

IBM Storage Protect for Enterprise Resource Planning: Data Protection for SAP for DB2

8.2.0



Contents

List of Tables	5
Who should read this guide	8
Publications	8
Getting started	11
Integration between SAP and DB2®	11
DB2® command line processor	11
DB2® Backup Object Manager utility	13
DB2® log manager	14
Backup objects and types of failures	15
Planning	17
Database server considerations	17
Network performance optimization	17
Backup server optimization	18
Store data on the server	18
Parallel backup paths and backup servers	19
Archive inactive data	20
Restore versus backup	20
Create multiple redo log copies	20
Planning for using IBM® HACMP™ for AIX®	21
HACMP™ impact	21
Digital signing of executable files (Windows™)	22
Installing	24
Preparing to install	24
Prerequisites	24
Installing in silent mode	25
Installing in a UNIX™ or Linux™ environment	26
Uninstalling older versions (UNIX™ and Linux™)	27
Installing in a Windows™ environment	28
Enabling ProLE to access configuration files on a remote share (Windows™)	29
Uninstalling older versions (Windows™)	29
Verify the installation or upgrade	30
Upgrading	31
Upgrading the base product	31
Migrate the Data Protection for SAP profile	32
Configuring	33
Changing configuration tasks for the Data Protection for SAP base product	33
Configuring profile tasks	33
DB2® tasks	36
Configuring distributed file system tasks	40
Configuring as an HACMP™ application	40
Configuring IBM Storage Protect™	42
IBM Storage Protect™ client tasks	42
IBM Storage Protect™ server tasks	45
Protecting data	52
Backing up SAP data	52
Schedule automated backup tasks	52
Windows™ scheduling example	53
Backups in a nonpartitioned database environment	54
Use DB2® single system view for backup	54
Creating multiple log file copies	54
Schedule batch sample	55
Full offline backup batch file sample	55
Full offline backup shell script sample	56
Segment large backup objects	56
Restoring SAP data	58
Start restores in a nonpartitioned database environment	58
Processing redirected restore in automatic mode	58
Tablespace definition information	59
Processing redirected restore in batch mode	61
Redirected restore in interactive mode	62

Sample work flow for redirected restore	63
Redirected restore plausibility checks	65
DB2® redirected restore using Backup Object Manager	65
Creating table space definition information	66
Redirected restore prerequisites	67
Tablespace definition information	67
Tuning performance	69
Server-related tuning	69
Manage data on the backup server	69
Alternate network paths and servers	70
Options	70
Performance options for Data Protection for SAP	70
Buffer copies	71
Buffer size	71
Compression of data for backup	71
Optimize for data deduplication in DB2	71
Automation options	72
Data transfer	72
Data throughput rate	73
Performance sensors	73
Performance tuning for data transfer	74
Multiple servers	74
Multiple sessions	75
Multiplexing	75
Multiple network paths	75
Storage space	76
Automated tablespace adaptations	76
Tablespace normalization	76
Scaling tablespace containers	77
Troubleshooting	79
Troubleshooting common problems	79
Reproducing problems	79
Internet Protocol version 6 (IPv6) support	80
Setup requirements	80
Information to collect for support	81
Log files that contain information and messages	82
Troubleshooting problems	82
Location of log files	84
DB2® vendor reason codes	84
Reference information	86
Backups and restores in partitioned database environments	86
Backup Object Manager	87
Backup Object Manager commands	88
Backup Object Manager command options	88
Backup command	90
Delete command	90
Password command	91
Query command	91
Restore commands	91
BACKOM command examples	92
Crontab example	93
Crontab file sample	93
Data Protection for SAP profile	94
Profile parameter descriptions	94
Sample profile file for UNIX™ or Linux™	98
Sample profile (Windows™)	100
Locating sample files	103
Client user options file sample (UNIX™, Linux™)	103
Client user options file sample (Windows™)	103
Client system options file sample (dsm.sys)	104
Include and exclude list sample (UNIX™, Linux™)	104
Include/exclude list sample (Windows™)	105
Client options files sample	106
Vendor environment file sample	106
Planning sheet for the base product	106
Network settings for IBM Storage Protect™	107
Networks with large bandwidth delay	108
SP switch (RISC 6000)	108
Accessibility features for the IBM Storage® Protect product family	109
Overview	109
Keyboard navigation	109

Interface information	109
Vendor software	109
Related accessibility information	109
Notices	110
Trademarks	111
Terms and conditions for product documentation	111
Privacy policy considerations	112
Glossary	113
Index	114

List of Tables

Table 1	9
Table 2: File Extensions for Shared Libraries	26
Table 3: SERVER statement and appropriate profile and option file settings	33
Table 4: Configuration parameters for DB2® database backup and restore, and log archive and retrieve.....	38
Table 5: Password handling for UNIX™ or Linux™	49
Table 6: Password handling for Windows™	50
Table 7: DB2® vendor reason codes	84
Table 8: Installation parameters for Data Protection for SAP	107
Table 9: Tuning IBM Storage Protect™ configuration file attributes.....	107
Table 10: Tuning of network settings	108
Table 11: Tuning of SP switch buffer pools	108

Note:

Before you use this information and the product it supports, read the information in “Notices” on page 110.

Second edition (5th December 2025)

This edition applies to version 8, release 2 of IBM Storage Protect™ for Enterprise Resource Planning (product number 5725-X03), available as a licensed program. It also applies to all subsequent releases and modifications until otherwise indicated in new editions.

About this publication

This publication documents how to use IBM Storage Protect™ 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 IBM Storage Protect™ server.

Who should read this guide

This guide is intended for system programmers and administrators who are responsible for implementing a backup solution in an SAP environment using the IBM Storage Protect™. It describes the procedures needed to install and customize Data Protection for SAP, the interface between SAP and the IBM Storage Protect™. The reader should be familiar with the documentation for SAP and IBM Storage Protect™.

Publications

The IBM Storage® Protect product family includes IBM Storage® Protect Plus, IBM Storage® Protect for Virtual Environments, IBM Storage® Protect for Databases, and several other storage management products from IBM®.

To view IBM® product documentation, see [IBM® Knowledge Center](#).

What's new for IBM Storage Protect™ for Enterprise Resource Planning

The update for IBM Storage Protect™ for Enterprise Resource Planning 8.2.0 is listed in the table. Review the release notes before you install the product.

New and changed information in this product documentation is indicated by a vertical bar (|) to the left of the change.

Release	New features and updates
8.2.0	<p>IBM Storage Protect for Enterprise Resource Planning now supports RHEL 9.X version</p> <p>IBM Storage Protect for Enterprise Resource Planning products including Oracle, DB2 and SAP HANA now supports Red Hat Enterprise Linux 9.x (RHEL). This currency upgrade helps you migrate and run their Oracle, DB2 and SAPHANA workloads on RHEL 9.x with improved compatibility and modernized infrastructure.</p> <p>SAPHANA version support</p> <p>IBM Storage Protect for Enterprise Resource Planning now supports SAP HANA version 2.00.077.00.1713529394 on Red Hat Enterprise Linux (RHEL) 9.x. This currency upgrade enables migration and operation of SAP HANA workloads on RHEL 9.x, improving compatibility and supporting infrastructure modernization.</p> <p>JAVA 21 Upgrade</p> <p>IBM Storage Protect for Enterprise Resource Planning including SAP HANA, Oracle, and DB2 previously used Java 8 and 17 for packaging. These versions introduced security vulnerabilities that affected product integrity. To address this, all products now use Java 21. This upgrade strengthens security and reduces vulnerability-related risks.</p> <p>IBM Storage Protect for Enterprise Resource Planning Rebranding</p> <p>In 8.2.0, the product name has changed from IBM Spectrum Protect for Enterprise Resource Planning to IBM Storage Protect for Enterprise Resource Planning.</p> <p>SAPHANA certification</p> <p>IBM Storage Protect for Enterprise Resource Planning is certified in collaboration with the external SAP team and holds the "SAP HANA Integration Certification". The certification validates interoperability through the following components:</p> <ul style="list-style-type: none"> • Backint SDK for SAP HANA • Backint Certification Test Suite • Backint API Version 1.0 • Backint API Version 1.5 <p>For details, refer to the official SAP Certified Solutions Directory.</p>
8.1.11	<p>Documentation updates</p> <p>The IBM Storage Protect™ for Enterprise Resource Planning Knowledge Center and User's Guides have been updated with entries from the 8.1 Documentation updates technote since the last full Knowledge Center update for 8.1.4.</p>
8.1.9	<p>Secure connection with TSL/SSL</p> <p>In 8.1.9, you can connect to your SAP HANA databases by using Transport Layer Security (TLS) or Secure Sockets Layer (SSL) connections.</p>
8.1.6	<p>Backup enhancement: backup size</p> <p>The accuracy of the estimated backup size that is sent to the IBM Storage Protect server has been improved. Instead of a fixed value that depended on the allocated memory size, the exact value of the backup is used.</p>

Release	New features and updates
8.1.4	<p>Take advantage of potential performance improvements enabled by parallel processing</p> <p>IBM Storage Protect™ for Enterprise Resource Planning Version 8.1.4 can process multiple redo log copies in dedicated threads when used with SAP HANA, IBM Db2®, or Oracle RMAN database technology. If the system offers sufficient CPU power and bandwidth, these redo logs are sent to IBM Storage Protect™ servers in parallel, which can improve system performance.</p> <p>Optimize the format of backup images for data deduplication</p> <p>You can optimize the format of backup images for data deduplication with Db2 by running the Backup Object Manager command, backom, with the command option -o. The new format is compressed and thus requires less storage space. For more information, see Optimize for data deduplication in Db2.</p>
8.1.1	The V8.1.1 release resolved defects, but did not introduce major new features.
8.1.0	<p>New product name</p> <p>IBM® Tivoli® Storage Manager for Enterprise Resource Planning is renamed to IBM Storage Protect™ for Enterprise Resource Planning in V8.1.0.</p>

Getting started

Data Protection for SAP and IBM Storage Protect™ provide a reliable, high performance, and production-oriented solution that enables back up and restore of SAP systems.

Data Protection for SAP is integrated with DB2® backup and recovery facilities and applies SAP backup and recovery procedures. Data Protection for SAP is optimized for SAP databases and therefore provides efficient management of large data volumes.

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

Other files, such as SAP and DB2® executable files, can be backed up using the IBM Storage Protect™ backup-archive client. This action is important for disaster recovery purposes, as all SAP and DB2® executable files must be available before you use Data Protection for SAP to restore and recover the database.

Integration between SAP and DB2®

Data Protection for SAP for DB2® operates as an unseen link between DB2® and the IBM Storage Protect™. A shared library is dynamically linked by DB2® backup/archive processes.

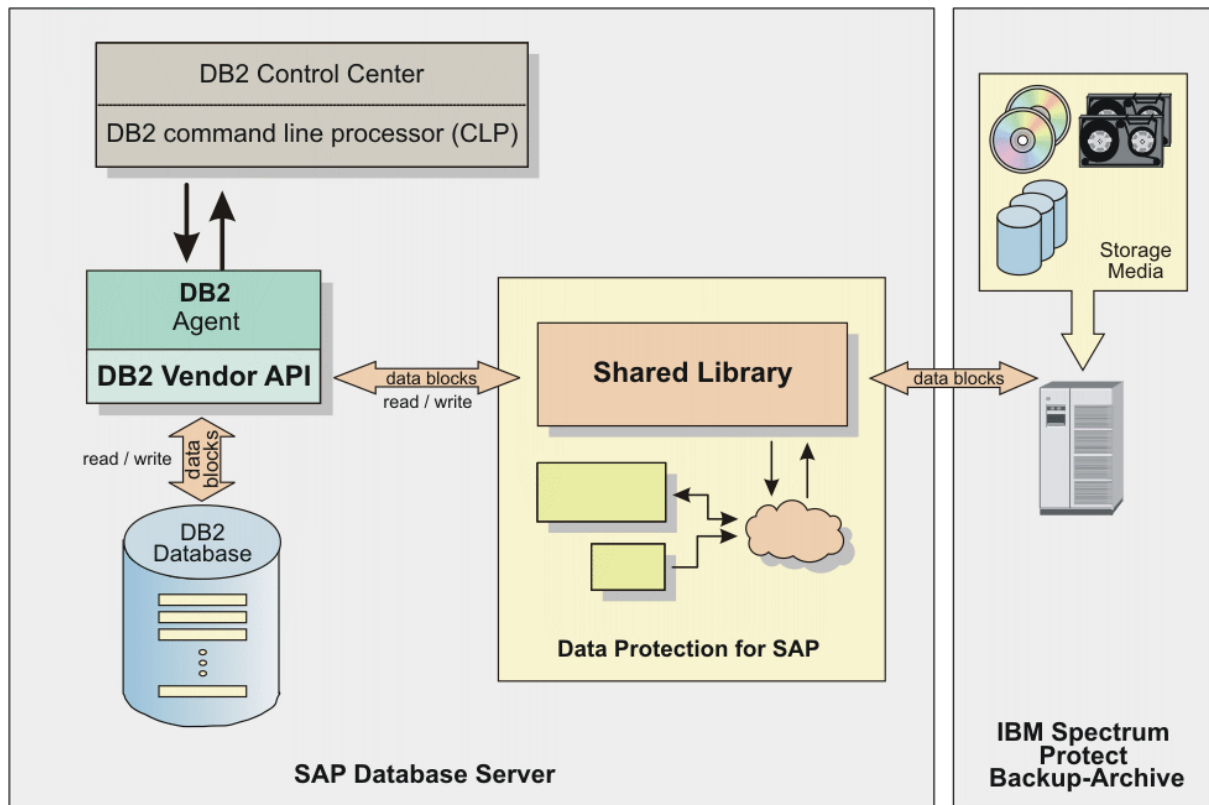


Figure 1: Integration of Data Protection for SAP with DB2®

DB2® command line processor

The DB2® Command Line Processor (CLP) interprets commands for the DB2® database and passes control to a DB2® Server Process.

For Data Protection for SAP for DB2®, the `LOADLIBRARYNAME` option instructs DB2® to start the Data Protection for SAP shared library. This process starts the backup or restore operation, dynamically loads the library, and communicates with Data Protection for SAP through the Vendor API.

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

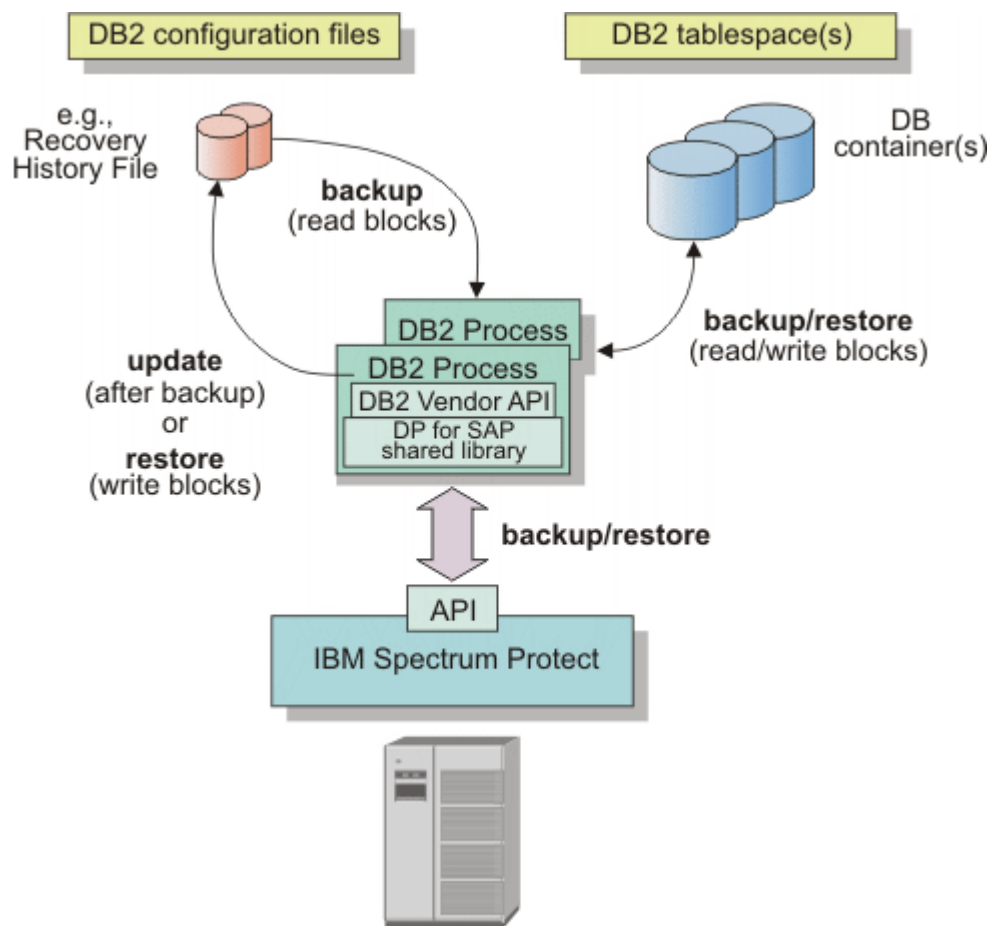


Figure 2: DB2® backup architecture

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

- Creates a unique timestamp to identify the backup.
- Loads Data Protection for SAP dynamically as a shared library.
- Reads the data from the database containers.
- Reads the DB2® configuration files.
- Creates data blocks that contain the backup image and passes these blocks to the data mover part of Data Protection for SAP.

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

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

- Loads Data Protection for SAP dynamically as a shared library.
- Requests the backup data from the shared library.

The Data Protection for SAP shared library:

- Checks with the IBM Storage Protect™ if the backup image is available.
- Retrieves the data blocks from IBM Storage Protect™.
- Passes the data blocks to the DB2® Server Process.

The DB2® Server Process

- Restores the DB2® data to the database containers.
- Logs the restore in the Recovery History File.

DB2® Backup Object Manager utility

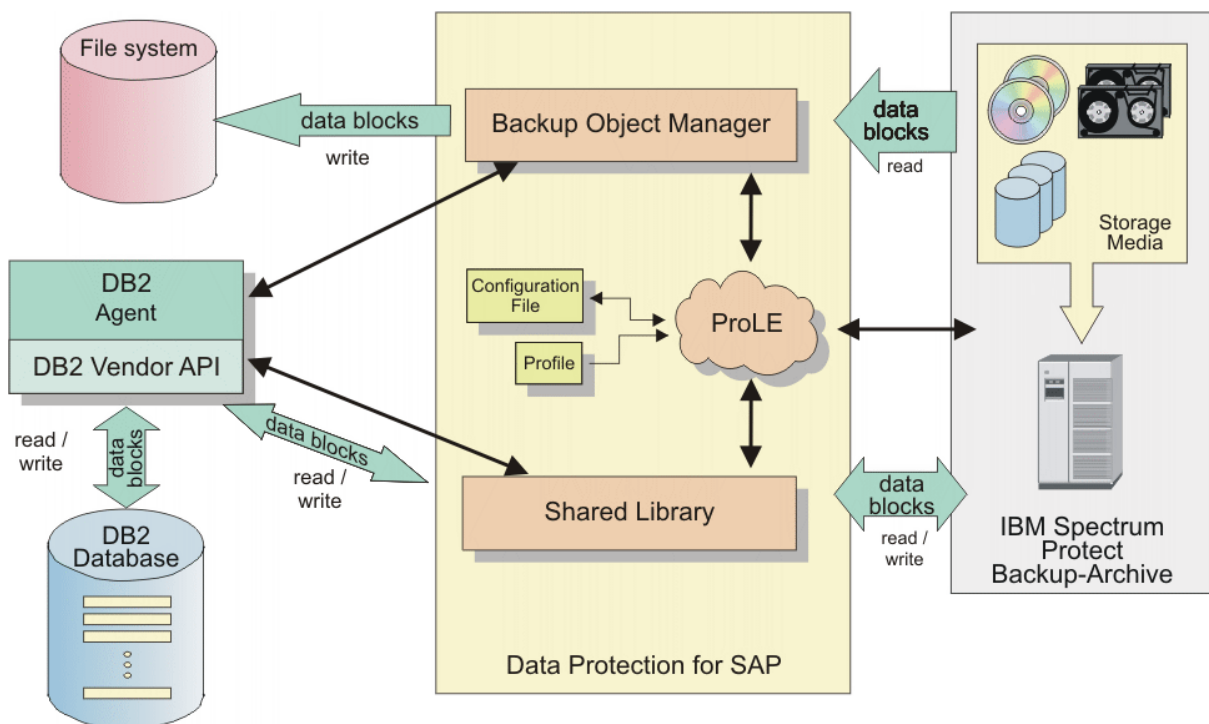
Backup objects, such as database or table space backups and DB2® log files, can be managed with the Data Protection for SAP for DB2® Backup Object Manager. Information about Backup Object Manager commands and options is provided.

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

- Verify and store a IBM Storage Protect™ password.
- Find backup objects in IBM Storage Protect™.
- Check the properties of the backup objects in IBM Storage Protect™.
- Remove any backup object from IBM Storage Protect™.
- Backup database and selected table spaces.
- Restore database and table space backups to the corresponding database.
- Retrieve files from IBM Storage Protect™ and restore them to the file system.
- Perform a redirected restore of databases (cloning).

The Backup Object Manager is designed to handle DB2® log files that are archived with Data Protection for SAP, the SAP tool BRARCHIVE, and those files that are archived with Data Protection for SAP and the DB2® Log Manager. No special Backup Object Manager customization or configuration is necessary.

This graphic displays how the Backup Object Manager interacts with the IBM Storage Protect™ server and the SAP database server:



SAP Database Server

Figure 3: Data Protection for SAP Backup Object Manager

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

- Requests to verify the IBM Storage Protect™ password are passed directly to IBM Storage Protect™.
- Requests to display or delete any data are answered by accessing the IBM Storage Protect™ server directly, thus working with the data that is available on IBM Storage Protect™.

- Requests to restore DB2® log files and raw files are also processed by using the IBM Storage Protect™ client.
- Requests to back up or restore any DB2® database data are routed to the DB2® agent. The DB2® agent employs the Data Protection for SAP shared library.

The Backup Object Manager is available for use upon successful installation and setup of Data Protection for SAP. The Backup Object Manager uses the settings in the Data Protection for SAP profile, configuration file, and the settings of the **XINT_PROFILE**, **TDP_DIR**, and **DB2_VENDOR_LIB** environment variables. As a result, extra installation and setup steps are not required.

DB2® log manager

Data Protection for SAP for DB2® is integrated with the built-in DB2® log manager.

When Data Protection for SAP is registered within the DB2® database configuration, the DB2® log manager uses Data Protection for SAP for archiving and retrieving log files.

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

Online active

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

Online retained

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

Offline retained

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

Archived

Filled or closed log files that were archived to IBM Storage Protect™ storage.

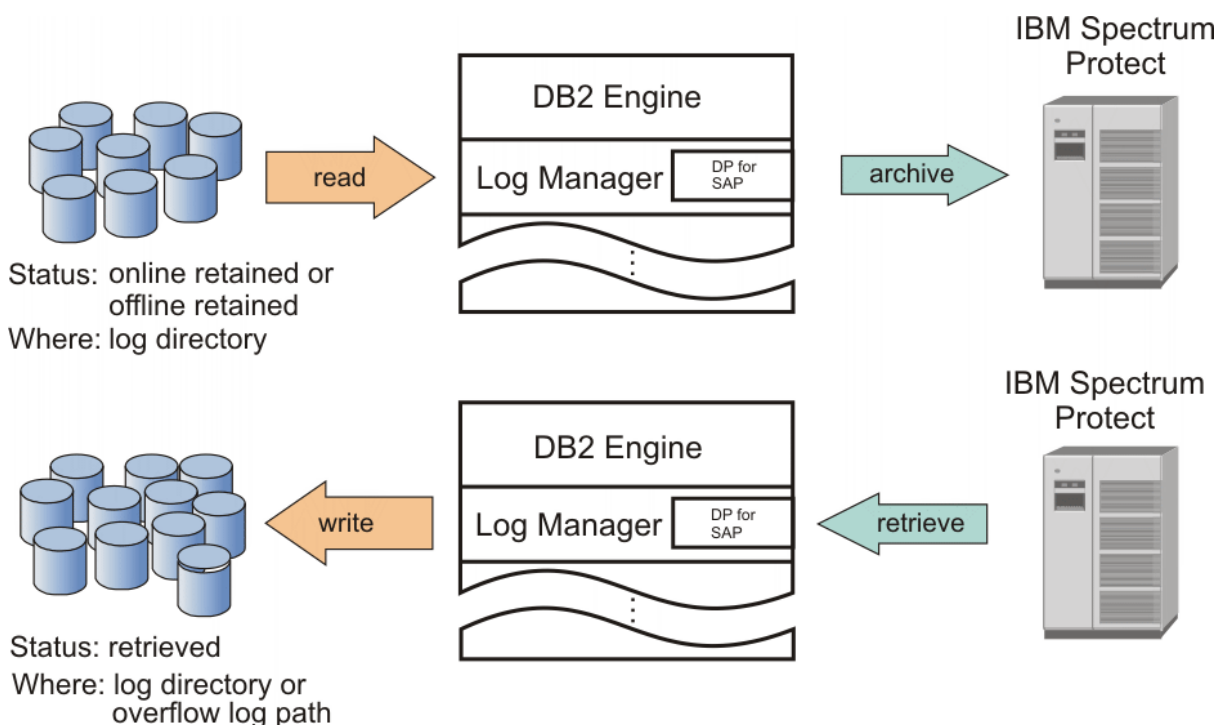


Figure 4: Log management with DB2® log manager and Data Protection for SAP

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

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

Detailed information about the DB2® log manager is available in your *DB2® Administration Guide*.

Backup objects and types of failures

Data Protection for SAP backs up and restores SAP® database objects.

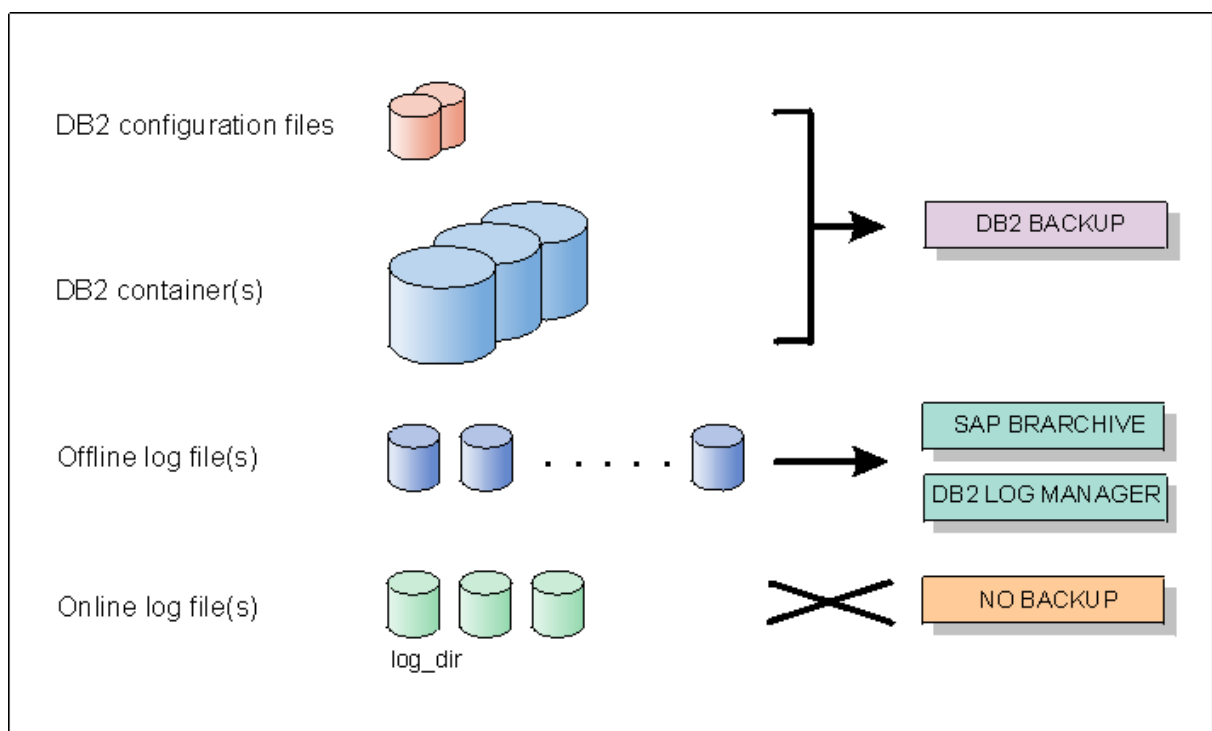


Figure 5: DB2® Backup Objects

Corrupt database

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

Hardware failure

In the event of a storage hardware failure, the database is typically restored to the most recent point in time. Thus, the most recent database image and DB2® log files are restored. However, the database executable files, SAP system data, and restoring user data might be necessary in the event of a hardware failure. To protect the system against the loss of SAP executable files, user data, or even operating system data, use the IBM Storage Protect™ backup-archive client incremental backup feature. You can use the client to define an include-exclude list of files that to be backed up during incremental backup operation. The include-exclude list is to exclude database container files and offline log files that were backed up or archived by Data Protection for SAP. Example include-exclude list files are provided in the Data Protection for SAP installation directory.

Disaster recovery

For a complete disaster recovery operation, all operating system data must be restored along with the database image, DB2® log files, database executable files, SAP system data, and user data. To help prevent a complete loss of the operating system, use operating system utilities (such as mksysb for AIX®) to run system backups. Such backups are to be done after you install, update, or upgrade the operating system. After these actions, you can start your system from the backup medium. A configured TCP/IP environment and IBM Storage Protect™ backup-archive client installation is to be included in a base backup to be able to restore all data. Since there is no provision for backing up online DB2® log files that are required for disaster recovery, place the DB2® log directory on a mirrored disk.

Planning

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

About this task

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

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

Database server considerations

In general, the production (SAP® database) server is the most critical component for data transfer, especially when parallelism is applied. As a result, special attention is given to the following items.

CPU power

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

I/O paths

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

Volume Manager settings

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

Disk layout

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

Disk layout

The manner in which the database files are laid out can affect data transfer rates. Since parallel is allowed, distribute data across several disks to take advantage of this feature.

Database size

The size of a database can be reduced by offloading inactive data to an external archive.

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 IBM Storage Protect™ 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

Data Protection for SAP uses the IBM Storage Protect™ archive function for all backup activities. When you are setting up the IBM Storage Protect™ server for use with IBM Storage Protect™ for Enterprise Resource Planning, the following considerations help you to optimize performance when you set up the IBM Storage Protect™ server.

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 IBM Storage Protect™ server CPU do increase when several clients access a single IBM Storage Protect™ 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 IBM Storage Protect™ 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 IBM Storage Protect™ 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 server

In SAP terminology, *backup* refers to the backup of data; *archive* refers to the backing up of log files. Data Protection for SAP uses the IBM Storage Protect™ archive function for backups and archives.

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

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 IBM Storage Protect™ 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 an IBM Storage Protect™ 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 IBM Storage Protect™ servers. Consider the location of all backup data before you remove an IBM Storage Protect™ 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 an IBM Storage Protect™ server if it contains valid backup data.

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 IBM Storage Protect™ 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 or DB2® partial and log file backups.

Parallel backup paths and backup servers

Data Protection for SAP can use several communication links between 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 IBM Storage Protect™ client-to-server communication and enable disaster recovery backup to a remote IBM Storage Protect™ server.

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

Each path in the `initSID.utl` profile is defined by a server statement and the corresponding definitions in the IBM Storage Protect™ client system option file `dsm.sys` (UNIX™ and Linux™) or `server.opt` (Windows™). The `SERVER <server 1..n>` statement denotes IBM Storage Protect™ servers that are defined in the Data Protection for SAP profile. This definition corresponds to the statement `SERVERNAME server 1..n` in the IBM Storage Protect™ 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 IBM Storage Protect™ 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 was unsuccessful, this path is skipped and Data Protection for SAP continues to the next path. This process continues until as many sessions are active as were specified in the total session number (`MAX_SESSIONS`). If this number is never reached (for example, because several paths were inactive), Data Protection for SAP ends the backup job.

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 IBM Storage Protect™ 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.

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

Compression

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

Multiple servers

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

Create multiple redo log copies

Data Protection for SAP can save a number of copies of each redo log by using different IBM Storage Protect™ 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 following Data Protection for SAP profile file keywords are important for creating multiple redo log copies:

- Keyword **BRARCHIVEMGTCLASS** denotes the IBM Storage Protect™ 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 IBM Storage Protect™ 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 IBM Storage Protect™ storage. This way, various redo log copies are on different backup media.
- When Db2 is used, Data Protection for Db2 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 possible.

- Data Protection for SAP cannot control the order in which IBM Storage Protect™ processes the requests. Therefore, an administrator cannot rely on sessions to be processed in the order they were started by Data Protection for SAP.
- When MAX_SESSIONS parameter is higher than SESSIONS value under server stanzas, Data Protection will enter a high performance mode and distribute the redo log copies across all usable server stanzas. However, when MAX_SESSIONS is equal or less than SESSIONS, Data Protection enters a high availability mode and redo log copies are distributed under first available stanza and defined archive management classes.

Planning for using IBM® HACMP™ for AIX®

Information is provided about Data Protection for SAP that is useful when you plan for HACMP™ failover configurations.

The following example uses the mutual takeover configuration (each node can take over the other node). If the application server and database server are installed on different hosts, the described actions must be taken on the database servers only.

This figure illustrates the takeover environment:




Figure 6: Sample environment for HACMP™ takeover

HACMP™ impact

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

Files

- The installation directory is `/usr/tivoli/tsm/tdp_r3`.
- Lock files are in `/var/tdp_r3`.
- There is only one ProLE running on each host (even after takeover).
- Each SAP® system has its own Data Protection for SAP configuration files (`initSID.ut1`, `initSID.bki`). These files are in a directory that is specified during the installation process.

Dependencies

- Both hosts must have the same level of IBM Storage Protect™ API installed.
- Both hosts must be Data Protection for SAP.
- On both hosts, the `dsm.sys` file (in `/usr/tivoli/tsm/client/api/bin/dsm.sys`) must contain all server names that are required for takeover.

Communication

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

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

Digital signing of executable files (Windows™)

Data Protection for SAP executable files (except .JAR files) for Windows™ systems have a digital signature.

The following files are affected:

- Passport Advantage® package for Windows™
- Data Protection for SAP installation files:
 - `version-TIV-TSMERPDB2-WinX64.exe`
- The Data Protection for SAP application files:
 - `backom.exe`
 - `prole.exe`
 - `tdpdb2.dll`

Code signing employs digital IDs, also known as certificates.

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

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

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

A warning is shown if the certificate is expired and if a time stamp is not present. A warning is also shown if the installation executable file is downloaded from a site that is not listed as a trusted site. The security warning is not related to the fact that executable files contain digital certificates. It is related to the security zone policy of the site you download the file from.

The executable file must be stored on an NTFS disk. The Internet Explorer Enhanced Security Configuration component (also known as Microsoft™ Internet Explorer hardening) reduces the server vulnerability to attacks from web content by applying more restrictive Internet Explorer security settings. As a consequence, Internet Explorer Enhanced Security Configuration might prevent some websites from displaying properly. It might also prevent users and administrators from accessing resources with Universal Naming Convention (UNC) paths on a corporate intranet. For more information about managing Internet Explorer Enhanced Security Configuration, see <http://www.microsoft.com/en-us/download/details.aspx?id=15013>. A security warning might be displayed whenever you run an executable file that is downloaded using the Internet Explorer from a URL or UNC that is not a member of the trusted security zone.

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

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

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

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

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

- Low risk: Opens the file with no warnings.

Warning messages do not prevent the file from being used. This is different from configuring the web server with a digital certificate.

Installing

This section provides installation instructions for a typical install. There are different procedures to install through the console, or to install the product in silent mode.

- Review the prerequisite information for the version before you start to install the software:
- Install the Data Protection for SAP product using the **InstallAnywhere** installation wizard.

Preparing to install

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

Before you begin

Be aware of differences between UNIX™ or Linux™, and Windows™ versions of Data Protection for SAP. For example, UNIX™ or Linux™ uses the path separator “/” and Windows™ uses the path separator “\” with a drive letter.

Procedure

1. Verify that the Data Protection for SAP package is complete. See the README .1ST file on each installation disk (or disk image) for a description of the contents.
2. Verify that the prerequisites are met as described in .
3. Review planning sheet information as described in the *Planning sheet for the base product* topic.
4. Install or upgrade Data Protection for SAP.

Prerequisites

Before you install Data Protection for SAP, review the hardware, software, and application requirements.

Requirements for Data Protection for SAP are published in the hardware and software requirements technote for each release. Review the technote for your version in the *IBM Storage Protect™ for Enterprise Resource Planning - all requirement documents* site, . From the page, follow the link to the technote for your release or update level.

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

These products must be installed before you install Data Protection for SAP:

- DB2®
- SAP R/3 or SAP e-business Solution
- IBM Storage Protect™ backup-archive client
For information about configuring the IBM Storage Protect™ API client, see the *Configure the IBM Storage Protect™ client options* topic. TCP/IP must be ready for communication between the IBM Storage Protect™ server and the IBM Storage Protect™ client.
- An operating system level that is supported by SAP and the IBM Storage Protect™ client

The release notes contain current information about Data Protection for SAP hardware, software, operating system, and maintenance levels.

When Data Protection for SAP is installed on a distributed file system, the root user requires read/write access to the file system during the installation.

An installation planning form for Data Protection for SAP is available in the `planning_sheet` (UNIX™ and Linux™) or `planning_sheet.txt` (Windows™) files in the installation directory. They are also available for printing in the *Planning sheet for the base product* topic. When prerequisites are met and installation planning information is completed, Data Protection for SAP is ready to be installed.

Installing in silent mode

You can install Data Protection for SAP for DB2® in silent mode using a response file. An installation that runs in silent mode suppresses the installation wizard. Instead, user data entry and status messages are displayed in the command line window.

Procedure

To run a silent or unattended installation complete the following steps.

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

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

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

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

```
./version-TIV-TSMERPDB2platform.bin -i silent -f properties file
```

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

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

Example

Sample properties file:

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

```
LOGGING_BOTH=1
LOGGING_NR=12
```

Lines starting with “#” are treated as comments.

Note: This example is a UNIX™ properties file. When you install IBM Storage Protect™ for ERP in silent mode for Windows™, use the corresponding Windows™ properties file.

Installing in a UNIX™ or Linux™ environment

Data Protection for SAP is delivered as a single executable file for each operating system. Use the executable file to start the installation wizard and to install the product.

About this task

Packages on the FTP server contain “FTP” before the operating system designation.

- For a disk or disk image, the name has the following format:

```
version-TIV-TSMERPDB2-platform
```

When the file is started, the IBM Storage Protect™ for ERP installation wizard guides you through the procedure. Read the descriptions carefully and follow the guidelines that are displayed on the windows.

Shared libraries have different file extensions on different UNIX™ or Linux™ operating systems. Within the following the section, the file extensions of shared libraries are represented as *ext*. Replace this text with the extension that applies to your operating system:

Table 2: File Extensions for Shared Libraries	
Operating System	Extension
AIX®	a
HP-UX	sl
Linux™	so
Solaris	so

Procedure

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

Result

These modifications are automatically done to your system during installation:

- An entry is created in */etc/inittab* that automatically starts the “ProLE” daemon on UNIX™ systems. If upstart is configured, */etc/init/prole_db2.conf* is created and upstart starts the “ProLE” daemon.
- An entry is created in */etc/inittab* that automatically starts the “ProLE” daemon on UNIX™ systems. If upstart is configured, */etc/init/prole_db2.conf* is created and upstart starts the “ProLE” daemon.
- The environment variable *XINT_PROFILE* specifies the IBM Storage Protect™ for ERP profile that is in the path that is specified for configuration files during installation. The file name is *initSID.utl* where *SID* is the DB2® database *SID* specified during installation.

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

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

If the DB2® instance was not running, you must complete these tasks manually, as described in the *Specifying the VENDOROPT parameter* and *Configuring the DB2 Log Manager* topics.

An entry is created in /etc/inittab that automatically starts the “ProLE” daemon on UNIX™ systems. The EN_US folder is created, which contains the message catalog file tsmerp.cat. The _uninst folder is also created, which contains more files.

These files are installed in the IBM Storage Protect™ for ERP installation directory:

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

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

These files are installed in the directory where the IBM Storage Protect™ for ERP configuration files are located:

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

Uninstalling older versions (UNIX™ and Linux™)

Follow the procedure to uninstall a previous version of IBM Storage Protect™ for Enterprise Resource Planning

Procedure

1. Log in as root user.
2. Make sure that the **DISPLAY** variable is set correctly as the uninstall procedure requires a graphical X-Window.
3. Make sure the previous version of IBM Storage Protect™ for Enterprise Resource Planning is not running.
4. Start the uninstall executable file as shown here:
AIX® 64-bit:

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

Other UNIX™ 64-bit or Linux™ 64-bit:

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

Follow the instructions of the uninstall dialog.

Installing in a Windows™ environment

IBM Storage Protect™ for ERP is delivered as a single executable file (.exe) for each operating system. Packages on the FTP server contain FTP before the operating system designation.

About this task

IBM Storage Protect™ for ERP for these operating systems is delivered as a single executable file for each operating system. The packages are named as follows:

- The package name on the disk or the disk image, which is shown in this example:

```
version-TIV-TSMERPDB2-platform
```

Procedure

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

Result

The following modifications are done on your system during installation:

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

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

If the DB2® instance is not running, these tasks must be run manually as described in the *Specifying the VENDOROPT parameter* and *Configuring the DB2 Log Manager* topics.

The following files are installed in the IBM Storage Protect™ for ERP installation directory:

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

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

These files are installed in the directory where the IBM Storage Protect™ for ERP profile is located:

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

Enabling ProLE to access configuration files on a remote share (Windows™)

When ProLE is started as a regular service, it operates under the ID of the local system account with Administrator privileges. However, a session opened on a remote system does not have credentials or permissions. You must grant access to ProLE to access the files on a remote share.

About this task

ProLE sessions on a remote system cannot access files that are on that remote share. This condition is true even when the share is mapped to a local drive letter or is accessed as a Uniform Naming Convention (UNC) notation (\\server\path\). Data Protection for SAP accepts UNC notation for the profile but not for all the files that are specified within the profile. These files are opened by ProLE, which by default has no permission to access remote shares.

Follow the procedure to enable ProLE to access all files on a remote share:

Procedure

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

Uninstalling older versions (Windows™)

Follow these steps to uninstall a previous version of Data Protection for SAP in a Windows™ environment.

Procedure

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

Verify the installation or upgrade

When you complete the installation or upgrade of the product, you can verify that the procedure was successful by running a backup task.

To verify the installation of Data Protection for SAP, do a full DB2® database backup. Then, restore with the DB2® Control Center or DB2® command line processor (CLP).

Before you start, you must plan to run a complete offline backup. Then you can run a complete restore or recovery of the entire SAP database for verification.

Upgrading

Follow the tasks to upgrade to Data Protection for SAP.

Upgrading the base product

Upgrade Data Protection for SAP from an earlier version.

Before you begin

If you are upgrading IBM Storage Protect™ for Enterprise Resource Planning on a busy SAP system where the software continuously starts log archives, it might be difficult to find a maintenance gap without any active log archiving processes. To alleviate this situation, you can stop the prole daemon. Each operating system has a different method to stop the prole daemon as follows.

- **AIX®** As a root user, run the following command: `rmmitab pd64`
- **RHEL 7 and later, SLES 12 and later** As root user, run the following command: `systemctl stop prole_db2`
- **Older Linux® versions and other Unix operating systems** As root user, comment out the line with `prole` in `/etc/inittab` and run command `init q`
- **Windows™** As a user with Administrator privileges, either run command `net stop prole`, or use the Services control panel to stop the prole service.

Tip: When you stop the prole daemon, redo log archive operations will not work. Typically, this is not an issue because the next archive run would pick up all the redo logs that previously failed to archive. But, if the system is generating a large amount of redo logs, the file system might run out of space.

About this task

The format of the configuration file (`.bki`) was changed with version 5.4. The software accepts the previous format and converts it automatically. If it is necessary to use a version earlier than 5.4, the old format can be recovered by overwriting the new file with the empty file. The previous version provides the empty file. The file must then be initialized by setting the IBM Storage Protect™ password. However, the information about the current backup number is lost. As a result, more backup versions must be retained for a longer time than is specified by the **MAX_VERSIONS** parameter.

Procedure

1. Verify that the Data Protection for SAP package is complete. The installation packages are provided on a disc or disc image (downloadable from Passport Advantage®), or the IBM® FTP server. See the release notes file in the IBM Storage Protect™ Knowledge Center for the most current release information.
2. Check the readme files and release notes for incompatibilities between the installed version and the new version. Make sure that data backed up with an older version of IBM Storage Protect™ for Enterprise Resource Planning can still be restored with the version to be installed.
3. Make sure that the requirements for the new version of Data Protection for SAP are met as described in .
4. Make sure that planning information is available as described in the *Prerequisites* topic.
5. A full backup of the SAP database must be performed before you upgrade to the new version.
6. Uninstall the old version as described in the *Uninstalling older versions* topics.
7. Install the new version of Data Protection for SAP as described in the *Prerequisites* topic.
8. Update the Data Protection for SAP profile as described in the *Migrate the Data Protection for SAP profile* topic.
9. Create the configuration file or files as described in the *Creating the configuration files* topic.

10. Perform the necessary configuration tasks as described in the *Configure the IBM Storage Protect™ client options* topic.
11. Verify the installation as described in the *Verify the installation or upgrade* topic.
12. A full backup must be performed after you upgrade to the new version.

Migrate the Data Protection for SAP profile

The license file, the profile, and the configuration files are not deleted when Data Protection for SAP is uninstalled.

These files can be used by the new version of Data Protection for SAP. To reuse the existing configuration and connection to the IBM Storage Protect™ server, choose not to change the profile when you are prompted during installation.

Configuring

In addition to configuring Data Protection for SAP, you need to configure other applications, for example, the IBM Storage Protect™ 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
- DB2® Log Manager and related DB2® files
- HACMP™
- Distributed File System
- IBM Storage Protect™ backup-archive client
- IBM Storage Protect™ server

Changing configuration tasks for the Data Protection for SAP base product

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

About this task

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

Configuring profile tasks

To configure the Data Protection for SAP profile file, you must set the server statement and in the IBM Storage Protect™ 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 IBM Storage Protect™ client option file.

There are corresponding keywords in the IBM Storage Protect™ 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 IBM Storage Protect™ client option file are established by using the logical server name. This logical server name is defined by the keywords SERVER or SERVERNAME.

Table 3: SERVER statement and appropriate profile and option file settings.		
Configuration possibilities	Data Protection for SAP profile <code>initSID.utl</code>	IBM Storage Protect™ client option file <code>dsm.sys</code> or <code>server.opt</code> ^[2]
single path; no password or manual password	<code>SERVER</code> <i>server</i> <code>ADSMNODE</code> <i>node</i> ^[1]	<code>SERVERNAME</code> <i>server</i> <code>TCPSERVERADDRESS</code> <i>address</i> <code>NODENAME</code> <i>do not</i> specify

Configuration possibilities	Data Protection for SAP profile <code>initSID.utl</code>	IBM Storage Protect™ client option file <code>dsm.sys</code> or <code>server.opt</code> [2]
single path; automatic password by IBM Storage Protect™	<pre> SERVER server ADSMNODE do not specify </pre>	<pre> SERVERNAME server NODENAME node TCPSERVERADDRESS address </pre>
several paths/servers; no password or manual password	<pre> SERVER server 1 ADSMNODE node 1 SERVER server 1 ADSMNODE node n </pre>	<pre> SERVERNAME server 1 NODENAME do not specify TCPSERVERADDRESS address 1 SERVERNAME server n NODENAME do not specify TCPSERVERADDRESS address n </pre>
several paths/servers; automatic password by IBM Storage Protect™ [3]	<pre> SERVER server 1 ADSMNODE do not specify SERVER server n ADSMNODE do not specify </pre>	<pre> SERVERNAME server 1 NODENAME do not specify TCPSERVERADDRESS address 1 SERVERNAME server n NODENAME do not specify TCPSERVERADDRESS address n </pre>
several paths/servers; automatic password by IBM Storage Protect™ [4]	<pre> SERVER server ADSMNODE do not specify TCP_ADDRESS address 1 SERVER server n ADSMNODE do not specify TCP_ADDRESS address n </pre>	<pre> SERVERNAME server NODENAME node TCPSERVERADDRESS address </pre>

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 IBM Storage Protect™ servers.
On Windows™, there is a separate client option file `server.opt` for each IBM Storage Protect™ server.

[3]

If two different physical systems have the same IBM Storage Protect™ 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 **dsmadm** 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]

If you are using IBM Storage Protect™ 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 IBM Storage Protect™ 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 paths

This example assumes that the IBM Storage Protect™ 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
  ADPMNODE         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
  ADPMNODE         C21
  SESSIONS         2
  PASSWORDREQUIRED YES
  BRBACKUPMGTCLASS mdb
  BRARCHIVEMGTCLASS mlog1 mlog2
# USE_AT          0 1 2 3 4 5 6
```

Example of SERVER statement with parallel servers

This example assumes the following configuration:

- Two IBM Storage Protect™ servers (each with two tape drives) with connections through two network paths:
 - *server_a* uses TCP/IP address xxx.xxx.xxx.xxx
 - *server_b* uses TCP/IP address yyy.yyy.yyy.yyy
- An SAP® database server that is connected to two networks.
- Daily backups are run on both systems.

The following is an example of the Data Protection for SAP profile that is used in this parallel configuration:

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

```

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

```

Example of SERVER statement with alternate servers

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 IBM Storage Protect™ 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
ADMSNODE      C21
SESSIONS      4
PASSWORDREQUIRED YES
BRBACKUPMGTCCLASS MDB
BRARCHIVEMGTCLASS MLOG1 MLOG2 MLOG3 MLOG4
USE_AT      1 2 3 4

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

```

DB2® tasks

When you are configuring Data Protection for SAP for DB2®, you must do some additional steps to complete the process.

- Review the vendor environment file `vendor.env`, and the Data Protection for SAP profile file.
- Set the `VENDOROPT` parameter for backup and restore commands.
- Configure parameters for DB2® Log Manager.
- Create configuration files for DB2® partitions.

DB2® and Data Protection for SAP configuration guidelines

Data Protection for SAP data transfer functions are implemented in a shared library that is accessed by DB2®. When back up or restore operations are started, and log archive or log retrieve commands are entered, the shared library is accessed. This library requires information about the path of the Data Protection for SAP profile and the path of the log files that are written by Data Protection for SAP.

BACKUP DATABASE and RESTORE DATABASE

If an action is initiated by using the DB2® commands **BACKUP DATABASE** or **RESTORE DATABASE**, the information that is required must be specified in a vendor environment file. The name of the vendor environment file is sent to DB2® through either the *OPTIONS* parameter of the **BACKUP DATABASE** or **RESTORE DATABASE** commands. It can also be stored persistently in the database configuration parameter *VENDOROPT*. For log archive or log retrieve, this file can be stored either in the database configuration parameter *LOGARCHOPT1* or in *LOGARCHOPT2*. For **BACKUP DATABASE** or **RESTORE DATABASE**, use of the *OPTIONS* keyword for this purpose is no longer necessary. It is advised that you keep the settings in the vendor environment file and in the system variables synchronized continually. For an example of a Data Protection for SAP vendor environment file, see [“Vendor environment file sample” on page 106](#). If **BACKUP DATABASE** or **RESTORE DATABASE** is triggered through the backom utility, the information that is required must be specified in the environment.

Vendor environment file

Consider the following adjustment rules for Data Protection for SAP:

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

Specifying the VENDOROPT parameter

You can select a default set of Data Protection for SAP for DB2® environment settings for commands **BACKUP DATABASE**, **RESTORE DATABASE**, and for the DB2® Log Manager.

About this task

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

- You do not need to restart the DB2® instance.
- You can define default values for the *OPTIONS* parameter of the **BACKUP DATABASE** command in the DB2® configuration.
- You can define default values for the *OPTIONS* parameter of the **RESTORE DATABASE** command in the DB2® configuration.
- The same settings can be applied to database backup and restore, and to log file archive and retrieve options.

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

Procedure

1. Modify the DB2® database configuration to denote a file that contains the following settings:

db2 update db cfg for *SID* using LOGARCHOPT1/2 *vendor environment file*, where *vendor environment file* is the fully qualified path of the file that contains Data Protection for SAP environment settings for DB2®.

2. Verify that the environment settings of your system match the settings in this file.

Configuring the DB2® Log Manager

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

Procedure

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

- (UNIX™ and Linux™):

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

- (Windows™):

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

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

- (UNIX™ and Linux™):

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

- (Windows™):

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

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

Result

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

Table 4: Configuration parameters for DB2® database backup and restore, and log archive and retrieve		
Parameter	Description	Default
LOGARCHMETH1	Media type of the primary destination for archived log files	Off
LOGARCHOPT1	Options field for the primary destination for archived log files (if required).	NULL
LOGARCHMETH2	Media type of the secondary destination for archived log files. If this path is specified, log files are archived to both this destination and the destination that is specified by LOGARCHMETH1.	Off
LOGARCHOPT2	Options field for the secondary destination for archived log files (if required).	NULL

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

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

Creating the configuration files

When you set the IBM Storage Protect™ password with the **backom** utility, the configuration files for all DB2® partitions are automatically created in the paths *path*/**%DB2NODE**/.

About this task

The *path* is the directory is the value of keyword **CONFIG_FILE** in the profile. **%DB2NODE** is replaced automatically by a DB2® partition name that is referenced in the DB2® configuration file **db2nodes.cfg**.

- If the directory denoted by the value of keyword **CONFIG_FILE** is not in the same network file system as the DB2® configuration file **db2nodes.cfg**, you must set the password for each system where a partition of the database is located.
- If the database is not partitioned, **NODE0000** is used as the only DB2® partition name.

Optional: Setting backup object segmentation

Environments that contain large databases that rapidly increase in size might encounter problems when data is transferred to the IBM Storage Protect™ server. For example, you might encounter the following problems when you back up or restore large databases:

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

To avoid potential problems that are related to transferring large objects, use the Data Protection for SAP **SEGMENTSZ** profile keyword. This keyword specifies the upper bound of the segments that are split from large backup objects during backup and restore processing.

Configuring distributed file system tasks

Configure Data Protection for SAP in a distributed file system. If the SAP systems are statically assigned to specific hosts, you do not need to configure in a distributed file system. If the root user ID has write access to the distributed file system, you do not need to configure in a distributed file system.

Configuring for a distributed file system

Configure IBM Storage Protect™ for ERP in a distributed file system with the following procedure.

Before you begin

For a single SID on a host, IBM Storage Protect™ for ERP sets the ProLE service to run with the db2SID user ID instead of root. Follow the procedure to set up the ProLE service to run with the db2SID user ID.

About this task

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

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

If these conditions exist, the standard installation process can be used as described in the *Preparing to install* topic.

Procedure

1. Enable root access to the distributed file system.
2. Install IBM Storage Protect™ for ERP by using the procedure that is described in the *Preparing to install* topic.
3. On a UNIX™ system, replace the following entry in the `/etc/inittab` file:

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

with this entry:

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

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

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

Result

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

Configuring as an HACMP™ application

Configure Data Protection for SAP for HACMP™. Data Protection for SAP must be defined as an application to HACMP™, and must be in a resource group that has a cascading or rotating takeover relationship. It does not support a concurrent access resource group.

Before you begin

A prerequisite for installation is a correct setup of the IBM Storage Protect™ client.

About this task

Although the *HACMP™ for AIX® Installation Guide* can be reviewed for detailed instructions, a high-level summary is provided here.

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

```
smit hacmp
```

2. Select **Cluster Configuration > Cluster Resources > Define Application Servers > Add an Application Server**.

3. Enter field values as follows:

Server Name

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

Stop Script

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

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

Adding Data Protection for SAP to a HACMP™ resource group

A final step in enabling Data Protection for SAP for HACMP™ failover is to define it to a cluster resource group.

Before you begin

Although the *HACMP™ for AIX® Installation Guide* can be reviewed for detailed instructions, a high-level summary is provided here. Perform these tasks to define the resources that are part of a resource group:

Procedure

1. From the Cluster Resources SMIT screen, select the **Change/Show Resources/Attributes for a Resource Group** option and press **Enter**. SMIT displays a picklist of defined resource groups.
2. Pick the wanted resource group. Press Enter and SMIT displays the **Configure a Resource Group** screen.
3. Enter values that define all the resources you want to add to this resource group.
4. After you enter field values, synchronize cluster resources.
5. Press F10 to exit SMIT or F3 to return to previous SMIT screens to run other configuration tasks or synchronize the changes that you just made. To synchronize the cluster definition, go to the Cluster Resources SMIT screen and select the **Synchronize Cluster Resources** option.

What to do next

The IBM Storage Protect™ client application must be added to the same resource group that contains the file systems it will back up. The file systems that are defined in the resource group are to also be the ones that are specified in the domain for this client instance in the client user options file. Both JFS and NFS file systems can be defined as cluster resources, although NFS supports only two node clusters in a cascading takeover relationship.

HACMP™ stop script example

A stop script that operates in an HACMP™ environment is illustrated.

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

The stop script is used in the following situations:

- HACMP™ is stopped.

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

The stop script is called by HACMP™ with the root user ID.

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

[illegible]

Configuring IBM Storage Protect™

Data Protection for SAP requires that you complete configuration tasks for the IBM Storage Protect™ backup-archive client and server.

IBM Storage Protect™ client tasks

Data Protection for SAP requires that configuration tasks be run for the IBM Storage Protect™ client as part of the overall product configuration.

Configure the IBM Storage Protect™ client options

The IBM Storage Protect™ clients must be configured after the IBM Storage Protect™ 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 IBM Storage Protect™ 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 IBM Storage Protect™ 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 IBM Storage Protect™ client options on UNIX™ or Linux™

IBM Storage Protect™ 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 IBM Storage Protect™ backup-archive client to operate in an SAP environment with the following procedure.

Procedure

1. Install the IBM Storage Protect™ 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 IBM Storage Protect™ to include or exclude the files that are listed in `inclexcl.list`.
You might want to exclude all database files that are processed by the DB2® database backup.
5. Throughput improves when tape drives attached to the IBM Storage Protect™ server provide hardware compression. However, combining hardware compression and IBM Storage Protect™ 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)
```

Result

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

The IBM Storage Protect™ provides two features for specifying the location of the IBM Storage Protect™ API Client error log: the environment variable **DSMI_LOG** and the IBM Storage Protect™ 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 IBM Storage Protect™ 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.

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

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

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.

Setting IBM Storage Protect™ client options

IBM Storage Protect™ clients on Windows™ are configured by setting options in the file **server_a.opt**, where **server_a** is the logical server name in the **initSID.utl** file. The **include/exclude** file is also used to define which files are included or excluded during backup, archive, or hierarchical storage processing.

About this task

To configure the IBM Storage Protect™ backup/archive clients to operate in an SAP environment, complete the following steps:

Procedure

1. Install the IBM Storage Protect™ client software on the SAP database server system.
2. For each logical IBM Storage Protect™ server, a corresponding client option file is needed. In this example, the file name must be **server_a.opt** since **server_a** is the logical server name:

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

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

3. Specify **TCPServeraddress 127.0.0.1** or loopback if the server and client are on the same system. This selection improves TCP/IP communication speed.

4. Specify `Inc1Exc1` if you want IBM Storage Protect™ to include or exclude the files that are listed in `inc1exc1.list`. You might want to exclude all database files that are processed by the DB2® database backup.
5. Throughput improves when tape drives attached to the IBM Storage Protect™ server provide hardware compression. However, combining hardware compression and IBM Storage Protect™ client software compression (`Compression ON`) is not advised.

Result

An IBM Storage Protect™ error log (required for each client) can be specified for each process regardless of the number of IBM Storage Protect™ client option files `server.opt` involved. The IBM Storage Protect™ error log is determined by these rules:

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

Set up the IBM Storage Protect™ client so that different processes write to separate error logs. The error log path must be defined in the **DSMI_LOG** environment variable if the client options files are shared among processes.

IBM Storage Protect™ server tasks

Data Protection for SAP requires configuration tasks to be done for the IBM Storage Protect™ server as part of the overall product configuration.

Configure the IBM Storage Protect™ server

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

Although the task examples use IBM Storage Protect™ commands, these tasks can also be run using the IBM Storage Protect™ web client GUI.

Consider the following performance-related guidelines before you install the IBM Storage Protect™ server.

IBM Storage Protect™ server host system

The IBM Storage Protect™ 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 IBM Storage Protect™ server is sufficient for a single SAP system landscape. If the IBM Storage Protect™ server is used to back up and restore other clients, consider installing the server on a large system or by using several IBM Storage Protect™ servers.

Network topology

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

In the AIX®: LPAR environment

An LPAR node can be used for an IBM Storage Protect™ server. The use of a High Performance Switch network can improve backup and performance.

These steps are considered complete when the IBM Storage Protect™ 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 IBM Storage Protect™ server is started.

The latest code fixes for IBM Storage Protect™ can be found at: <ftp://public.dhe.ibm.com/storage/tivoli-storage-management/maintenance>

Specifying an IBM Storage Protect™ server

To configure Data Protection for SAP, you need to specify an IBM Storage Protect™ server in the profile file.

About this task

Follow these steps to add an IBM Storage Protect™ server:

Procedure

1. Add a server statement to the Data Protection for SAP profile.
2. Adapt the IBM Storage Protect™ options files as described in the *Verifying the IBM Storage Protect™ server name* topic.
3. Set and save the IBM Storage Protect™ password for the new server as described in the *Setting the IBM Storage Protect™ password* topic.

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 IBM Storage Protect™ 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 containers: `define stgpool sap_db device_class_name maxscr=20`
3. Define a storage pool for the offline log files: `define stgpool sap_log1 device_class_name maxscr=3`

Result

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 IBM Storage Protect™ 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 an IBM Storage Protect™ 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 log files and copies of offline log files :

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

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

3. Define a copy group:

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

Data Protection for SAP uses *version control* for managing SAP database backups by backing up all data to only those management classes for which an archive copy group is defined (**typearchive**). To prevent backed up files within IBM Storage Protect™ server storage from being deleted due to expiration dates (IBM Storage Protect™ deletes expired files), the copy group parameter **retver**, which specifies the number of days a file is to be kept, must be set to unlimited (9999ornolimit).

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 completing the configuration. The IBM Storage Protect™ 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 IBM Storage Protect™ 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 IBM Storage Protect™ server.

Determining the IBM Storage Protect™ password method

Specify how Data Protection for SAP manages the IBM Storage Protect™ 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 IBM Storage Protect™ server. Each user that is connected to the backup server can access IBM Storage Protect™ data without a password. This method is advised only if adequate security measures are established.

For example, no password might be acceptable when the IBM Storage Protect™ 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 IBM Storage Protect™.

Manual handling of password

A password is required for each connection to the IBM Storage Protect™ 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 IBM Storage Protect™ server, Data Protection for SAP automatically uses the stored password when it connects to IBM Storage Protect™. This method provides password security and can be set up easily. Whenever the password expires or is changed, the new password must be set with this command:

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

If you are setting the password to be automated (such as in a script), enter this command:

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

where *passwordA* is the password for IBM Storage Protect™ node *nodeA* on IBM Storage Protect™ 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 IBM Storage Protect™ passwords in a script.

Automatic handling of password

A password is required for each connection to the IBM Storage Protect™ server. After the first connection, the password is managed by IBM Storage Protect™. The IBM Storage Protect™ client stores the current password locally. When the password expires, the password is changed and stored automatically.

If you schedule your backups or restore from a system user different from the database owner, you must grant access permissions to your data files on disk for this user. Specify the IBM Storage Protect™ password in use before you start by using Data Protection for SAP to connect to the server. Whenever the password is changed manually on the IBM Storage Protect™ server, again connect to the server and update the password with the command **update node**. Use the following command for automatic password handling:

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


This method is advised for an automated production environment.

Setting the IBM Storage Protect™ password

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

About this task

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

Provide Data Protection for SAP with the password for the IBM Storage Protect™ 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 IBM Storage Protect™, 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.

Password handling parameters and profile actions in a UNIX™ or Linux™ environment.

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

Note:

1. See appropriate IBM Storage Protect™ 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 IBM Storage Protect™ server for the node.
4. **ADSMNODE** must not be set when **PASSWORDACCESS** generate is set.
5. The users *SIDadm* and *db2SID* must have read and write permission for the path specified.
6. This step is only necessary if the password is expired (manual-handling only) or must be changed on the IBM Storage Protect™ server.
7. A password must be entered for each server statement in the Data Protection for SAP profile.
8. When you use **PASSWORDACCESS GENERATE**, the operations must always be used with the same user ID provided in step 5 (setting of passwords).

Password configuration matrix (Windows™)

When the preferred method of password-handling is determined, review the following steps to set the keywords and parameters in the various profiles.

Detailed information about password-handling methods is available in the *Determining the IBM Storage Protect™ password method* topic.

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

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

Note:

1. See IBM Storage Protect™ documentation.
2. If you are using manual password generation during testing, make sure that the expiration period is set to an appropriate time.
3. For an initial setup, this password must be the same password that is specified when the node was registered to IBM Storage Protect™. The password must be changed first on the IBM Storage Protect™ server and then on Data Protection for SAP.
4. **ADSMNODE** must not be set when **PASSWORDACCESS** generate is set.

5. The users *SIDadm* and *sapserviceSID* must have read and write permission for the path specified .
6. This step is only necessary if the password is expired (manual-handling only) or must be changed on the IBM Storage Protect™ 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 *initSID.utl* file.

- Review the IBM Storage Protect™ client options files to make sure that the server name matches the name that is specified in the server statement of the *initSID.utl* file.
- Review that other parameters are set correctly. These settings depend on the password method selected.
- (UNIX™ or Linux™) Define the IBM Storage Protect™ server in the IBM Storage Protect™ client system options file (*dsm.sys*). The server stanza that is specified in *dsm.sys* must match the entry in *initSID.utl*.
- (Windows™) Define a client options file *servername.opt*. This file must be in the directory that contains *dsm.opt*. The value of *servername* is the server name that is specified in *initSID.utl*.

Protecting data

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

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.

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**.

Scheduler (Windows™) or Crontab (UNIX™ or Linux™)

Automating backups at the database server level is available by using either the Schedule Services feature (on Windows™) or the **crontab** command (for UNIX™ or Linux™).

IBM Storage Protect™ scheduler

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

Sample IBM Storage Protect™ schedule

This sample procedure is flexible because you can define a command file with any set of commands you choose. This allows you to use the same command file to define schedules on other IBM Storage Protect™ servers.

1. Enter the following command on the server console or from an administrative client to define the schedule. The administrative client does not have to be running on the same system as the IBM Storage Protect™ server.

```
def sched PolicyDB daily_db_bkup desc="Daily Online DB Backup"
  action=command objects="/home/admin/sched/schedbkdb.scr"
  starttime=21:00 duration=15 duru=minutes period=1 perunits=day
  dayofweek=any
```

IBM Storage Protect™ displays this message:

```
ANR2500I Schedule daily_db_bkup defined in policy domain PolicyDB.
```

2. To associate Data Protection for SAP to this backup schedule, issue the following command:

```
define association PolicyDB daily_db_bkup NodeA1
```

IBM Storage Protect™ displays this message:

```
ANR2510I Node NodeA1 associated with schedule
daily_db_bkup in policy domain PolicyDB.
```

A backup event (schedule) is now defined on the IBM Storage Protect™ server. The schedule runs a command file called schedbkdb.scr located in the /home/admin/sched directory. The backup starts around 9:00 PM., runs once a day, and can start on any day of the week. Use the IBM Storage Protect™ administrative commands query schedule or query association to confirm that you set the schedule and association correctly.

IBM® Workload Scheduler

The IBM® 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.

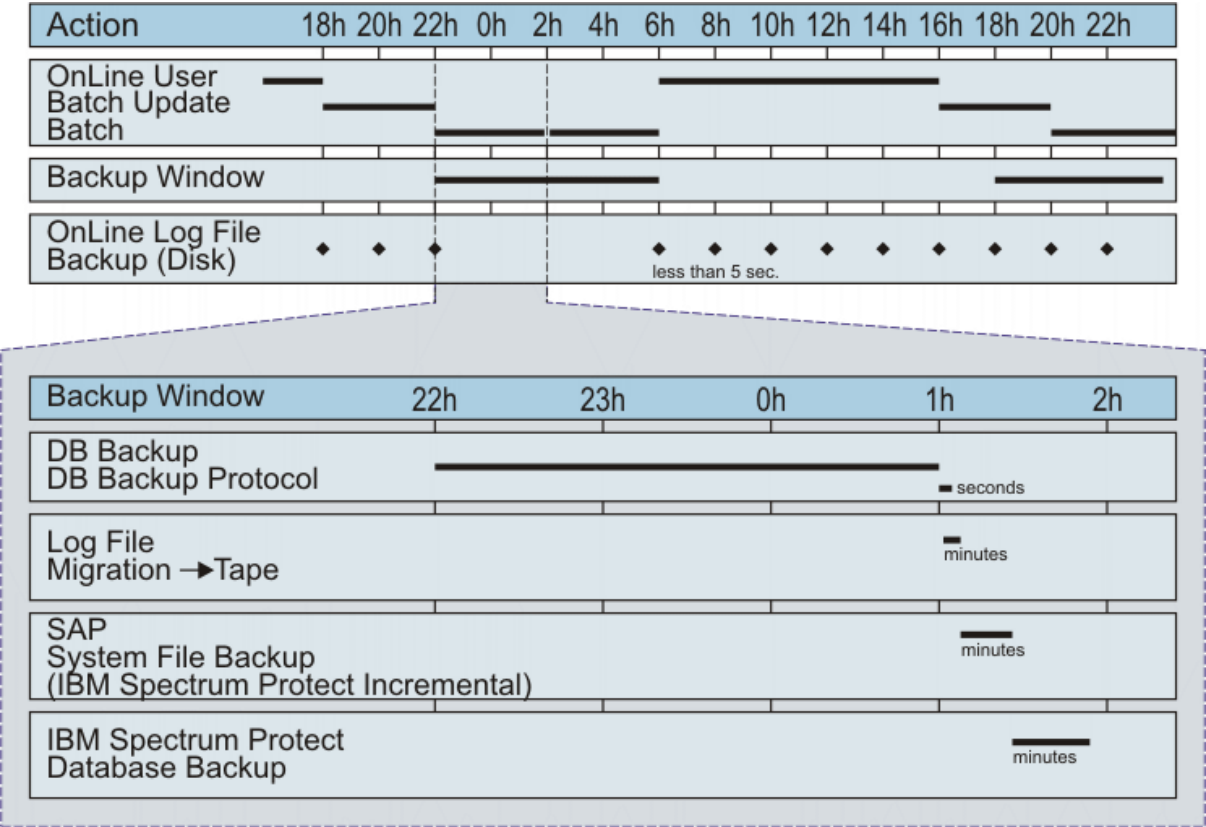


Figure 7: Production Backup Example

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

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

Backups can be moved to disk storage and to tape media. The IBM Storage Protect™ server manages the data regardless of the storage media. However, backing up the SAP database directly to tape is the preferred media.

Windows™ scheduling example

An example of a batch file schedule is shown.

About this task

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

```
net start schedule
```

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

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

Backups in a nonpartitioned database environment

The following examples show how you can start DB2® database or table space backups from the command line by using DB2® CLP.

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

- Full online backup (database parameter **LOGRETAIN** must be activated):

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

- Online table space backup (database parameter **LOGRETAIN** must be activated):

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

Use DB2® single system view for backup

DB2® Version 9.5 and later provides the single system view (SSV) function, which allows backing up multiple database partitions immediately.

In earlier releases, partitioned databases that are needed to be backed up one partition at a time that can be time-consuming and prone to errors. Backing up a partitioned database one partition at a time also failed to include the log files in the backup image. These log files are required to restore and recover the data. Restoring multiple partitions that were backed up individually is complicated as well because the backup timestamp for each partition is slightly different. Identifying all database partitions belonging to the same backup is difficult. Also, determining the minimum recovery time for the backup that contains these partitions is difficult. Use of **db2_all** simplifies the backup of partitioned databases. However, backup and restore operations restrictions still exist that complicate these tasks.

With DB2® Version 9.5 and later, when you do a backup operation of a partitioned database, you can specify which partitions to include in the backup. If necessary, you can include all the database partitions. The specified partitions are backed up simultaneously and the backup timestamp that is associated with all specified database partitions is the same. Also, by default, database logs are included in an SSV backup image. Finally, when you restore from an SSV backup image, you can specify to roll forward to end of logs, which is the minimum recovery time that is calculated by the database manager. For more information, see the *DB2® Command Reference*.

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 IBM Storage Protect™ volumes or on different IBM Storage Protect™ 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.
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 IBM Storage Protect™ 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 IBM Storage Protect™ 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.

Schedule batch sample

Example

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

Full offline backup batch file sample

Example

```
@echo off
rem Full Offline Backup batch file:
rem -----
rem file name: backup.cmd
rem -----
rem Sample DB2 backup batch file for 3264bit environments
rem -----
rem Task:
rem Invokes a DB2 backup in order to perform a full offline backup of
rem all DB2 tablespaces
rem -----
rem ***** NOTE ***** NOTE ***** NOTE *****
rem
rem This script is intended only as a model and should be
rem carefully tailored to the needs of the specific site.
rem
```

```

rem ***** NOTE ***** NOTE ***** NOTE *****
rem -----
rem
rem For the following examples, the system ID of the DB2 database
rem is assumed to be "C21".
rem
rem -----COMMAND-----
db2 backup db C21 load 'C:\Program Files\tivoli\tsm\tdp_r3\db264\tdpdb2.dll'

```

Full offline backup shell script sample

Example

```

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

```

Segment large backup objects

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

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

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

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

When IBM Storage Protect™ for Enterprise Resource Planning initiates a change of IBM Storage Protect™ objects, the segment number, for the new backup object segment, increases by one.

For integrity check processing of the backup segment group, an extra zero-byte backup object, the so-called commit object, is generated. This object is used by IBM Storage Protect™ for Enterprise Resource Planning to check the integrity of the related backup segment group. The naming convention of the commit object is as follows:

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


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

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

IBM Storage Protect™ for Enterprise Resource Planning log entries

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

DB2® Backup Object Manager

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

Segmentation and IBM Storage Protect™ server

If segmentation is used for backup operations, the IBM Storage Protect™ server might issue the error message, "ANS0326E This node exceeds its maximum number of mount points. This situation happens because there might be a short delay before IBM Storage Protect™ actually closes client sessions. To overcome this problem, the IBM Storage Protect™ server MAXNUMMP parameter for the IBM Storage Protect™ node is set to twice the number of IBM Storage Protect™ sessions that are used for the backup.

The number of active parallel sessions for IBM Storage Protect™ for Enterprise Resource Planning to the IBM Storage Protect™ server is limited by the IBM Storage Protect™ for Enterprise Resource Planning parameters SESSION, in the SERVER stanza, and MAX_SESSIONS.

For example, if two IBM Storage Protect™ sessions are needed for the database backup, the MAXNUMMP parameter for the used node is set to four in the IBM Storage Protect™ server. In this example, IBM Storage Protect™ for Enterprise Resource Planning sends the data by using two IBM Storage Protect™ sessions only.

Segmentation and backup processing

Review the following backup characteristics before you apply segmentation to your DB2® backup operations:

- The data stream that is sent from DB2® is segmented during a DB2® database backup.
- When implemented, segmentation is applied to every participating DB2® backup session.
- Back up and restore sessions are isolated from each other. As a result, segments that are generated by IBM Storage Protect™ for Enterprise Resource Planning are isolated on a per session basis. Therefore, segments cannot be mixed between different sessions. All segments that are backed up within the same session are restored in the same session.
- DB2® logs are not partitioned into multiple segments.

Segmentation and restore processing

Review the following restore characteristics before you apply segmentation to your DB2® restore operations:

- Metadata that is associated with the backup object indicates whether the object is part of a segmented data stream. If the backup object is part of a segmented data stream, IBM Storage Protect™ for Enterprise Resource Planning automatically joins the segments to the object DB2® expects to receive from IBM Storage Protect™ during the restore operation.
- Back up and restore sessions are isolated from each other. As a result, segments that are generated by IBM Storage Protect™ for Enterprise Resource Planning are isolated on a per session basis. Therefore, segments cannot be mixed between different sessions. All segments that are backed up within the same session are restored in the same session.
- Do not use segmentation into two or more segments for a backup that is to be restored to DB2® by using the Backup Object Manager command `backom -c r_raw ...`. This backup can be restored to the destination directory, but not into DB2® if two or more segments were created. If the backup was created

by using a single segment, it can be restored to DB2® from the destination directory after retrieval from IBM Storage Protect™. There is no limitation that concerns segmentation for other restore methods.

Restoring SAP data

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

Start restores in a nonpartitioned database environment

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

Every successful backup run generates a timestamp that is required for later restore operations. These timestamps are written to the DB2® Recovery History file (RHF), which can be queried with DB2® commands. The timestamps of backup images that are currently stored on the IBM Storage Protect™ server can be queried by using the Backup Object Manager query commands. If no timestamp is specified in a restore command, the latest backup image that is found on IBM Storage Protect™ is restored.

DB2® database and table space restores are performed as follows:

- Full restore to a certain point in time:

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

or

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

- Online table space restore

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

or

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

- Recovery History File restore

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

or

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

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

Processing redirected restore in automatic mode

Backup Object Manager provides an automatic cloning function, which creates an exact copy of the original SAP database in a different location.

About this task

The physical database layout of the target database is identical to that of the source system. The physical database layout consists of table spaces, table space number, and size of the table space containers. The path names of the new table space containers are constructed by replacing the original SID with the SID of the target system. In addition, modifications to the sizes of all or selected table space containers of the target database can be made to optimize the I/O performance. Backup Object Manager provides automated table space resizing and automated table space normalizing for these modifications. The Backup Object Manager automatic mode redirected restore function can be used to resize table space containers of the source database. This action is

accomplished by performing a redirected restore in automatic mode with the same SID set as both the original and the target SID and requesting scaling or normalizing (or both) during the operation.

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

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

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

1. Backup Object Manager retrieves the TDI for the requested backup from IBM Storage Protect™ into memory.
2. Backup Object Manager replaces the source database alias with the target database alias. If no target database alias is specified, Backup Object Manager uses the original database alias as the target database alias.
3. Using the modified TDI, Backup Object Manager performs basic plausibility checks.
4. Backup Object Manager uses the modified TDI to create the necessary table space containers on the target system. If the target database alias is the same as the original database alias, the database is restored to the original database alias and SID. When Backup Object Manager restores to the original system, it attempts to overwrite the original database. Overwriting the original database requires approval by the administrator.
5. Backup Object Manager calls the DB2® redirected restore function.

Tablespace definition information

To automate a redirected restore operation, Backup Object Manager requires information about the table spaces and the table space containers that are used in the original database.

The following information is used to create the table space containers of the target database, and is required for each table space:

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

The following information must be available for each table space container:

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

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

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

- Add or remove of table space containers from dedicated table spaces

- Modify names (locations) of table space containers
- Modify the size of a DMS table space container, whereby the sum of container sizes must have at least the number of pages that are used plus $((\text{number of containers} + 1) * \text{extent})$, where extent is the extent size in pages.
- Add an automatic storage path, if at least one automatic storage path is already present
- Change the location of an existing automatic storage path
- Remove one or more existing automatic storage paths, whereby in any case at least one automatic storage path must exist

tablespace

The following is a sample TDI file:

```
; IBM Storage Protect™ for Enterprise Resource Planning
; Data Protection for SAP(R) for DB2
; - Tablespace Definition Information (TDI) -
;
; The following TDI sections can be modified manually:
; - Automatic_Storage_Path
; - Container
;
; An automatic storage path section consists of the following format:
;
; Automatic_Storage_Path = path#1
; ...
; Automatic_Storage_Path = path#n
;
; It is possible to add or remove an automatic storage path entry. For already existing
; automatic storage path entries the assigned path can be updated.
;
; A tablespace section consists of the following format:
;
; [Tablespace ID "tbsp. name" type page size extent size in pages
;                                     used pages yes|no]
; Container[ID 1] = definition ;
; ...
; Container[ID n] = definition
;
; where the definition of a container statement is characterized by its tablespace:
; - SMS tablespace: "path"
; - DMS tablespace: file | "path/container name" | size in pages
;
; If the tablespace containers are modified manually (add or remove container,
; adjust container path or size) at least the following conditions have to be
; guaranteed for ensuring the TDI integrity:
; 1) Any new container specified requires empty brackets '[]'. The ID is calculated
;     internally.
;
; 2) Each tablespace block has to have at least one container specification
;
; 3) The sum of container sizes of a DMS tablespace has to have at least the number
;     of used pages plus  $((\text{number of containers} + 1) * \text{extent})$ .
;
```

```
; !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
; ! DO NOT EDIT ANYTHING ELSE EXCEPT THE SECTIONS !
; ! - Automatic_Storage_Path (if present) !
; ! - Container !
; !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
[TDI]
Version = 1.1
Generator = Data Protection for SAP(R) 6400

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

```
[Tablespace 0 "SYSCATSPACE" dms 16384 4 9264 yes]
Container[0] = file | "/db2/T01/sapdata1/db2t01/NODE0000/T01/T00000000/C0000000.CAT" | 8192
Container[1] = file | "/db2/T01/sapdata2/db2t01/NODE0000/T01/T00000000/C0000001.CAT" | 8192

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

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

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

The following details are related to the TDI file:

- The [TDI] header block is used to identify the data as TDI and holds some meta-information about it. The *Version* key holds the version of the TDI syntax. The *Generator* key denotes some product information.
- The [Backup] block holds various kinds of information about the database backup the TDI is associated with. This information must be kept within the TDI file so that it is available even when the file is renamed. [Backup] additionally includes the database path where database metadata is stored, and all automatic storage paths the database provides for table spaces supporting automatic storage. It is possible to add or remove an automatic storage path entry in that section. Optionally, for automatic storage path entries that already exist, the assigned path can be updated.
- The [Tablespace] block marks the start of the container definitions of a specific table space.
- The block header contains the following items in exactly this order: the ID of the table space, its name, its type, the page size in bytes, the extent size in pages and the number of used pages in the table space. Do not change any data within the table space block header.
- Each container statement defines one container of a table space according to the following rules:
 - The ID is denoted in square brackets if the line was written by the system. If a new container is to be added to a table space, the ID is not yet known. Therefore, the administrator specifies a new container without an ID, just entering consecutive brackets.
 - For an SMS table space, only the fully qualified path is specified.
 - For a DMS table space, the type, location, and size of the container are specified, in this order, and separated by a vertical bar (|). The type is given by one of the strings *file* or *device*. The size is interpreted as a number of pages unless a unit is specified. In this case, the unit is used.
 - Names of table spaces and paths must be quoted strings.

Processing redirected restore in batch mode

Backup Object Manager provides a redirected restore batch mode function where the TDI for the target database is modified before it starts the redirected restore.

About this task

The TDI image to be used must be available as an ASCII file in the file system. For example, a TDI image that is created during an interactive redirected restore can be used as target TDI for a redirected restore in batch mode. Batch mode can also be used for multiple redirected restores to different locations with identical changes of the physical database structure. As with the interactive mode, the original TDI is used to test whether the changes of table space container sizes and locations that are made are valid. In addition, modifications to the sizes of all or selected table space containers of the target database can be made to optimize the I/O performance. Backup Object Manager provides automated table space resizing and automated table space normalizing for these modifications.

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

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

This original TDI image can be renamed and modified.

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

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

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

1. Backup Object Manager replaces the alias that is specified in the target TDI file with the alias of the target database. If no target database alias is specified, Backup Object Manager uses the original database alias as the target database alias.
2. Backup Object Manager retrieves the original TDI from IBM Storage Protect™ and verifies whether the target TDI defines table space containers that are sufficient to replace the original table space containers.
3. Backup Object Manager uses the target TDI and the original TDI to run basic plausibility checks.
4. Backup Object Manager uses the target TDI to create the necessary table space containers on the target system. If the target database alias is the same as the original database alias, the database is restored to the original database alias and SID. When Backup Object Manager restores to the original system, it attempts to overwrite the original database. Overwriting the original database requires approval by the administrator.
5. Backup Object Manager calls the DB2® redirected restore function.

Redirected restore in interactive mode

Backup Object Manager interactive mode is a menu-driven dialog where the table space container layout is redefined by adding, deleting, moving, or resizing items.

About this task

Backup Object Manager compares the table space definitions that are entered in the menu dialog with the original database layout (as documented in the original TDI) and provides immediate feedback about potential configuration problems. In addition, modifications to the sizes of all or selected table space containers of the target database can be made to optimize the I/O performance. Backup Object Manager provides automated table space resizing and automated table space normalizing for these modifications.

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

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

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

1. Backup Object Manager retrieves the TDI for the requested backup from IBM Storage Protect™ into memory.
2. Backup Object Manager replaces the source database alias with the target database alias. If no target database alias is specified, Backup Object Manager uses the original database alias as the target database alias.
3. Backup Object Manager determines whether specific containers must be redefined.
4. Backup Object Manager displays the main menu, which shows a list of sorted table spaces for the database to be restored. A “!” mark in front of a table space or table space container indicates a warning about a potential problem. Although the redirected restore can still begin, the problem is to be resolved before proceeding. A “!!” character in front of a table space or table space container indicates an error was detected, such as a problem about their location or size. The redirected restore does not succeed until the error is first resolved.
5. The administrator can select table spaces or table space containers to be changed by using their IDs. When all modifications of the physical database layout are completed and no more errors (“!!”) are displayed. The redirected restore can be started by entering c from the main menu. The administrator can also end the redirected restore from any menu dialog by entering a.
6. When the -foption is specified during the redirected restore, the modified physical database layout of the target database is stored in an ASCII file in the file system. This file can be used later as input for a

redirected restore in batch mode at another location, where the same physical changes to the restored database must be applied.

7. Backup Object Manager uses the modified TDI and the original TDI to perform basic plausibility checks.
8. Backup Object Manager uses the modified TDI to create the necessary table space containers on the target system. If the target database alias is the same as the original database alias, the database is restored to the original database alias and SID. When Backup Object Manager restores to the original system, it attempts to overwrite the original database. Overwriting the original database requires approval by the administrator.
9. Backup Object Manager calls the DB2® redirected restore function.

Sample work flow for redirected restore

A sample work flow for a redirected restore with Data Protection for SAP Backup Object Manager is given.

About this task

To clone the SAP production database (PRD) to a test system (TST) on a different system, apply the following procedure:

1. Make sure that the administrator account to be used has the appropriate rights on the target system. An example of such rights is permission to allocate files of a size greater than 2 GB.
2. Verify that the source database PRD meets the prerequisites for a redirected restore operation.
3. Set up Data Protection for SAP on the target system. Verify that these environment variables specify these values:
 - **XINT_PROFILE** specifies the Data Protection for SAP profile.
 - **DB2_VENDOR_LIB** specifies the Data Protection for SAP shared library.
 - **TDP_DIR** specifies the path for the Data Protection for SAP process log files.
4. For the restore process, customize the Data Protection for SAP profile (`initTST.utl`) on the test system with these settings:
 - Use **BACKUPIDPREFIX** as specified on the source system: `PRD_---`
 - Use the IBM Storage Protect™ server that is specified on the source system. You might include adding the appropriate IBM Storage Protect™ server stanza to the client system options file (`dsm.sys`) on the test system.
 - Use the **ADSMNODE** specified on the source system.
 - Use **BRBACKUPMGTCLASS** as specified on the source system.
5. Issue the following command to record the password of the appropriate node on the IBM Storage Protect™ server:

```
backom -c password
```

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

6. Issue the following command to check the Data Protection for SAP database backup images on IBM Storage Protect™:

```
backom -c q_db
```

Verify that the TDI flag is set to yes for the backup image to be restored.

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

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

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

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

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

9. If the database is in rollforward pending mode and must be recovered, there are two possibilities for retrieving the required logs.
- Automatically by the DB2® Log Manager during the recovery process, or
 - Manually with **BackOM** before the DB2® rollforward process is started.

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

The following are the configuration parameters:

- **SRC_DB_INSTANCE**
- **SRC_DB_ALIAS**

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

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

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

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

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

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

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

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

10. To retrieve the log files, issue:

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

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

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

- **BRBACKUPMGTCCLASS**
- **BRARCHIVEMGTCLASS**

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

Attention: Do not back up the test system with the **BACKUPIDPREFIX** of the production system.

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

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

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

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

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

Redirected restore plausibility checks

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

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

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

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

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

DB2® redirected restore using Backup Object Manager

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

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

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

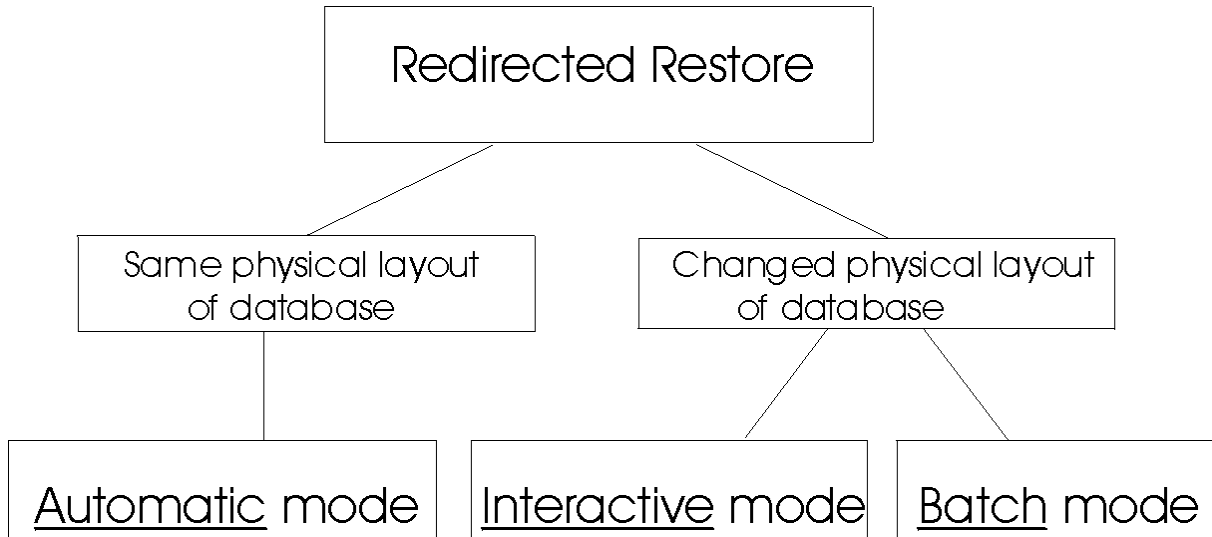


Figure 8: Redirected restore overview

Backup Object Manager provides these redirected restore modes:

Automatic

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

Batch

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

Interactive

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

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

Creating table space definition information

A table space definition image (TDI) is created after a full database backup completes successfully. It is stored on the IBM Storage Protect™ server with the database backup image. Only database backups for which a corresponding TDI is available can be used for redirected restore with Backup Object Manager.

About this task

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

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

A TDI is not created when you use the DB2® command-line interface to run an offline database backup.

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

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

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

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

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

The DB2® backup command can then be entered on the DB2® command-line interface:

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

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

Redirected restore prerequisites

The following requirements must be met for the Backup Object Manager to successfully run a redirected restore.

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

Tablespace definition information

To perform a redirected restore, Backup Object Manager requires information about the physical layout of the original database, such as the tablespace containers used by the original database.

The size of database managed containers (DMS) must be available to create new tablespace containers with sufficient space. Backup Object Manager keeps information about all tablespaces of a database that is backed up for every backup image on IBM Storage Protect™. The following information is collected for each tablespace:

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

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

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

The TDI is stored in ASCII format to allow for read and edit usability. For example, the number of used pages that are recorded in the TDI image can help identify the correct sizes to request when resizing containers. Backup Object Manager also calculates the number of total pages and used pages from the data that is stored for each tablespace container. Editing might be necessary when you request a redirected restore in batch mode.

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 9: A balanced configuration

- 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.

Server-related tuning

You can manage the data that is stored on the IBM Storage Protect™ server for IBM Storage Protect™ for ERP. You can manage which servers are used to store data.

Manage data on the backup server

The Data Protection for SAP Backup Object Manager can search for backup objects on the IBM Storage Protect™ server to restore or delete them.

For more information, see the *DB2® Backup Object Manager utility* topic.

ADDING INFO TO TEST CHECK IN

Lorem ipsum dolor sit amet, consectetur adipiscing elit. Integer in consequat ex. Aenean eget orci quis velit convallis sodales vitae vitae velit. Ut nulla leo, facilisis ac est a, laoreet placerat nunc. Praesent turpis orci, maximus posuere enim nec, ultrices tincidunt lectus. Aliquam in elit lacus. Quisque at erat vitae dolor suscipit faucibus. Donec tristique mollis tellus, eget cursus libero vulputate vel. Aliquam bibendum purus non fringilla vulputate.

Alternate network paths and servers

Multiple network paths and multiple backup servers can be specified. When the number of available sessions to multiple servers exceeds the number of sessions that are allowed, Data Protection for SAP uses the first sessions that it can establish.

Data Protection for SAP continues to use the number of sessions that are allowed as defined by the MAX_SESSIONS keyword. This setup allows data to be backed up even when a resource (such as an IBM Storage Protect™ server or its network interface) is unavailable. The servers that are used for the backup must be available to restore the data. The days of the week that a server is used can also be specified using the USE_AT keyword. For more information, see the *Profile parameter descriptions* topic.

Options

Use Data Protection for SAP options to tune performance.

Tune performance for Data Protection for SAP by using multiple sessions, network paths, servers, or through multiplexing and other options.

Performance options for Data Protection for SAP

Data transfer rates are impacted by the types of disks that are used for the databases, the network capabilities that are accessed by the database host and the IBM Storage Protect™ server, and the type of storage device that contains the backup.

Data Protection for SAP provides the following options to help optimize the data transfer rate for these components.

Parallel and multiple sessions

Data Protection for SAP can back up or restore data to multiple tape drives in parallel. Parallelism is achieved by using more than one session to send data to a backup server.

Multiple and parallel network paths and servers

Improve performance by configuring Data Protection for SAP to distribute a database backup across two or more IBM Storage Protect™ servers. In addition, you can balance network traffic by providing two (or more) separate network connections between the SAP database host and the IBM Storage Protect™ server.

Incremental and delta backups

Data Protection for SAP supports incremental and delta backups of DB2® databases. Depending on the system environment, incremental backups might decrease backup processing time.

RL_COMPRESSION

The RL_COMPRESSION profile keyword compresses a partially filled database. This process can result in reduced network traffic and fewer tapes that are required for backup.

Adjustments for improving performance of data transfer

Data Protection for SAP is configured to send uncompressed data to the IBM Storage Protect™ server that uses a single session.




Figure 10: Data transfer for a backup and restore

A single configuration that is best for all environments is not possible or realistic. However, the information that is provided here can help in determining which configuration is best for your environment.

Buffer copies

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

Data Protection for SAP uses internal buffers to store and exchange data with the IBM Storage Protect™ 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.

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 IBM Storage Protect™ 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 IBM Storage Protect™ server by compressing zero-byte blocks (RL_COMPRESSION profile keyword).




Figure 11: Null Block Compression

Data Protection for SAP can decrease the amount of data that is sent to the IBM Storage Protect™ 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.

Optimize for data deduplication in DB2

You can request DB2® to optimize the format of backup images for target storage devices that support data deduplication.

When backup data is stored on locations that perform data deduplication, for example IBM Storage Protect™ container storage pools, streams of unchanged data must be created in the same way for each backup. When command line option **-o** is passed to **backom**, DB2 is requested to optimize the format of the backup images for target storage devices that support data deduplication.

Automation options

You can improve administrative productivity by using the Data Protection for SAP automation options.

Selectable management classes

Specify different IBM Storage Protect™ 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 Db2® log files. For more information about specifying management classes, see the *Profile parameter descriptions* topic.

Retain backups by version

Retaining backups by version, limits the number of full backups that are retained on the IBM Storage Protect™ server. When the number of full backups on the server exceeds the specified number, the oldest version is deleted. Retaining backups provides a trace of all Db2® log files, and all incremental and delta backups, that are associated with a full backup. All these objects are removed together with the full backup.

Important: If a backup is created when the profile parameter **MAX_VERSIONS** is set to zero, this backup is excluded from the backup versions processing. It is not considered when counting the number of backup generations, and it is not deleted when it becomes older than the backups that are retained.

Multiple redo log copies

Backing up multiple copies of a log file in a single archive operation helps protect data in the event of tape defects or disaster recovery situation. These copies can be on different physical IBM Storage Protect™ volumes or different IBM Storage Protect™ 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 IBM Storage Protect™ servers. Also using multiple network connections to the IBM Storage Protect™ server can help. 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 IBM Storage Protect™ 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 IBM Storage Protect™ 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 12: 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 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 sensors

Data Protection for SAP uses sensors that observe incoming and outgoing data streams. They measure throughput and the idle time of the I/O threads in comparison to the duration of the backup. This function provides a way to determine whether the streams of incoming and outgoing Data Protection for SAP data are balanced.

The method of transferring data packages depends on how IBM Storage Protect™ is configured. In a standard configuration, the data packages are sent from the IBM Storage Protect™ API client through the network to the backup server. In an environment that is configured for LAN-free operations, the data packages are processed by the IBM Storage Protect™ API client and the IBM Storage Protect™ Storage Agent.

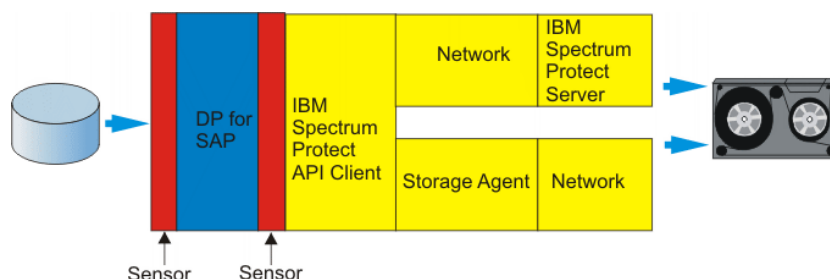


Figure 13: Performance optimizing by using sensors

When a backup operation begins, filling the buffers is necessary before the effects of a bottleneck are viewable.

Performance tuning for data transfer

During data transfer, a continuous stream of data is generated between the SAP database server, the network, and the IBM Storage Protect™ 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 IBM Storage Protect™ server, which is also referred to as a backup server.




Figure 14: Data Protection for SAP data transfer

Multiple servers

Data Protection for SAP supports multiple servers, which can distribute backup data among two (or more) backup servers. This feature helps eliminate constraints that are frequently encountered among backup servers.




Figure 15: Multiple servers

A server statement must be entered in the Data Protection for SAP profile for each adapter of the backup server. A description for the SERVER keyword is given in the *Profile parameter descriptions* topic. The value of the MAX_SESSIONS keyword is not greater than the sum of all SESSION values specified for the SERVER statements that are used concurrently.

Multiple sessions

You can use multiple tape drives simultaneously to increase the transfer rate to or from the IBM Storage Protect™ server. Several backup sessions access the database in parallel on the database server, and the data is written simultaneously to several tape drives.




Figure 16: Parallel (multiple) sessions

The Data Protection for SAP DB2® keyword MAX_SESSIONS is used to define the number of parallel sessions to be established with the IBM Storage Protect™ server for database backup, archive (backup of log files) and restore. When you run a database backup, the data is typically written directly to tape drives on the IBM Storage Protect™ server. The parameter that is specified in the MAX_SESSIONS keyword must match the number of tape drives that are used simultaneously.

These must be available to the management class defined as BRBACKUPMGTCLASS in the Data Protection for SAP profile. When you set up the IBM Storage Protect™ server, make sure not to activate collocation in the (tape) storage pool that is defined for the management class that is chosen as BRBACKUPMGTCLASS. In addition, ensure that as many tape drives for this management class are available as the number of sessions that are defined in MAX_SESSIONS. Multiple access to the same tape might slow down data transfer.

These must be available to the management class defined as BRARCHIVEMGTCLASS in the Data Protection for SAP profile. If you are using tape pools as primary pools for this management class, this consideration for database backups also applies to disk storage pools. Several DB2® log archive sessions can simultaneously use one or two independent disk storage pools.

The number of storage pools that are required depends on the number of backup copies that are requested for a DB2® log file.

Multiplexing

Multiplexing uses parallel access points to data on the database server. This configuration is useful when tape drives are used during database backup operations on the backup server.




Figure 17: Multiplexing

This feature is provided by the **PARALLELISM** parameter available with the **BACKUP DATABASE** and **RESTORE DATABASE** commands. Refer to your *DB2® Command Reference* for details about these commands and the **PARALLELISM** parameter.

Multiple network paths

Data Protection for SAP allows multiple network connections (paths) for data transfer between the database server and the backup server. Parallel paths can be used to eliminate network points of constraint.




Figure 18: Parallel (multiple) paths

For each additional path, more network adapters are required on both the production and the backup server. A server statement must be entered in the Data Protection for SAP profile for each adapter of the backup server. This scenario is described for the SERVER keyword in the *Profile parameter descriptions* topic. The value of the MAX_SESSIONS keyword is not greater than the sum of all SESSION values specified for the SERVER statements that are used concurrently. Detailed information about setting up multiple parallel network paths is described in the *Parallel backup paths and backup servers* topic.

Storage space

You can manage aspects of storage space to tune the performance.

Automated tablespace adaptations

Backup Object Manager can adapt the sizes of tablespace containers when it creates the containers of the target databases during a redirected restore operation.

For example, tablespace container sizes might be increased to provide more space or decreased to use storage more efficiently. Tablespaces can also be allocated with similar sizes to make parallel I/O operations more efficient. These features are supported by Backup Object Manager resizing and normalizing functions.

Tablespace normalization

To achieve optimal parallel I/O operation performance for a database, all containers of a tablespace are to be the same size. During tablespace maintenance, containers might be added or extended which creates different container sizes. As a result, data is unevenly distributed among the containers, which can result in decreased parallel I/O operation performance during table scans (sequential prefetching).

Backup Object Manager provides an automated tablespace normalizing function that allows the location and size of tablespace container to be redefined. This function also helps prevent I/O-intensive tablespace rebalancing that can be detected by DB2®.




Figure 19: Tablespace Normalizing

This graphic shows that the original system tablespace consists of different-sized containers. Although sequential prefetching allows three processes to simultaneously read the data during table scans, the different container sizes and uneven data distribution prevent parallel I/O during part of the scan. The degree of parallelism decreases over time. To counteract this situation, adjust the containers for each tablespace so that they are the same size. This type of adjustment, in combination with a redirected restore operation, requires no further resizing tasks for the tablespace containers after the restore completes.

Backup Object Manager simplifies the tablespace container resizing process by providing an automatic tablespace normalizing function. It can be used in combination with any mode of redirected restore facility by specifying the `-N` option. Issue this command to resize each container of a tablespace to the average size of all containers within the same tablespace during a redirected restore:

```
backom -c rr_db_type -aDB2 source alias,DB2 target alias -t timestamp -N
```

After the redirected restore completes successfully, all containers of a tablespace of the target database are the same size. As a result, the continuous parallel I/O performance of the physical layout of the restored database is optimized.

Scaling tablespace containers

Sufficient free space must be available in a tablespace for the database to function properly. Backup Object Manager provides an automated table space scaling function that allows the location and size of table space container to be redefined. This function also helps prevent I/O-intensive table space rebalancing that can be detected by DB2®.

About this task

This graphic shows that 98 percent of the SYSCATSPACE table space of the original system is being used. Disk 1 has 4 percent free space while 100 percent of the second container is used. The free space available in a table space can be increased as part of the redefinition feature during a redirected restore. The goal is to achieve an overall filling rate for the target side of 70 percent. This rate can be achieved by manually increasing the free space that the first container must provide at 20 percent and 40 percent for the second container. This type of adjustment, in combination with a redirected restore operation, requires no further resizing tasks for the table space containers after the restore completes.

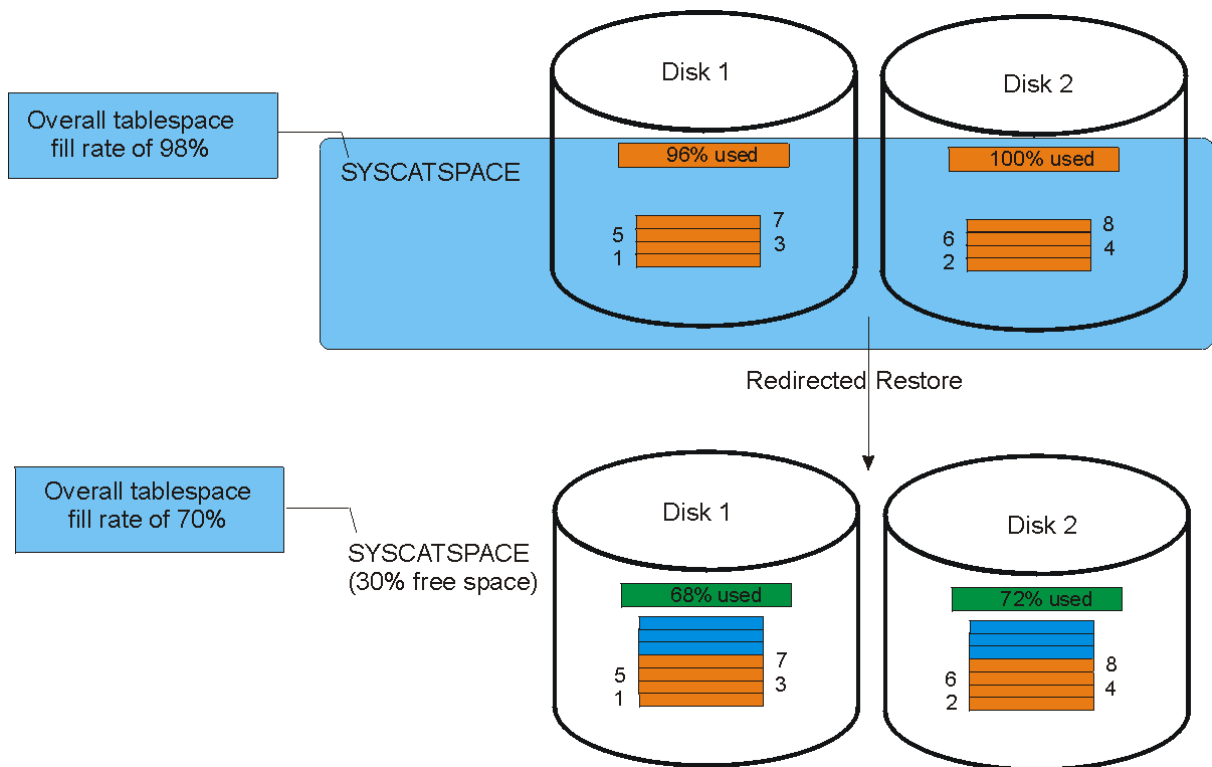


Figure 20: Tablespace scaling

Backup Object Manager simplifies the table space resizing process by providing an automatic table space scaling function. It can be used in combination with any mode of redirected restore by specifying the `-S` option with a floating-point sizing factor. Consider these factors when you resize table spaces:

- A value of '1' indicates that the target table space is 100 percent the size of the original (nothing is changed).
- A value greater than one increases the target table space. For example, a value of 1.1 increases the target table space by 10 percent to a target value of 110 percent of the original.
- A value less than 1 decreases the target table space. For example, a value of 0.9 decreases the target table space by 10 percent to a value of 90 percent of the original.

Therefore, manual adaptation of the table space containers can be replaced by the following procedure by using the Backup Object Manager redirected restore function:

1. Issue this Backup Object Manager query to determine the original fill rate of the table space:

```
backom -c q_tdi -a DB2 source alias -t timestamp -m detailed
```

2. Calculate the table space scaling factor by using this formula:

```
scaling factor = original fill rate / new fill rate
```

For example:

```
scaling factor = 0.98 / 0.7 = 1.4
```

3. Issue this command to begin the Backup Object Manager redirected restore:

```
backom -c rr_db_type -a DB2 source alias,DB2 target alias -t timestamp ...  
...-T SYSCATSPACE -s 1.4
```

After the redirected restore completes successfully, the SYSCATSPACE table space on the target side is increased by 40 percent during table space container redefinition. The new overall fill rate of the SYSCATSPACE in the target database is now 70 percent.

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

Investigate and compare the log files of the application and the IBM Storage Protect™ server activity log to find out the differences between successful and unsuccessful operations. The log file names are provided for reference:

- `tdpdb2.SID.nodename.log`
- `db2diag.log`

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.

Troubleshooting common problems

Compare the log files of the application in question (`tdpdb2.SID.nodename.log`, `db2diag.log`, and the IBM Storage Protect™ server activity log) to find out the differences between successful and unsuccessful operations.

About this task

When problems occur, look for one of the following patterns:

- The problem always occurs at the same time. If this condition is true, view the appropriate log files to determine whether there are any scheduled processes that occur 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 check 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 changed.
- One or more of the operating system, network, database, or hardware components changed.
- Patches or updates to one or more of the components were applied.
- Changes occur that originate by the system:
 - Check whether the disks are running full with the UNIX™ `df` command.
 - If network performance decreases, check whether additional hosts, additional applications, or defects in software or hardware occurred. Compare the log files. The log file names are provided for reference:
 - `tdpdb2.SID.nodename.log`

- db2diag.log
- If IBM Storage Protect™ server processing decreases, check whether more clients or more operations were added. Information is also available in the IBM Storage Protect™ server activity log.

If none of these changes caused the problem, view the last modified time stamp of the configuration files (vendor.env, initSID.utl, dsm.sys, dsm.opt, /etc/services, /etc/inittab, ...) command lists all files in the /etc directory, which are modified during the previous five days:

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

If you are able to 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 supports both IPv4 and IPv6 for internal communication in that it runs in IPv4, IPv6, and mixed environments on AIX® and Linux™. However, these products do not use IPv6. In a mixed environment, the communication depends on the adapter network settings. There is no option to enforce the use of a specific protocol other than by network configuration. Specifically, the ProLE or acsd service 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 to other systems 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.

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.

Review these considerations to better understand the installation setup on UNIX™ or Linux™ systems:

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

```
pd64:2:respawn:/usr/tivoli/tsm/tdp_r3/db264/prole -p tdpr3db264
Server component hostname 5126
```

```
pd64:2:respawn:/usr/tivoli/tsm/tdp_r3/db264/prole
-p tdpr3db264 Server component hostname 5126
```

Review these considerations to better understand the installation setup on Windows™ systems:

- Make sure that all files are installed.
- Verify that service ProLE Service is running and set to automatic startup. If this service is not running, Data Protection for SAP does not function properly.
- The installer adds lines to the %SYSTEMROOT%\system32\drivers\etc\services file similar to the example:
(Data Protection for SAP for DB2)

```
tdpr3db264      57324/tcp
```

The example shows the Data Protection for SAP 64-bit port.

- Make sure the Data Protection for SAP configuration file initSID.utl is in the directory pointed to by the TDP_DIR environment variable.
- Make sure the Data Protection for SAP configuration file initSID.utl is in the directory pointed to by the TDP_DIR environment variable (on DB2®).

- The vendor environment file `vendor.env` must contain the fully qualified path and file name of the `initSID.utl` file for Data Protection for SAP for DB2®.
- The vendor environment file `vendor.env` is to contain the path of the location where the Data Protection for SAP for DB2® runs logs that are written. If this location is not specified, a temporary directory of the system is used.
- The names of the IBM Storage Protect™ servers that are specified in `initSID.utl` must match the names in the `dsm.sys` file. If the IBM Storage Protect™ API or IBM Storage Protect™ 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.

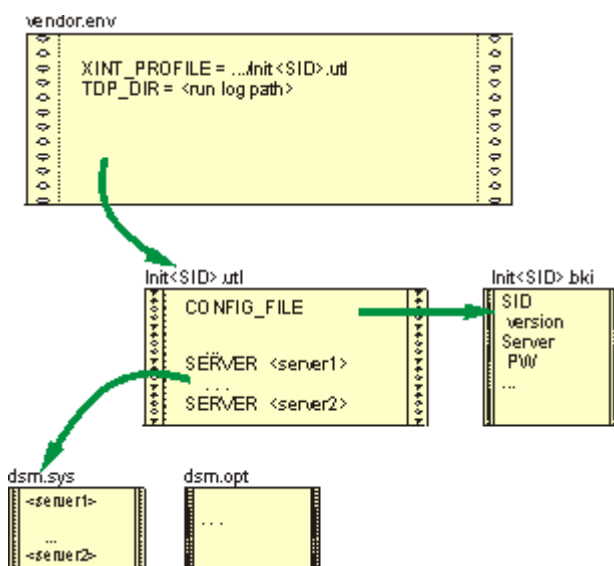


Figure 21: SAP and Data Protection for SAP for DB2® configuration files on UNIX™ or Linux™

On UNIX™ or Linux™ systems, the names of the IBM Storage Protect™ servers that are specified in `initSID.utl` must match the names in the `dsm.sys` file. If the IBM Storage Protect™ API or IBM Storage Protect™ 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 `initSID.bki` file as this file is the only one to which Data Protection for SAP writes persistent information.

On Windows™ systems, the `dsm.opt` file is used instead of the `dsm.sys` file. However, the content of this file is not relevant to Data Protection for SAP. The directory that contains the `dsm.opt` file must also contain a `server.opt` file for each server that is specified in the `initSID.utl` file. The environment variable `DSMI_CONFIG` must specify an option file within this directory. `DSMI_CONFIG` is to specify the `dsm.opt` file in this directory. The `DSMI_DIR` environment variable must also specify the directory where the IBM Storage Protect™ API message text file is in. This example shows a typical the `c:\Program Files\Tivoli\tsm\api64` directory.

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 version level.
- The IBM Storage Protect™ server version.
- The IBM Storage Protect™ server operating system level.
- Data Protection for SAP configuration file (`vendor.env`, `initSID.utl`) including IBM Storage Protect™ 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.

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.

The `backint.log` log file contains the IBM Storage Protect™ for ERP data for all database and redo log file backup and restore operations that complete successfully or fail.

These files are in the path indicated by the `TDP_DIR` environment variable. After the installation, `TDP_DIR` points to the subdirectory `tdplog` of the path for the Data Protection for SAP configuration files. If `TDP_DIR` is not set, or if a log file cannot be created in the path pointed to by `TDP_DIR`, the log files are created in path `/tmp` (UNIX™ or Linux™) or in the path pointed to by environment variable `TEMP` (Windows™).

The Data Protection for SAP shared library writes to the `tdpdb2.SID.node name.log` log file.

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

These files are in the following paths:

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

Windows:

- `%SAPDATA_HOME%\sapbackup` for backup and restore runs
- `%SAPDATA_HOME%\saparch` for redo log archive runs

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

- `backSID.log`
- `restSID.log`
- `archSID.log`

Troubleshooting problems

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

About this task

The following figure helps to isolate problems that occur when you back up or restore your DB2® database.

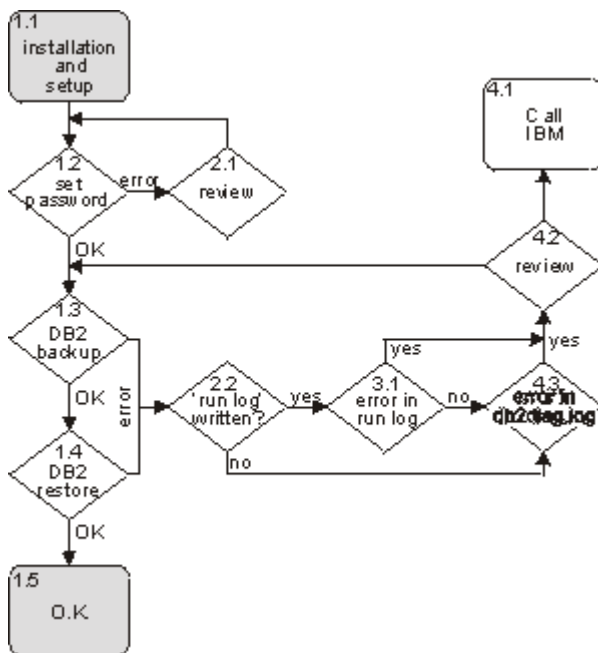


Figure 22: General problem isolation

After installation is completed (Step 1.1) and manual password handling is specified, set the password (Step 1.2). When the operation completes successfully, the informational messages BKI0051I: Password successfully verified for node *NODENAME* on server *SERVERNAME* and BKI0024I: Return code is: 0. display for each server that is configured within the *initSID.utl* file. An error message displays when a problem occurred.

These errors are frequently encountered at Step 1.2:

BKI2001E: Socket error while connecting to ProLE at IP-Address:PORT: Connection refused

On Windows™, verify that the ProLE Service is running by viewing the Computer Management Services screen or issue this command:

```
net start
```

A list of all running services displays. On UNIX™ or Linux™, verify that the background daemon is running by issuing this command:

```
ps -ef | grep prole
```

Check the entry in */etc/services* (UNIX™ or Linux™) and *%SYSTEMROOT%\system32\drivers\etc\services* (Windows™). Compare the port number from the error message with the port number within */etc/services*. Also, check the entry in */etc/inittab* (UNIX™ or Linux™). If another port was set by using the option *-pPORT*, check this port as well. If all of this effort does not help, start the ProLE from another shell on UNIX™ or Linux™ with this command:

```
prole -p PORT
```

Issue this command on Windows™:

```
prole -console -p PORT
```

Attempt to start **backom** again.

BKI5001E: IBM Storage Protect™ Error: Server not found in configuration file

On UNIX™ or Linux™, the IBM Storage Protect™ server that is defined in the *initSID.utl* file does not match the server that is specified in the *dsm.sys* file. On Windows™, the *server.opt* file might be missing.

BKI5001E: IBM Storage Protect™ Error: ANS1353E (RC53) Session rejected: Unknown or incorrect ID entered

This message can display when the node in the server stanza of the UTL file is not valid on the server.

HANG

If **backom** hangs after the password is entered, the server IP address that is specified in the UNIX™ or Linux™ `dsm.sys` file might be incorrect.

When Step 1.2 (setting the password) is successful, proceed to Step 1.3 and do a backup by using the DB2® backup command. Verify that the settings are correct. If the backup was successful, a message from DB2® is displayed:

```
Backup successful. The timestamp for this backup image is: timestamp
```

If an error message displays, find the message and information about how to resolve the error.

When an error occurs, always view the Data Protection for SAP run log `tdpdb2.SID.nodename.log` first. This log file is in the directory that is specified by the **TDP_DIR** environment variable. If the variable is not specified, the log file is in the system temporary directory. If the `tdpdb2.SID.nodename.log` file does not exist (Step 2.2), then two cases are possible. The first case is that DB2® was unable to load the shared library that contains the DB2® connector for Data Protection for SAP. The second case is that an error was encountered before the Data Protection for SAP library was called. In both cases, a DB2® error message is displayed on the command line. It begins with the SQL prefix and is also written in the DB2® diagnostic log `db2diag.log` (Step 4.3). DB2® provides detailed error descriptions by issuing this command:

```
db2 ? SQLnumber
```

Replacenumber with the appropriate message number. Try to resolve this problem by using the DB2® documentation.

If the `tdpdb2.SID.nodename.log` file exists, search for a message beginning with BKIXXXXY where XXXX is a four-digit number, and Y is the letter I, W, or E. When such a message occurs, the DB2® connector for Data Protection for SAP that loaded correctly is called by DB2®. In Step 3.1, the `tdpdb2.SID.nodename.log` file is created and an error message that starts with BKI is recorded.

Location of log files

Text that is displayed on the screen during DB2® backup, DB2® restore, and BackOM operations are typically written to log files.

DB2® also writes messages of internal operations, events, or status in the administration notification log file (`db2SID.nfy`) and diagnostic log file (`db2diag.log`). These log files are in the directory that is specified with the DB2® database management configuration parameter **DIAGPATH**. Query the DB2® database management configuration with this command:

```
db2 get dbm cfg
```

DB2® vendor reason codes

Data Protection for SAP uses these reason codes, which might also be displayed or logged by DB2® in the case of problems.

Table 7: DB2® vendor reason codes		
Reason code	Explanation	User response
1	The library that is specified could not be loaded.	Check the DB2® diagnostic log for further details.
2	Communication error between shared library and ProLE	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.
6	Object that is specified cannot be found on IBM Storage Protect™.	There is no backup image on IBM Storage Protect™ matching the search criteria.
10	Invalid options that are specified with the options parameter of the DB2® backup/restore command.	Check the options string that is specified and check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.

Reason code	Explanation	User response
11	Initialization procedure for shared library failed.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> and the DB2® diagnostic log file for further details.
17	During end processing of either backup/archive or restore/retrieve session or sessions, an error occurred.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.
18	An error occurred during reading or writing data from or to IBM Storage Protect™.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.
26	An error occurred during deleting data from IBM Storage Protect™.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> for further details.
29	A terminate request from DB2® could not be handled correctly.	Check the Data Protection for SAP run log file <code>tdpdb2.SID.nodename.log</code> and the DB2® diagnostic log file for further details.
30	A severe error occurred.	Check the DB2® diagnostic log file for further details.

Reference information

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

Backups and restores in partitioned database environments

DB2® provides the Single System View (SSV) function, which allows the backup of all partitions to be triggered with a new dedicated DB2® backup command option **ON ALL DBPARTITIONNUMS**. Furthermore, partitioned database backups and restores can be made by using the already established DB2® framework for partitioned databases called **db2_a11**.

Partitioned database backups and restores can be made by using the already established DB2® framework for partitioned databases called **db2_a11**. A requirement of DB2® is to back up and restore the catalog partition separately from all other DB2® partitions. Thus, backup and restore operations of a partitioned database are two-step scenarios, whereby the first step is to backup/restore the catalog partition and the second step to back up and restore all other partitions in parallel. For more information about DB2® backup operations, see [“Use DB2 single system view for backup” on page 0](#).

Data Protection for SAP uses the DB2® command **db2_a11** for separation and parallelization of the backup and restore commands. The **db2_a11** command provides special characters for handling partitions and for running commands in parallel or sequentially. The DB2® Administration Guide contains further information.

The DB2® **db2_a11** command supports the following special characters or character combinations:

- **<<+N<** Runs a command only on partition N
- **<<-N<** Runs a command on all partitions except on partition N
- **"** Substitute occurrences of () by the machine index and substitute occurrences of ## by the partition number
- **;** Runs the commands in parallel in the background and finishes the command after all remote commands are completed.

By using these characters, each partition of the database can be backed up or restored with its special adapted environment.

This list presents possible partitioned database backups and restores with Data Protection for SAP for DB2®:

Full offline backup

Full offline backup of all database partitions with two sessions, starting with the catalog partition (shown as partition 0) followed by all other partitions in parallel:

```
db2_a11 '<<+0< db2 backup db SID load shared library open 2 sessions'  
db2_a11 '<<-0<; db2 backup db SID load shared library open 2 sessions'
```

Full restore of latest backup

Full restore of latest backup starting with the catalog partition (shown as partition 0) followed by all other partitions (shown with three partitions) in parallel using different number of sessions for some of the partitions. Before you start restoring, a temporary environment script (this example uses `/db2/SID/EEEenv.sh`) must be created. The `/db2/SID` directory must be an NFS share between all hosts where the partitions are located.

```
export SESSION0=2  
export SESSION1=2  
export SESSION2=4  
export SESSION3=4
```

Partitions 0 and 1 are restored with two sessions, and partitions 2 and 3 is restored with four sessions. As a result, the following restore command must be used:

```
db2_a11 '<<+0<" ./db2/SID/EEEenv.sh; db2 restore db SID load shared library open  
$SESSION##'  
db2_a11 '<<-0<" ./db2/SID/EEEenv.sh; db2 restore db SID load shared library open
```

```
$SESSION##'
```

The string '\$SESSION##' in the **db2_a11** command is replaced while running it with the value provided by the environment variables SESSION0 to SESSION3.

Full online backup

Full online backup of all database partitions, starting with the catalog partition (shown as partition 1) followed by all other partitions in parallel using different Data Protection for SAP profiles for each partition. To support this scenario, one Data Protection for SAP profile (`initSID.ut1`) must be created and maintained for each partition. Each profile can have different settings for the IBM Storage Protect™ node and management class. The following profiles are needed for this example:

- `initSID.ut1.1` (profile for partition 1)
- `initSID.ut1.2` (profile for partition 2)
- `initSID.ut1.3` (profile for partition 3)
- `initSID.ut1.4` (profile for partition 4)

In addition, a Data Protection for SAP vendor environment file (`vendor.env`) must be created and maintained for each partition. Each vendor environment file has an `XINT_PROFILE` entry that refers to the corresponding Data Protection for SAP profile. For example:

```
vendor.env.1 (environment for partition 1 ->  
XINT_PROFILE=/db2/SID/vendor.env.1)
```

```
vendor.env.2 (environment for partition 2 ->  
XINT_PROFILE=/db2/SID/vendor.env.2)
```

```
vendor.env.3 (environment for partition 3 ->  
XINT_PROFILE=/db2/SID/vendor.env.3)
```

```
vendor.env.4 (environment for partition 4 ->  
XINT_PROFILE=/db2/SID/vendor.env.4)
```

As a result, the following backup command must be used:

```
db2_a11 '<<+1<" db2 backup db SID online load shared library options  
/db2/SID/vendor.env.##'  
db2_a11 '<<-1<," db2 backup db SID online load shared library options  
/db2/SID/vendor.env.##'
```

The string `vendor.env.##` in the **db2_a11** command is replaced while running it with the values `vendor.env.1` - `vendor.env.4`, whereby the special characters '`##`' is substituted by the corresponding partition number.

Backup Object Manager

The Backup Object Manager manages DB2® log files that are archived with Data Protection for SAP, BR*Tools BRARCHIVE, and files that are archived with Data Protection for SAP and the DB2® Log Manager.

The following shows the Backup Object Manager syntax:

`backom [-?]` displays the syntax help.

Note: For the C shell, enclose the option string in single quotation marks (`backom '-?'`).

`backom -h [password|query|backup|restore|delete]` displays the command online help.

```
backom -c command [ command option ...]
```

Where '`command`' is one of the following choices:

Select for Password: password

Select for Query: q_all | q_db | q_ts | q_log | q_tdi | q_raw

Select for Backup: b_db

Select for Restore: r_db | r_ts | r_log | r_hfile | r_raw | r_tdi | rr_db_interactive | rr_db_batch | rr_db_clone

Select for Delete: d_db | d_ts | d_log | d_raw

Where 'command option' is one of the following choices:

Example

```
-i instance
-a alias name
-n node number
-u userid
-p password
-t timestamp | timerange
-l log number | log number range
-k log chain | log chain range
-f file name
-d destination directory
-e execution profile
-b buffer size
-s scaling factor
-N
-S sessions
-B number of buffers
-P parallelism
-D target database
-T tablespace
-R full | incremental | delta
-O
-L
-C
-x
-v
-m output mode
```

Backup Object Manager commands

There are six types of Backup Object Manager commands:

- Password command
- Query commands
- Backup command
- Restore commands
- Delete commands

Optional command options are listed in brackets [], parameter descriptions, that must be replaced, are listed in angle brackets < >.

Backup Object Manager command options

Use the following options together with Backup Object Manager commands.

- *-a database alias* or *-a original database alias,target database alias*. Denotes the name of the database for which an operation is requested. For a redirected restore to a different database or of database cloning, the database aliases of both the original and the target databases must be specified and separated by a comma. When a redirected restore is requested that specifies a single database alias, the database is restored to the original database.
- *-b buffer size* Denotes the size of DB2® backup or restore buffers, in 4-KB allocation units (pages). The minimum is eight units. The buffer size is limited by the memory available.
- *-B number of buffers* Denotes the number of DB2® buffers to be used for backup or restore. The minimum number is 2. The number of buffers is limited by the available memory.

- **-C** If specified for a redirected restore, this option indicates that the Backup Object Manager is to run only a test of the setup but not start copying data.
- **-d destination directory** Denotes the destination path for restoring a file to the file system.
- **-D target database directory** For a redirected restore, this option denotes the fully qualified name of the target database directory. This command option is ignored when the database alias of the target database is the same as the database alias of the original database.
- **-e execution profile** Denotes the complete path of the IBM Storage Protect™ for Enterprise Resource Planning for DB2® profile to be used with the Backup Object Manager. This option overrides the profile name that is set in the **XINT_PROFILE** environment variable.
- **-f file name** Denotes the name of a file in the file system. Unless the file denotes a TDI image, the following wildcard characters are accepted: ??? denotes any single character. * denotes any number of any characters.
- **-i DB2 instance** Used in query commands to limit the database or table space data to be displayed to a specific DB2® instance. With all other commands, this command option is used to override the DB2® instance name that is defined in the **DB_INSTANCE** environment variable.
- **-k log chain | log chain range** where *log chain range = chain1 - chain2*. Denotes the log chain number or numbers of DB2® log file or files. DB2® log chains can be specified either in the format *Cnnnnnnn*, where *nnnnnnn* is a string of seven decimal digits or in the format *mmmmmmm*, where *mmmmmmm* is a string of up to seven decimal digits that denote the log chain number.
- **-l log number | log number range** where *log number range = log number 1 - log number 2*. Denotes the log serial numbers of DB2® log files. DB2® log numbers can be specified either in the format *Snnnnnnn .log* (DB2® log file name), where *nnnnnnn* is a string of seven decimal digits, or in the format *mmmmmmm*, where *mmmmmmm* is a string of up to seven decimal digits that denote the log serial number.
- **-L** If specified for a database backup, the DB2® log files are saved to IBM Storage Protect™ with the database backup.
- **-m output mode** where *output mode = short | normal | detailed* denotes the detail of information that is requested with a query command. The default is “short” for information that is related to DB2® log files, “normal” for all other kinds of information.
If you must override the default values generally, you might set environment variables **FULL_OUTPUT** (for information about database backups), **TABLESPACE_OUTPUT** (for information about table space backups) and **LOG_OUTPUT** (for information about DB2® log file backups) to the values wanted.
- **-N** If specified, this command option causes all containers of a table space to be allocated with the same size during a redirected restore.
- **-n node number** Denotes the DB2® node number. For the password command in a DB2® partitioned environment: If for only one DB2® node/partition the password must be set/changed, specify the command option *-nnode number*. If the node/partition number is not specified, the new password is saved to all available node/partitioned based Data Protection for SAP configuration files. For all other commands: If the node number is not specified, node **NODE0000** is assumed.
- **-o** If specified for a backup, DB2 is requested to optimize the format of the backup images for target storage devices that support data deduplication.
- **-O** If specified when you request a backup or table space restore operation, an online backup or an online table space restore is run.
- **-p password** The password of the user ID specified in option **-u**.
- **-P parallelism** Denotes the degree of parallelism within DB2®, which is the number of buffer manipulator processes reading from or writing to table spaces at the same time. The minimum parallelism is 1. The maximum is 1024.
- **-R backup type** where *backup type = full | incremental | delta*. Denotes the type of backup requested. If no backup type is specified, a full backup is done.
- **-S sessions** Denotes the number of I/O sessions that are to be started by DB2®. The value of this command option must be less than or equal to the value of the keyword **MAX_SESSIONS** in the Data Protection for SAP profile.

- `-s scaling_factor` Denotes the positive floating point factor to be used for resizing all containers of a table space during redirected restore. The default is 1, indicating that the new table space is exactly the size of the original.
- `-t timestamp | time range` where `time range = timestamp1-timestamp2` Denotes the time when a backup object was created. For database and table space backups, this time stamp matches the time stamp that is listed in the DB2® Recovery History File. It consists of 14 decimal digits and has the format: `yyyymmddhhmmss` where `yyyy` is the year; `mm` is the month of the year, 01 - 12; `dd` is the day of the month, 01 - 31; `hh` is the hour of the day, 00 - 23; `mm` is the minute of the hour, 00 - 59; `ss` is the second of the minute, 00 - 59. For **restore** commands, an exact time stamp must be given. For **query** and **delete** commands, a time range can be specified, or the time stamp can contain wildcard characters. The following wildcard characters are accepted:
 - `?`denotes any single digit
 - `*`denotes any number of any digits.

If a time stamp is not specified for a query, the result contains all eligible backup object available on the IBM Storage Protect™ server. If a time stamp is not specified for a restore, the newest object is retrieved from IBM Storage Protect™.

- `-T tablespace list` where `tablespace list = tablespace[,tablespace list]` For a backup request, denotes the names of the table space or spaces to be backed up. Table space names are separated by commas. If there is no table space list that is specified, a full database backup is done.
- `-u userid` Denotes the DB2® user ID used for backing up or restoring a DB2® database, table space, or recovery history file if it is different from the current login user ID.
- `-v`If set, all log messages are displayed on STDOUT.
- `-x`If specified, this option suppresses all confirmation requests. Otherwise, confirmation requests are sought for restore commands that would overwrite existing data, and for delete requests. With the `-c passwordoption`, `-x`causes the password to be changed on all database partitions.

Backup command

Use the Backup Object Manager backup command, to back up a complete database or selected table spaces of a database.

For more information about command options, see [“Backup Object Manager command options” on page 0](#).

- Back up the database data with these command options:
`backom -c b_db -a database alias [-T tablespace list][-R backup type][-i instance] [-n node number] [-u userid][-p password] [-b buffer size][-B number of buffers] [-S sessions][-P parallelism] [-e execution profile] [-o] [-O] [-L] [-v]`

Delete command

Use the Backup Object Manager delete commands to remove backup objects from IBM Storage Protect™ that were sent by Data protection for DB2®.

For information about the command options, see [“Backup Object Manager command options” on page 0](#).

- Delete the database backup or backups that are specified by the command options from IBM Storage Protect™:
`backom -c d_db -a database alias -t timestamp | time range[-i instance] [-n node number] [-e execution profile] [-x] [-v]`
- Delete the table space backup or backups that are specified by the command options from IBM Storage Protect™:
`backom -c d_ts -a database alias -t timestamp/time range [-i instance] [-n node number] [-e execution profile] [-x] [-v]`
- Delete the DB2® log file backup or backups that are specified by the command options from IBM Storage Protect™:
`backom -c d_log -a database alias -l log number/log number range[-n node number] [-t timestamp/time range] [-e execution profile][-x] [-v]`
- Delete the file or files that are specified by the `-f`command option from IBM Storage Protect™:
`backom -c d_raw -f file name [-e execution profile] [-x] [-v]`

Password command

The password command connects to the backup server, prompts for a new password, and verifies the password that is entered with the backup server. If the verification is successful, the new password is encrypted and stored in the Data Protection for SAP configuration file.

Successful password verification requires that the password entered must be the active password for the corresponding node on the IBM Storage Protect™ server. Issue this command to verify and save a IBM Storage Protect™ password:

```
backom -c password [-x] [-a DB2 alias name] [-n DB2 node number]
[-e execution profile]
```

Query command

Use the query command to list backup objects that were sent to IBM Storage Protect™ by Data Protection for DB2®.

The objects that are displayed can be filtered with command options. For more information about command options, see [“Backup Object Manager command options” on page 0](#).

- List all backup objects that are related to DB2® (database or table space backups and DB2® log file backups):
backom -c q_all [-i instance] [-a database alias] [-n node number] [-t timestamp | time range] [-l log number | log number range] [-e execution profile] [-m output mode] [-v]
- List database backups:
backom -c q_db [-i instance] [-a database alias] [-n node number] [-t timestamp | time range] [-e execution profile] [-m output mode] [-v]
- List table space backups:
backom -c q_ts [-i instance] [-a database alias] [-n node number] [-t timestamp | time range] [-e execution profile] [-m output mode] [-v]
- List table space definition information (TDI) images that are related to a full DB2® database backup:
backom -c q_tdi -a database alias -t timestamp [-i instance] [-n node number] [-e execution profile] [-m output mode] [-v]
- List DB2® log file backups:
backom -c q_log [-a database alias] [-n node number] [-t timestamp | time range] [-l log number | log number range] [-e execution profile] [-m output mode] [-v]
- List backup objects available on IBM Storage Protect™ (database or table space backups, DB2® log file backups, and file backups):
backom -c q_raw [-f file name] [-e execution profile] [-m output mode] [-v]

Restore commands

Use the restore commands to restore any backup object that was created by Data Protection for DB2®.

- Restore the database with these command options:
backom -c r_db -a database alias [-n node number] [-u userid] [-p password] [-t timestamp] [-b buffer size] [-B number of buffers] [-S sessions] [-P parallelism] [-R restore type] [-O] [-e execution profile] [-x] [-v]
- Restore the database with these command options to a different location (redirected restore) in automatic mode:
backom -c rr_db_clone -a original database alias,target database alias [-i instance] [-n node number] [-u userid] [-p password] [-t timestamp] [-b buffer size] [-B number of buffers] [-S sessions] [-P parallelism] [-D target database directory] [-e execution profile] [-s scaling factor] [-N] [-C] [-v]
- Restore the database with these command options to a different location (redirected restore) in batch mode:
backom -c rr_db_batch -a original database alias,target database alias -f TDI image [-i instance] [-n node number] [-u userid] [-p password] [-t timestamp] [-b buffer size] [-B number of buffers] [-S sessions] [-P parallelism] [-D target database directory] [-e execution profile] [-s scaling factor] [-N] [-C] [-v]
- Restore the database with these command options to a different location (redirected restore) in interactive mode:
backom -c rr_db_interactive -a original database alias,target database alias [-i

instance] [-n *node number*] [-u *userid*] [-p *password*] [-t *timestamp*] [-f *modified TDI image*] [-b *buffer size*] [-B *number of buffers*] [-S *sessions*] [-P *parallelism*] [-D *target database*] [-e *execution profile*] [-s *scaling factor*] [-N] [-C] [-v]

- Restore the table spaces with these command options: `backom -c r_ts -a database alias [-n node number] [-u userid] [-p password] [-t timestamp] [-b buffer size] [-B number of buffers] [-S sessions] [-P parallelism] [-R restore type] [-O] [-e execution profile] [-x] [-v]`
- Restore the DB2® Recovery History File of the database with these command options: `backom -c r_hfile -a database alias [-n node number] [-t timestamp] [-u userid] [-p password] [-b buffer size] [-B number of buffers] [-S sessions] [-P parallelism] [-e execution profile] [-x] [-v]`
- Restore the table space definition information (TDI) denoted by the command options: `backom -c r_tdi -t timestamp -a database alias [-d destination directory] [-e execution profile] [-x] [-v]`
- Retrieve the DB2® log files with these command options: `backom -c r_log -a database alias -l log number|log number range -d destination directory [-n node number] [-t timestamp|time range] [-e execution profile] [-x] [-v]`
- Retrieve the file or files that are specified by the -f command option to the path specified by command option -d: `backom -c r_raw -f file name -d destination directory [-e execution profile] [-x] [-v]` This command can restore data to the destination directory. If a single segment was used during the backup, the data can be restored to DB2® from the destination directory after retrieval. If two or more segments were used during the backup, the data can be restored to the destination directory but cannot be restored to DB2®.

BACKOM command examples

Commands for certain tasks such as verifying and saving a IBM Storage Protect™ password are shown in the following examples.

- Use the following command to verify and save a IBM Storage Protect™ password:

```
backom -c password
```

- Use the following command to create a list of all available backup objects sent to IBM Storage Protect™ by IBM Storage Protect™ for Enterprise Resource Planning for DB2®:

```
backom -c q_all
```

- Use the following command to create a list of all DB2® log files for the SAMPLE database. The log number is greater than 123 and it is created in November 2002 with normal output detail level:

```
backom -c q_log -a SAMPLE -l 124-9999999 -t 200211* -m normal
backom -c q_log -a SAMPLE -l S0000124.log-S9999999.log -t 200211* -m normal
```

- Use the following command to create a list of DB2® log files for log chains 5 - 15 for the SAMPLE database. The log numbers range is 98 - 180 and it is archived from 4 p.m. to 8.30 p.m.:

```
backom -c q_log -a SAMPLE -k C0000005-C0000015 -l 98-180 -t ????????16*-????????2030*
```

- Use the following command to create a list of all table space backups. The list is for partition NODE0001, in the SAMPLE database SAMPLE, and for the backups that were created in November 2002 from 4 p.m. to 5 p.m.:

```
backom -c q_ts -a SAMPLE -n NODE0001 -t 200211??16*
```

- Use the following command to backup online database SAMPLE by using two I/O sessions and four backup buffers:

```
backom -c b_db -a SAMPLE -S 2 -B 4 -O
```

- Use the following command to backup the table spaces SYSCATSPACE and USERSPACE1 of database SAMPLE, by using the execution profile `initSAMPLE.utl` at `/db2/SAMPLE/config`:

```
backom -c b_db -a SAMPLE -T SYSCATSPACE,USERSPACE1 -e /db2/SAMPLE/config/initSAMPLE.utl
```

- Use the following command to restore a table space of database SAMPLE with the table space backup created on November 27, 2002, at 6:32:15 p.m.:

```
backom -c r_ts -a SAMPLE -t 20021127183215
```

- Use the following command to restore database SAMPLE with the latest backup:

```
backom -c r_db -a SAMPLE
```

- Use the following command to delete all DB2® log files for database SAMPLE that were created before June 2002:

```
backom -c d_log -a SAMPLE -t 1900*-20020601000000
```

- Use the following command to delete all versions of files that contain “tmp” in their path or file names that were sent to IBM Storage Protect™ by IBM Storage Protect™ for Enterprise Resource Planning:

```
backom -c d_raw -f *tmp*
```

Crontab example

UNIX™ or Linux™ cron jobs can be scheduled with the **crontab** command. This command starts an editing session that creates a crontab file. The cron jobs and the appropriate times are defined within the crontab file.

The file can be customized with this command:

```
crontab -e
```

In this example, a cron job starts the shell script `backup.ksh` at 11:30 p.m. Monday through Friday and uses DB2® backup to back up the SAP database. This example shows the entry in the crontab file that starts the script for this scenario:

```
30 23 * * 1,2,3,4,5 /usr/bin/su - db2c21 -c "/db2/C21/sapscripts/backup.ksh"
```

Crontab file sample

The following sample output, shows the root crontab jobs.

Example

```
# -----
# crontab.sample:
# Sample crontab file to be included in the root crontab jobs.
# -----
# Task:
# Submits backup/archive commands at regularly scheduled intervals
# using two simple shell scripts containing backup/archive commands
# and IBM Storage Protect™ commands.
# -----
#          *****      NOTE      *****      NOTE      *****      NOTE      *****
#
#          This file is intended only as a model and should be
#          carefully tailored to the needs of the specific site.
#
#          *****      NOTE      *****      NOTE      *****      NOTE      *****
# -----
#
# Remarks on the crontab file format:
#
# Each crontab file entry consists of a line with six fields, separated
# by spaces and tabs, that contain, respectively:
#   o The minute (0 through 59)
#   o The hour (0 through 23)
#   o The day of the month (1 through 31)
#   o The month of the year (1 through 12)
#   o The day of the week (0 through 6 for Sunday through Saturday)
```

```
# o The shell command
# Each of these fields can contain the following:
# o A number in the specified range
# o Two numbers separated by a dash to indicate an inclusive range
# o A list of numbers separated by commas
# o An * (asterisk); meaning all allowed values
#
# -----
#
# For the following examples, the system id (alias) of the DB2 database
# is assumed to be 'C21' and the username 'db2c21'.
#
# -----
# Full database backup, scheduled every Friday at 8:00 p.m.
#
0 20 * * 5
# /usr/bin/su - db2c21 -c "/db2/C21/sqllib/scripts/backup.ksh"
#
# -----
# Save redo logs, scheduled twice a day at 11:30 a.m. and at 5:30 p.m.
# Monday through Friday
#
30 11,17 * * 1,2,3,4,5
# /usr/bin/su - db2c21 -c "/db2/C21/sqllib/scripts/archive.ksh"
```

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.

During installation on Windows™ systems, the sample profile (along with all other files) is placed in the C:\Program Files\Tivoli\TDP4SAP directory.

The profile is copied to the profile path during installation if no other profile exists there. Data Protection for SAP reads the profile pointed to by environment variable `XINT_PROFILE` (shared library, **BackOM**) or sent as a parameter (**BackOM**) immediately before a backup or restore operation.

These rules apply to the keyword syntax:

- Each line is analyzed separately.
- Keywords can start in any column of the line.
- Keywords must not be preceded by any string, except blanks.
- If a keyword is encountered several times, the last one is used.
- File processing ends when the END keyword is encountered or the end of file is reached.
- The comment symbol is the 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 node_name

Specifies a *node_name* that is registered to the IBM Storage Protect™ server as an IBM Storage Protect™ 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 IBM Storage Protect™ server. This keyword must not be set when automated password handling is selected. It is to be set for manual password-handling.

ASNODE

Specifies a node name that is registered to the IBM Storage Protect™ server as an IBM Storage Protect™ node. This parameter must be defined with the respective SERVER statement, as shown in the sample profile. When automated password handling is selected and the node is accessed from multiple different SAP systems, for example, HANA scale-out or IBM Storage Protect™ Snapshot offload operations, this parameter avoids storing the encrypted password on multiple hosts (which would cause the password update to fail on all but the first host). This parameter must not be set when manual password handling is selected.

BACKEND pgmname [parameterlist]

Specifies a program *pgmname* that is called by IBM Storage Protect™ for ERP after the backup function completed and before program control is returned to DB2®. If *pgmname* is not a fully qualified path, the default search path is used to locate the program. If not specified, no backend processing is finished. Example for UNIX™ or Linux™:

```
BACKEND write operator@remotesite Backup of SAP database object completed.
```

This process sends a message to a remote user when the backup is finished.

BACKUPIDPREFIX 6-charstring | SAP____

Specifies a six-character prefix that is used to create a backup identifier for each archived object. If not specified, the default value is SAP____. All partitions of a partitioned DB2® database have the same **BACKUPIDPREFIX**.

BRARCHIVEMGTCLASS management_class [management_class...]

Specifies the IBM Storage Protect™ management classes that IBM Storage Protect™ for ERP uses to back up offline DB2® log files. Each parameter string can consist of up to 30 characters. Specify a separate management classes for each log file copy requested. As a result, make sure the number of different management classes that are specified must be greater than or equal to the number of log file copies. This parameter must be defined with the respective SERVER statement, as shown in the sample profile. To use different IBM Storage Protect™ 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 management classes only.

BRBACKUPMGTCLASS management_class [management_class...]

Specifies the IBM Storage Protect™ management classes that IBM Storage Protect™ for ERP uses to back up the DB2® database. The parameter string can consist of up to 30 characters. This parameter must be defined with the respective SERVER statement, as shown in the sample profile.

BUFFCOPY SIMPLE|PREVENT|AUTO

This optional parameter controls how IBM Storage Protect™ 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 IBM Storage Protect™ components. This option is the default. If set to PREVENT, the original data buffers are sent between IBM Storage Protect™ components.

For this mode, **BUFFSIZE** is restricted to a maximum of 896 KB. Furthermore, it cannot be selected when the IBM Storage Protect™ client encryption or client compression features are activated. If set to AUTO, IBM Storage Protect™ for ERP runs in PREVENT mode whenever the configuration supports it. Otherwise, SIMPLE mode is automatically selected. This parameter has no effect on restore operations.

BUFFSIZE n|131072

This parameter specifies the block size (in bytes) for the buffers that are used when it interacts with DB2®. The size of the buffers that are sent to the IBM Storage Protect™ 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. If not specified, the default value is 131072 (128 KB) for UNIX™ or Linux™ systems and 32768 (32 KB) for Windows™ systems. In most cases, these values are appropriate. If you plan to increase the size of internal buffers, make sure that sufficient storage is available. The number of buffers that are acquired by IBM Storage Protect™ for ERP correlates to the number of sessions (keyword **SESSIONS**). By activating **RL_COMPRESSION**, the number of buffers is doubled.

CONFIG_FILE path/SID.bki

Specifies the configuration file *initSID.bki* for IBM Storage Protect™ for ERP to store all variable parameters such as passwords and the date of the last password change. During processing, the string %DB2NODE is replaced automatically by the current DB2® node of a partitioned database or by “NODE0000” otherwise. This parameter is required.

END

Specifies the end of the parameter definitions. IBM Storage Protect™ for ERP stops searching the file for keywords when **END** is encountered.

FRONTEND pgmname [parameterlist]

Specifies a program *pgmname* that is called by IBM Storage Protect™ for ERP in a backup run before the connection to the IBM Storage Protect™ 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 IBM Storage Protect™ 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 IBM Storage Protect™ for ERP messages to be logged in the IBM Storage Protect™ 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 IBM Storage Protect™ servers.

MAX_SESSIONS n

Specifies the maximum number of parallel IBM Storage Protect™ client sessions that IBM Storage Protect™ for ERP establishes. For a direct backup or restore on tape drives, the number of sessions must be less than or equal to the number of tape drives available for the backup. Make sure that the **MOUNTLIMIT** (*mount1*) 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.

MAX_VERSIONS n|0

The *n* value defines the maximum number of database backup versions to be kept in backup storage. The default setting for this value is 0, meaning that backup version control is disabled. Every time a full backup completes successfully, the version count is increased by an increment of 1 and stored in the IBM Storage Protect™ for ERP configuration file. This value is also assigned to the table space files and to all subsequent DB2® log file backups. If the number of versions that are kept in backup storage is larger than the specified maximum number of backup versions (stored by the parameter **MAX_VERSIONS**), the oldest version is deleted, together with the corresponding table space, incremental, and log file backups until only the specified maximum number of most recent versions remain. For partitioned DB2® databases, a backup version control is done on a partition basis. Therefore, full backups must always be initiated for all partitions at the same time, for example by the DB2® script *db2_a11*. For details on the *db2_a11* script, see your DB2® documentation. Also, consider these characteristics:

- When IBM Storage Protect™ for ERP deletes an old full backup, all partial backups older than this full backup are also deleted.
- If the backups are distributed over multiple IBM Storage Protect™ servers and one of the servers is temporarily unavailable at the time of a new full backup, it is not possible to find all the backup versions. This situation might result in retaining a backup that would otherwise delete it.
- Every database partition needs its own configuration file. Partitions of a partitioned database must have the same **BACKUPIDPREFIX**.

IBM Storage Protect™ uses the value of the **RETVER** parameter (specified when a copy group is defined) to give files an expiration date. Use only one of these methods to control how long you keep backups:

- If you use IBM Storage Protect™ for ERP backup version control, you must bypass this expiration function. Set the IBM Storage Protect™ parameter **RETVER=9999** so that the files are not considered expired and are not deleted by IBM Storage Protect™.
- If you use the IBM Storage Protect™ expiration function, turn off the maximum number of full database backups versions control. Deactivate IBM Storage Protect™ for ERP backup version control by setting **MAX_VERSIONS=0**.

PASSWORDREQUIRED NO|YES

Specifies whether IBM Storage Protect™ requires a password to be supplied by the IBM Storage Protect™ client. This situation depends on the IBM Storage Protect™ 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 IBM Storage Protect™ for ERP stores for each processed DB2® log file. The valid range is 1 - 9. If not specified, IBM Storage Protect™ for ERP stores one copy of each log file. The number of different management classes for archived logs (keyword **BRARCHIVEMGTCLASS** specified must

be greater than or equal to the number of log file copies specified. The number of different management classes 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, IBM Storage Protect™ 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 IBM Storage Protect™ 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.

SEGMENTSIZ size[GB|TB]

This keyword specifies the maximum size of the segments that are split from large backup objects. The required *size* value must be a positive integer equal to or greater than 1. Consider these characteristics when you specify this parameter:

- The scale units (GB or TB) are not required. GB is the default value.
- When you specify the scale units (GB or TB), you can use lowercase letters, uppercase letters, or a combination of both cases. However, you cannot specify single-character abbreviations (G or T).
- When this parameter is not specified, one backup object per DB2® backup session is transferred to IBM Storage Protect™.
- If the specified segment size is less than the DB2® block size specified during the backup, the specified segment size is ignored at run time. The specified DB2® backup block size is used instead.

The following example sets the maximum size of the backup objects segments on IBM Storage Protect™ to 100 GB:

```
SEGMENTSIZ 100 GB
```

You can also specify the command in the following way:

```
SEGMENTSIZ 100
```

SERVER servername

This keyword specifies the name of the IBM Storage Protect™ server to which IBM Storage Protect™ for ERP backups are to be stored. This statement begins a server section in the IBM Storage Protect™ 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:

- ADMSNODE
- BRARCHIVEMGTCLASS
- BRBACKUPMGTCLASS
- PASSWORDREQUIRED
- SESSIONS
- TCP_ADDRESS
- USE_AT

The server name must be defined in the IBM Storage Protect™ profiles `dsm.sys` (UNIX™, Linux™) or `servername.opt` (for Windows™). 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. Do NOT use any profile keywords, ADSM, or TSM as the server name.

SESSIONS n|1

The *n* value specifies the number of parallel sessions IBM Storage Protect™ 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.

TCP_ADDRESSIP address of server

Specifies the IP address of the IBM Storage Protect™ server in dotted decimal notation. This parameter overrides the value for the parameter **TCPSERVERADDRESS** in the IBM Storage Protect™ client system options file (*dsm.sys*) on UNIX™ or Linux™ or in the client options file (*servername.opt*) on Windows™. This parameter must be defined 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 IBM Storage Protect™ for ERP support. Using it can significantly deteriorate the performance of IBM Storage Protect™ for ERP.

TRACEFILE path

Specifies the name and location of the trace file for IBM Storage Protect™ 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 IBM Storage Protect™ 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 IBM Storage Protect™ server on all days. This parameter must be defined with the respective **SERVER** statement. The parameter **USE_AT** has no effect on actions other than on a backup.

Sample profile file for UNIX™ or Linux™

A sample profile file (*initSID.utl*) is included in the IBM Storage Protect™ for ERP installation package.

```
#-----
#
# Data Protection for SAP (R) interface for DB2 UDB
#
# Sample profile for Data Protection for SAP (R) Version 7.1
#
#-----
#
# See the 'Data Protection for SAP (R) Installation &
# User's Guide' for a full description.
#
# For a comment symbol the character '#' can be used.
# Everything following this character will be interpreted as comment.
#
# Data Protection for SAP (R) accesses its profile
# in "read only" mode. All variable parameters like passwords, date of
# last password change, current version number will be written into the file
# specified with the CONFIG_FILE parameter. The passwords will be encrypted.

#-----
# Prefix of the 'Backup ID' which is stored in the description field of
# the IBM Storage Protect™ archive function.
# Maximum 6 characters.
# Default: none.
#-----
BACKUPIDPREFIX      SID___

#-----
# Number of parallel sessions to be established.
# Note: This number must not exceed the number of tape drives simultaneously
# available to the node on the IBM Storage Protect™ servers to be accessed.
# The valid range of MAX_SESSIONS is from 1 and 32.
# Default: none.
```

```

#-----
MAX_SESSIONS          1 # IBM Storage Protect™ client sessions

#-----
# Number of backup copies of the DB2 log files.
# The valid range of REDOLOG_COPIES is from 1 to 9.
# Default: 1.
#-----
#REDOLOG_COPIES          2

#-----
# Specifies the block size for disk I/O (in bytes).
# The default values have been chosen from our performance experiments in
# standard hardware environments.
# The valid range of BUFFSIZE is from 4KB to 32MB.
# Default: 131072 (128 KB) on UNIX, 32768 (32 KB) on Windows.
#-----
BUFFSIZE              131072          # block size in bytes

#-----
# This optional parameter controls how Data Protection for SAP(R) uses
# the internal buffers for transferring data during a backup.
# Valid values:  SIMPLE | PREVENT | AUTO
# Default: SIMPLE
#-----
#BUFFCOPY                AUTO

#-----
# Name of a program to be called before the backup task is started.
# Default: none.
#-----
#FRONTEND                pgmname parameterlist

#-----
# Name of a program to be called after the backup task is completed.
# Default: none.
#-----
#BACKEND                pgmname parameterlist

#-----
# Maximum number of data base backup versions to be kept.
# Note: Version control by Data Protection for SAP (R) is only activated
# only activated if the parameter MAX_VERSION is not 0.
# The valid range of MAX_VERSIONS is from 0 to 9999.
# A value of 0 means no versioning.
# Default: 0, no versioning.
#-----
#MAX_VERSIONS           4

#-----
# Specifies whether a null block compression of the data is to be performed
# before transmission to IBM Storage Protect™.
# Although RL compression introduces additional CPU load, throughput can be
# improved when the network is the bottleneck. RL compression in Data
# Protection for SAP(R) should not be used together with
# IBM Storage Protect™ API compression.
# Default: NO
#-----
#RL_COMPRESSION          YES          # NO is default

#-----
# Controls generation of a trace file.
# Note: We recommend using the trace function only in cooperation with
# Data Protection for SAP (R) support.
# Default: OFF.
#-----
#TRACE                   OFF
#TRACEFILE                /db2/C21/sqllib/log/tdpr3.trace

#-----
# Denotes the maximum size of the trace file in KB.

```

```

# If not specified, the trace file size is unlimited.
#-----
#TRACEMAX          max. size          # trace file size in KB

#-----
# Specify the full path of the configuration file.
# Default: none.
#-----
CONFIG_FILE          /db2/C21/sqlllib/%DB2NODE/initSID.bki

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

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

SERVER              server_a          # Servername, as defined in dsm.sys
SESSIONS            2                 # Maximum number of sessions
                                     # to server_a
PASSWORDREQUIRED    YES               # Use a password
ADSMNODE            NODE              # IBM Storage Protect™ Nodename
BRBACKUPMGTCCLASS   MDB              # Mgmt-Classes for database backup
BRARCHIVMGTCCLASS   MLOG1 MLOG2       # Mgmt-Classes for redo log backup
# TCP_ADDRESS        192.168.1.1      # IP address of network interface
                                     # on server_a
# USE_AT             0 1 2 3 4 5 6    # Overrides IP address of dsm.sys
                                     # Days when server_a is used for
                                     # backup
*****
# USE_AT : 0=Su 1=Mo 2=Tu 3=We 4=Th 5=Fr 6=Sa
# The valid range of USE_AT is from 0 to 6.
# Default: all days
*****

#SERVER              server_b          # Servername, as defined in dsm.sys
# SESSIONS            2                 # Maximum number of sessions
                                     # to server_b
# PASSWORDREQUIRED    YES               # Use a password
# ADSMNODE            NODE              # IBM Storage Protect™ Nodename
# BRBACKUPMGTCCLASS   MDB              # Mgmt-Classes for database backup
# BRARCHIVMGTCCLASS   MLOG1 MLOG2       # Mgmt-Classes for redo log backup
# TCP_ADDRESS        192.168.1.1      # IP address of network interface
                                     # on server_b
# USE_AT             0 1 2 3 4 5 6    # Overrides IP address of dsm.sys
                                     # Days when server_b is used for
                                     # backup
*****
# USE_AT : 0=Su 1=Mo 2=Tu 3=We 4=Th 5=Fr 6=Sa
# Default: all days
*****

#-----
# End of profile

END

```

Sample profile (Windows™)

The sample profile file (initSID.utl) is included in the installation package.

```

#-----
#
# Data Protection for SAP (R) interface for DB2 UDB
#
# Sample profile for Data Protection for SAP (R)
# Version x.x for Windows
#
#-----
#
# See the 'Data Protection for SAP (R) Installation & User's Guide' for
# a full description.
#
# For a comment symbol the character '#' can be used.
# Everything following this character will be interpreted as comment.
#
# Data Protection for SAP (R) accesses its profile in "read only" mode.
# All variable parameters like passwords, date of last password
# change, current version number will be written into the file specified
# with the CONFIG_FILE parameter. The passwords will be encrypted.

#-----
# Prefix of the 'Backup ID' which is used for communication with the
# SAP® BR*Tools and stored in the description field of the
# IBM Storage Protect™ archive function.
# Must be 6 characters.
# Default: none.
#-----
BACKUPIDPREFIX          SID___

#-----
# Number of parallel sessions to be established.
# Note: This number must not exceed the number of tape drives simultaneously
# available to the node on the IBM Storage Protect™ servers to be accessed.
# The valid range of MAX_SESSIONS is from 1 and 32.
# Default: none.
#-----
MAX_SESSIONS           1 # IBM Storage Protect™ client sessions

#-----
# Number of backup copies of the DB2 log files.
# The valid range of REDOLOG_COPIES is from 1 to 9.
# Default: 1.
#-----
#REDOLOG_COPIES        2

#-----
# Specifies the block size for disk I/O (in bytes).
# The default values have been chosen from our performance experiments in
# standard hardware environments.
# The valid range of BUFFSIZE is from 4KB to 32MB.
# Default: 131072 (128 KB) on UNIX, 32768 (32 KB) on Windows.
#-----
BUFFSIZE               32768          # block size in bytes

#-----
# This optional parameter controls how Data Protection for SAP(R) uses
# the internal buffers for transferring data during a backup.
# Valid values:  SIMPLE | PREVENT | AUTO
# Default: SIMPLE
#-----
#BUFFCOPY              AUTO

#-----
# Name of a program to be called before the backup task is started.
# Default: none.
#-----
#FRONTEND              pgmname parameterlist

#-----
# Name of a program to be called after the backup task is completed.
# Default: none.
#-----
#BACKEND               pgmname parameterlist

```

```

#-----
# Maximum number of data base backup versions to be kept.
# Note: Version control by Data Protection for SAP (R) is
# only activated if the parameter MAX_VERSION is not 0.
# The valid range of MAX_VERSIONS is from 0 to 9999.
# Default: 0
#-----
#MAX_VERSIONS          4

#-----
# Specifies whether a null block compression of the data is to be performed
# before transmission to IBM Storage Protect™.
# Although RL compression introduces additional CPU load, throughput can be
# improved when the network is the bottleneck. RL compression in Data
# Protection for SAP(R) should not be used together with
# IBM Storage Protect™ API compression.
# Default: NO
#-----
#RL_COMPRESSION        YES

#-----
# Controls generation of a trace file.
# Note: We recommend using the trace function only in cooperation with
# Data Protection for SAP (R) support.
# Default: OFF
#-----
#TRACE                 OFF
#TRACEFILE              c:\sqllib\tdp_r3\log\tdpr3.trace

#-----
# Denotes the maximum size of the trace file in KB.
# If not specified, the trace file size is unlimited.
#-----
#TRACEMAX              max. size          # trace file size in KB

#-----
# Specify the full path of the configuration file.
# Default: none.
#-----
CONFIG_FILE            c:\sqllib\tdp_r3\%DB2NODE\initSID.bki

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

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

SERVER                server_a            # Servername, as defined in dsm.sys
SESSIONS              2                  # Maximum number of sessions
                                # to server_a
PASSWORDREQUIRED      YES                # Use a password
ADSMNODE              NODE               # IBM Storage Protect™ Nodename
BRBACKUPMGTCLASS      MDB                # Mgmt-Classes for database backup
BRARCHIVEMGTCLASS     MLOG1 MLOG2        # Mgmt-Classes for redo log backup
# TCP_ADDRESS         192.168.1.1        # IP address of network interface
                                # on server_a
# USE_AT              0 1 2 3 4 5 6      # Overrides IP address of dsm.sys
                                # Days when server_a is used for
                                # backup
*****
# USE_AT : 0=Su 1=Mo 2=Tu 3=We 4=Th 5=Fr 6=Sa
# The valid range of USE_AT is from 0 to 6.
# Default: all days

```

```

#*****
#SERVER          server_b          # Servername, as defined in dsm.sys
# SESSIONS       2                 # Maximum number of sessions
#                                     # to server_b
# PASSWORDREQUIRED YES             # Use a password
# ADMSNODENAME   NODE              # IBM Storage Protect™ Nodename
# BRBACKUPMGTCCLASS MDB            # Mgmt-Classes for database backup
# BRARCHIVEMGTCLASS MLOG1 MLOG2    # Mgmt-Classes for redo log backup
# TCP_ADDRESS    192.168.1.1       # IP address of network interface
#                                     # on server_b
#                                     # Overrides IP address of dsm.sys
# USE_AT         0 1 2 3 4 5 6     # Days when server_b is used for
#                                     # backup
#*****
# USE_AT : 0=Su 1=Mo 2=Tu 3=We 4=Th 5=Fr 6=Sa
# Default: all days
#*****

#-----
# End of profile

END

```

Locating sample files

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

- 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 user options file sample (UNIX™, Linux™)

```

*****
* IBM Storage Protect™                                     *
*                                                         *
* Sample Client User Options file for Unix platforms      *
*****

Servername      server_a
Replace         On
Tapeprompt      No
DOM             /usr/sap /sapmnt/C21 /usr/sap/trans /db2/C21

```

Client user options file sample (Windows™)

Data Protection for SAP requires a client options file `dsm.opt` to be present in the location indicated by environment variable **DSMI_CONFIG**. The specific options that are used by Data Protection for SAP for each server are taken from file `server.opt` in the same path.

Example

```

*****
*
* DSM.OPT (for Data Protection for SAP)
*
* This file is intentionally left empty. It must be present in the location
* indicated by environment variable DSMI_CONFIG. The specific options used
* by Data Protection for SAP for each server however are taken from files
* server.opt residing in the same path.
*
* Please note: This client options file is not meant to be used by other
*               IBM Storage Protect™ clients.
*

```

```
*****
```

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.

Example

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

SErvername server_a
COMMmethod      TCPip
TCPport         1500
TCPserveraddress your_ITSM_server_1
TCPbuffsize     32
TCPwindowsize   24
Compression     Off
InclExcl        /usr/lpp/adsm/bin/inclexcl.list

SErvername server_b
COMMmethod      TCPip
TCPport         1500
TCPserveraddress your_ITSM_server_2
TCPbuffsize     32
TCPwindowsize   24
Compression     Off
InclExcl        /usr/lpp/adsm/bin/inclexcl.list
```

Include and exclude list sample (UNIX™, Linux™)

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

Example

```
* -----
* incl excl.list:
* Sample include/exclude list
* -----
* Task:
* Include/Exclude list of files and directories for IBM Storage Protect™
* incremental backups
* -----
*          *****      NOTE          *****      NOTE          *****      NOTE          *****
*
*          This file is intended only as a model and should be
*          carefully tailored to the needs of the specific site.
*
*          *****      NOTE          *****      NOTE          *****      NOTE          *****
* -----
*
* For all UNIX systems
*
* exclude /unix
* exclude ../../core
* exclude /u/../../.sh_history
* exclude /home/../../.sh_history
*
* Note: It is recommended to perform system backups on a regular
*       basis. Consequently, you can exclude at least the following
*       directories:
*
* exclude /usr/games/../../*
* exclude /usr/bin/../../*
```



```

exclude /usr/lbin/.../*
exclude /usr/sbin/.../*
exclude /usr/sbin/.../*
* -----
*
* For those using AFS, exclude the cache filesystem or file
*
* exclude /usr/vice/cache/*
* exclude /var/vice/cache/*
* or
* exclude /afscfs
* -----
*
* This stuff is either not worthwhile to be included or should be backed up
* using DB2 backup techniques and the SAP utility brarchive.
*
exclude /db2/C21/log_archive/C21/*
* exclude /db2/C21/sapreorg/.../* (There may be important scripts
*                               located, check it out and decide.)
exclude /db2/C21/sapdata*/.../*
exclude /db2/C21/sapraw*/.../*
* -----
*
* With the above include/exclude list we implicitly include everything not
* excluded above. Especially for DP for SAP, this means including:
*
* /sapmnt/C21      > 300 MB
* /usr/sap         > 50 MB
* /db2/C21         > 200 MB
* and UNIX related > 350 MB
* -----

```

Include/exclude list sample (Windows™)

The include/exclude list is intended for the standard client user option file. The purpose is to exclude files that are easy to restore. Also, exclude files that are already saved by Data Protection for SAP from routine IBM Storage Protect™ incremental backups. Typically, such files are Windows™ system files and database files.

Example

```

*****
* This Include-Exclude list is used for incremental backups of file
* systems by the IBM Storage Protect™ command-line backup client.
* Therefore the name of this file has to be set under the keyword InclExcl
* in the standard IBM Storage Protect™ client user option file "dsm.opt".
*
* Since the backup of the DB2 database is done by
* Data Protection for SAP(R) and not by IBM Storage Protect™
* command-line backup client, the DB2 database should be excluded
* from backups by the IBM Storage Protect™ command-line backup client.
*
* Note 1:
* The environment variable DSM_CONFIG contains the full file name of
* the IBM Storage Protect™ client user option file "dsm.opt".
* Note 2:
* This Include-Exclude is not used by Data Protection for SAP(R)
*
*****
Exclude *:...\*.swp
Exclude *:...\*.obj
Exclude *:...\*.csm
Exclude *:...\*.dsk
Exclude *:...\*.bak
Exclude *:...\win386.swp
Exclude *:...\386spart.par
Exclude *:...\pagefile.sys
Exclude *:...\*.par
Exclude *:...\SYSTEM32\CONFIG\*.
Exclude *:...\SYSTEM32\CONFIG\...\*
Exclude *: \IBMBIO.COM
Exclude *: \IBMDOS.COM
*
*Exclude the following DB2 database files:
*
Exclude *:db2\C21\log_archive\C21\...\*

```

```
Exclude *:\db2\C21\sapreorg\...\*
Exclude *:\db2\C21\sapdata*\...\*
```

Client options files sample

Data Protection for SAP requires a corresponding client option file *server.opt* for each IBM Storage Protect™ server. These files must be in the same directory. This directory must also contain the client options file *dsm.opt*, which is specified in the environment variable *DSMI_CONFIG*.

Example

```
*****
*
* SERVER.OPT
*
* Data Protection for SAP (R) obtains the necessary information about
* an IBM Storage Protect™ server 'server' from a client option file
* called 'server.opt'. For each IBM Storage Protect™ server a
* corresponding client option file is required.
*
* Note: This file contains the client options for the IBM Storage Protect™
* server called 'server_a'.
*
* Please see the IBM Storage Protect™ documentation for details.
*
*****
COMMethod      TCPIP
COMPression    OFF
*NODEname      C21
TCPPort        1500
TCPServeraddress xxx.xxx.xxx.xxx
PASSWORDACCESS PROMPT
TCPBUFFSIZE    31
TCPWINDOWSIZE  32
```

Vendor environment file sample

A DB2® vendor environment file (*vendor.env*) is created from the information that is entered in the installation dialog panels during installation. A sample DB2® vendor environment file is included in the Data Protection for SAP for DB2® installation package.

Ensure that there are no blanks within the paths that are specified for the vendor-specific environment variables of the vendor environment file. DB2® is unable to handle embedded blanks. For a standard Windows™ installation, the Data Protection for SAP profile is at

```
c:\Program Files\Tivoli\tsm\tdp_r3\db264\initSID.utl
```

Sample DB2® vendor environment file for UNIX™ or Linux™:

```
XINT_PROFILE=/db2/C21/tdpr3/initC21.utl
TDP_DIR=/db2/C21/tdpr3/tdplog
BACKOM_LOCATION=/usr/tivoli/tsm/tdp_r3/db264/backom
```

Sample DB2® vendor environment file for Windows™:

```
XINT_PROFILE=c:\db2\C21\tdpr3\initC21.utl
TDP_DIR=c:\db2\C21\tdpr3\tdplog
BACKOM_LOCATION=c:\tivoli\tsm\tdp_r3\db264\backom.exe
```

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.

This table is also provided in file form as planning_sheet_db2 for UNIX™ and Linux™ and planning_sheet_db2.txt for Windows™.

Table 8: Installation parameters for Data Protection for SAP		
UNIX™ or Linux™	Windows™	Installation parameter
X	X	DB2® database SID
X	X	IBM Storage Protect™ server name or IP address:
X	X	IBM Storage Protect™ node name: IBM Storage Protect™ node that is configured on the IBM Storage Protect™ server that is named for the backup of the SID previously listed.
X	X	IBM Storage Protect™ management classes for database and log file backups. Management classes that are configured for the database backup and for the backup of log files. Default: MDB for database backups, MLOG1 and MLOG2 for log file backups.
	X	Path where the IBM Storage Protect™ 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 IBM Storage Protect™ (contents of environment variable DSMI_CONFIG).
	X	Path to IBM Storage Protect™ log files (contents of environment variable DSMI_LOG): The IBM Storage Protect™ API creates the file dserror.log in this path. Default: C:\temp
	X	Installation path for Data Protection for SAP executable files: C:\Program Files\Tivoli\TSM\tdp_r3\db264

Network settings for IBM Storage Protect™

When you are using IBM Storage Protect™ with Data Protection for SAP, you can improve performance by making updates to the configuration files. Before you edit configuration files, save a backup copy.

The performance adjustments for IBM Storage Protect™ are completed by editing the following files:

- IBM Storage Protect™ server option file `dsmserv.opt`
- IBM Storage Protect™ backup-archive client option file `dsm.sys` (UNIX™ and Linux™ systems), or `server.opt` (Windows™ systems).

This table shows the corresponding IBM Storage Protect™ configuration file attributes with the values.

Table 9: Tuning IBM Storage Protect™ configuration file attributes		
Attributes	Value	Description
TCPBuffsize	32	Specifies the size, in KB, of the buffer that is used for TCP/IP send requests. This option affects whether IBM Storage Protect™ sends the data directly from the session buffer or copies the data to the TCP buffer. A buffer size of 32 KB forces IBM Storage Protect™ to copy data to its communication buffer and flush the buffer when it fills.
TCPNODElay	YES	Specifies whether the server is to send small amounts of data or allow TCP/IP to buffer the data. Disallowing buffering might improve throughput but more packets are sent over the network.
TCPWindowSize	640 (AIX) 63 (others)	Specifies the size, in KB, which is used for the TCP/IP sliding window for the client node. This size is used when data is sent or received. The range of values is 0 - 2048.

Networks with large bandwidth delay

For networks with a large bandwidth-delay, activate the TCP enhancements as specified in RFC1323.

For example, the network on an AIX® system can be configured with the **no** command. This command sets or displays current network attributes in the kernel. For more information, see the man page.

This table shows the network attributes with their advised values:

Table 10: Tuning of network settings		
Attributes	Value	Description
rfc1323	1	Enables TCP enhancements as specified by RFC 1323, TCP Extensions for High Performance. The default is 0. A value of 1 specifies that all TCP connections attempts to negotiate the RFC enhancements.
sb_max	131072	Specifies the maximum buffer size that is allowed for a socket. The default is 65536 bytes. From the point of view of performance recommendations, the sb_max value is to be twice the TCPWindowsize set within the IBM Storage Protect™ configuration file dsm.sys.

Set these values by issuing these commands by the root user on the appropriate system:

```
no -o rfc1323=1
no -o sb_max=131072
```

The **no** command does not do range checking. It accepts all values. If used incorrectly, the command might cause the system to become inoperable. These changes are lost at system restart. To permanently change the values, edit the `/etc/rc.net` file.

SP switch (RISC 6000)

If an SP switch (RISC 6000) is used, the `rpoolsize` and `spoolsize` values must be set as shown in the following table.

Table 11: Tuning of SP switch buffer pools		
Attributes	Value	Description
rpoolsize	1048576	The receive pool is a buffer pool for incoming data. The size for values is in bytes.
spoolsize	1048576	The send pool is a buffer for outgoing data. The size for values is in bytes.

The buffer pool settings can be changed by using the **hgcss** command. After the changes are made, restart the node.

Accessibility features for the IBM Storage® Protect product family

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Storage® Protect family of products includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Storage® Protect family of products uses the latest W3C Standard, [WAI-ARIA 1.0 \(www.w3.org/TR/wai-aria/\)](http://www.w3.org/TR/wai-aria/), to ensure compliance with [US Section 508 \(www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards\)](http://www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and [Web Content Accessibility Guidelines \(WCAG\) 2.0 \(www.w3.org/TR/WCAG20/\)](http://www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

The product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the [Accessibility section of the IBM Knowledge Center help \(www.ibm.com/support/knowledgecenter/about/releasenotes.html?view=kc#accessibility\)](http://www.ibm.com/support/knowledgecenter/about/releasenotes.html?view=kc#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

User interfaces do not have content that flashes 2 - 55 times per second.

Web user interfaces rely on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

Web user interfaces include WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Storage® Protect 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 accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see [IBM Accessibility \(www.ibm.com/able\)](http://www.ibm.com/able).

Notices

This information was developed for products and services offered in the US. This material might be available from IBM® in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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
US

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

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 jurisdictions 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 wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM® Director of Licensing
IBM® Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
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.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

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: © (your company name) (year). Portions of this code are derived from IBM® Corp. Sample Programs. © Copyright IBM® Corp. _enter the year or years_.

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.

The registered trademark Linux® is used pursuant to a sublicense from the Linux® Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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.

Red Hat®, OpenShift®, Ansible®, and Ceph® are trademarks or registered trademarks of Red Hat®, Inc. or its subsidiaries in the United States and other countries.

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

VMware, VMware vCenter Server™, and VMware vSphere™ are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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>.

Glossary

Index

A

accessibility features [109](#)

B

backup

- full offline [86](#)
- full online [86](#)
- incremental [15](#)
- objects [87](#)

Backup Object Manager

- commands [87](#)

backup-archive client [15](#), [15](#)

C

command syntax

- Backup Object Manager [87](#)

configuration files

- creation of [39](#)

D

DB2 Version 8.2

- rules when using [37](#)
- vendor environment file [37](#)

delete command [90](#)

disability [109](#)

E

environment variable

- DB2_INSTANCE [88](#)

F

full offline backup [86](#)

full online backup [86](#)

full restore [86](#)

I

IBM Knowledge Center [8](#)

incremental backup [15](#)

K

keyboard [109](#)

Knowledge Center [8](#)

M

multiple copies of redo logs [20](#)

multiple redo log copies [20](#)

O

offline log file [15](#)

P

partition [39](#), [39](#), [88](#), [94](#), [86](#), [86](#), [86](#), [86](#), [92](#)

ProLE [26](#), [28](#)

publications [8](#)

Q

query command [91](#)

R

Recovery History File [88](#), [88](#)

restore

- full restore [86](#)

S

sessions

- multiple (parallel) [86](#), [86](#)

SID [86](#), [86](#), [82](#)

T

tablespace [88](#), [88](#), [88](#), [88](#), [90](#), [91](#), [91](#), [91](#), [92](#), [92](#)

tablespace

- definition information [91](#)

V

vendor environment file [36](#), [36](#), [36](#), [36](#), [86](#)

X

XINT_PROFILE [88](#)

© Copyright International Business Machines Corporation 2014, 2025

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

