IBM Tivoli Storage Manager for Enterprise Resource Planning
Data Protection for SAP HANA
Version 7.1.3

Installation and User's Guide

IBM Tivoli Storage Manager for Enterprise Resource Planning
Data Protection for SAP HANA
Version 7.1.3

Installation and User's Guide

Note:
Before you use this information and the product it supports, read the information in "Notices" on page 57.

Second edition (2015)

This edition applies to version 7, release 1, modification 3 of IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP HANA (product number 5608-E05), available as a licensed program. It also applies to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Tables

About this publication

This publication documents how to use IBM® Tivoli® Storage Manager for Enterprise Resource Planning: Data Protection for SAP HANA. It describes the procedures that are needed to install, configure, and protect your SAP HANA data with Data Protection for SAP HANA.

The Data Protection for SAP HANA product is the interface between SAP HANA and the Tivoli Storage Manager server.

Who should read this publication

This publication is intended for system programmers and administrators who are responsible for implementing a backup solution in a SAP HANA environment with the Tivoli Storage Manager server.

It describes the procedures that are needed to install and customize IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP HANA, the interface between SAP HANA and Tivoli Storage Manager.

The reader must be familiar with the documentation for SAP HANA and the Tivoli Storage Manager server.

Publications

The Tivoli Storage Manager product family includes IBM Tivoli Storage FlashCopy® Manager, IBM Tivoli Storage Manager for Space Management, IBM Tivoli Storage Manager for Databases, and several other storage management products from IBM.

To view IBM product documentation, see http://www.ibm.com/support/knowledgecenter.

What's new

Read about features that are new in this version of IBM Tivoli Storage Manager for Enterprise Resource Planning.

Data Protection for SAP HANA

This software update provides support for SAP HANA® on POWER®. There are no specific instructions related to this additional platform support.

All components

From the Tivoli Storage Manager Operations Center, you can view front-end and back-end capacity usage to estimate compliance with license entitlements.

Important: The Tivoli Storage Manager client software needs to be at the latest level for the Operations Center to present complete data.

Chapter 1. Getting started

Data Protection for SAP HANA operates as a link between SAP HANA and the Tivoli Storage Manager server.

The Data Protection for SAP HANA `hdbbackint` process is used by SAP HANA for backing up full databases and redo log files. The configuration of the `hdbbackint` process is stored in the `initSID.utl` profile file. This file contains information that describes how to run backup and restore operations, and can be customized for your SAP HANA environment.

The `hdbbackint` process communicates with the Tivoli Storage Manager server through an API that is shared with other IBM data protection products. It requires that the Data Protection for SAP HANA `ProLE` process is running. The `ProLE` process coordinates multiple `hdbbackint` instances in a full database backup. The process ensures that all backup objects that belong to the same full database backup get assigned to the same backup id. The full database backup is handled as a single entity even it consists of numerous single objects.

In a SAP HANA scale-out environment that consists of multiple SAP HANA nodes, Data Protection for SAP HANA is running on each node.

Depending on the number of SAP HANA services that are on a node, multiple instances of `hdbbackint` are started by SAP HANA for data transfer.

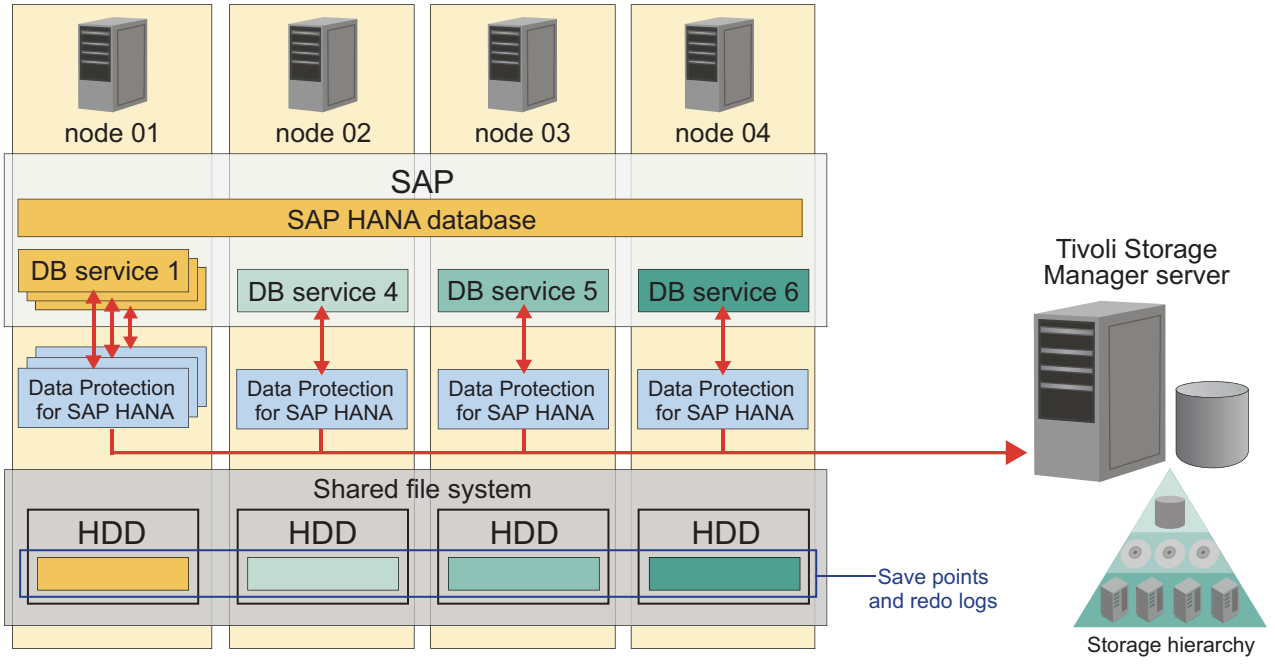

Figure 1. An example of how Data Protection for SAP HANA, Tivoli Storage Manager and SAP HANA are integrated.

The backint interface

The **backint** interface communicates between SAP HANA, Data Protection for SAP HANA, and Tivoli Storage Manager to run full online and offline backups of SAP HANA databases and redo log files. The **backint** interface communicates directly with SAP.

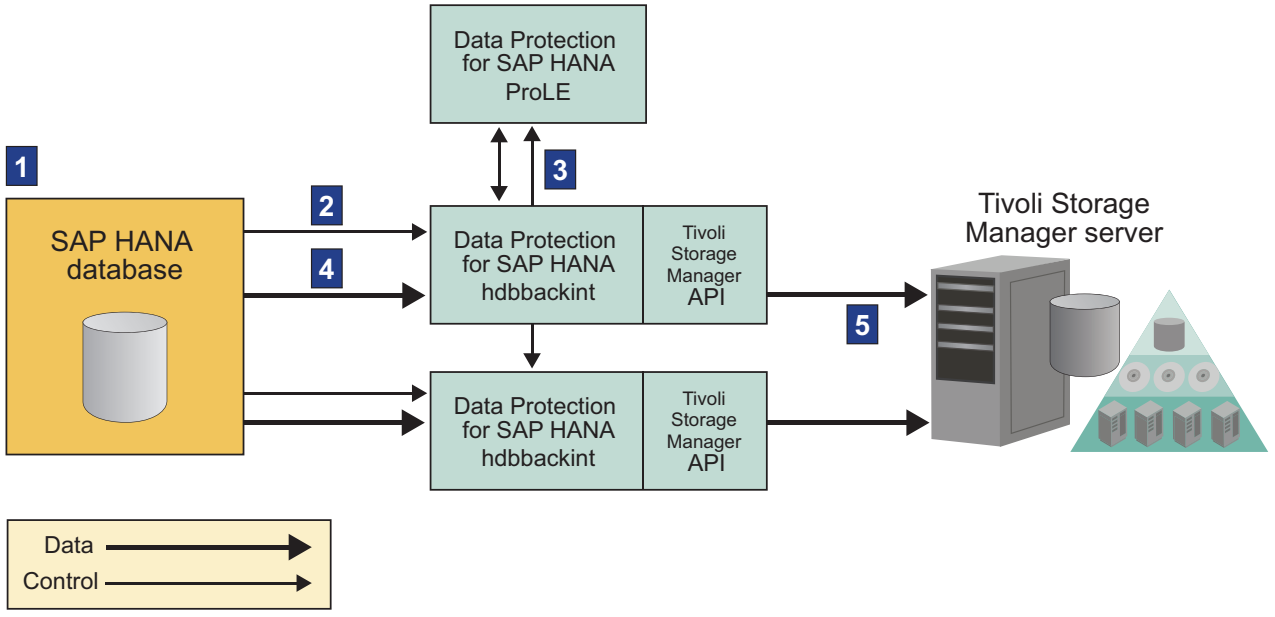

Figure 2. The interactions between SAP HANA and the Tivoli Storage Manager through the Data Protection for SAP HANA **hdbbackint** *interface.*

A Data Protection for SAP HANA backup operation proceeds in the following order:

1. You start a backup operation using the SAP HANA Studio or through the **hdbsql** command-line interface.
2. A number of SAP HANA **hdbbackint** processes are started.
3. The **hdbbackint** processes connect to the ProLE to get the configuration information.
4. SAP HANA sends data to the **hdbbackint** processes.
5. The data is sent to the Tivoli Storage Manager server through the Tivoli Storage Manager API.

All the database and redo log backup information is stored by SAP HANA. The backup history can be viewed in the SAP HANA Studio in the backup tab or by querying the database view M_BACKUP_CATALOG. For more information about **hdbbackint**, see the *SAP HANA Administration Guide*.

Chapter 2. Planning

Planning information regarding strategies and components is provided.

Backup strategy for Data Protection for SAP HANA

To avoid data loss from SAP HANA databases, you must run full backups and redo log file backups regularly.

Full database backup

The full SAP HANA database is backed up to Tivoli Storage Manager for Enterprise Resource Planning. If needed, the database can be restored without extra redo log files. The frequency of full backups is controlled by the Tivoli Storage Manager recovery time objective (RTO) and recovery point objective (RPO). For more information about full backups, see the *SAP HANA Administration Guide*.

Redo log file backups

Back up redo log files so that if data is lost in between two full backups, the data can be restored to a particular point in time. Set the **SAP HANA LOG_BACKUP_USING_BACKINT** parameter to `TRUE` to enable backing up redo log files. Edit the frequency of the backup operations with the **LOG_BACKUP_TIMEOUT_S**. The default value is for a backup every 15 minutes. For more information about redo log file backups, see the *SAP HANA Administration Guide*.

Backing up data with the backup-archive client

Use the Tivoli Storage Manager backup-archive client include/exclude function to back up any files that are not backed up by the full database backup, such as configuration files. A scheduled incremental backup would ensure that the configuration files are backed up periodically to be used if data was lost.

```
*
* Sample include/exclude list for SAP HANA appliances
*
* first exclude everything
exclude /.../*
*
* now include relevant files and directories only
include /usr/sap/C21/SYS/profile/.../*
include /usr/sap/C21/SYS/global/hdb/custom/config/.../*
```

Backup strategy for the SAP HANA appliance

If you experience a total loss of the SAP HANA appliance, use the SAP HANA recovery procedure to recover the system. For information about hardware recovery, see the SAP documentation.

Optimization of Data Protection for SAP HANA performance

When SAP HANA backs up data, it sends data for all SAP HANA services, such as **nameserver** and **indexserver** in parallel from all nodes. If data is intended to go directly to tape, the number of allowed mount points for the Tivoli Storage Manager node must be adjusted.

When you are planning to store data in a disk storage pool and then migrate it to tape, be aware that SAP HANA data is restored in the order that is determined by HANA. Use either disk storage pools or tapes with enough mount points to optimize the restore by reducing the number of tape-seeks and mounts. If the data to be restored for all nodes is on a single tape drive, it can lead to increased restore times or deadlocks.

Network performance optimization

When you are setting up the network, there are some items to consider that can improve network performance.

Consider these items when you set 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 you use LAN-free backup, ensure that Fibre Channel adapter capacity to the SAN can accommodate the data transfer rates of the disk reads and tape writes.

Network bandwidth
In general, 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
Increase the overall throughput rate to the backup server by providing a way to specify multiple network paths.

Backup server optimization

When you are setting up the Tivoli Storage Manager server, there are some items to consider to optimize performance.

Consider these items when you set up the Tivoli Storage Manager server. Data Protection for SAP 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 usage.

CPU power
For a specific data throughput, the CPU load on the backup server is

approximately 60% of the load on the database server. Therefore, backup server CPU power is not 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 is to be directed to tape to achieve the highest transfer rates. If disks must be used, use one disk pool per session. Small files such as log files, are to be directed to disk storage first and then moved to tape collectively to avoid excessive tape mounts.

Parallel sessions
The Tivoli Storage Manager server allows the use of several tape drives in parallel to store data. This setup can increase overall data throughput. To fully use this feature, two conditions must exist. 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.

Store data on the Tivoli Storage Manager server

In SAP terminology, *backup* (**BRBACKUP**) refers to the backup of data; *archive* (**BRARCHIVE**) refers to the backing up of log files. Data Protection for SAP uses the Tivoli Storage Manager archive function for backups and archives.

Tape storage is the preferred media for storing database contents as it provides the best data throughput for backup and restore operations. For a large scale-out system, the number of required tape drives might become too large. In this case, use a virtual tape library (VTL). A disk-tape storage hierarchy is used for backing up redo log files. This action provides the best protection against data loss, and eliminates the need to mount a tape for each redo log file.

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 that is 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 that are required for a complete restore.

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

- Deactivate collocation for Data Protection for SAP backups when you enable parallel sessions for use with multiple tape drives in parallel.
- Deactivate collocation when you use the multiple log copy function.

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 you remove a Tivoli Storage Manager server from the Data Protection for SAP profile.

Because Data Protection for SAP accesses only those servers that are defined in the profile, be cautious when you remove a Tivoli Storage Manager server if it contains valid backup data.

Database backups are retained for a specified period and then become obsolete. Manage backup storage space by deleting obsolete backups and automating archive retention period with Tivoli Storage Manager options. Alternatively the obsolete backups can be deleted manually in the SAP HANA Studio.

Database backups are retained for a specified period and then become obsolete. Manage backup storage space efficiently, by deleting obsolete backups in one of the following ways:

- 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 that are specified by this function is exceeded, entire backup generations are deleted. The backups that can be deleted are full backups and all related redo log file backups.
- Use the Data Protection for SAP backup version control function. When the number of backup versions that are specified by this function is exceeded, entire backup generations are deleted. The backups that can be deleted are full backups and all related redo log files backups.

Planning for space required for SAP HANA backups

Before Data Protection for SAP HANA sends data to Tivoli Storage Manager, it notifies the Tivoli Storage Manager server of the amount of data that is going to be sent. This enables the Tivoli Storage Manager server to select the appropriate storage pool that accommodates this specific amount of data.

About this task

Through the `backint` SAP HANA interface, Data Protection for SAP HANA does not know the amount of data that is due to be sent by SAP HANA. It is assumed that half of the memory size is to be transferred. If the system has 1 TB RAM, then the estimated size for the backup object is assumed to be 512 GB. For the data that is stored by the SAP HANA `indexserver` this number is close to the value for a fully used system. For other SAP HANA services like `nameserver` or `statisticserver`, the value is much smaller. Typically the backups of these services are several MB only, and a backup can result in a storage pool that is intended for large objects being used.

Procedure

- To avoid backing up small items to the next storage pool in line, the required size of the first storage pool should be at least half the size of the appliance memory multiplied with the number of sessions (SESSIONS, MAX_SESSIONS) that are simultaneously sending files. If you use the **`maxsize`** parameter for the first storage pool, make sure it is at least half the size of the appliance memory.
- Another option is to use Virtual Tape Library for backups.

Parallel backup paths and backup servers

Data Protection for SAP can use several communication links between Tivoli Storage Manager clients 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 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`. The *server 1..n* statement denotes Tivoli Storage Manager servers that are defined in the Data Protection for SAP profile. This definition corresponds to the statement `SERVERNAME` *`server 1..n`* in the Tivoli Storage Manager client option file or files. These servers are identified by their `TCPSERVERADDRESS` and can be on one

system (multiple paths) or several systems (multiple servers). SESSIONS denotes the number of parallel sessions that Data Protection for SAP schedules for the 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 by using the first path in the profile. If this attempt is successful, Data Protection for SAP starts the number of parallel sessions as specified for this path.

If the attempt is 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 ends the backup job.

Backup backint processes

When SAP HANA runs a database backup operation through Data Protection for SAP HANA, a dedicated backint process for each backup object is started. It is possible to use multiple sessions in a single backint process, and each backint process transfers a single object.

You do not need to configure multiple sessions for database backup operations. Multiple backint processes run in parallel, and all objects are backed up in parallel. The number of objects depends on the number of SAP HANA nodes in the environment. All nodes host an index server that holds the payload of the database. The backup objects from the index servers are typically the largest objects. The master node hosts a name server, a statistic server, and other services. Backup objects from these additional services are typically smaller than the index server. If the backups must go directly to tape, then the number of mount points must include the additional services.

Redo log files

Configure multiple sessions for a single backint process for backing up redo log files. When you are using the redo log copy feature of Tivoli Storage Manager for Enterprise Resource Planning, each redo log file is saved simultaneously in multiple storage pools on the Tivoli Storage Manager servers. Ensure that you have the sane number of sessions opened as redo log copies that are created to be able to duplicate and transfer data.

Archive inactive data

Data Protection for SAP creates a database image that is stored at the bit-level and can be used for routine backup operations.

To restore an outdated backup, you must restore it into the same environment it was originally taken from. This process requires you to maintain older versions of SAP, the operating system, database, and Tivoli Storage Manager data to enable a rebuild of the 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.

Restore versus backup

Configuration changes and infrastructure problems affect backup and restore operations.

Changes that support a fast backup while you are using resources can be considered applicable to the restore operation. Tune the backup operation and then run a restore to verify that the restore operation works in a satisfactory manner.

If backups are compressed during a restore operation, the data must be decompressed before you can use it.

Create multiple redo log copies

Data Protection for SAP 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 corrupted.

The Data Protection for SAP profile file keywords important for creating multiple redo log copies:

- Keyword **BRARCHIVEMGTCLASS** denotes the Tivoli Storage Manager server management classes to be used when it saves redo logs. With the use of different management classes, the backup media that is targeted for redo logs is separated from the backup media that is 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 databases in an SAP environment. The restore and recover can occur even if a backup tape becomes corrupted 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 as many archive management classes that are defined by **BRARCHIVEMGTCLASS** 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. This way, various redo log copies are on different backup media.
- When SAP HANA is used, Data Protection for SAP HANA requires that the maximum number of sessions that are defined by **MAX_SESSIONS** is greater than or equal to the number of redo log copies that are requested. A setup with a smaller number of sessions is not advised 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.

Chapter 3. Installing

Install Data Protection for SAP HANA using the install wizard, through the console, or in silent mode using a response file.

Before you begin

Requirements for Data Protection for SAP HANA are available in the hardware and software requirements technote for each release. For requirements, review the *Hardware and Software Requirements* technote for your version. See the technote at http://www.ibm.com/support/docview.wss?uid=swg21219410. From the page, follow the link to the technote for your release or update level.

Before you install Data Protection for SAP HANA, verify that your system meets the following prerequisites:

- SAP HANA SPS 05 revision 45 is installed.
- IBM Tivoli Storage Manager API Version 5.5 or later is installed and configured on all SAP HANA nodes where you are going to install and configure Data Protection for SAP HANA.
- The SAP HANA database is configured on the system where you are going to install and configure Data Protection for SAP HANA.
- The SAP HANA HDB client is installed on the system.
- During the installation and configuration of Data Protection for SAP HANA, root access to the appliance host operating system is required.

Procedure

1. Log in to the SAP HANA host with the root user ID, and choose where you want to install the package.
2. Mount the DVD and navigate to the Data Protection for SAP HANA installation package. If you are installing from a file share, to ensure that the installer file has adequate permissions to run, enter the following command:

   ```
   chmod +x 7.1.3.0-TIV-TSMERP-HANA-Linux.bin
   ```
3. To start the installation process, enter the following command:

   ```
   ./7.1.3.0-TIV-TSMERP-HANA-Linux.bin
   ```

   If you are logged on to the SAP HANA host with an X Window System or X terminal, follow the instructions to complete the installation.

   If you are installing the product from the command line, a console mode installation is started.
4. After you accept the license agreement, click **Next** and read the pre-installation summary that lists details about the installation folder, and the required disk space. Click **Install** to begin the installation, and click **Done** when the process finishes.

Results

Data Protection for SAP HANA is installed in the following directory: `/opt/tivoli/tsm/tdp_hana`.

Installing Data Protection for SAP HANA in silent mode

You can install Data Protection for SAP HANA in silent mode by using a response file. After you create the response file, you can install the product in silent mode without monitoring the process or inputting any details.

About this task

To install Data Protection for SAP HANA in silent mode, you must first create a response file.

Procedure

1. Create a response file for Data Protection for SAP HANA with the following command:

   ```
   ./6.4.1.0TIVTSMERP-HANA-Linux.bin -i console -r responsefile
   ```

   This command runs the installation process in console mode and all user input is recorded in the response file.
2. Run the following command to install Data Protection for SAP HANA in silent mode:

   ```
   ./6.4.1.0TIVTSMERP-HANALinux.bin -i silent -f responsefile
   ```

   This command runs the installation process automatically without requiring any user input, Values for options are read from the response file.

Replication environments

An environment that has a number of SAP HANA database instances that are synchronized with a primary database instance is a *replication environment*. Replication is possible on the database level, for example with SAP HANA system replication, or on the storage level, for example with GPFS™ storage replication. The replicated database instances are typically not online.

In a replicated SAP HANA environment the standard installation procedure is not applicable. Installation of Data Protection for SAP HANA is done through one of the following methods:

Installation as part of takeover testing.

Manual installation on each SAP HANA node.

Installing to test in a replication environment

In a replication environment, when the primary system is shut down to verify a failover procedure, Data Protection for SAP HANA can be installed as part of the test.

About this task

When you are setting up a replication environment and you are running a takeover test, the primary system is shut down. Install Data Protection for SAP HANA when the SAP HANA database instance in the replicated environment is online.

Procedure

- Install Data Protection for SAP HANA by following the procedure at this link Chapter 3, “Installing,” on page 9.

- Configure Data Protection for SAP HANA by following the procedure at this link "Configuring Data Protection for SAP HANA" on page 21.

Installing manually in a replication environment

In a replicated environment, you need to install Data Protection for SAP HANA manually on each SAP HANA node for a particular database instance. When you are installing in large scale-out environments, you can choose to create a response file to install on each node in silent mode.

Before you begin

For more information about installing Data Protection for SAP HANA in silent mode, see "Installing Data Protection for SAP HANA in silent mode" on page 10.

Procedure

- Install Data Protection for SAP HANA on each SAP HANA node in a database instance.
- Install the Tivoli Storage Manager client API on each of the SAP HANA nodes.
- Copy the Data Protection for SAP HANA profile from the primary database instance, and use this profile in the replicated environment.
- Configure the Tivoli Storage Manager password.
  - If automatic password handling is used, edit **NODENAME** in the server stanza of the `dsm.sys` file so that each SAP HANA node uses a unique Tivoli Storage Manager node name.
  - If manual password handling is used, the Tivoli Storage Manager node password must be stored locally by entering the following command for each SAP HANA node: `/opt/tivoli/tsm/tdp_hana/hdbbackint -p` *`full path to profile`*`/initSID.utl -f password`

Uninstalling Data Protection for SAP HANA

Uninstall Data Protection for SAP HANA, and remove all of the associated files from your system.

Before you begin

If you are using `.rpm` files to uninstall Data Protection for SAP HANA, the uninstallation program remains on the disk. The `.rpm` packages that were installed during the setup on other SAP HANA nodes remain on these nodes.

Procedure

1. Change directory to the `uninstall` subdirectory in the `install` folder.
2. Enter the following command: **/opt/tivoli/tsm/tdp_hana/uninstall/uninstaller.bin**
3. To remove `.rpm` packages that were installed during the setup, enter the following command:

   rpm e TIVTSMERPHANA

Results

The uninstallation program removes Data Protection for SAP HANA from your system. Services that were installed and used by Data Protection for SAP HANA are stopped and removed.

Chapter 4. Configuring

In addition to configuring Data Protection for SAP, you need to configure other applications, for example, the Tivoli Storage Manager backup-archive client.

About this task

Data Protection for SAP requires certain configuration tasks to be run for the following applications.

- Data Protection for SAP base product
- Tivoli Storage Manager backup-archive client
- Tivoli Storage Manager server

Tivoli Storage Manager server tasks

Data Protection for SAP HANA requires configuration tasks to be done for the Tivoli Storage Manager server as part of the overall product configuration.

Configure the Tivoli Storage Manager server

When you are configuring Data Protection for SAP HANA you must set up the Tivoli Storage Manager server, and run general and specific server configurations such as setting up storage devices.

Although the task examples use Tivoli Storage Manager commands, these tasks can also be run using the Tivoli Storage Manager web client GUI.

Consider the following performance-related guidelines before you install the Tivoli Storage Manager server.

Tivoli Storage Manager server host system
The Tivoli Storage Manager server must be installed on an exclusive system. The tasks that are presented here 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 is used to back up and restore other clients, consider installing the server on a large system or by 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. Use fast network topologies 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.

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

- Recovery log volume is allocated and initialized.
- Recovery log mirror volume is allocated and initialized.
- Database volume is allocated and initialized.

- Database mirror volume is allocated and initialized.
- Extra labeled volumes for the backup and archive storage pools are allocated and initialized (disks, tapes, or combinations).
- Licenses are registered.
- The Tivoli Storage Manager server is started.

The latest code fixes for Tivoli Storage Manager can be found at:
`ftp://public.dhe.ibm.com/storage/tivoli-storage-management/maintenance`

Specifying a Tivoli Storage Manager server

To configure Data Protection for SAP HANA, you need to specify a Tivoli Storage Manager in the profile file.

About this task

Follow these steps to add a Tivoli Storage Manager server:

Procedure

1. Add a server statement to the Data Protection for SAP HANA profile.
2. Adapt the Tivoli Storage Manager options files.
3. Set and save the Tivoli Storage Manager password for the new server.

Specifying a storage device

A storage device needs to be added when you are configuring. A storage device defines a device class, which handles the type of media. The default device class that is defined for disks is `DISK` and is considered sufficient.

About this task

Verify that the following 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
```

Defining a storage pool

A storage pool needs to be added when during the configuration. 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.

Procedure

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 : `define stgpool sap_db device_class_name maxscr=20`
3. Define a storage pool for the : `define stgpool sap_log1 device_class_name maxscr=3`

Results

When a library tape device is associated, the maximum *scratch volumes* (labeled volumes that are empty or contain no valid data) that this storage pool is 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.

Defining a policy

A server policy needs to be specified when you are configuring Tivoli Storage Manager policies. Specify how files are backed up, archived, moved from client node storage, and 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*.

About this task

After you set definitions, a default policy set must be assigned, validated, and activated. For the policy definition, log on as a Tivoli Storage Manager Administrator by using the *Admin Command Line* or the *Web Admin* and run the following commands.

Procedure

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 log files and copies of offline redo 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=90
define copygroup sap_c21 p_c21 mlog1 type=archive destination=sap_log1 retver=90
define copygroup sap_c21 p_c21 mlog2 type=archive destination=sap_log2 retver=90
```

   Data Protection for SAP HANA stores data in the archive copy group of the management classes. This data expires after a defined number of days. In the example, a **retver** value of 90 days is used. When no backup is being done within this time frame, all backup data expires and is no longer available for restore. As an alternative the copy group parameter **retver**, which specifies the

number of days a file is to be kept, can be set to unlimited (*9999* or *nolimit*). Obsolete backups can be deleted manually using the SAP HANA Studio.

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
```

Registering a node

The node must be registered when you are configuring The Tivoli Storage Manager server views its registered clients, application clients, host servers, and source servers as nodes.

About this task

To register a node, log on as the Tivoli Storage Manager administrator by using the *Admin Command Line* or the *Web Admin*, run the following command

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

When you use 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 must use more than one mount point, the parameter must be set to the wanted number of mount points. This parameter is not to be set higher than the total number of drives available on the Tivoli Storage Manager server.

Setting 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.

About this task

To avoid sessions timing out, set the parameter **IdleTimeOut** to a value higher than the time required for sending the largest table space file to the Tivoli Storage Manager. For example:

```
setopt IdleTimeOut 60
```

Determining the Tivoli Storage Manager password method

Specify how Data Protection for SAP manages the Tivoli Storage Manager password. There are three options.

About this task

There are three methods of password handling:

No password is required

No authentication is completed on the Tivoli Storage Manager server. Each user that is connected to the backup server can access Tivoli Storage Manager data without a password. This method is advised only if adequate security measures are established.

For example, no password might be acceptable when the Tivoli Storage Manager is only used for SAP, and authentication and authorization is done at the operating system level. This scenario is valid when no other clients are registered to the Tivoli Storage Manager.

Manual handling of password

A password is required for each connection to the Tivoli Storage Manager server. In this method, Data Protection for SAP stores the encrypted password in its configuration files.

While 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 it connects 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:

If you are setting the password to be automated (such as in a script), enter this command: 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 when you record Tivoli Storage Manager passwords in a script.

Automatic handling of password

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.

Setting the Tivoli Storage Manager password:

Data Protection for SAP is to be installed after the Tivoli Storage Manager installation is completed. Tivoli Storage Manager provides different password methods to protect data.

About this task

Data Protection for SAP must use the same method as specified in Tivoli Storage Manager. The default password method during Data Protection for SAP installation is PASSWORDACCESS prompt.

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

```
backom -c password
```

The default parameters for Data Protection for SAP are set according to this default value. If a different password method is set in Tivoli Storage Manager, adjust the Data Protection for SAP parameters.

Password configuration matrix:

After you select the suitable password-handling method, follow this configuration matrix to set the password keywords and parameters.

Proceed as indicated by the step number.

Table 1. 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*
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 profile (`init`*`SID`*`.utl`)	For each SERVER statement, specify:PASSWORDREQUIRED ADSMNODE	NO *nodename*	YES *nodename*	NO (see note 4)
6	Command line	`backom -c password`	Unavailable	*password* (See notes 3, 7)	*password* (See notes 3, 7)

Note:

1. See appropriate Tivoli Storage Manager documentation.
2. If you are using manual password generation during testing, make sure that the expiration period is set to an appropriate 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 must have read and write permission for the `path` specified.
6. This step is only necessary if the password is expired (manual-handling only) or must 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.

Verifying the server name

You must verify that the server name and the parameters are correct in the `init*SID*.utl` file.

Procedure

- Review the Tivoli Storage Manager client options files to make sure that the server name matches the name that is specified in the server statement of the `init*SID*.utl` file.
- Review that other parameters are set correctly. These settings depend on the password method selected.
- (UNIX or Linux) Define the Tivoli Storage Manager server in the Tivoli Storage Manager client system options file (`dsm.sys`). The server stanza that is specified in `dsm.sys` must match the entry in `init*SID*.utl`.

Deleting backups with the `RETVER` parameter

Use the **RETVER** profile parameter to automatically delete obsolete backups.

About this task

Control Data Protection for SAP HANA backup file expiration with the **RETVER** parameter. **RETVER** is the parameter of the archive copy group that is used to store the backup data on the Tivoli Storage Manager server.

For example, to keep 30 daily backups, set **RETVER** to 30.

Tip: Backups are expired even if no backups are run on that day.

Tivoli Storage Manager client tasks

Data Protection for SAP requires that configuration tasks be run for the Tivoli Storage Manager client as part of the overall 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 is used 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.

The clients must be installed on all nodes that interface with the Tivoli Storage Manager server. In a SAP® system landscape, the backup-archive client must be installed on every system that is scheduled for a file system backup. Examples of these systems are SAP application servers and the SAP database server. The Tivoli Storage Manager API client must be installed only on the SAP database server system to enable backup and restore operations of the SAP database by using Data Protection for SAP.

Setting 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 used to define which files are included or excluded during backup, archive, or hierarchical storage processing.

About this task

Configure the Tivoli Storage Manager backup-archive client to operate in an SAP environment with the following procedure.

Procedure

1. Install the Tivoli Storage Manager client software on the SAP database server system.
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`. If the server and client are on the same system, select loopback. This selection improves TCP/IP communication speed.
4. Specify `InclExcl` if you want Tivoli Storage Manager to include or exclude the files that are listed in `inclexcl.list`.
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 advised.
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)
```

Results

When the Tivoli Storage Manager API client is installed on a UNIX or Linux system, ensure that a link exists that points to the Tivoli Storage Manager API installation directory, `/usr/tivoli/tsm/client/api/bin64`.

```
/usr/lib/libApiDS.so
```

The Tivoli Storage Manager provides two features for specifying the location of the Tivoli Storage Manager API Client error log: the environment variable **`DSMI_LOG`** and the Tivoli Storage Manager system client option `ERRORLOGName` in `dsm.sys`. For **`DSMI_LOG`**, a directory is specified to which a file named `dsierror.log` is written. For `ERRORLOGName` a path and user-defined file name are defined.

To achieve conclusive logical linking of the environment, configuration and log files in your SAP backup-archive system, you must use the Tivoli Storage Manager system client option ERRORLOGName rather than the environment variable **DSMI_LOG**.

When you use ERRORLOGName, you can include the SID in the file name. This information can speed up problem determination by simplifying identification of the correct error log file. You can match log file names to the active user client options file name, which must also contain the SID and be stored in environment variable **DSMI_CONFIG**. This information is especially useful on systems with several SIDs.

With this 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** exists in your environment from an earlier setup, it is overridden by dsm.sys option **ERRORLOGName**. However, 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.

Configuring Data Protection for SAP HANA

Configure Data Protection for SAP HANA with the setup.sh script. The setup.sh is stored in the installation directory during the installation process.

Configuring Data Protection for SAP HANA with the setup script

Data Protection for SAP HANA must be configured using the setup script before it can work with the Tivoli Storage Manager. The setup script setup.sh is stored in the Data Protection for SAP HANA installation directory /opt/tivoli/tsm/tdp_hana.

Before you begin

Ensure that the Tivoli Storage Manager node name is specified in the server stanza of the dsm.sys file.

About this task

The Data Protection for SAP HANA configuration files are stored in the configuration directory, /usr/sap/*SystemID*/SYS/global/hdb/opt/hdbconfig.

Procedure

1. Log in to Data Protection for SAP HANA using the root user ID, and change to the installation directory, /opt/tivoli/tsm/tdp_hana.
2. Run the Data Protection for SAP HANA setup.sh script, by entering the following command:

   ```
   ./setup.sh
   ```

3. When prompted, enter the SAP HANA system ID, or accept the proposed value.
4. Enter the instance number or accept the default value of *00*.
5. Enter the password for the `SYSTEM user`.
6. Choose to configure the Tivoli Storage Manager server, or defer the server configuration. For information about manually configuring the Tivoli Storage Manager server, see "Manually configuring the Tivoli Storage Manager for Data Protection for SAP HANA" on page 23.
7. Choose one of the following Tivoli Storage Manager server password handling methods:
   - Automatic password handing: to store the Tivoli Storage Manager node password in the Tivoli Storage Manager API. When the password expires on the server, the Tivoli Storage Manager client and server generates a new password. The API updates the password on the client.
   - Manual password handling: to store the node password in the Tivoli Storage Manager for ERP configuration file. When the password expires, you must update it by using the **hdbbackint -f password** command.
8. Enter the Tivoli Storage Manager server name as defined in the `dsm.sys` file. The following parameters must be entered to enable backup and restore of databases and redo log files to the Tivoli Storage Manager
   a. Enter the Tivoli Storage Manager node name for the **ADSMNODE** parameter.
   b. Enter the Tivoli Storage Manager management class for the **BRBACKUPMGTCLASS** parameter.
   c. Enter the Tivoli Storage Manager management class for the **BRARCHIVEMGTCLASS** parameter.
9. Enter a password for the Tivoli Storage Manager node. After you enter the password, it is verified. For a scale-out system with automatic password handling, this step is repeated for each SAP HANA node that belongs to the scale-out system.

Protecting multiple SAP HANA databases

Multiple SAP HANA database instances that are installed on a single Tivoli Storage Manager host can be protected with Data Protection for SAP HANA.

About this task

Run the setup script on each database to ensure that each is protected, and has a dedicated `init`*`SID`*`.utl` file.

Procedure

Log in with the root user ID. For each SAP HANA database instance, run the `setup.sh`. For more information about the setup process, see "Configuring Data Protection for SAP HANA with the setup script" on page 21.

Results

Running the setup script on each database creates a dedicated profile `init`*`SID`*`.utl` file for each instance. Parameters are adapted to each database.

When you are deleting a database, all corresponding backup data that includes stored objects for the Tivoli Storage Manager node, is deleted.

Manually configuring the Tivoli Storage Manager for Data Protection for SAP HANA

If you chose to defer the configuration of the Tivoli Storage Manager during the setup script for Data Protection for SAP HANA, you must configure it manually. Configure the Tivoli Storage Manager by manually editing the Data Protection for SAP HANA init*SID*.utl file.

Before you begin

You must create the server stanza for Tivoli Storage Manager in the dsm.sys file before you proceed to the Data Protection for SAP HANA configuration tasks.

Procedure

1. After the setup.sh script finishes, open the Data Protection for SAP HANA init*SID*.utl file that is in the configuration directory.

   /usr/sap/*SID*/SYS/global/hdb/opt/hdbconfig/*SID*init.utl
2. With details listed in the dsm.sys file, update the following Tivoli Storage Manager parameters in the server section of the init*SID*.utl file.

   SERVER server # Servername Tivoli Storage Manager server name

   ADSMNODE NODE Tivoli Storage Manager node name

   BRBACKUPMGTCLASS MDB Management classes for database backups

   BRARCHIVEMGTCLASS MLOG Management classes for redo log backups
3. From the command line, enter the following command

   /opt/tivoli/tsm/tdp_hana/hdbbackint -p *full path to profile*/init*SID*.utl -f password
4. If the password for the Tivoli Storage Manager node is not set, you are prompted to enter the password.

Data Protection for SAP HANA profile file

The Data Protection for SAP HANA profile file, init*SID*.utl is automatically created when the setup script is run during the configuration process. The file is used for backup and restore operations.

The Data Protection for SAP HANA profile file is named init*SID*.utl, where the system identifier for the SAP HANA database instance is stored. During the configuration of Data Protection for SAP HANA, the profile file is created in the following directory, /usr/sap/*SID*/SYS/global/hdb/opt/hdbconfig. The SAP HANA configuration is adjusted to use the profile file for backup and recovery operations.

Configuring profile tasks

To configure the Data Protection for SAP profile file, you must set the server statement and in the Tivoli Storage Manager client options file.

Set the SERVER statement in the Data Protection for SAP profile:

The SERVER statement is specified in the Data Protection for SAP profile, and in the Tivoli Storage Manager client option file.

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 by using the logical server name. This logical server name is defined by the keywords SERVER or SERVERNAME.

Table 2. 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` do not specify
single path; automatic password by Tivoli Storage Manager	`SERVER` *server* `ADSMNODE` do not specify	`SERVERNAME` *server* `NODENAME` *node* `TCPSERVERADDRESS` *address*
several paths/servers; no password or manual password	`SERVER` *server 1* `ADSMNODE` *node 1* `SERVER` *server 1* `ADSMNODE` *node n*	`SERVERNAME` *server 1* `NODENAME` do not specify `TCPSERVERADDRESS` *address 1* `SERVERNAME` *server n* `NODENAME` do not specify `TCPSERVERADDRESS` *address n*
several paths/servers; automatic password by Tivoli Storage Manager[3]	`SERVER` *server 1* `ADSMNODE` do not specify `SERVER` *server n* `ADSMNODE` do not specify	`SERVERNAME` *server 1* `NODENAME` do not specify `TCPSERVERADDRESS` *address 1* `SERVERNAME` *server n* `NODENAME` do not specify `TCPSERVERADDRESS` *address n*
several paths/servers; automatic password by Tivoli Storage Manager [4]	`SERVER` *server* `ADSMNODE` do not specify `TCP_ADDRESS` *address 1* `SERVER` *server n* `ADSMNODE` do not specify `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 or Linux, `dsm.sys` is the single client option file for all Tivoli Storage Manager servers.

[3] If two different physical systems have the same Tivoli Storage Manager node name or if multiple paths are defined on one node by using several server stanzas, `passwordaccess generate` might work only 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. 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. The updated copy is generated by the first stanza that is used after password expiration. To avoid this situation, update the passwords before they expire. When the passwords are expired, run 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 valid entry.
3. Run `dsmc -servername=stanza2` and use the new password to generate a valid entry.

[4] You must use Tivoli Storage Manager API 5.5 (or later), you can use the **TCP_ADDRESS** parameter in the Data Protection for SAP profile. This parameter eliminates the requirement to set multiple stanzas in the Tivoli Storage Manager client option file for multiple paths. The parameter also eliminates the problem when it updates the password (see [3]).

Example of SERVER statement with alternate servers:

Data Protection for SAP profile is used in certain disaster recovery configurations.

This example assumes the following configuration for two servers `a` and `b`:

- 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 that is connected to this FDDI network.
- Normal backups are processed with server a, which is local to the SAP database server.
- A disaster recovery backup is stored on remote server b every Friday.

The following is an example of the Data Protection for SAP profile that is 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
```

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 processed through network path 1 (**SERVER** statement 1). If network path 1 is unavailable, the backup is processed by using network path 2 (**SERVER** statement 2). If path 1 is active, Data Protection for SAP 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 two sessions on path 2. Since **MAX_SESSIONS** specifies 2, the backup is processed by using path 2.

The Data Protection for SAP profile that is used in this alternate path configuration is shown in the following example:

```
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
```

Protection of a scale-out solution

When Data Protection for SAP HANA is used to protect a scale-out solution, backup and restore operations run simultaneously on all SAP HANA nodes. SAP requires that each SAP HANA node has access to all backups that are run by any SAP HANA node within the cluster.

All data must be stored on a single Tivoli Storage Manager server.

For manual password handling, all SAP HANA nodes must have identical Tivoli Storage Manager configurations. Ensure that the stanzas in the `dsm.sys` file that are referenced by the Tivoli Storage Manager for ERP profile are identical. When you are manually handling passwords, the parameter **nodename** in the `dsm.sys` file is commented out.

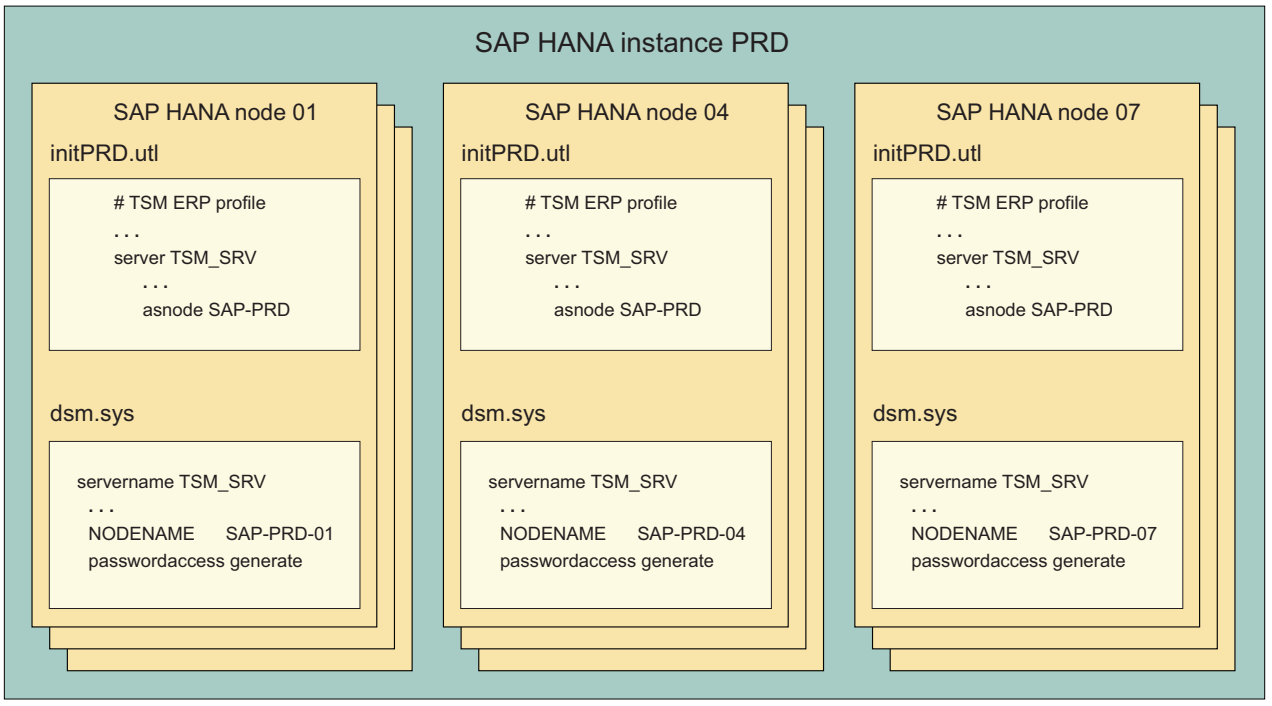

Figure 3. An example of a Data Protection for SAP HANA scale-out solution with manual password handling selected.

If automatic password handling is used, the stanzas in the `dsm.sys` file that are referenced by the Tivoli Storage Manager for ERP profile must have unique Tivoli Storage Manager node names for each SAP HANA node. Otherwise, the stanzas must be identical.

The nodes are used for authentication purposes. By using the Tivoli Storage Manager proxy node feature, the nodes gain access to a single Tivoli Storage Manager node that holds the data for the entire SAP HANA instance.

In the following example, the Tivoli Storage Manager node `SAPPRD` is used to store the backup of the entire SAP HANA database.

The Tivoli Storage Manager nodes `SAPPRD01`, `SAPPRD04`, and `SAPPRD07` are used by SAP HANA nodes `01`, `04`, and `07` to authenticate with the Tivoli Storage Manager server.

These Tivoli Storage Manager nodes must have proxy authority to theTivoli Storage Manager node `SAPPRD`. Similar configurations must be applied to all other nodes within the SAP HANA scale-out system.

The following example is provided for reference. In this sample scenario, there is a SAP HANA instance called PRD which is distributed over two HANA nodes *hana01* and *hana02*. There is a Tivoli Storage Manager server called *tsmsrv*. The `/opt/tivoli/tsm/client/api/bin64/dsm.sys` file on node *hana01* has an entry like the following sample:

```
SERVERNAME tsmsrv
   COMMMETHOD       TCPIP
   TCPPORT          1500
   TCPSERVERADDRESS tsmsrv.domain.local
   nodename         hana01
   passwordaccess   generate
```

While the /opt/tivoli/tsm/client/api/bin64/dsm.sys on node hana02 file would have an entry like the following sample:

```
SERVERNAME tsmsrv
   COMMMETHOD       TCPIP
   TCPPORT          1500
   TCPSERVERADDRESS tsmsrv.domain.local
   nodename         hana02
   passwordaccess   generate
```

The Data Protection for SAP HANA profile /usr/sap/PRD/SYS/global/hdb/opt/hdbconfig/initPRD.utl (which is located in the shared file system) would have the following server stanza (among other parameters):

```
SERVER     tsmsrv      # Servername, as defined in dsm.sys
 SESSIONS      2      # Maximum number of sessions to this server
 PASSWORDREQUIRED   NO     # Use a password
 ASNODE       hana_tpr  # Tivoli Storage Manager Nodename
 BRBACKUPMGTCLASS   mdbdisk1  # Mgmt-Classes for database backup
 BRARCHIVEMGTCLASS  mdbdisk1  # Mgmt-Classes for redo log backup
```

This example requires three Tivoli Storage Manager nodes. Nodes *hana01* and *hana02* are used for authentication only. There is no data stored on these nodes. Node *hana_tpr* is the node that is used to store all the data from the entire SAP HANA instance TPR.

Chapter 5. Protecting data

Information that is needed to back up, restore, and clone your SAP data is provided.

About this task

Backing up SAP data

Plan a daily backup strategy with scheduled and automated backups for the system.

About this task

Follow the tasks to put the backup strategy in place. Use the samples to help you for your operating system.

Backing up log and data files

During the setup procedure, the SAP HANA configuration is adjusted to use the BACKINT for SAP HANA window to back up the redo logs. SAP HANA regularly backs up the redo logs to the Tivoli Storage Manager.

About this task

The frequency of backups can be adjusted with the SAP HANA configuration parameter **log_backup_timeout_s** that is stored in the persistence section of the `global.ini` file. For more information about how to adjust this parameter, see the SAP HANA documentation.

Procedure

1. In the SAP HANA Studio, select the database instance to be backed up.
2. Right-click the database, and select **Back Up**. Alternatively, in the Specify Backup Settings window, select **Backint** for the destination type.
3. Optional: On the command line, use the SAP HANA SQL client `hdbsql`. For example, the following sample connects to the instance with number `53`, on host `vhana06`, as user *system* with password manager and runs a complete database backup:

   ```
   hdbsql -i 53 -n vhana06 -u system -p manager "backup data using backint (DAILY)"
   ```

Schedule automated backup tasks

Scheduling and automating backup and archive operations helps to ensure that data is backed up regularly at a specified time. Products that are used to schedule backup operations can be used to automate these 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**.

Crontab (Linux)
Automating backups at the database server level is available by using the `crontab` command.

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

IBM Tivoli Workload Scheduler
The IBM Tivoli Workload Scheduler provides event-driven automation, monitoring, and job control for both local and remote systems.

Sample backup strategy for daily backup processing

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

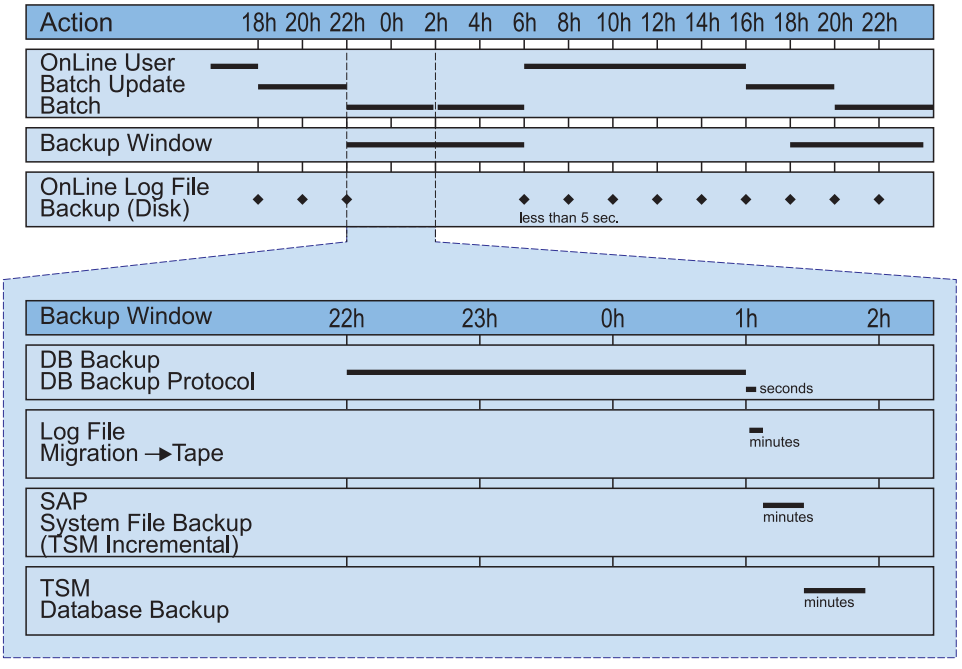

Figure 4. Production Backup Example

The automated backup example shown in the graphic displays these common tasks:

- A full database backup (offline or without application load) runs each night.
- The Tivoli Storage Manager server moves 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 backup must always be done.

Backups can be moved to disk storage and 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.

Creating multiple log file copies

Backing up multiple copies of a log file in a single archive operation helps protect against this data in the event of a storage hardware failure or disaster recovery situation. These copies can be on different physical Tivoli Storage Manager volumes or on different Tivoli Storage Manager servers.

When a log file copy is unavailable at restore time, the software switches to another copy, and continues to restore the log file from that copy. To create multiple backup copies of a log file, complete the following steps:

1. Open the Data Protection for SAP profile. The default directory and profile name is `/usr/sap/<SID>/SYS/global/hdb/opt/hdbconfig/init<SID>.utl`
2. Enter the keyword `REDOLOGS_COPIES`, and specify the number of backup copies that are required for the redo logs.
3. (Optional) Adjust the number in the `MAX_SESSIONS` keyword. This keyword specifies the maximum number of sessions that a single Data Protection for SAP instance has on the Tivoli Storage Manager server.
4. In the server stanza, search for the `BRARCHIVEMGTCLASS` keyword, and ensure that there are as many archive management classes specified as there are redo log copies that are requested.

If you distribute the redo log copies to more than one Tivoli Storage Manager server, the management classes for all server stanzas must be greater than or equal to the number of redo log copies. Data Protection for SAP requires that the maximum number of sessions, which are defined by `MAX_SESSIONS`, is greater than or equal to the number of redo log copies that are requested. A setup with a smaller number of sessions is not advised with the backint interface.

Restoring SAP data

Use the Data Protection for SAP file manager for managing restore operations.

Preparing to restore SAP HANA data

The actions that you take to restore SAP HANA data depend on the state of the SAP HANA application.

Before you begin

Determine the health of the SAP HANA application before you proceed to restoring your SAP HANA data. If you have an SAP HANA hardware failure that cannot be recovered with hardware components, you must restore SAP HANA with the application recovery DVD. Ensure that the SAP HANA operating system and software are running before you restore any SAP HANA databases.

When a logical error arises, SAP HANA data can be restored with Data Protection for SAP HANA.

Procedure

- Recover the SAP HANA configuration files from the Tivoli Storage Manager if they are backed up with the backup-archive client.
- Create the `dsm.sys` and `dsm.opt` configuration files from the backed up copies that are obtained from the Tivoli Storage Manager.

Restoring and recovering data

To restore or recover data with Data Protection for SAP HANA, start the SAP HANA Studio.

Procedure

1. In the SAP HANA Studio Navigator tab, right-click to select the instance. Select **Recovery**. If the database instance is still running, a message might be displayed. The message states that the system must be shut down. When you confirm the message, SAP HANA Studio automatically shuts down the instance.
2. In the Specify Recovery Type window, choose the type of recovery that you want to run.
3. In the Locate Log Backups panel the default location can be used.
4. Choose the backup to restore in the Select Data Backups panel. The list of backups is generated from the SAP HANA recovery catalog. This catalog contains every backup that is created. By selecting a specific backup and pressing **Check Availability** it can be verified that the backup exists on the Tivoli Storage Manager server.
5. Advance through the panels, and select the options that you require.

What to do next

For more information about restoring and recovering data, see the SAP HANA documentation.

Chapter 6. Tuning performance

Information needed to tune Data Protection for SAP performance is provided. A system is considered balanced when the threads on the disk and the network sides are similarly busy throughout the backup, and when resource usage is good. To improve overall throughput, consider adding more resources to create a balanced system.

About this task

In an optimum setup, a slight network bottleneck is preferred. Under certain conditions, the degree of imbalance cannot be determined from the graphical presentation. Depending on your system characteristics that include system buffering and buffer sizes, usage might reduce to almost zero in the graphical presentation although the system is 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. A balanced system, however, does not necessarily mean that the data throughput cannot be improved further. Adding new resources can improve the throughput rate.

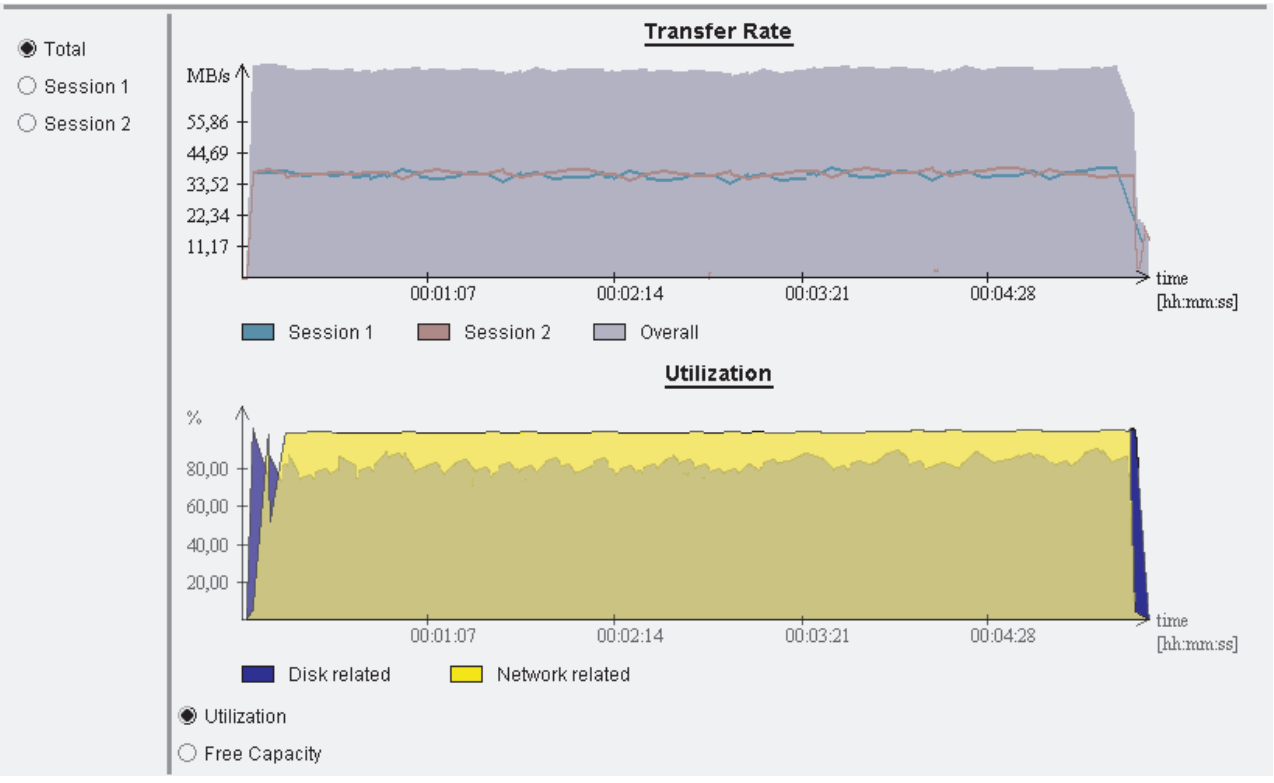

Figure 5. A balanced configuration

Procedure

- Maintain an optimum setup by ensuring tapes are maintained in streaming mode.

- Ensure that there is no network idle time, and that the network is at least as fast as the tape.
- Consider adding new resources to improve the throughput rate.

Options

Use Data Protection for SAP options to tune performance.

Performance tuning for Data Protection for SAP can be addressed by reviewing buffer size and copies, compression of backup data, and automation of backup options.

Buffer copies

You can change the Data Protection for SAP options to prevent copying data buffers, the original data buffers are sent between Tivoli Storage Manager components. This change can improve performance.

Data Protection for SAP uses internal buffers to store and exchange data with the Tivoli Storage Manager server. When data is sent from one component to another, data buffers are copied by default. Data Protection for SAP can prevent copying data buffers by sending the original data buffers. This process reduces the CPU load of the database server.

If client compression or client encryption are specified in the Tivoli Storage Manager options file (`dsm.sys` or `dsm.opt`, the original data buffers are sent.

Buffer size

Adjust buffer size disk I/O to improve transfer rates.

The internal data buffer size can be adjusted for Data Protection for SAP. These buffers are used for reading the disk and sending data to the Tivoli Storage Manager client API. The default values typically produce acceptable performance.

Optimize the buffer size for disk I/O to improve transfer rates. For disk subsystems, the best transfer rates are achieved when the buffer size is set equal to the stripe size. Before you increase the size of internal buffers, however, ensure that sufficient storage is available for the number of buffers that are specified by Data Protection for SAP. This number correlates to the number of sessions requested. The number of buffers doubles when compression is specified.

Compression of data for backup

You can adjust the amount of data that is being sent to the Tivoli Storage Manager server by compressing zero-byte blocks (RL_COMPRESSION profile keyword).

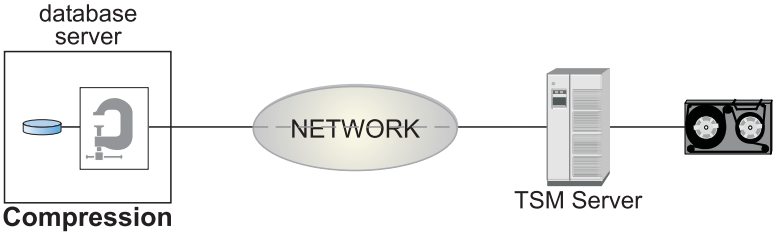

Figure 6. Null Block Compression

Data Protection for SAP can decrease the amount of data that is sent to the Tivoli Storage Manager server by compressing zero-byte blocks. Compression can increase the CPU load on the database server and can improve performance in situations when the network is at 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, in the *Profile parameter descriptions* topic, for details on how to activate Data Protection for SAP compression.

Automation options

Administrative productivity can be improved by using the Data Protection for SAP automation options.

Selectable management classes

Specify different Tivoli Storage Manager management classes for backup data and archive data. 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 be specified to use with multiple redo log files. For more information about specifying management classes, see the *Profile parameter descriptions* topic.

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 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. It continues restoring the log file from that copy. The description of the profile keyword REDOLOG_COPIES, in the *Profile parameter descriptions* topic, provides detailed information about creating and by 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 does 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 and 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 the *Profile parameter descriptions* topic.

Data transfer

When you use Data Protection for SAP, data is passed from disk through to the network and finally to tape. A balanced configuration can help to prevent bottlenecks and to ensure optimized performance.

Data throughput rate

Throughput rates differ for different environments because of different disk, network bandwidth, server systems, number of tapes, and configuration settings. When you are moving data, certain elements that are used in the movement of data can be tuned to improve data throughput.

Throughput rates differ widely among various environments because of different disk, network bandwidth, server systems, number of tapes, and configuration settings. The information that is provided here concentrates on selected elements that are involved in the movement of data. This information determines how to use existing resources to their maximum efficiency and provide insight as to how throughput can be improved.

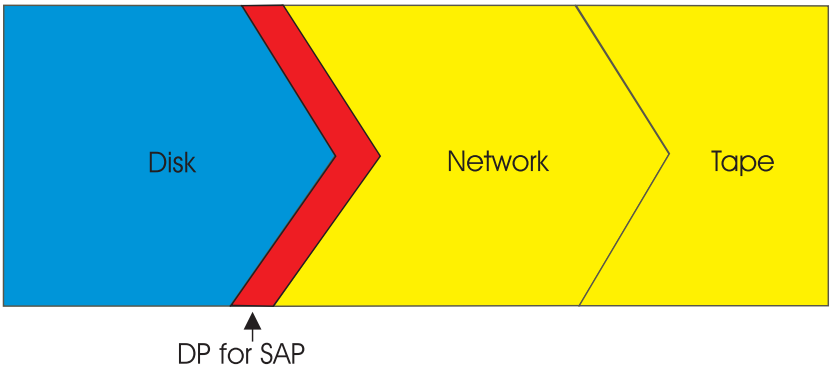

Figure 7. High-level view of the data flow during backup

From a high-level view, the data packages must send these elements when it does a backup with Data Protection for SAP: Data is read from disk that is 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. This situation 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 the logfile `backint.log` as the average transmission rate. However, identifying bottlenecks that are derived from log file messages is difficult. For this analysis effort, Data Protection for SAP provides performance sensors that indicate a bottleneck. These bottlenecks are located either in the elements that are

represented in blue (for disk) or in yellow (for network and tape respectively) in the graphic.

Performance tuning for data transfer

During data transfer, a continuous stream of data is generated between the SAP database server, the network, and the Tivoli Storage Manager server. The weakest component in this stream decreases the overall data transfer rate.

There are three main components that are involved during a Data Protection for SAP data transfer:

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

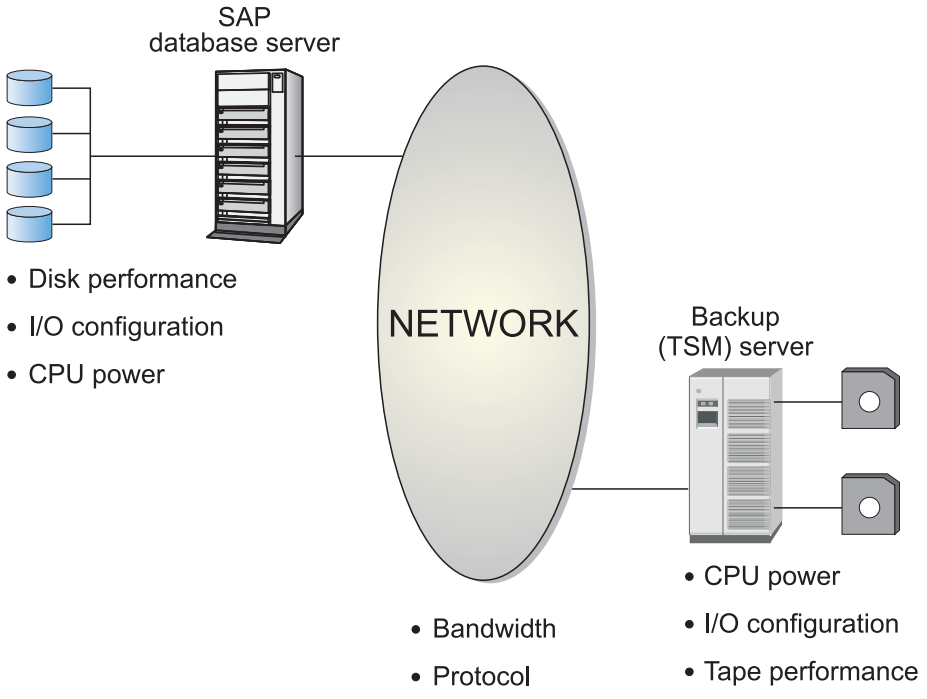

Figure 8. Data Protection for SAP data transfer

Chapter 7. Troubleshooting

To assist with troubleshooting and problem determination, diagnostic files and system information are displayed in a centralized location. Investigating the details in log files helps to troubleshoot problems.

About this task

Look for one of these patterns when a problem occurs:

- The problem always occurs at the same time. If this condition is true, view the appropriate log files to determine if scheduled processes are occurring simultaneously. Examples of such processes are virus checker, automatic updates, or batch jobs.
- The problem always occurs after another operation is done or the same operation is done.
- The problem occurs when another application or process is processed in parallel.

Investigate the log files for the SAP HANA application, the Tivoli Storage Manager server activity log, and the `backint` log files 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 condition is true, view the appropriate log files to determine whether any scheduled processes are occurring simultaneously. Examples of such processes are virus checker, automatic updates, or batch jobs.
- The problem always occurs after another operation is done or the same operation is done.
- The problem occurs when another application or process is processed in parallel.

Reproducing problems

Use the checklist to che\ck what caused the problem, and then attempt to reproduce the problem.

About this task

When you encounter a problem that occurs during an operation that previously ran successfully, review this list to determine the root cause of the problem.

- The setup has changed.
- A component changed such as the SAP HANA, Tivoli Storage Manager, operating system, network, or hardware components.
- Patches or updates to one or more of the components were applied.
- Changes occurred that originated from the system:
  - Check whether the disks are running full with the Linux `df` command.
  - If network performance decreases, check whether there are new hosts, or new applications. Check whether defects in software or hardware have occurred.
  - If Tivoli Storage Manager server processing decreases, check whether more clients or more operations were added. Information is also available in the Tivoli Storage Manager server activity log.

If none of these changes caused the problem, view the last modified time stamp of the following configuration files:

- `backint.log`
- `backup.log`
- `dsm.sys`
- `dsm.opt`

Run the following command to list all files in the `/etc` directory, that have been modified in the previous five days:

```
find /etc -type f -ctime 5 -print
```

If you can identify changes that are 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 supports both IPv4 and IPv6 for internal communication.

Data Protection for SAP runs in IPv4, IPv6, and mixed environments on Linux. 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 service listens for both IPv4 and IPv6 connection requests if the system is configured accordingly. Connection requests to ProLE are made for the addresses that are returned by the system for the respective port on the local host. Connection requests are made for the addresses that are 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 must be changed to *service or port*@*IP address* if the IP address is specified in the IPv6 notation. If 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.

Log files that contain information and messages

Data Protection for SAP processes are recorded in log files. Information about backup operations can be used to determine which backup should be used to restore your data.

Data Protection for SAP records data in two log files that can be used during troubleshooting:

`backup.log`

`backint.log`

Both log files can be opened through the SAP HANA Studio menu options **Open Perspectives** > **Administration Console**.

The backup.log log file records the start and finish of backup and restore operations. The success or failure of the operations is also recorded. All SAP HANA node details in a scale-out environment are stored in a single backup.log log file.

The backint.log log file contains the Tivoli Storage Manager for ERP data for all database and redo log file backup and restore operations that complete successfully or fail.

The Backup Object Manager writes to the backom.log log file.

Setup requirements

When you are troubleshooting issues while using Data Protection for SAP software there are items that you can check to ensure that the setup completed correctly.

Ensure that the Data Protection for SAP installation setup is correct by reviewing the following list:

- Make sure an entry similar to this example is defined in the /etc/inittab file:

  ```
  tdph:2345:respawn:/opt/tivoli/tsm/tdp_hana/prole -p tdphana
  ```

  The purpose of this entry is to start a daemon process for ProLE, and to verify that the process is running. This process listens to activity on the Data Protection for SAP port. The name of the port must match the name in the /etc/services file as shown in this example:

  ```
  tdphana 57321/tcp       #TDP for SAP HANA
  ```

  The lines are added to the /etc/services file during the installation process. If there are conflicts with existing entries, the port number must be changed to any unused number.
- Make sure that all the files are installed by running this command:

  ```
  rpm -V TIV-TSMERP-HANA
  ```

  If the command returns no output, all files are found. Otherwise, a list of missing files is returned in the output.
- On one host, make sure that the Data Protection for SAP profile init*SID*.utl and configuration file init*SID*.bki are in the /usr/sap/SID/SYS/global/hdb/opt/hdbconfig/ directory.
- Make sure that /usr/sap/SID/SYS/global/hdb/opt/hdbbackint link exists and points to /opt/tivoli/tsm/tdp_hana/hdbbackint.

The names of the Tivoli Storage Manager servers that are 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 are installed into their default locations, then it is not necessary to set the DSMI_* variables. If the variables are set, however, make sure that 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 that are specified by these variables. Also, verify that write permissions exist for the init*SID*.bki file as this file is the only one to which Data Protection for SAP HANA writes persistent information.

Information to collect for support

When you contact support, you must be able to provide the following information.

- The Data Protection for SAP version level.
- The operating system level and patches that were applied.
- The SAP HANA version level.
- The Tivoli Storage Manager server version.
- The Tivoli Storage Manager server operating system level.
- Data Protection for SAP configuration file `init`*`SID`*`.utl` including Tivoli Storage Manager client configuration files (`dsm.sys`, `dsm.opt`)
- Data Protection for SAP profile (`initSID.utl`)
- The change history of the system components (if the process worked previously).

More information might also be requested from the service representative.

Chapter 8. Reference information

Reference information, such as versioning and profile information, is provided.

Version numbers

When Tivoli Storage Manager for ERP backup version control is active, version information is stored on the Tivoli Storage Manager server.

The number of Tivoli Storage Manager for ERP backup versions is defined by the MAX_VERSIONS keyword. The version number is increased only after successful backups.

Manage 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.

BRARCHIVE might process redo logs while a database backup is still processing or several BRARCHIVE processes might run simultaneously. These combined sessions might exceed the number of available tape drives. To avoid this situation, save redo logs to disk storage pools and then move them to tape storage.

Crontab file sample

The following sample output, shows the root crontab jobs.

```
# crontab.sample:
# Sample crontab file to be included in the root crontab jobs.
# ------------------------------------------------------------------------
# Task:
# Submits backup commands at regularly scheduled intervals
# using the SAP HANA command line interface hdbsql.
# ------------------------------------------------------------------------
#      *****     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 SAP HANA database
# is assumed to be 'GT3', the instance number 0 and the username
# of the database instance owner 'gt3adm'.
#
# ------------------------------------------------------------------------
# Full database backup, scheduled every weekday at 8:00 p.m. providing the
# database user name (system) and password (manager) on command line
#
0  20  * * 1-5
/bin/su - gt3adm -c "hdbsql -i 0 -u system -p manager
\"backup data using backint
('/usr/sap/GT3/SYS/global/hdb/backint/full_weekday_$(date +'m%d')')\""
#
#
# Full database backup, scheduled at weekends at 9:00 a.m. using a key
# store entry named TSM_BACKUP to connect to the database (please refer
# to the SAP HANA administration guide for details how to create a key
# store entry)
#
0  9  * * 0,6
/bin/su - gt3adm -c "hdbsql -i 0 -U TSM_BACKUP
\"backup data using backint
('/usr/sap/GT3/SYS/global/hdb/backint/full_weekend_$(date +'m%d')')\""
#
```

Data Protection for SAP profile

The Data Protection for SAP profile provides keyword parameters that customize how Data Protection for SAP operates. A sample profile `initSID.utl` is provided on the product media.

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 number sign (#). Scanning of the current line stops when the comment symbol is encountered. No comment is allowed between the keyword and the value or values. 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.

Profile parameter descriptions

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

ADSMNODE

Specifies a node name that is registered to the Tivoli Storage Manager server as a Tivoli Storage Manager node. This parameter must be defined 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 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 is to be set for manual password-handling.

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___`.

BRARCHIVEMGTCLASS *management_class [management_class...]*

Specifies the Tivoli Storage Manager management classes 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 that are specified must be greater than or equal to the number of redo log copies . This parameter must be defined with the respective SERVER statement, as shown in the sample profile.

To use different Tivoli Storage Manager servers for backup and archive data, the value ":SKIP:" can be used to define a server stanza with no archive management classes. This value is allowed for the parameter **BRARCHIVEMGTCLASS** only.

BRBACKUPMGTCLASS *management_class [management_class...]*

Specifies the Tivoli Storage Manager management classes that Tivoli Storage Manager for ERP uses 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 option 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**

The size of the buffers that are sent to the Tivoli Storage Manager API is the value of **BUFFSIZE** increased by approximately 20 bytes. The valid range is 4096 (4 KB) - 32 MB. Inappropriate values are adjusted automatically. If **BUFFCOPY** is set to `PREVENT`, the value of **BUFFSIZE** must not exceed 896 KB.

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.

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, front-end processing is not done.

Example for UNIX or Linux:

```
FRONTEND write operator@remotesite Backup of SAP database
object is starting.
```

This process sends a message to a remote user before backup begins.

LOG_SERVER *servername [verbosity]*
The *servername* value specifies the name of the Tivoli Storage Manager server to which log messages are sent. The *servername* must match one of the servers that are 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 parallel Tivoli Storage Manager client sessions that Tivoli Storage Manager for ERP establishes for backup, archive redo logs and restore. 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.

PASSWORDREQUIRED NO | **YES**
Specifies whether Tivoli Storage Manager requires a password to be supplied by the Tivoli Storage Manager client. This situation 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 with the respective **SERVER** statement, as shown in the sample profile.

REDOLOG_COPIES *n* | **1**
Specifies the number of copies Tivoli Storage Manager for ERP stores for each processed redo log file. The valid range is 1 - 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 that are specified must be greater than or equal to the number of log file copies specified.

RL_COMPRESSION NO | YES
If set to YES, Tivoli Storage Manager for ERP runs a null block compression of the data before they are sent over the network. Although RL compression introduces more CPU load, throughput can be improved when the network is the bottleneck. It is not advised 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 done. **RL_COMPRESSION** is only run if a full database backup was started. The offline log files 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 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 profile dsm.sys. 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 profiles specify the same TCP/IP address for these server names. 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 profiles specify different TCP/IP addresses for the different servers.

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 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 that are separated by spaces. A trace is written only 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 used. For example, specifying `/tmp/%BID.trace` yields a trace file `/tmp/myBackup.trace` for backup ID myBackup. A trace is written only if both **TRACE** and **TRACEFILE** are specified.

TRACEMAX *n*
Specifies the maximum size of the trace file in KB. The valid range is 4096 (4 MB) - 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 in the range 0 (Sunday) - 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.

Sample profile file for UNIX or Linux

A sample profile file (`initSID.utl`) is included in the Tivoli Storage Manager for ERP installation package.

```
#----------------------------------------------------------------------
#
# IBM Tivoli Storage Manager for Enterprise Resource Planning
#
# Data Protection for SAP HANA (R)
#
# Sample profile for Data Protection for SAP HANA (R)
#
#----------------------------------------------------------------------
#
# See the 'Data Protection for SAP HANA (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 HANA (R) accesses its profile
# in "read only" mode. All variable parameters will be written into the file
# specified with the CONFIG_FILE parameter. The passwords will be encrypted.


#----------------------------------------------------------------------
# Prefix of the 'Backup ID' which will be stored in the description field
# of the Tivoli Storage Manager archive function.
# If this parameter is not specified then the SID of the SAP HANA (R)
# instance will be used to prefix the backup ID by default. The value of
# this parameter does overrule the default behaviour.
# Must be exactly 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 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 HANA (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)
#-------------------------------------------------------------------------
BUFFSIZE  131072             # block size in bytes


#-------------------------------------------------------------------------
# This optional parameter controls how Data Protection for SAP (R) HANA 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


#-------------------------------------------------------------------------
# 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) HANA 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. Furthermore the current hostname, a time stamp and
# the process name will be appended. (.../backup_%BID.trace changes to
# .../backup_SAP___9809182300.trace.mizar.20130731134735.4226.backint).
# Default: none.
#-------------------------------------------------------------------------
#TRACEFILE  /usr/sap/SID/home/backup.trace
#TRACEFILE  /usr/sap/SID/home/backup_%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  CONFIGDIR/initSID.bki


#-------------------------------------------------------------------------
# Denotes if Data Protection for SAP (R) HANA shall send error/status
# information to a Tivoli Storage Manager server.
# The servername must match one of the servers listed in a SERVER statement.
# Valid values for verbosity are ERROR | WARNING | DETAIL.
# Default: none.
#-------------------------------------------------------------------------
#LOG_SERVER             servername    [verbosity]
#LOG_SERVER             server_a      ERROR


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


#*************************************************************************
# Statement for servers and paths.
# Multiple servers may be defined.
#*************************************************************************


SERVER         SED_SERVER                  # Servername, as defined in dsm.sys
  SESSIONS            2                  # Maximum number of sessions to this server
  PASSWORDREQUIRED    YES                # Use a password
  ADSMNODE            SED_NODE              # Tivoli Storage Manager Nodename
  ASNODE              SED_ASNODE            # Tivoli Storage Manager Nodename
  BRBACKUPMGTCLASS    SED_MDB               # Mgmt-Classes for database backup
  BRARCHIVEMGTCLASS   SED_MLOG              # 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
```

Locating sample files

Use the file samples to assist you with Data Protection for SAP operations.

Procedure

- Review the out put samples for `dsm.opt`, the include/exclude statement, and `dsm.sys`.
- Use the planning sheet to help you plan the installation parameters for Data Protection for SAP.

Client system options file sample (`dsm.sys`)

The system options file lists information that includes the **`buffersize`** and compression status. The following sample shows the typical output.

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


SErvername  server_a
  COMMmethod         TCPip
  TCPPort            1500
  TCPServeraddress   192.168.1.1
  TCPBuffsize        32
  TCPWindowsize      24
  Compression        Yes
  InclExcl           /opt/tivoli/tsm/client/ba/bin/hana_inclexcl.list

SErvername  server_b
  COMMmethod         TCPip
  TCPPort            1500
  TCPServeraddress   192.168.1.2
  TCPBuffsize        32
  TCPWindowsize      24
  Compression        Yes
  InclExcl           /opt/tivoli/tsm/client/ba/bin/hana_inclexcl.list
```

Include and exclude list sample (UNIX, Linux)

The include and exclude list shows the files and directories that are included or excluded for backup operations.

```
*
* Sample include/exclude list for SAP HANA appliances
*
* first exclude everything
exclude /.../*
*
* now include relevant files and directories only
include /usr/sap/C21/SYS/profile/.../*
include /usr/sap/C21/SYS/global/hdb/custom/config/.../*
```

Client user options file sample (UNIX, Linux)

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

SErvername       server_a
```

Planning sheet for the base product

Use the planning sheet to assist you when you are installing and configuring Data Protection for SAP.

Collect the information in this planning sheet before you install Data Protection for SAP.

Table 3. Installation parameters for Data Protection for SAP

Linux	Installation parameter
X	Database SID.
X	Database instance number.
X	Password of database user SYSTEM.
X	Tivoli Storage Manager server name or IP address.
X	Tivoli Storage Manager node name: Tivoli Storage Manager node that is configured on the Tivoli Storage Manager server that is named for the backup of the SID previously listed. In a scale-out environment, there can be multiple Tivoli Storage Manager node names required.
X	Tivoli Storage Manager management classes for database and redo log backups. Management classes that are configured for the database backup and for the backup of redo logs.
X	Path where the Tivoli Storage Manager API are in (contents of environment variable **DSMI_DIR**): Default: `C:\Program Files\Common Files\tivoli\TSM\api64`
X	Path to client option file of Tivoli Storage Manager (contents of environment variable **DSMI_CONFIG**).
X	Path to Tivoli Storage Manager log files (contents of environment variable **DSMI_LOG**): The Tivoli Storage Manager API creates the file `dsierror.log` in this path. Default: `C:\temp`

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 IBM Tivoli Storage Manager family of products includes the following accessibility features:

- Keyboard-only operation using standard operating-system conventions
- Interfaces that support assistive technology such as screen readers

The command-line interfaces of all products in the product family are accessible.

Tivoli Storage Manager Operations Center provides the following additional accessibility features when you use it with a Mozilla Firefox browser on a Microsoft Windows system:

- Screen magnifiers and content zooming
- High contrast mode

The Operations Center and the Tivoli Storage Manager server can be installed in console mode, which is accessible.

The Operations Center help system is enabled for accessibility. For more information, click the question mark icon on the help system menu bar.

Vendor software

The Tivoli Storage Manager product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

IBM and accessibility

See the IBM Human Ability and Accessibility Center (http://www.ibm.com/able) for information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
United States of America

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.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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
US

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 document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

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:

```
Portions of this code are derived from IBM® Corp. Sample Programs.
© Copyright IBM® Corp. _enter the year or years_. All rights reserved.
```

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 www.ibm.com/legal/copytrade.shtml.

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

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights Except as expressly granted in this permission, no other permissions,

licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://www.ibm.com/privacy/details in the section entitled “Cookies, Web Beacons and Other Technologies,” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at http://www.ibm.com/software/info/product-privacy.

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, MD-NC119
Armonk, NY 10504-1785
United States of America

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.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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
US

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 document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

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:

```
Portions of this code are derived from IBM® Corp. Sample Programs.
© Copyright IBM® Corp. _enter the year or years_. All rights reserved.
```

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 www.ibm.com/legal/copytrade.shtml.

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

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at http://www.ibm.com/privacy and IBM’s Online Privacy Statement at http://www.ibm.com/privacy/details in the section entitled “Cookies, Web Beacons and Other Technologies,” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at http://www.ibm.com/software/info/product-privacy.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at http://www.ibm.com/privacy and

IBM's Online Privacy Statement at http://www.ibm.com/privacy/details in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Glossary

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

See Tivoli Storage Manager glossary (http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.3/tsm/glossary.html).

To view glossaries for other IBM products, see http://www.ibm.com/software/globalization/terminology/.

Index

Product Number: 5608-E05

Printed in USA