Storage Manager for ERP options to tune performance.

Performance Options of Data Protection for SAP for Oracle

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 Oracle 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 101.

Multiplexing
Multiplexing simultaneously transfers data from different files through one session (MULTIPLEXING) in order to maximize tape performance. Multiplexing is useful for tape storage since tape drives often have higher data transfer rates than the disks. Combining multiplexing and parallel sessions can optimize overall backup and restore performance. See the description of the MULTIPLEXING option on page "Tivoli Storage Manager for ERP for Oracle profile parameter descriptions" on page 120.

Disk Sorting
Data Protection for SAP uses Adaptive File Sequencing during backup processing. This feature sorts database files in sequential order to avoid simultaneously reading files located on the same disk. As a result, to reduce backup processing time is reduced.

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 102 and "Multiple Servers" on page 100.

Incremental Backup
Data Protection for SAP supports incremental RMAN backup of a SAP databases. Depending on the system environment, incremental backups might decrease backup processing time.

Individual Tablespace Locking
Data Protection for SAP provides a backup profile parameter (`util_file_online`) that minimizes the number of archived redo logs backed up during online backup operations. This parameter informs the SAP database utilities of the files (and related table spaces) to be backed up. The SAP utilities then switch those table spaces into backup mode. After the files are backed up, the table spaces are released and a new cycle starts. See page "The Data Protection for SAP for Oracle Profile" on page 119 for detailed information.

RL_COMPRESSION
The RL_COMPRESSION profile keyword is compresses a partially filled database. This can result in reduced network traffic and fewer tapes required for backup. See "Compression" on page 96 for complete details.

Adjustments to Data Protection for SAP for Oracle for Improving Performance of Data Transfer

Data Protection for SAP for Oracle is configured (by default) to send uncompressed data to the Tivoli Storage Manager server using a single session.

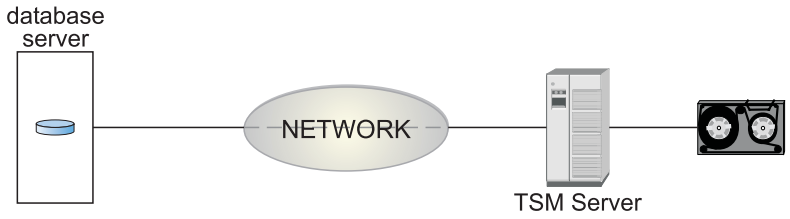

Figure 25. Data Transfer for a Backup / Restore

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 83.

Buffer Copies

Data Protection for SAP for Oracle 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 Oracle profile parameter descriptions" on page 120 for more information.

Buffer Size

Data Protection for SAP for Oracle 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 Oracle profile parameter descriptions" on page 120 for more information.

Compression

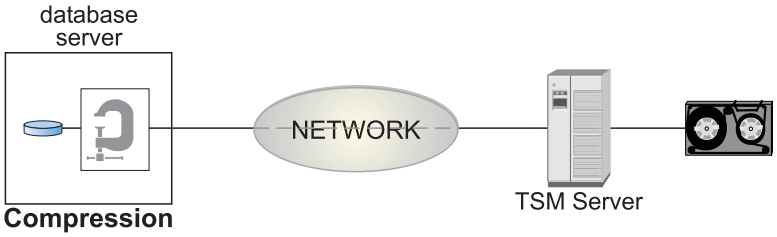

Figure 26. Null Block Compression

Data Protection for SAP for Oracle 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 Oracle profile parameter descriptions" on page 120 for details on how to activate Data Protection for SAP compression.

Automation Options for Data Protection for SAP for Oracle

Administrative productivity can be improved by using these Data Protection for SAP for Oracle 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 log files to a disk storage pool. Multiple management classes can also be specified to use in conjunction with multiple redo log copies. The profile keywords BRARCHIVEMGTCLASS and BRBACKUPMGTCLASS in "Tivoli Storage Manager for ERP for Oracle profile parameter descriptions" on page 120 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 value of the MAX_VERSIONS parameter, the oldest versions are deleted. Retaining backups keeps track of all redo log files, database control files, partial and incremental 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 Oracle profile parameter descriptions" on page 120 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 FRONTEND in "Tivoli Storage Manager for ERP for Oracle profile parameter descriptions" on page 120.

Data transfer

Review information about ways to improve performance for data transfer.

Observations on the Data Protection for SAP for Oracle 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.

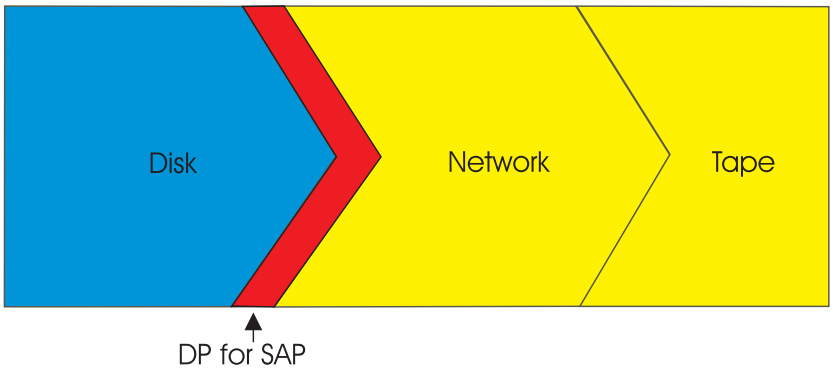

Figure 27. High-level View of the Data Flow During Backup

From a high-level view, the data packages need to send these elements when doing a backup with Data Protection for SAP for Oracle: 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 Oracle" on page 94. Additional performance issues are available in "General Performance Considerations" on page 99.

Data Protection for SAP for Oracle 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.

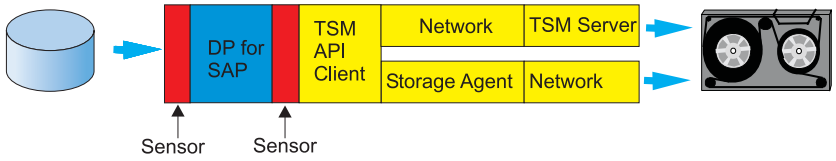

Figure 28. Performance Optimizing by Using Sensors

Data Protection for SAP for Oracle 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 29 on page 100 provides a high level overview of these three main components involved during a Data Protection for SAP for Oracle data transfer:

- The SAP® database server.
- The network.
- The Tivoli Storage Manager server which is also referred to as a backup server.

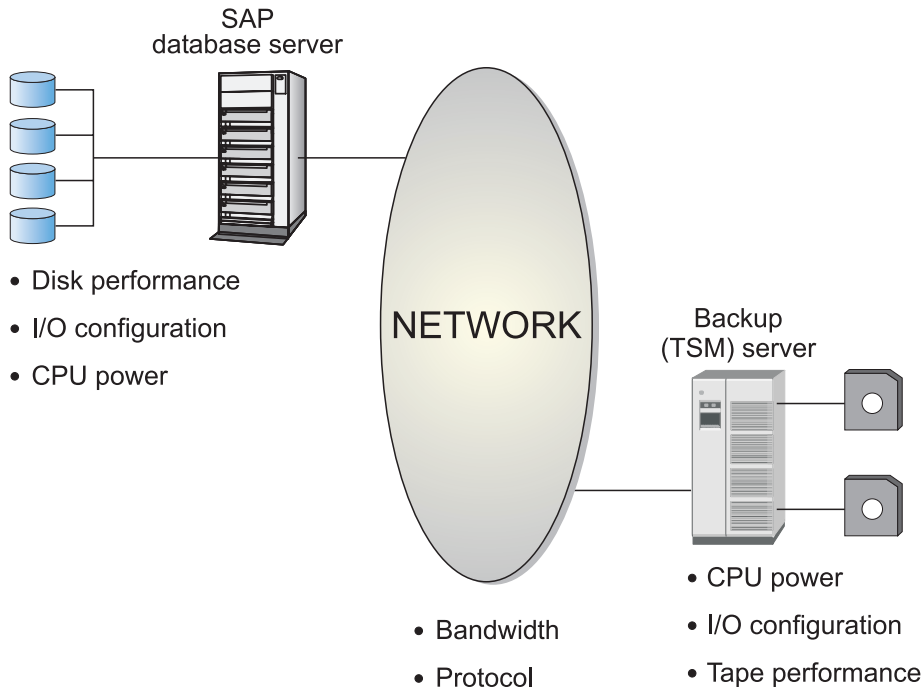

Figure 29. Data Protection for SAP Data Transfer

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 Oracle 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.

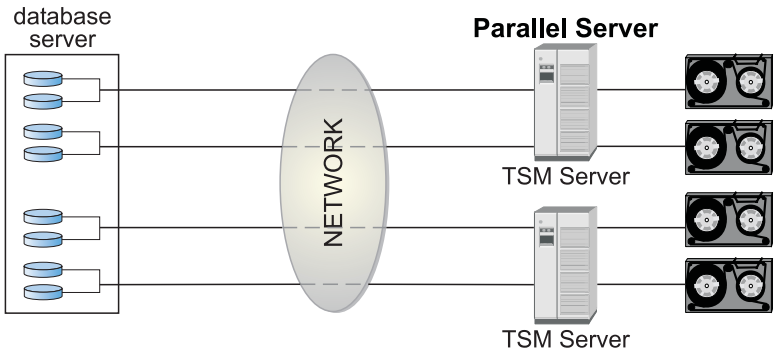

Figure 30. Multiple Servers

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 Oracle profile parameter descriptions" on page 120. The value of the MAX_SESSIONS keyword is not greater than the sum of all SESSION values specified for the SERVER statements used concurrently.

When RMAN is used, the number of SESSIONS configured for each SERVER must be greater than or equal to the number of sessions configured for the MAX_SESSIONS keyword specified during restore operations. This prevents RMAN from requesting a number of objects (in parallel from the same server) that exceeds the number of sessions that are available for that server. Detailed information regarding parallel servers is available in “Alternate or parallel backup paths and backup servers” on page 17.

Multiple Sessions

Data Protection for SAP for Oracle allows use of multiple tape drives simultaneously in order to increase the transfer rate to or from the Tivoli Storage Manager server.
The keywords MAX_SESSIONS, MAX_BACK_SESSIONS, MAX_ARCH_SESSIONS

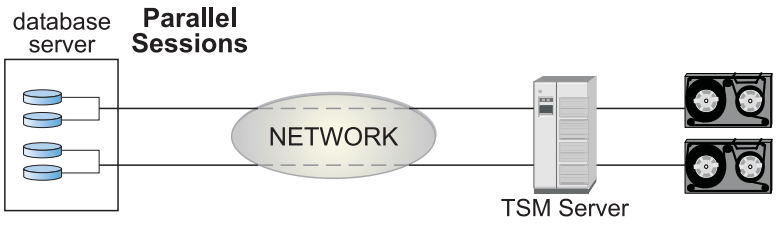

Figure 31. Parallel (Multiple) Sessions

and MAX_RESTORE_SESSIONS are used for defining the number of parallel 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 these keywords, see page “Tivoli Storage Manager for ERP for Oracle profile parameter descriptions” on page 120. 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 Oracle profile parameter descriptions” on page 120.

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 running BRARCHIVE for log file backups , 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, these considerations for database backups also apply to disk storage pools:

- Several sessions of one BRARCHIVE operation can utilize one or two independent disk storage pools.
- Several sessions of BRARCHIVE operations of different databases can simultaneously utilize one or two independent disk storage pools.

The number of storage pools that are required depends on the number of backup copies requested for a log file. See keyword REDOLOG_COPIES in “Tivoli Storage Manager for ERP for Oracle profile parameter descriptions” on page 120.

BRARCHIVE details are available in “Tivoli Storage Manager for ERP for Oracle profile parameter descriptions” on page 120.

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.

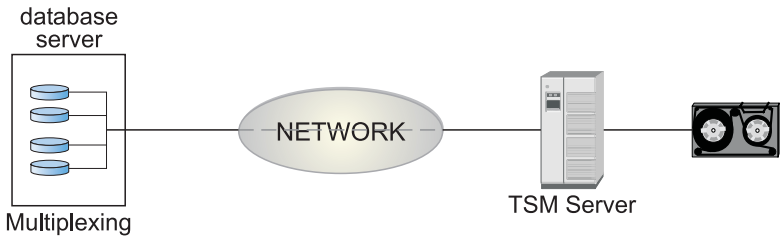

Figure 32. Multiplexing

The value of keyword MULTIPLEXING defines the number of files read in parallel within a single session. Appropriate MULTIPLEXING values are expected in the range of 1 to 4. The best value for your environment depends on the I/O rate of your disks, the location of your data on the disks, the network capacity, the throughput rate of the storage media, and the compression setting. If the MULTIPLEXING value is too high, a thread management overhead may be occur that might offset any performance gain. details regarding this keyword is available on page “Tivoli Storage Manager for ERP for Oracle profile parameter descriptions” on page 120.

Multiple Network Paths

Data Protection for SAP for Oracle allows you to use multiple network connections (paths) for data transfer between the database server and the backup server.

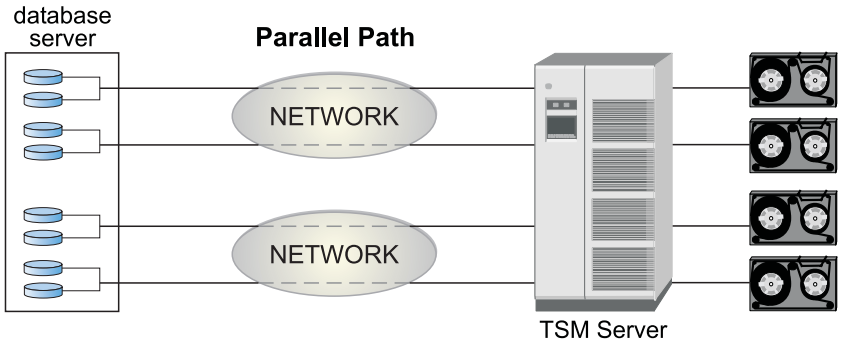

Figure 33. Parallel (Multiple) Paths

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 Oracle profile parameter descriptions” on page 120. 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 17.

Storage space

You can manage aspects of storage space to tune the performance of Tivoli Storage Manager for ERP.

Disk Sorting

Data Protection for SAP for Oracle uses Adaptive File Sequencing which attempts to read operations from disks that are truly parallel. For example, when MULTIPLEXING set to 5 and MAX_SESSIONS is set to 3, Data Protection for SAP backs up fifteen files to three tapes. Due to disk sorting, the fifteen files are selected from different disks whenever possible. However, some storage subsystems only provide the information required for disk sorting to users with administrator authority. In such a situation, the administrator can perform a manual sort file to retrieve the locations of the files. See the description of the SORT_FILE keyword on page "Tivoli Storage Manager for ERP for Oracle profile parameter descriptions" on page 120 for more information.

The information required for sorting is determined by the UNIX or Linux createinfo program. Beginning with Data Protection for SAP version 5.4, createinfo is no longer run automatically on UNIX or Linux. due to the wide use of storage subsystems. If performance suffers or you want to retain the previous functionality (because of directly attached disks), this can be achieved by starting the 'prole' process with the 'c' option:

1. As user root, modify the entry in /etc/inittab (add '-c'). For example:

   ```
   .../tivoli/Tivoli Storage Manager/tdp_r3/ora/prole -p tdpr3ora -c ....
   ```

2. Activate the change with 'init q'.

This simplifies, or even eliminates the need for, the start/stop scripts for HACMP™ takeover.

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 (`brbackup / brrestore log`, `sbtio.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 Oracle setup changed.
- One (or more) of the Oracle, 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 or Linux `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 `brbackup / brrestore log files`.
  - 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 (`init`*`SID`*`.utl`, `init`*`SID`*`.sap`, `dsm.sys`, `dsm.opt`, `/etc/services`, `/etc/inittab`, ...). This UNIX or Linux 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 Oracle 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:

```
po64:2:respawn:/usr/tivoli/tsm/tdp_r3/ora64/prole -p tdpr3ora64
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 Oracle port `tdpr3ora64` for backint and RMAN connections 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:

```
tdpr3ora64      57323/tcp
```

  These lines are added to the `/etc/services` file by the installer.
- Make sure the Data Protection for SAP configuration file `init`*SID*`.utl` is located in the `$ORACLE_HOME/dbs` directory.

- When using the BR*Tools, modify the `init*SID*.sap` file by setting `backup_dev_type = util_file` and variable `util_par_file` to the fully qualified path and file name of `init*SID*.utl`.

See Figure 34 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:

```
tdpr3ora64      57323/tcp
```

- Make sure the Data Protection for SAP configuration file `init*SID*.utl` is located in the `%ORACLE_HOME%\database` directory.
- When using the BR*Tools, modify the `init*SID*.sap` file by setting `backup_dev_type = util_file` and variable `util_par_file` to the fully qualified path and file name of `init*SID*.utl`.

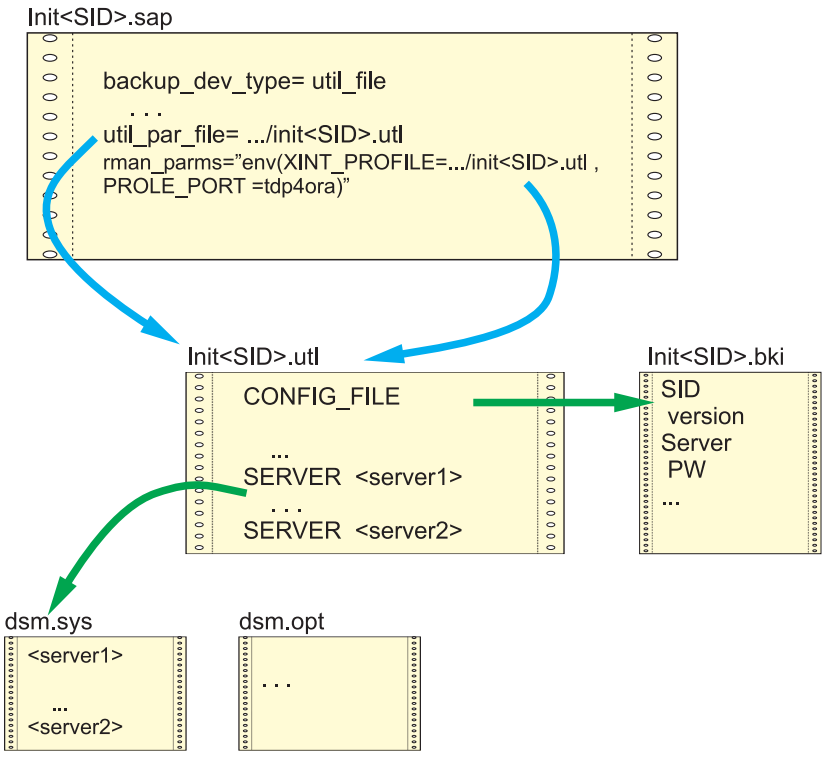

Figure 34. SAP® and Data Protection for SAP configuration files on UNIX or Linux

On UNIX or Linux systems, the names of the Tivoli Storage Manager servers specified in `init*SID*.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 init*SID*.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 init*SID*.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 Oracle version.
- The operating system level and patches that were applied.
- The Oracle version
- The Tivoli Storage Manager server version.
- The Tivoli Storage Manager server operating system level.
- Data Protection for SAP configuration file (initSID.utl) including Tivoli Storage Manager client configuration files (dsm.sys, dsm.opt)
- Data Protection for SAP profile (initSID.utl)
- BR*Tools output for the operation in question (brarchive, brrestore)
- 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 Oracle problems

Information on how to resolve errors that might occur during Data Protection for SAP for Oracle operations is provided.

Location of log files

Text displayed on the screen during brbackup, brrestore, and SAP® Tools operations are typically written to a log file. Oracle also writes internal operations in the alert log and core files that reside in the directory specified in the Oracle control files, for example $SAPDATA_HOME/saptrace/background/alert_*SID*.log. Information about how to locate these log files is available in "How to find files containing message output (log files)."

How to find files containing message output (log files)

Data Protection for SAP for Oracle process results are logged in files.

These files are located in the following paths:

- UNIX or Linux: $SAPDATA_HOME/sapbackup for backup and restore runs
- UNIX or Linux: $SAPDATA_HOME/saparch for redo log archive runs

Windows:

- %SAPDATA_HOME%\sapbackup for backup and restore runs

- %SAPDATA_HOME%\saparch for redo log archive runs

All log files written during a backup, restore or archive operation are listed in summary log files with start and end timestamps. The summary log files are located in the same directory as the log files themselves and have the following names:

- back*SID*.log
- rest*SID*.log
- arch*SID*.log

If you are running Oracle RMAN, the log file sbtio.log (which is specified by user_dump_dest in the Oracle control files) might also need to be viewed. For most installations, this file is defined as $SAPDATA_HOME/saptrace/usertrace/sbtio.log. This file contains all messages issued by the Data Protection for SAP RMAN connector during operation of Oracle RMAN.

Messages

During BR*Tools processing, logs that contain all issued messages are written to paths /oracle/*SID*/sapbackup (for BRBACKUP) or /oracle/*SID*/saparch (for BRARCHIVE). The message prefix indicates the issuing components. Refer to the documentation for the component that issued the message for detailed information. However, the following prefixes are used when employing BR*Tools with Data Protection for SAP for Oracle:

*Table 13. Prefixes when using BR*Tools*

Prefix	Issuing Component
ANS / ANR	Tivoli Storage Manager
BKI	Data Protection for SAP
BR	BR*Tools
ORA	Oracle database kernel
RMAN	RMAN

File Manager

The most important requirement for File Manager is that Data Protection for SAP for Oracle is set up correctly. This is especially true in regard to the backint executable file, as this file must be able to connect to the Tivoli Storage Manager server without errors. If this call fails, the File Manager displays an error message but does not analyze the reason for the failure. To analyze the error, invoke backint manually with the inquire function and check the output for error messages.

BACKINT problem resolution

Figure 35 on page 110 displays how to isolate the problem once the settings performed by the installer are verified. Make sure the BACKINT interface is working correctly before examining the RMAN interface.

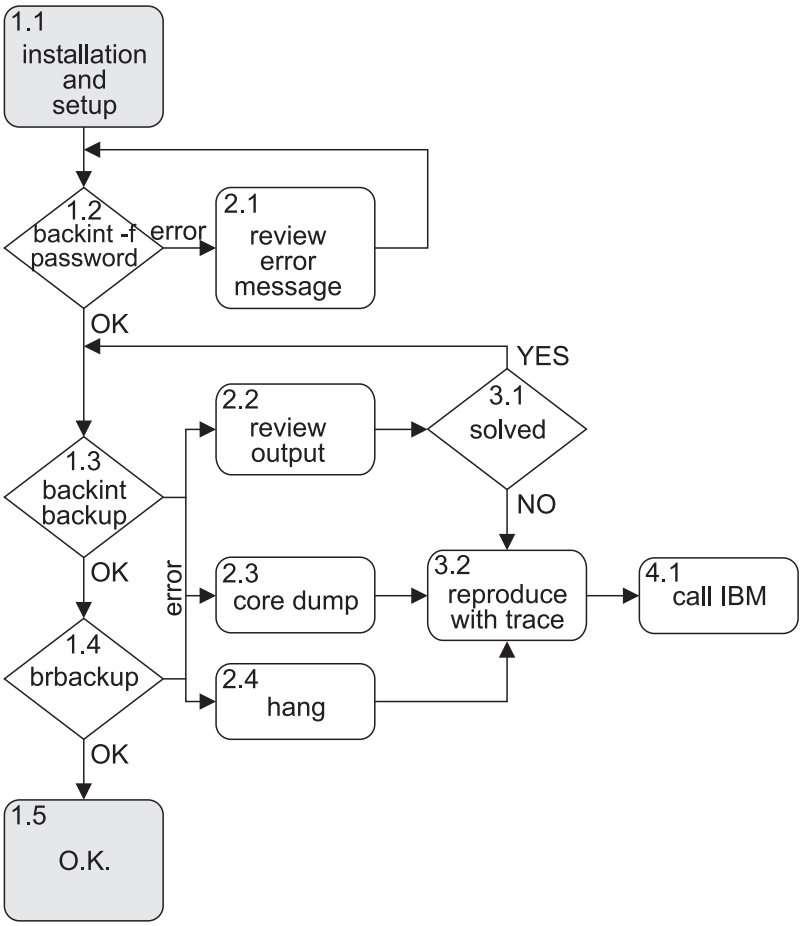

Figure 35. Problem Isolation for Backint

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 61. 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 `init`*`SID`*`.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 `-p`*PORT*, 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 backint again.

BKI5001E: Tivoli Storage Manager Error: Server not found in configuration file
On UNIX or Linux, the Tivoli Storage Manager server defined in the `init`*SID*`.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 backint 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 backint to verify the settings are correct as described in "Backup function" on page 114. When the backup completes successfully, the message `#SAVED`*BIDFILENAME* displays for each saved file and `BKI0024I: Return code is: 0` also displays. If an error message displays, view the error description in *IBM Tivoli Storage Manager for Enterprise Resource Planning Data Protection for SAP Messages* for information regarding how to resolve it. At this point, the primary Data Protection for SAP for Oracle setup is almost complete. The BR*Tools and Oracle (when using RMAN) must also be configured correctly. Proceed to Step 1.3 and start `brbackup` as described in "Verify the installation" on page 39.

RMAN problem resolution

The following graphic (Figure 36 on page 112) will help you to isolate problems that occur when using RMAN.

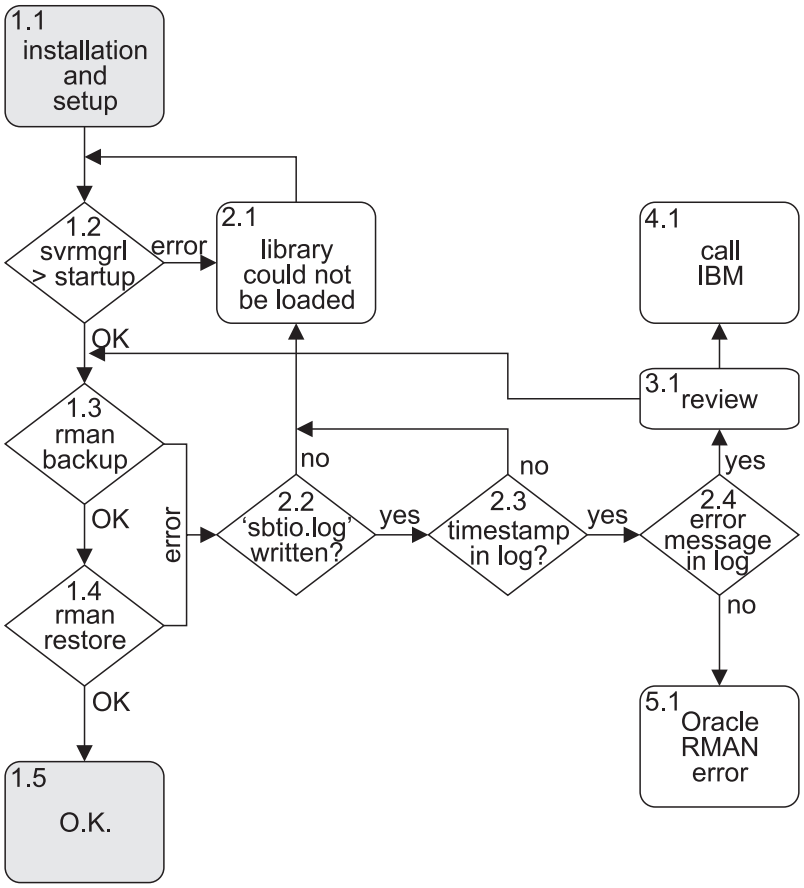

Figure 36. Problem Isolation for RMAN

After Data Protection for SAP for Oracle and Oracle are configured to work together (Step 1.1 in Figure 36), attempt to start Oracle using the server manager `svrmgrl` (on UNIX or Linux) or `svrmgr30` (on Windows) with Oracle 8.x. Use SQL Plus (sqlplus) with Oracle 9.x. When an error occurs while using RMAN, always view the `sbtio.log` file first. This file is located in the directory specified by the `user_dump_dest` keyword in the Oracle `init`*`SID`*`.ora` profile (located at `$ORACLE_HOME/saptrace/usertrace/sbtio.log` by default). If the `sbtio.log` file does not exist (Step 2.2), then either Oracle was unable to load the shared library that contains the RMAN connector for Data Protection for SAP or an error was encountered before the Data Protection for SAP library was called. In both cases, an Oracle error message should exist in the `brbackup` log file that begins with ORA-, PLS-, or RMAN-. Try to resolve this problem using the Oracle and SAP® documentation. If the `sbtio.log` file exists, search for a message beginning with `BKI`*`XXXXY`* where *`XXXX`* is a four digit number and *`Y`* is the letter I, W, or E. When such a message occurs, the RMAN connector for Data Protection for SAP loaded correctly and was called by RMAN. This should be the first message for every new session in Step 2.3:

```
BKI7060I: Data Protection for SAP version and build number session: 764
```

If this message is not available, Oracle loaded an incorrect library.

Perform these tasks on Windows when an incorrect library is loaded by Oracle:

1. Remove or rename all occurrences of the file `orasbt.dll` except the one in the Data Protection for SAP installation directory. Then copy this one to `%ORACLE_HOME%\bin`.
2. Stop the `OracleService`*`SID`* and restart it.

Several factors must be considered when an incorrect library is loaded by Oracle on UNIX or Linux. For example, the RMAN library `libtdp_r3.`*`ext`* is not located by the Oracle executable. Oracle suggests to use the SBT_LIBRARY variable to specify the library. However, do not use this variable for a version of Oracle prior to Oracle 9.2. Be aware that Oracle recommends not to issue the `make` command as described in "Required installation tasks" on page 23. However, this recommendation is not applicable for all combinations of operating system and Oracle combinations. As a result, issuing the `make` command on any UNIX or Linux system with Data Protection for SAP is acceptable. When issued correctly, this command can confirm that the library and the Oracle executable are compatible. Also, make sure the library and soft link entered during the command exists and that the soft link is valid:

```
make -f ins_rdbms.mk ioracle LLIBMM=lib or link
```

It might be helpful to add the location of the link or library to the LIBPATH environment variable (on AIX) or to the LD_LIBRARY_PATH environment variable (on other UNIX or Linux systems).

On Windows based systems, the location of `orasbt.dll` must be in the PATH. Also, ensure that you have only one `orasbt.dll` in your system's PATH. It will be helpful as well to review the setup procedure according to "Required installation tasks" on page 23 and the information given in that chapter. Also, check if a core file is written or if Oracle has written a trace within the `saptrace/usertrace` directory.

In (Step 2.4) the file `sbtio.log` is written and you find an error message starting with BKI. Find the message and information about how to resolve the error in *Tivoli Storage Manager for ERP Messages*. Using the backint executable file to determine any problems may make it easier because you can see all messages on the screen. Also, you will not disturb Oracle if something goes wrong. If backint is working as expected, return to problem determination with RMAN.

When isolating a problem with Data Protection for SAP and RMAN, you can follow the same steps as in "BACKINT problem resolution" on page 109. There must be a connection established to ProLE and the Tivoli Storage Manager server, and a password must be set (using backint) as well. If some of these steps fail, you will get exactly the same error messages with RMAN as you get with backint. Find the messages and information about how to resolve the errors in *Tivoli Storage Manager for ERP Messages*.

Manually invoke Data Protection for SAP for Oracle

Information about how to invoke Data Protection for SAP for Oracle from the command line to assist with troubleshooting efforts is provided.

Data Protection for SAP for Oracle is typically invoked by the BR*Tools utilities with a set of appropriate parameters. For troubleshooting purposes, call Data Protection for SAP directly from the command line:

```
backint -?
```

This command displays a list of all possible Data Protection for SAP command line options. It enables you to manually perform data protection operations which can assist with correcting errors. For the C shell, enclose the option string in quotes (`backint '-?'`).

Backup function

Data Protection for SAP for Oracle is typically invoked by the BR*Tools utilities with a set of appropriate parameters. For troubleshooting purposes, call Data Protection for SAP directly from the command line:

```
backint -?
```

This command displays a list of all possible Data Protection for SAP command line options. It enables you to manually perform data protection operations which can assist with correcting errors. For the C shell, enclose the option string in quotes (`backint '-?'`).

The backup function is typically invoked by the SAP® database utilities BRBACKUP and BRARCHIVE. These programs provide an input file (in the case of backup and inquire) to Data Protection for SAP that contain the names and paths of the database files to be processed. For troubleshooting purposes, however, it might be necessary to directly call this Data Protection for SAP function directly in order to back up individual files as shown in these examples. Issue this command on UNIX or Linux systems:

```
backint -p /oracle/SID/dbs/initSID.utl -f backup
```

Issue this command on Windows systems:

```
backint -p drive: or UNC name\orant\database\initSID.utl -f backup
```

The Data Protection for SAP profile `initSID.utl` has to be specified with the path and file name statement as shown in the examples. The program prompts you to enter one (or more) file names. Every successful backup operation (collection of one or more files) is allocated its own backup ID within Tivoli Storage Manager. Remember to press CTRL + D (on a UNIX or Linux system) or CTRL + Z (on a Windows system) after the file name to backup has been entered.

Delete function

Data Protection for SAP for Oracle is typically invoked by the BR*Tools utilities with a set of appropriate parameters. For troubleshooting purposes, call Data Protection for SAP directly from the command line:

```
backint -?
```

This command displays a list of all possible Data Protection for SAP command line options. It enables you to manually perform data protection operations which can assist with correcting errors. For the C shell, enclose the option string in quotes (`backint '-?'`).

The delete function is used as part of the Data Protection for SAP version control mechanism and can only be called by Data Protection for SAP or by a user. This

function can be invoked from the command line as shown in these examples. Issue this command on UNIX or Linux systems:

```
backint -p /oracle/SID/dbs/initSID.utl -f delete
```

Issue this command on Windows systems:

```
backint -p drive: or UNC name\orant\database\initSID.utl -f delete
```

You will be prompted to enter the backup ID to be deleted. It is not possible to delete single files within a backup ID. You can only delete complete backup IDs.

Inquire function

Data Protection for SAP for Oracle is typically invoked by the BR*Tools utilities with a set of appropriate parameters. For troubleshooting purposes, call Data Protection for SAP directly from the command line:

```
backint -?
```

This command displays a list of all possible Data Protection for SAP command line options. It enables you to manually perform data protection operations which can assist with correcting errors. For the C shell, enclose the option string in quotes (`backint '-?'`).

The inquire function (typically invoked by BR*Tools and BRRESTORE) is used to query the Tivoli Storage Manager server for backup IDs or files which belong to a particular backup ID. For troubleshooting purposes, however, it might be necessary to invoke this function manually as shown in these examples. Issue this command on UNIX or Linux systems:

```
backint -p /oracle/SID/dbs/initSID.utl -f inquire
```

Issue this command on Windows systems:

```
backint -p drive: or UNC name\orant\database\initSID.utl -f inquire
```

Data Protection for SAP prompts you to enter the inquiry in one of these four formats:

- **#NULL**: Display all backup IDs that have been saved to this point. A typical line of the response could be `#BACKUP JE0___A0DNE9Z74C`. The backup ID in this case is `JE0___A0DNE9Z74C` (`#BACKUP` does not belong to the backup ID). The first six characters are the user defined prefix (see BACKIDPREFIX as described in "Tivoli Storage Manager for ERP for Oracle profile parameter descriptions" on page 120). The next ten characters after this represent a unique ID of the backup.
- **BackupID**: Display all of the files which belong to that backup ID. A typical result could be `##BACKUP JE0___A0DNE9Z74C /oracle/C21/dbs/initC21.utl`.
- **#NULL `filename`**: Display all of the backup IDs corresponding to the specified file. *Filename* requires an input consisting of path and name of the file.
- **BackupID `filename`**: Verify whether a particular file has been saved under a certain backup ID. *Filename* requires an input consisting of path and name of the file.

Restore function

Tivoli Storage Manager for ERP for Oracle is typically started by the BR*Tools utilities with a set of appropriate parameters. For troubleshooting purposes, call Data Protection for SAP directly from the command line:

```
backint -?
```

This command displays a list of all possible Data Protection for SAP command-line options. It enables you to manually perform data protection operations which can assist with correcting errors. For the C shell, enclose the option string in quotes (`backint '-?'`).

The restore function is typically started by the SAP® database utility BRRESTORE. For troubleshooting purposes, however, it might be necessary to directly call this Data Protection for SAP function directly in order to restore individual files as shown in these examples. Issue this command on UNIX or Linux systems:

```
backint -p /oracle/SID/dbs/initSID.utl -f restore
```

Issue this command on Windows systems:

```
backint -p drive: or UNC name\orant\database\initSID.utl -f restore
```

You will be prompted to enter the backup ID and the full file names of the files to be restored. If the files are to be restored to another directory, it is necessary to specify the path to the input files. If a file is restored directly, any existing file with the same name will be overwritten without warning. Thus, it is recommended that you restore database files directly only in a controlled manner, when it is necessary in order to remove an error. In normal operation, a database should never be restored directly because the SAP database might become corrupted.

Loading the message catalog

When Tivoli Storage Manager for ERP fails to load the message catalog check the following completed successfully:

1. Verify that the installation was successful and the language files are contained in *DP for SAP install path*>/lang .
2. If the install path is not the default and you are using a backup interface library like RMAN, then you must set the environment variable XINT_NLS_CATALOG_PATH to the new install path before you run any functions. For Tivoli Storage Manager for ERP for Oracle this variable is set in the parameter `rman_parms` in the profile `initSID.sap`.

Chapter 9. Data Protection for SAP for Oracle reference information

Data Protection for SAP for Oracle reference information is provided here.

Commands used with Data Protection for SAP for Oracle

A list of various commands that are used with Data Protection for SAP for Oracle operations is provided.

Cooperation of Data Protection for SAP for Oracle with BRARCHIVE

BR*Tools are a collection of utilities that simplify Oracle database system administration within an SAP environment. Details regarding setting up and using BR*Tools is available in SAP® documentation such as the *SAP Database Guide: Oracle*.

BRARCHIVE backs up redo logs. To save each redo log to Tivoli Storage Manager, BRARCHIVE calls Data Protection for SAP for Oracle either through the BACKINT interface or through RMAN. BRARCHIVE maintains a list of redo logs to be saved. Redo logs that were successfully saved by Data Protection for SAP may be deleted from the file system immediately by BRARCHIVE. However, BRARCHIVE deletes redo log files only in the order of the list. As a result, if the requested number of backup copies cannot be saved for a redo log, this redo log and all subsequent redo logs are maintained. When BRARCHIVE starts again, these redo logs are saved again even if some were successfully saved earlier. Data Protection for SAP informs BRARCHIVE about the redo logs that were saved successfully to Tivoli Storage Manager. If a problem occurs, Data Protection for SAP makes several attempts to save the redo log. When a redo log cannot be saved to the number of copies requested, Data Protection for SAP terminates with an error. Data Protection for SAP does not try to save redo logs with a higher sequence number because they will be saved in a later BRARCHIVE run.

Managing Tivoli Storage Manager Sessions

When redo logs are saved directly to a tape pool, the number of Tivoli Storage Manager sessions must not exceed the number of available tape drives. Be aware that BRARCHIVE might process redo logs while a database backup is still processing or several BRARCHIVE processes may run simultaneously. As a result, these combined sessions might exceed the number of available tape drives. To avoid this situation, save redo logs to disk storage pools and eventually have Tivoli Storage Manager migrate them to tape storage.

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.

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 the SAP® database utility BRBACKUP 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 - ora<SID> -c "/oracle/SID/sapscripts/backup.ksh"
```

The content of `backup.ksh` is available in "Full Offline Backup Shell Script Sample" on page 68.

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 SAP backup/archive 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 of the ORACLE database
# is assumed to be 'C11' and the username 'oraC11'.
#
# ------------------------------------------------------------------------
# Full database backup, scheduled every Friday at 8:00 p.m.
#
0  20    * * 5
# /usr/bin/su - oraC11 -c "/oracle/C11/sapscripts/backup/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 - oraC11 -c "/oracle/C11/sapscripts/backup/archive.ksh"
```

The Data Protection for SAP for Oracle Profile

The Data Protection for SAP for Oracle 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.

During installation on UNIX and Linux systems, this file is copied and renamed to `$ORACLE_HOME/dbs/init$ORACLE_SID.utl`, where *$ORACLE_HOME* is the Oracle home directory and *$ORACLE_SID* is the Oracle System ID (for example, `/oracle/`*`SID`*`/dbs/init`*`SID`*`.utl`).

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.
- The `backint` program on Windows systems accepts the value of the profile name ('-p' option) in Universal Naming Convention (UNC) format as shown here : `\\SERVER_A\profiles\init`*`SID`*`.utl`. However, any file specifications within the profile must use the `drive:path` syntax.
- 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 Oracle profile parameter descriptions

The default value is underlined in these descriptions and applies if the parameter is not specified.

ADSMNODE *ORACLE_sid*
Specifies an *ORACLE_sid* 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 Oracle 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 59.

BACKEND *pgmname [parameterlist]*
Specifies a program *pgmname* that is called by Tivoli Storage Manager for ERP for Oracle after the backup function has completed and before program control is returned to the SAP backup utility. If *pgmname* is not fully qualified, the default search path is used to find 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___.

BATCH YES|NO
Specify NO if Tivoli Storage Manager for ERP is running with an operator standing by. Specify YES if Tivoli Storage Manager for ERP is running in unattended mode. In unattended mode, Tivoli Storage Manager for ERP terminates the run if operator intervention is required. The default for the BATCH parameter is YES for the backup run and NO for the restore run if the BATCH parameter is not present or is commented out in the Tivoli Storage Manager for ERP profile. This parameter has no effect if an RMAN backup or restore is started.

BRARCHIVEMGTCLASS *management_class [management_class...]*
Specifies the Tivoli Storage Manager management class(es) that Tivoli Storage Manager for ERP uses when called from BRARCHIVE. 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 redo log copies (keyword REDOLOG_COPIES on page "Tivoli Storage Manager for ERP for Oracle profile parameter descriptions." This parameter must be defined with the respective SERVER statement, as shown in the sample profile.For more detailed information about implementing and using BRARCHIVEMGTCLASS see "Cooperation of Data Protection for SAP for Oracle with BRARCHIVE" on page 117.

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 when called using BRBACKUP. 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 for disk I/O. 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 files multiplexed in a data stream (keyword MULTIPLEXING) multiplied by the number of sessions (keyword SESSIONS). By activating RL_COMPRESSION, the number of buffers is doubled.

COMPR_INFO *path*
Specifies the file where Tivoli Storage Manager for ERP stores information about the compressed size of files. The *path* value specifies the full path and name of the file. When multiplexing is used, Tivoli Storage Manager for ERP attempts to optimize performance by putting files of the same size in one multiplexing stream. If RL_COMPRESSION is used in addition to multiplexing, the file sizes of the compressed files can differ very much from the original file sizes. Tivoli Storage Manager for ERP can collect information about the compressed file sizes and use it for further file sorting. This file size information is stored in the file specified by the COMPR_INFO parameter. If backups shall serve as a basis for simulations (see "Simulating Backup and Restore" on page 78), COMPR_INFO must denote a valid file and RL_COMPRESSION must be set to YES in order to get meaningful simulation results for compression. When the parameter RL_COMPRESSION is set to NO, this parameter has no effect. If specified, the information file is written after each backup and the information is

used by the following backups and simulations. If there is no compression information about a file because of a database extension, the uncompressed file size is used for file sorting.

CONFIG_FILE *path/initSID.bki*
Specifies the configuration file `init`*`SID`*`.bki` for Tivoli Storage Manager for ERP to store all variable parameters such as passwords and the date of the last password change. 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.

EXITONERROR YES | NO | *n*
This keyword specifies whether or not Tivoli Storage Manager for ERP exits on a backup or restore error during a BRBACKUP/BRRESTORE run. NO means do not exit if an error occurs. YES means exit if one file cannot be backed up. If a number is specified as an argument, Tivoli Storage Manager for ERP counts the number of errors (not warnings or retries) and exits after the specified number of errors. This keyword works only for the BRBACKUP/BRRESTORE runs. BRARCHIVE and RMAN runs always exit after the first error. This parameter is ignored if the BATCH parameter is set to NO.

FCS_FILE *path*
Specifies the profile for Data Protection for Snapshot Devices version 5.4. If Tivoli Storage Manager for ERP and Data Protection for Snapshot Devices version 5.4 are used together, this parameter is required. See the Data Protection for Snapshot Devices documentation for details. For a stand-alone installation of Tivoli Storage Manager for ERP, this parameter must not be used.

FILE_RETRIES *n* | **3**
This parameter specifies the number of retries when a file could not be saved or restored. This parameter has no effect if an RMAN backup/restore is started.

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.

INCREMENTAL NO/CUMULATIVE/DIFFERENTIAL
This parameter specifies if a backup is performed by Oracle RMAN. If it is set to CUMULATIVE or DIFFERENTIAL, incremental backups are performed by using Oracle RMAN. The default value is NO. All the other INCREMENTAL* parameters have no effect if INCREMENTAL is set to NO.

Note: This parameter can be used only when IBM Tivoli Storage FlashCopy Manager is used to offload backups of an SAP Oracle database from the production system and the backups on the backup server are

executed by Oracle RMAN. To perform backups using Oracle RMAN on the production system, see "Oracle Recovery Manager (RMAN)" on page 4.

INCREMENTAL_CATALOG_CONNECT_STRING *string*
This parameter specifies the name of the catalog that is passed to RMAN to connect to the catalog database. This catalog is the name of the listener for the catalog database. There is no default value. If INCREMENTAL is enabled and this value is missing an error message is displayed.

INCREMENTAL_CATALOG_USER *string*
This parameter specifies the name of the user that is passed to RMAN to connect to the catalog database. There is no default value. If INCREMENTAL is enabled and this value is missing an error message is displayed.

INCREMENTAL_CHANNELS *integer in the range 1 or higher*
Specifies the number of parallel RMAN channels that can transfer the data. The default is 1.

INCREMENTAL_LEVEL *integer 0 or 1 [USE_AT days of week][From time TO time*
The RMAN incremental level. The default for the incremental level is 0. You can optionally limit the specified incremental level to a specific time or day. If you specify a time or day, multiple occurrences of this parameter are valid as long as the time specification does not overlap. Time must specified in 24 hour format. Days can be specified by weekday abbreviations like 'Mon, Tue, ...' or by numerical values 0, 1, ..., 6 where 0 stands for Sunday and 6 for Saturday.

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* | **1**
Specifies the maximum number of parallelTivoli Storage Manager client sessions that Tivoli Storage Manager for ERP establishes for backup, archive (redo logs) and restore. Each session transfers one database object or, in the case of an RMAN backup or restore, a set of data blocks to or from the Tivoli Storage Manager server by using the Tivoli Storage Manager API client functions. This keyword is required. Tivoli Storage Manager for ERP optimizes the data transfer regarding the physical location of the Oracle objects. Files stored on different volumes are backed up in parallel if multiple sessions are configured. A maximum of 32 parallel sessions can be configured. 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 `maxnummp` 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. For more detailed information about implementing and using MAX_SESSIONS see "Cooperation of Data Protection for SAP for Oracle with BRARCHIVE" on page 117.

MAX_ARCH_SESSIONS, MAX_BACK_SESSIONS, MAX_RESTORE_SESSIONS, MAX_CONTROL_SESSIONS

These parameters provide the same function as the MAX_SESSIONS parameter but they also provide a more specific use:

- MAX_ARCH_SESSIONS defines the number of parallel sessions used for archive (backup of log files). Usually archive does not need as many sessions as (data file) backups since the volume is much smaller with log files. This value overrides the value of MAX_SESSIONS for the backup of database files.
- MAX_BACK_SESSION defines the number of parallel sessions used for (data file) backup. This value overwrites the value of MAX_SESSIONS for the backup of database files.
- MAX_CONTROL_SESSIONS defines the number of parallel sessions used for backing up the control files after a database or redo log backup. If MAX_CONTROL_SESSIONS is not specified the number of sessions used for the control file backup is the same as for the corresponding database or redo log backup. Typically, for a control file backup, the number of sessions can be reduced in order to avoid unnecessary tape mounts. This value overwrites the value of MAX_ARCH_SESSIONS or MAX_BACK_SESSIONS for the backup of control files.
- MAX_RESTORE_SESSIONS defines the number of parallel sessions used for restore. For restore, more tape drives may be available than for backup. Using additional tape drives may speed up the data transfer for restore if the backup was written to a sufficiently large number of tapes. This value overwrites the value of MAX_SESSIONS for restore.

If MAX_SESSIONS is specified with one or more of these parameters, these specific parameters override the MAX_SESSIONS parameter. You must specify them all if you do not specify the MAX_SESSIONS parameter. For the valid range as well as the rules, see keyword MAX_SESSIONS.

MAX_VERSIONS *n* | **0**

The *n* value defines the maximum number of full database backup versions to be kept in backup storage. The default setting for this value is 0, meaning that backup version control is disabled. If the number of versions found in backup storage is larger than the specified maximum number of backup versions (as specified by the parameter MAX_VERSIONS), the oldest versions are deleted (together with the corresponding tablespace and redo log files) until only the specified maximum number of most recent versions remain. 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.

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 Tivoli Storage Manager for ERP 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 58.

MULTIPLEXING *n* | **1**
Specifies the number of files which are multiplexed into one data stream. The allowed range is from 1 to 8. The optimal value depends on the actual hardware environment. Multiplexing is most effective when fast tape access exists, fast networks are available, database files are compressed, and the CPU load is moderate. Optimal values are in the range from 1 to 4. If not specified, the default value of 1 means multiplexing is not used. This parameter has no effect if an RMAN backup or restore operation is started.

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 59.

REDOLOG_COPIES *n* | **1**
Specifies the number of copies Tivoli Storage Manager for ERP stores for each processed Oracle redo log. The valid range is from 1 to 9. If not specified, Tivoli Storage Manager for ERP stores one copy of the redo logs. 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. For more detailed information about implementing and using REDOLOG_COPIES see "Cooperation of Data Protection for SAP for Oracle with BRARCHIVE" on page 117.

REPORT NO | YES | 2
If set to YES, Tivoli Storage Manager for ERP produces additional information such as information about transferred files. If set to 2, Tivoli Storage Manager for ERP generates an additional summary report containing detailed backup and restore performance statistics. This summary is displayed at the end of the complete operation. The output is sent to stdout, which is typically the console. If not specified, the default is REPORT NO. This keyword has no effect if an RMAN backup or restore operation is started.

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 (BRBACKUP) was started. The offline log files (BRARCHIVE) are not compressed.

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:

- ADSMNODE
- 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, ADSM, 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.

SORT_FILE

To perform manual sorting, a file must be created (*sortfile*). This is an example of the sortfile contents:

```
/path/filename1 disknumbers
/path/filename2 disknumber
    .
     .
     .
/path/filenameN disknumber
```

The disk numbers are counted from 1 to n. They do not have any relation to the physical disks. You only have to specify the same number for the files on the same physical disk.

TCP_ADDRESS
Specifies the IP address of the Tivoli Storage Managerserver 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. The parameter TCP_ADDRESS 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. Make sure that the same Tivoli Storage Manager server is used for a simulation and its corresponding basis production backup. See “Simulating Backup and Restore” on page 78 for details on simulations. The parameter USE_AT must be defined with the respective SERVER statement as shown in the sample profile. The parameter has no effect on actions other than backup.

Sample Tivoli Storage Manager for ERP for Oracle Profile for UNIX or Linux

The sample profile (initSID.utl) is included in the Tivoli Storage Manager for ERP for Oracle installation package. Although the UNIX, Linux, and Windows versions are similar, all example versions are provided.

```
#----------------------------------------------------------------------
#
# Data Protection for SAP (R) interface for ORACLE
#
# Sample profile for Data Protection for SAP (R) Version 6.2
# for UNIX
#
#----------------------------------------------------------------------
#
# This file should be renamed to $ORACLE_HOME/dbs/init$ORACLE_SID.utl
# where $ORACLE_HOME is the home directory of the Oracle database and
# $ORACLE_SID is the system ID of the Oracle database.
#
# 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) V6.2 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.
# 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 parallel sessions to be established for the database backup.
# Note: This number must not exceed the number of tape drives simultaneously
# available to the node for a database backup on the Tivoli Storage Manager
# servers to be accessed.
# The valid range of MAX_BACK_SESSIONS is from 1 to 32.
# Default: MAX_SESSIONS.
#----------------------------------------------------------------------
#MAX_BACK_SESSIONS 1 # Tivoli Storage Manager client sessions for backup


#----------------------------------------------------------------------
# Number of parallel sessions to be established for the redo log backup.
# Note: This number must not exceed the number of tape drives simultaneously
# available to the node for a redo log backup on the Tivoli Storage Manager
```

```
# servers to be accessed.
# The valid range of MAX_ARCH_SESSIONS is from 1 to 32.
# Default: MAX_SESSIONS.
#------------------------------------------------------------------------
#MAX_ARCH_SESSIONS 1 # Tivoli Storage Manager client sessions for archive


#------------------------------------------------------------------------
# Number of parallel sessions to be established for the backup of control
# files. This number is typically used to reduce the number of sessions
# to be used for the control file backup after another backup operation.
# The valid range of MAX_CONTROL_SESSIONS is from 1 to 32.
# Default: MAX_BACK_SESSIONS or MAX_ARCH_SESSIONS, depending on the type of
#          the control file backup.
#------------------------------------------------------------------------
#MAX_CONTROL_SESSIONS 1 # Tivoli Storage Manager client sessions for control
                        # file backup.


#------------------------------------------------------------------------
# Number of parallel sessions to be established for the restore of files.
# Note: This number must not exceed the number of tape drives simultaneously
# available to the node for restore processing backup on the Tivoli Storage
# Manager servers to be accessed.
# The valid range of MAX_RESTORE_SESSIONS is from 1 to 32.
# Default: MAX_SESSIONS.
#------------------------------------------------------------------------
#MAX_RESTORE_SESSIONS 1 # Tivoli Storage Manager client sessions for restore


#------------------------------------------------------------------------
# Number of backup copies of redo logs.
# The valid range of REDOLOG_COPIES is from 1 to 9.
# Default: 1.
#------------------------------------------------------------------------
#REDOLOG_COPIES  2


#------------------------------------------------------------------------
# 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


#------------------------------------------------------------------------
# Specifies how many files are read simultaneously and are multiplexed into
# one data stream to a Tivoli Storage Manager server. Multiplexing is usefull
# when the data rate to a Tivoli Storage Manager server is higher (fast
# tapes, fast network) than the I/O rate of a single disk.
# The valid range of MULTIPLEXING is from 1 to 8.
# Default: 1 (meaning no multiplexing)
#------------------------------------------------------------------------
#MULTIPLEXING  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
# if the R/3 release is 3.0C and higher and the parameter MAX_VERSIONS 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


#------------------------------------------------------------------------
# Indicates whether processing is to be done unattended or whether human
# intervention is allowed.
# Default:
# YES for backup processing
# NO  for restore processing
#------------------------------------------------------------------------
#BATCH                  YES                # unattended automated operation
#BATCH                  NO                 # manual operation


#------------------------------------------------------------------------
# Control of error situations: Indicates whether and when database backups
# and restore operations should be ended when an error occurs during
# unattended processing.
# Valid values:
# YES: Exit if a single file cannot be backed up or restored.
# NO:  Do not exit when an error occurs.
# the number of errors resulting in exiting the processing.
# The valid range of EXITONERROR is from 0 to 100.
# Default: NO.
#------------------------------------------------------------------------
#EXITONERROR  3                  # exit after 3 errors


#------------------------------------------------------------------------
# Control of information for reporting purposes, e.g. messages, statistics.
# Default: NO (no additional data will be reported).
```

```
#-------------------------------------------------------------------------
#REPORT   NO                  # no additional messages
#REPORT   YES                 # all additional messages
#REPORT   2                   # all additional messages + summary


#-------------------------------------------------------------------------
# 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


#-------------------------------------------------------------------------
# The full path of the trace file.
# Note: for an actual trace the string '%BID' will be replaced by
# the current backupid.
# (.../backint_%BID.trace changes to .../backint_SAP___9809182300.trace).
# Default: none.
#-------------------------------------------------------------------------
#TRACEFILE  /oracle/C21/dbs/backint.trace
#TRACEFILE  /oracle/C21/dbs/backint_%BID.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  /oracle/C21/dbs/initSID.bki


#-------------------------------------------------------------------------
# Number of times to retry saving/restoring a file in case an error occurs.
# The valid range of FILE_RETRIES is from 0 to 100.
# Default: 3.
#-------------------------------------------------------------------------
#FILE_RETRIES  3


#-------------------------------------------------------------------------
# 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


#-------------------------------------------------------------------------
# Denotes if Data Protection for SAP (R) shall use a manual sorting file
# for disk sorting.
# Default: none.
#-------------------------------------------------------------------------
#SORT_FILE  /oracle/C21/dbs/manual_sort_file


#-------------------------------------------------------------------------
```

```
# Denotes if Data Protection for SAP (R) shall use a compressed filesize
# sorting file for disk sorting.
# For backup simulations with compression (see manual) this parameter must
# be set to a valid file.
# Default: none.
#------------------------------------------------------------------------
#COMPR_INFO  /oracle/C21/dbs/initSID.cfi


#------------------------------------------------------------------------
# If Tivoli Storage FlashCopy Manager is used to offload backups to
# another host and Oracle RMAN should be utilized for these backups
# then the following parameters need to be activated. This is not
# required for Oracle RMAN backups on the production system
#------------------------------------------------------------------------
# Type of RMAN backup to perform (CUMULATIVE, DIFFERENTIAL, NO).
# Default: NO (disables the function)
#------------------------------------------------------------------------
#INCREMENTAL CUMULATIVE
#------------------------------------------------------------------------
# Number of RMAN channels to establish.
# Default: 1
#------------------------------------------------------------------------
#INCREMENTAL_CHANNELS 2
#------------------------------------------------------------------------
# Incremental level for the backup (0 or 1).
# Default: 0
# Optional time specifications can be defined with the USE_AT clause.
#------------------------------------------------------------------------
#INCREMENTAL_LEVEL 1 USE_AT MON TUE WED Thu Fri Sat
#INCREMENTAL_LEVEL 1 USE_AT SUN FROM 00:00 TO 06:00
#INCREMENTAL_LEVEL 0 USE_AT SUN FROM 06:01 TO 23:59
#------------------------------------------------------------------------
# Name of the recovery catalog database.
#------------------------------------------------------------------------
#INCREMENTAL_CATALOG_CONNECT_STRING catdb
#------------------------------------------------------------------------
# Name of the user that is used to connect against the recovery catalog database.
#------------------------------------------------------------------------
#INCREMENTAL_CATALOG_USER rman
#*************************************************************************
# 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
  BRBACKUPMGTCLASS    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_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
#  BRBACKUPMGTCLASS    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 Oracle Profile for Windows

```
#-----------------------------------------------------------------------------
#
# Data Protection for SAP (R) interface for ORACLE
#
# 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 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 parallel sessions to be established for the database backup.
```

```
# Note: This number must not exceed the number of tape drives simultaneously
# available to the node for a database backup on the Tivoli Storage Manager
# servers to be accessed.
# The valid range of MAX_BACK_SESSIONS is from 1 to 32.
# Default: MAX_SESSIONS.
#------------------------------------------------------------------------
#MAX_BACK_SESSIONS 1 # Tivoli Storage Manager client sessions for backup


#------------------------------------------------------------------------
# Number of parallel sessions to be established for the redo log backup.
# Note: This number must not exceed the number of tape drives simultaneously
# available to the node for a redo log backup on the Tivoli Storage Manager
# servers to be accessed.
# The valid range of MAX_ARCH_SESSIONS is from 1 to 32.
# Default: MAX_SESSIONS.
#------------------------------------------------------------------------
#MAX_ARCH_SESSIONS 1 # Tivoli Storage Manager client sessions for archive


#------------------------------------------------------------------------
# Number of parallel sessions to be established for the backup of control
# files. This number is typically used to reduce the number of sessions
# to be used for the control file backup after another backup operation.
# The valid range of MAX_CONTROL_SESSIONS is from 1 to 32.
# Default: MAX_BACK_SESSIONS or MAX_ARCH_SESSIONS, depending on the type of
#          the control file backup.
#------------------------------------------------------------------------
#MAX_CONTROL_SESSIONS 1 # Tivoli Storage Manager client sessions for control
                        # file backup.

#------------------------------------------------------------------------
# Number of parallel sessions to be established for the restore of files.
# Note: This number must not exceed the number of tape drives simultaneously
# available to the node for restore processing backup on the Tivoli Storage
# Manager servers to be accessed.
# The valid range of MAX_RESTORE_SESSIONS is from 1 to 32.
# Default: MAX_SESSIONS.
#------------------------------------------------------------------------
#MAX_RESTORE_SESSIONS 1 # Tivoli Storage Manager client sessions for restore


#------------------------------------------------------------------------
# Number of backup copies of redo logs.
# The valid range of REDOLOG_COPIES is from 1 to 9.
# Default: 1.
#------------------------------------------------------------------------
#REDOLOG_COPIES  2


#------------------------------------------------------------------------
# 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


#------------------------------------------------------------------------
# Specifies how many files are read simultaneously and are multiplexed into
# one data stream to a Tivoli Storage Manager server. Multiplexing is usefull
# when the data rate to a Tivoli Storage Manager server is higher (fast
# tapes, fast network) than the I/O rate of a single disk.
```

```
# The valid range of MULTIPLEXING is from 1 to 8.
# Default: 1 (meaning no multiplexing)
#------------------------------------------------------------------------
#MULTIPLEXING  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 SAP R/3 release is 3.0C and higher and the parameter
# 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


#------------------------------------------------------------------------
# Indicates whether processing is to be done unattended or whether human
# intervention is allowed.
# Default:
# YES for backup processing
# NO  for restore processing
#------------------------------------------------------------------------
#BATCH                 YES                # unattended automated operation
#BATCH                 NO                 # manual operation


#------------------------------------------------------------------------
# Control of error situations: Indicates whether and when database backups
# and restore operations should be ended when an error occurs during
# unattended processing.
# Valid values:
```

```
# YES: Exit if a single file cannot be backed up or restored.
# NO:  Do not exit when an error occurs.
# the number of errors resulting in exiting the processing.
# The valid range of EXITONERROR is from 0 to 100.
# Default: NO.
#------------------------------------------------------------------------
#EXITONERROR  3                  # exit after 3 errors


#------------------------------------------------------------------------
# Control of information for reporting purposes, e.g. messages, statistics.
# Default: NO (no additional data will be reported).
#------------------------------------------------------------------------
#REPORT   NO                  # no additional messages
#REPORT   YES                 # all additional messages
#REPORT   2                   # all additional messages + summary


#------------------------------------------------------------------------
# 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


#------------------------------------------------------------------------
# The full path of the trace file.
# Note: for an actual trace the string '%BID' will be replaced by
# the current backupid.
# (...\backint_%BID.trace changes to ...\backint_SAP___9809182300.trace).
# Default: none.
#------------------------------------------------------------------------
#TRACEFILE  x:\oracle\C21\database\backint.trace
#TRACEFILE  x:\oracle\C21\database\backint_%BID.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  x:\oracle\C21\database\initSID.bki


#------------------------------------------------------------------------
# Number of times to retry saving/restoring a file in case an error occurs.
# The valid range of FILE_RETRIES is from 0 to 100.
# Default: 3.
#------------------------------------------------------------------------
#FILE_RETRIES  3


#------------------------------------------------------------------------
# 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


#-----------------------------------------------------------------------
# Denotes if Data Protection for SAP (R) shall use a manual sorting file
# for disk sorting.
# Default: none.
#-----------------------------------------------------------------------
#SORT_FILE  x:\oracle\C21\database\manual_sort_file


#-----------------------------------------------------------------------
# Denotes if Data Protection for SAP (R) shall use a compressed filesize
# sorting file for disk sorting.
# For backup simulations with compression (see manual) this parameter must
# be set to a valid file.
# Default: none.
#-----------------------------------------------------------------------
#COMPR_INFO  x:\oracle\C21\database\initSID.cfi


#************************************************************************
# 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
  BRBACKUPMGTCLASS    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_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
#  BRBACKUPMGTCLASS    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
```

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 14. 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)
programid	Specifies the program that handles the result of the SQL statement. • programid 0: Administration Assistant

Table 14. Contents of the Custom SQL File (continued)

Tag	Definition
actionid	Defines the way the result will be handled, depending on the programid (currently, the only value for actionid is 0): • (programid 0: Administration Assistant): Send e-mail when threshold exceeded (SQL statement returns data)
displaygroup	List of display group IDs separated by commas, or "ALL" for all display groups.
system	List of system IDs separated by commas, or "ALL" for all systems.
backuptype	List of backup types separated by commas, or "ALL" for all backup types. • 0: Archive • 1: Partial backup • 2: Incremental backup • 3: Full backup
executionmode	executionmode sets the time the entry will be performed (i.e., the SQL statement issued): • 0: Entry will be performed after each backup run • 1: Entry will be performed periodically
param	Parameters needed by the programs. The number of parameters depends on the selected program and action. Multiple parameters are coded using repeating *param/param* tag pairs. • (programid 0: Administration Assistant): – 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 15. Tags for Defining Thresholds in the Custom SQL File

Tag	Value
sql	An SQL statement that will return data when the threshold is exceeded.
programid	0 (Administration Assistant)

Table 15. Tags for Defining Thresholds in the Custom SQL File (continued)

Tag	Value
actionid	0 (send e-mail when threshold exceeded)
executionmode	1 (run periodically)
param	(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 *param/param* 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 Oracle files and samples

Use these file samples to assist with Data Protection for SAP for Oracle operations.

Save and Delete Redo Logs Batch File Sample

```
@echo off
rem ------------------------------------------------------------------------
rem file name: archive.cmd
rem ------------------------------------------------------------------------
rem Sample BRArchive batch file
rem ------------------------------------------------------------------------
rem Task:
rem Invokes the SAP utility BRArchive in order to save ORACLE's archived
rem redo logs (using Data Protection for SAP (R) ) and deletes the redo
rem logs from their original location. After completing this, the BRArchive
rem protocol is saved separately.
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 Remarks on the parameters of BRArchive:
```

```
rem
rem -u system/manager ORACLE username/password
rem -sd save and delete archived redo logs
rem -c run BRArchive in quiet mode
rem (-n number of redo logs to be saved,
rem default is 10000,
rem which means all available)
rem
rem The following should be configured within the SAP profile
rem initC21.sap:
rem
rem backup_dev_type = util_file
rem causes BRBACKUP to use the external program
rem Data Protection for SAP (R)
rem util_par_file = %ORACLE_HOME%\database\initC21.utl
rem Data Protection for SAP (R) profile
rem ------------------------------COMMAND----------------------------------
brarchive -u system/manager -sd -c
```

Save and Delete Redo Logs Shell Script Sample

```
#!/bin/ksh
# ------------------------------------------------------------------------
# archive.ksh:
# Sample BRARCHIVE shell script
# ------------------------------------------------------------------------
# Task:
# Invokes the SAP utility brarchive in order to save ORACLE's archived
# redo logs (using Data Protection for SAP (R) ) and deletes the redo
# logs from their original location. After completing this, the brarchive
# protocol is saved separately.
# ------------------------------------------------------------------------
#      *****     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     *****
# ------------------------------------------------------------------------
#
# Remarks on the parameters:
#
# -u system/manager     Oracle username/password
# -sd                   save and delete archived redo logs
# -c                    run BRARCHIVE in unattended mode
# (-n                   number of redo logs to be saved, default is 10000,
#   which means all available)
#
# The following should be configured within the SAP profile initC11.sap:
#
# backup_dev_type = util_file
#  causes brbackup to use the external program backint
# util_par_file =  initC11.utl
#  Data Protection for SAP profile
#
# --------------------------------COMMAND----------------------------------
brarchive -u system/manager -c -sd
```

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 Oracle 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 Oracle
#
# 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 Oracle 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 Oracle
#
# 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

```
************************************************************************
* IBM Tivoli Storage Manager                                           *
*                                                                      *
* Sample Client User Options file for Unix platforms                   *
************************************************************************

SErvername       server_a
Tapeprompt       No
DOM              /usr/sap /sapmnt/C11 /usr/sap/trans /oracle/C11
```

Client User Options File Sample (dsm.opt) Windows

Data Protection for SAP for Oracle 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)

```
* ------------------------------------------------------------------------
* inclexcl.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 AIX 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 (e.g. using 'smit mksysb'). Consequently, you can exclude
*       at least the following directories (which make up about 30 MB).
*
exclude /usr/games/.../*
exclude /usr/bin/.../*
exclude /usr/lbin/.../*
exclude /usr/mbin/.../*
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 SAP's BR*Tools utilities brbackup/brarchive.
*
exclude /oracle/C11/saparch/.../*
* exclude /oracle/C11/sapbackup/.../*
* exclude /oracle/C11/sapreorg/.../* (There may be important scripts
*                                     located, check it out and decide.)
exclude /oracle/C11/sapdata*/.../*
exclude /oracle/C11/sapraw*/.../*
* ------------------------------------------------------------------------
*
* With the above include/exclude list we implicitly include everything not
* excluded above. Especially for DP for SAP (R), this means including:
*      /sapmnt/C11     > 270 MB
*      /usr/sap        > 14 MB
*      /oracle/stage   > 89 MB
*      /oracle/C11     > 90 MB
* and  OS  related     > 220 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 Oracle from routine Tivoli Storage Manager incremental backups. Typically, such files are Windows system files and Oracle 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 ORACLE database is done by
* Data Protection for SAP (R) and not by Tivoli Storage
* Manager command-line backup client, the ORACLE 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 ORACLE database files:
*
Exclude *:\oracle\C21\saparch\...\*
Exclude *:\oracle\C21\sapbackup\...\*
Exclude *:\oracle\C21\sapreorg\...\*
Exclude *:\oracle\C21\sapdata*\...\*
```

Client Options Files Sample (*server*.opt)

Data Protection for SAP for Oracle 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.
*
************************************************************************
COMMmethod         TCPIP
COMPression        OFF
*NODEname           C21
TCPPort            1500
TCPServeraddress   xxx.xxx.xxx.xxx
PASSWORDACCESS     PROMPT
TCPBUFFSIZE        31
TCPWINDOWSIZE      32
```

Data Protection for SAP for Oracle planning sheets

Uses these planning sheets to assist with installing and configuring Data Protection for SAP for Oracle.

Data Protection for SAP for Oracle (base product) planning sheet

Collect the information in this planning sheet before installing Data Protection for SAP for Oracle. This table is also provided in file form as `planning_sheet_oracle` for UNIX and Linux and `planning_sheet_oracle.txt` for Windows.

Table 16. Installation Parameters for Data Protection for SAP

UNIX or Linux	Windows	Installation Parameter
X	X	Oracle database SID:
X	X	Path where the SAP BR*Tools reside: Default: /usr/sap/*SID*/SYS/exe/run or C:\oracle*SID*\sapmnt\SYS\exe\run
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 59.
X	X	Tivoli Storage Manager management classes for database and redo log backups. Management classes configured for the database backup and for the backup of redo logs. For details, refer to "4. Define a policy" on page 58. Default: MDB for database backups, MLOG1 and MLOG2 for redo log 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

Table 16. Installation Parameters for Data Protection for SAP (continued)

UNIX or Linux	Windows	Installation Parameter
	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 `dsierror.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\ora64
X	X	Path for Data Protection for SAP configuration files (directory for SAP configuration files). During the installation, the Data Protection for SAP configuration files will be saved to this path. If old configuration files are found, they are renamed to *`filename`*`.nnn`, where `nnn` is a three-digit decimal number.This path must not contain blanks. Default: /oracle/*SID*/dbs or C:\orant\database
X	X	Options: • Use of Oracle RMAN. • Use of the Administration Assistant (see "Administration Assistant function for Data Protection for SAP" on page 6 and Table 17). The Administration Assistant should be installed prior to Data Protection for SAP so that the interface between the two can be automatically established.

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 17. 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

Table 17. Installation Parameters for the Administration Assistant function for Data Protection for SAP (continued)

Installation Option	Installation Parameter
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
Parameters applying to the Administration Assistant Server component	Hostname or IP address: Default: Hostname of current system Port number for Data Protection for SAP for Oracle (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 Port number for client connect in non-secure mode (HTTP). Default: 80 Port number for client connect in secure mode (HTTPS). Default: 443 RMI registry port number Default: 1099 Port number for performance data from Database Agent Default: 5129 Port number for communication with Database Agent Default: 5128

Table 17. Installation Parameters for the Administration Assistant function for Data Protection for SAP (continued)

Installation Option	Installation Parameter
Parameters applying to the Administration Assistant Database Agent component	Hostname or IP address: Default: Hostname of current system Port number for Data Protection for SAP (ProLE) connection Default: 5125 Port number for communication with Administration Assistant Server component Default: 5127
Parameters applying to the Administration Assistant Database component (Apache Derby)	Hostname or IP address: Default: Hostname of current system Port number for database connect Default: 1527 User ID and password to access internal database.
Parameters applying to the Administration Assistant Database component (IBM DB2)	Hostname or IP address: Default: Hostname of current system Port number for database connect Default: 50000 User ID and password of the system user for which the DB2 instance should be installed that the internal database should access.
Installation directory	Installation directory (on each host) Default: `/opt/tivoli/tsm/tdp_r3_assist` on UNIX and Linux, or `C:\Program Files\tdpr3assi` on Windows.
Product Support	Location of mail.jar (Java Mail)
Product Support	Location of activation.jar (Java Beans Activation Framework):
History file	History file directory (on Server component host) Default: `history` (in installation directory) History file retention time (days). Can be changed via the Administration Assistant client. Default: 14

Table 17. Installation Parameters for the Administration Assistant function for Data Protection for SAP (continued)

Installation Option	Installation Parameter
Secure Communication	Information on the public key infrastructure (PKI): • *Keystore name*. 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. • *Keystore password.* Password ensuring the consistency of the keystore. The server's key pair must be protected by the same password. • *Truststore name*. 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. • *Truststore password*. 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. • *Certificate file*. Path of the certificate file in case you already have a server certificate issued by a certificate authority. • *Certificate creation information.* 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. • *New certificate file name*. 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. • *CSR file name*. 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	DB2 JDBC Universal Driver. The corresponding packages are bundled with your IBM DB2 installation. • db2jcc.jar location:Default: None • db2jcc_license_cu.jar location:Default: None 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. Default: None The name of the internal database is predefined and cannot be changed. Default: `AADB`

Table 17. 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 18. 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 19. Tuning of Network Settings

Attributes	Value	Description
`rfc1323`	`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.
`sb_max`	`131072`	Specifies the maximum buffer size allowed for a socket. The default is 65536 bytes. From the point of view of performance recommendations, the `sb_max` value should be twice the `TCPWindowsize` set within the Tivoli Storage Manager configuration file `dsm.sys`.

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 20. Tuning of SP Switch Buffer Pools

Attributes	Value	Description
`rpoolsize`	`1048576`	The receive pool is a buffer pool for incoming data. The size for values is in bytes.
`spoolsize`	`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 `chgcss` 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 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 (ACSLS). 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 **dsmattr** 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

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

R

S

T

U

V

W

Product Number: 5608-E05

Printed in USA

SC33-6340-12