



SAP HANA Database – Backup and Recovery Guide

- SAP HANA Appliance Software SPS 04

Target Audience

- Consultants
- Administrators
- SAP Hardware Partner
- Others

Public

Document version 1.0 – April, 2012



SAP AG
Dietmar-Hopp-Allee 16
69190 Walldorf
Germany
T +49/18 05/34 34 34

© Copyright 2012 SAP AG. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors.

Microsoft, Windows, Excel, Outlook, and PowerPoint are registered trademarks of Microsoft Corporation.

IBM, DB2, DB2 Universal Database, System i, System i5, System p, System p5, System x, System z, System z10, System z9, z10, z9, iSeries, pSeries, xSeries, zSeries, eServer, z/VM, z/OS, i5/OS, S/390, OS/390, OS/400, AS/400, S/390 Parallel Enterprise Server, PowerVM, Power Architecture, POWER6+, POWER6, POWER5+, POWER5, POWER, OpenPower, PowerPC, BatchPipes, BladeCenter, System Storage, GPFS, HACMP, RETAIN, DB2 Connect, RACF, Redbooks, OS/2, Parallel Sysplex, MVS/ESA, AIX, Intelligent Miner, WebSphere, Netfinity, Tivoli and Informix are trademarks or registered trademarks of IBM Corporation.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Adobe, the Adobe logo, Acrobat, PostScript, and Reader are either trademarks or registered trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Oracle is a registered trademark of Oracle Corporation.

UNIX, X/Open, OSF/1, and Motif are registered trademarks of the Open Group.

Citrix, ICA, Program Neighborhood, MetaFrame, WinFrame, VideoFrame, and MultiWin are trademarks or registered trademarks of Citrix Systems, Inc.

HTML, XML, XHTML and W3C are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Massachusetts Institute of Technology.

Java is a registered trademark of Sun Microsystems, Inc.

JavaScript is a registered trademark of Sun Microsystems, Inc., used under license for technology invented and implemented by Netscape.

SAP, R/3, xApps, xApp, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP Business ByDesign, and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world. All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

These materials are subject to change without notice. These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.



Disclaimer

Some components of this product are based on Java™. Any code change in these components may cause unpredictable and severe malfunctions and is therefore expressly prohibited, as is any decompilation of these components.

Any Java™ Source Code delivered with this product is only to be used by SAP's Support Services and may not be modified or altered in any way.

Documentation in the SAP Service Marketplace






You can find this documentation at the following Internet address:

`service.sap.com/hana`

Typographic Conventions

Type Style	Description
<i>Example Text</i>	Words or characters quoted from the screen. These include field names, screen titles, pushbuttons labels, menu names, menu paths, and menu options. Cross-references to other documentation
Example text	Emphasized words or phrases in body text, graphic titles, and table titles
EXAMPLE TEXT	Technical names of system objects. These include report names, program names, transaction codes, table names, and key concepts of a programming language when they are surrounded by body text, for example, SELECT and INCLUDE.
Example text	Output on the screen. This includes file and directory names and their paths, messages, names of variables and parameters, source text, and names of installation, upgrade and database tools.
Example text	Exact user entry. These are words or characters that you enter in the system exactly as they appear in the documentation.
<Example text>	Variable user entry. Angle brackets indicate that you replace these words and characters with appropriate entries to make entries in the system.
EXAMPLE TEXT	Keys on the keyboard, for example, F2 or ENTER.

Icons

Icon	Meaning
	Caution
	Example
	Note
	Recommendation
	Syntax

Additional icons are used in SAP Library documentation to help you identify different types of information at a glance. For more information, see *Help on Help* → *General Information Classes and Information Classes for Business Information Warehouse* on the first page of any version of *SAP Library*.

Contents

SAP HANA DATABASE BACKUP AND RECOVERY GUIDE.....	1
About This Guide.....	1
SAP HANA Guides.....	2
Prerequisites	2
Points to Note	3
ABOUT SAP HANA DATABASE BACKUP	5
About Data Area Backup	6
Default Location of the Data Backup Files	6
Backup File Names	6
Checking Available Disk Space for Backups.....	7
About Log Area Backup	8
Location of the Log Backup Files	8
Log Backup Options.....	9
Avoiding LOG FULL (File System Full) Situations	11
About the SAP HANA Configuration Files.....	12
Location of the Configuration Files	12
Configuration Files Created During Installation	13
About the Backup Catalog.....	14
Monitoring Views for the Backup Catalog	16
Identifying Old Log Backups	17
BACKING UP THE SAP HANA DATABASE.....	18
Performing a Data Backup Using SAP HANA Studio.....	18
Performing a Data Backup Using SQL Commands	19
Performing a Data Backup in Batch Mode	20
Setting Up a Data Backup Using Cron	20
Canceling a Running Data Backup.....	21
Prerequisites	21
Canceling a Running Data Backup Using SAP HANA Studio	21
Canceling Other Running Data Backups.....	21
ABOUT SAP HANA DATABASE RECOVERY.....	22
Recovery Scenarios	23
Data Area is Unusable	23



Log Area is Unusable	24
Logical Error – Point in Time Recovery	24
Copying a Database.....	25
RECOVERING THE SAP HANA DATABASE.....	26
Prerequisites	26
Constraints	26
Recovering the Database	26
Procedure.....	27
COPYING A DATABASE USING BACKUP AND RECOVERY	29
Prerequisites	29
Procedure.....	29
Result.....	30
BACKUP CHECKS	31
SQL SYNTAX FOR BACKUP AND RECOVERY.....	34
Prerequisites	34
Commands for Data Backup.....	34
Examples of Commands for Data Backup	35
Data Backup Check.....	36
Example of Commands for Data Backup Check	36
Data Backup Cancel.....	36
Example of Commands for Backup Cancel.....	37
Data Snapshot	37
Examples of Commands for Data Snapshot.....	37
Database Recovery and Data Recovery	37
Examples of Commands for Database Recovery and Data Recovery	39

SAP HANA Database Backup and Recovery Guide

To ensure optimal performance, the SAP HANA database holds the bulk of its data in memory. However, it still uses persistent storage to provide a fallback in case of failure.

During normal database operation, data is automatically saved from memory to disk at regular savepoints. Additionally, all data changes are recorded in the redo log. The redo log is saved from memory to disk with each committed database transaction. After a power failure, the database can be restarted like any disk-based database, and it returns to its last consistent state by replaying the redo log since the last savepoint.

While savepoints and log writing protect your data against power failures, savepoints do not help if the persistent storage itself is damaged. To protect against data loss due to disk failures, backups are required. Backups save the payload (the actual data) of the data area and log area to different locations. Currently only backups to the file system are supported.

Backups are performed while the database is running. The impact of backups on system performance is negligible, and users can continue to work normally while the backup is running.

About This Guide

This Guide describes how to implement and administer the backup and recovery features of the SAP HANA database.

This Guide describes the steps to perform the following tasks:

- Create a data backup of the SAP HANA database.
- Configure when log backups are created.
- Perform a recovery of the SAP HANA database.

Note

This Guide deals with the SAP HANA database only. It does not describe how to perform a backup and recovery of all the components that can be part of a SAP HANA installation.

Note

This Guide does not replace the documentation that customers are recommended to use for their specific production operations.

SAP HANA Guides

For more information about SAP HANA landscape, installation, deployment and security, see the resources listed in the following table:

Topic	Guide/Tool	Quick Link
SAP HANA Landscape, Deployment & Installation	SAP HANA Knowledge Center on SAP Service Marketplace	https://service.sap.com/hana → SAP HANA Master Guide → SAP HANA Installation Guide
SAP HANA Administration & Security	SAP HANA Knowledge Center on SAP Help Portal	http://help.sap.com/hana → SAP HANA Technical Operations Manual → SAP HANA Security Guide

Prerequisites

To perform backup operations, the following requirements must be met:


- Backups can only be executed by database users that have the following system privileges:

BACKUP ADMIN	Authorization to perform backups.
CATALOG READ	Authorization to collect the information needed by the backup wizard.

More information:

[SAP HANA Database - SQL Reference Guide](#)

[SAP HANA Security Guide](#)

-  During installation, the SYSTEM user is created. The SYSTEM user has all system privileges, but is not allowed to see the data owned by other users. We do not recommend that the SYSTEM user be used for normal database operation. Instead, we recommend that you create dedicated administration users with only the privileges necessary for backup and recovery.

- The recovery process is executed as operating system user <sid>adm.

Points to Note

Below is an overview of important information to consider when planning your backup and recovery strategy with SAP HANA database. You can find more information on the individual points in the subsequent sections of this guide.

- Data and logs can only be backed up when the SAP HANA database is online (when all configured services are running).
- Only the database payload is backed up.
- Until an initial data backup has been completed, the logs are written in overwrite mode.
- During data and log backup, the system is available as usual.
- Backup only supports files as backup media.
- The configuration path for data and log backup must be valid throughout the whole system, and not only for specific hosts.
- Backup and recovery always applies to the whole database.
It is not possible to backup and recover individual database objects.
- To recover the database, you need at least one data backup.
- At the beginning of a recovery, all the data and log backups to be recovered must be available.
- To recover the SAP HANA database, the database needs to be shut down.
For this reason, during recovery, the database cannot be accessed by end users or applications.
- If a recovery fails, the complete recovery must be repeated.
- Shared storage must be used.
This is to ensure that the nameserver process can access the backup files at the time of recovery.
- The SAP HANA database software version used during the recovery must always be the same or higher than the version of the software used to create the backup.

Below is an overview of the backup and recovery capabilities for SAP HANA database:



Capability	Available
Data Backup	YES More information: About Data Area Backup
Log Backup	YES More information: About Log Area Backup
Backup of Configuration Files (.ini files)	NO More information: About the SAP HANA Configuration Files
Point-in-Time Recovery	YES More information: About SAP HANA Database Recovery
Recovery to a specific data backup (without using log information)	YES More information: About SAP HANA Database Recovery

About SAP HANA Database Backup

The database backup consists of the following:

- Backup of the data area
This is done from persistent storage to external backup locations. Backups are done either manually or scheduled.
- Backup of the log area
By default, logs are backed up automatically. You can configure whether and how the logs are backed up.

We recommend that you perform a data backup in the following situations:

- After the initial load.
- At regular intervals.
This ensures that a database recovery is possible.
- Before the database software is upgraded to a new version.
If a software upgrade fails, it is then possible to recover the database to its state before the upgrade.
- If a new host is added to the landscape.
 When a new host is added to the landscape, the redo log history is interrupted because the landscape has changed (“History Lost”). For this reason, it is no longer possible to use the redo log to recover the database. After a new host has been added, it will be possible to use the existing database and log backups to recover the database only to the point before which the host was added.
If you add a new host to the system, you should also stop production operation until the host has been added and a new complete data backup has been made.
- After any history lost alert.
For example:
 - Immediately after logging is turned on again.
 - After a data import.
 An IMPORT BINARY is not recorded in the redo log. For this reason, if you restore the database using a backup that was made before the import, the imported data will not be restored. In this situation, a history lost is recorded in the log.

About Data Area Backup


When the data area is backed up, all the payload data from all the servers is backed up. This happens in both single-host and multihost environments.

 SAP HANA database currently only supports backups to files.

Default Location of the Data Backup Files

The data backup files are written to the location specified by the parameter *basepath_databackup* in the *persistence* section of the *global.ini* configuration file. By default, the location for data backup files is $\$(DIR_INSTANCE)/backup/data$.

To use a different location, you can specify a different path when you perform the backup. If you need to, you can specify a different path for each backup. Alternatively, you can change the value of *basepath_databackup*. Go to the *Configuration* tab in the SAP HANA studio, choose *global.ini* -> *persistence*.

 If you change the backup location in *basepath_databackup*, the change takes effect immediately.

To ensure data security, we recommend that you specify a path to an external backup location. The backup location should never be on the same file system as the data or log areas.

 Note

All the files for a particular data backup are written to the same location. The files belonging to the same data backup cannot be written to multiple locations. Different data backups can be written to different locations, but all the files belonging to one particular data backup are written to the same location. We recommend that you create the directory structures before the backup is started.

Backup File Names

Each backup file name comprises the following elements: $\langle path \rangle \langle prefix \rangle \langle suffix \rangle$

The $\langle path \rangle$ and $\langle prefix \rangle$ are optional. If no complete path is specified, the default backup location is used.

You can specify a prefix for the backup file name, or you can use the prefix proposed by the system.

The system adds a unique suffix to each backup file name. As this is done for each service that is included in the backup, you only need to specify one file name prefix for all the backups on the different hosts.

The suffix that is appended to a file name prefix is only unique for each service. Consequently, the next time you back up a service, the system assigns the same backup suffix to the backup file for that service. If you do not change the file name, the existing backup file for that service will be overwritten by the new backup.

We therefore recommend that you copy a data backup to a new location as soon as it is created. Alternatively, specify a different file prefix or location when starting the next backup.

During the backup process, a backup file for each service is created in the backup location.

Example

Below is an example of a set of backup files from one data backup created with SAP HANA studio. The files can have different names.

```
COMPLETE_DATA_BACKUP_databackup_0_1 (name server topology)  
COMPLETE_DATA_BACKUP_databackup_1_1 (name server)  
COMPLETE_DATA_BACKUP_databackup_2_1 (for example, index server)  
...
```

In the above example, COMPLETE_DATA_BACKUP is the file name prefix; databackup_0_1 is the suffix.

Note

You can use older data backups for recovery, provided that the log backups that were created after the older data backups are available. If more log backups have to be replayed, the recovery will take longer. For that reason, to recover the database, we recommend that you use the most recent data backup and subsequent log backups.

The more frequently a database is backed up, the faster the recovery will be.

Checking Available Disk Space for Backups

The Administrator has to ensure that sufficient free space for the backup files is available. The amount of free space that will be needed in the backup directory needs to be calculated.

To estimate the size of a backup, you can use the system table `M_CONVERTER_STATISTICS` in the SQL Editor in the SAP HANA studio. This system table contains information about the used blocks.

To estimate the size of the next complete data backup, you can use either of the following commands:


- *select sum(allocated_page_size) from M_CONVERTER_STATISTICS*
The result is a single value that gives the sum of the sizes of all services in bytes. The size may differ between SELECT statement and DATA BACKUP execution. For this reason, it is advisable to include a reserve of free space.
- *select volume_id, sum(allocated_page_size) from M_CONVERTER_STATISTICS*
Group by volume_id
This displays a list of the volumes (index server, name server, statistics server), with the size of each volume in bytes.

About Log Area Backup

The system can perform regular log backups to allow the reuse of log segments. During a log backup, the payload of the log segments is copied from the log area to service-specific log backup files.

A log segment is backed up in the following situations:


- The log segment is full
- The log segment is closed after exceeding the configured time threshold
- The database is started


 If you do not regularly move the log backup files to an external destination, you run the risk of the file system becoming full. Log segments can only be overwritten by the system after they have been backed up.

Location of the Log Backup Files

The log backup files are written to the location specified by the parameter *basepath_logbackup* in the *persistence* section of the *global.ini* configuration file. By default, the location for log backup files is $\$(DIR_INSTANCE)/backup/log$.

To use a different location, change the value of *basepath_logbackup*. Go to the *Configuration* tab in the SAP HANA studio, choose *global.ini -> persistence*.

 If you change the backup location in *basepath_logbackup*, the change takes effect immediately.

 Note

You can release free log segments explicitly using the SQL command `ALTER SYSTEM RECLAIM LOG`. The `ALTER SYSTEM RECLAIM LOG` command physically deletes log segments that are no longer needed.

Log Backup Options

You can configure the log backup behavior using the following parameters in the `global.ini` configuration file:

`enable_auto_log_backup`


This parameter enables or disables automatic log backup.

Default: `enable_auto_log_backup = yes`


 Note


In the default `log_mode normal`, if automatic log backup is disabled, the log area grows until the file system is full. At that stage, the database will freeze.

`log_mode`


 After changing the log mode parameters, you must restart the database system to activate the changes. We also recommend that you create a full data backup of the database.

This parameter can have the following values:

- `log_mode = normal`
Normal mode (default).
Log segments are automatically backed up if parameter `enable_auto_log_backup` is enabled.
 `log_mode = normal` is recommended to provide support for point-in-time recovery.
- `log_mode = overwrite`
Overwrite mode.
Log segments are freed by savepoints and no log backup is performed. For example, this can be useful for test installations that do not need to be backed up or recovered.

 `log_mode = overwrite` is not recommended for production systems. No point-in-time recovery is possible. For recovery, only data backups are used; the logs must be disregarded.

- `log_mode = legacy`
Legacy mode.

 `log_mode = legacy` is still supported, but not recommended. After an upgrade from SAP HANA SPS 02 to SPS 03, it is recommended that you use `log_mode = normal`.

Log segments are retained until a full data backup is performed. This is to allow recovery from the most recent full backup + log in the log area. No log backup is performed. This was the default setting for HANA 1.0 SPS2.

Switching from `log_mode = legacy` to `normal`

When you switch `log_mode` from `legacy` to `normal`, check that the following parameters have the following values:

- `enable_auto_log_backup = yes`
- `log_backup_timeout_s > 0`
(for example, 900s)

A value of 900s ensures that if you need to recover the database, you will lose no more than 900s of data.

In SAP HANA studio, you can search for occurrences of each parameter. Go to the *Configuration* tab and use the *Filter* field.

Note

When you switch to `log_mode = normal`, log backups are created on a continuous basis. This requires a steadily increasing amount of disk space that could lead to a full disk if not handled properly. For this reason, we recommend that you change the paths where log and data backups are located to a separate storage medium with sufficient disk space. In addition to avoiding space issues, using a separate storage medium also leads to improved safety in case of disk failures, and better performance.

More information: SAP Note [1645183](#) - Changing log mode for SAP HANA DB after SPS2 to SPS3 upgrade

log_backup_timeout_s

This parameter forces log backups at a fixed time interval, specified in seconds. Log backups triggered by `log_backup_timeout_s` are performed in addition to the log backups that are performed when a log segment becomes full.

Recommended: Specify a time interval. For example, 900s.

(0 = Log backups are only made when a log segment is full and when services are restarted.)

Specifying an appropriate time interval for log backups enables you to recover a SAP HANA database with minimum data loss. For example, if you need to recover the database in a situation where the log area is unusable, and only the data and log backups are available.

Note

The `log_backup_timeout_s` parameter only takes effect if `enable_auto_log_backup` is set.

Avoiding LOG FULL (File System Full) Situations

If the log area becomes full and no more log segments can be created on the disk, a log full situation arises. When the log area is full, no more log entries can be written until the log backup has been completed.

1. When a log segment becomes full, the system internally tags it for backup.
2. After the system has backed up the full log segment, the system can reuse the space that the full log segment occupied on the log area (overwrite with new log entries).

In this way, automatic log backup can prevent log full situations from arising.

Caution

Do not delete log segments on operating system level, as the log area will become unusable and the database may stop working immediately.

More information: ALTER SYSTEM RECLAIM LOG in the [SAP HANA Database - SQL Reference Guide](#)

More information:

Location of the Log Backup Files

About the SAP HANA Configuration Files

The configuration files (.ini files) contain the SAP HANA database configuration settings.

The configuration files are not backed up as part of the database backup. Configuration files that contain customer-specific changes can be backed up manually in order to more easily identify and restore customer-specific changes in a recovery situation.

- ❗ The configuration files are not essential to perform a recovery. If you want to use the customer-specific configuration, you need to reconfigure the system using SAP HANA studio. To display the configuration values, go to the *Configuration* tab in SAP HANA studio.

More information:

Copying a Database Using Backup and Recovery

Location of the Configuration Files



The configuration files (.ini files) are located by default in the following directories:

Example Directory Paths	Example Configuration Files
For global configuration settings: <code>\$(DIR_INSTANCE)/../SYS/global/hdb/ custom/config</code>	global.ini indexserver.ini nameserver.ini
For host-specific configuration settings: <code>\$(SAP_RETRIEVAL_PATH)</code>	daemon.ini saprofile.ini

- ❗ **Note**
Configuration files are only created in these directories if customer-specific changes are made to them after installation. If no customer-specific changes have been made, these directories may be empty.

Configuration Files Created During Installation

When SAP HANA database is installed, the following customer-specific configuration files are created during installation:

saprofile.ini	<p>Contains system identification information, such as the system name (SID) or the instance number.</p> <p>saprofile.ini is not changed again after installation.</p> <p>saprofile.ini is not displayed in the <i>Configuration</i> tab in SAP HANA studio, but it is located in the following directory: <code>\$(DIR_INSTANCE)/../SYS/global/hdb/custom/config</code></p> <p> saprofile.ini contains information that is specific to each host. In a recovery situation, the saprofile.ini file must not be copied manually to a different host, as it will not be compatible with a new landscape.</p>
daemon.ini	<p>Contains information about which database services to start.</p>
nameserver.ini	<p>The nameserver.ini file contains global information for each installation. The landscape section contains the system-specific landscape ID and assignments of hosts to roles MASTER, WORKER, and STANDBY.</p> <p>If the system landscape is changed using hdbaddhost, the landscape section of nameserver.ini is also changed.</p> <p> In a recovery situation, the nameserver.ini file must not be copied manually to a different host, as it will not be compatible with a new landscape. If the landscape ID in the backup is different from the landscape ID in the nameserver.ini file, the recovery process will fail and it will not be possible to restart the database successfully.</p>

About the Backup Catalog

The backup catalog provides information about the execution of backups and their history. It enables the system to do the following:

- To determine whether a recovery is possible
- To choose which data and log backup to use to recover the database
- To determine which backup files are no longer needed

The backup catalog includes the following information:

- Backups and recoveries performed for a database
- The start and finish times of the backups
- Whether a backup is still running
- Whether a backup was successful or not
- Volumes that were backed up
- Log backups and which part of the log they contain
- Backup destinations and their sizes
- Whether the redo log history was interrupted


Note

The backup catalog does not show the progress of a backup.


Information in the backup catalog is stored as XML, making it readable by users and software. Once entries have been recorded, they cannot be changed.

The backup catalog is stored at the following location:

```
$ (DIR_INSTANCE) / ../SYS/global/hdb/metadata/BackupCatalog.xml
```

-  The backup catalog does not need to be backed up, as the system recreates it during a recovery using the accessible data and log backups.

More information: SAP Note [1705945](#) (Determining the files needed for a recovery)

 **Note:** backup.log

The backup.log file records information about the data and log backups.

To open backup.log:


1. From the SAP HANA studio, choose *Diagnosis Files*.
2. Search for backup.log in the *Filter* field.
3. Open the file.

Monitoring Views for the Backup Catalog

The monitoring views M_BACKUP_CATALOG and M_BACKUP_CATALOG_FILES can be used to display information about the backup catalog.

 **Note**

To use the M_BACKUP_CATALOG and M_BACKUP_CATALOG_FILES views, the system privilege BACKUP ADMIN or CATALOG READ is needed.

Monitoring View	Description
M_BACKUP_CATALOG	<p>Provides an overview of information about backup and recovery activities.</p> <p>Each row in the view provides information for a separate catalog entry that is identified by an entry ID. This information includes the type (complete data backup, log backup and log missing, recovery), and start and completion times.</p> <p> Note The backup ID is unique. The entry ID is unique, but can change when the backup catalog is rebuilt.</p>
M_BACKUP_CATALOG_FILES	<p>Provides information about the backup files created, and the backup destinations that are used by data and log backups.</p> <p>Each row in the view corresponds to a row in the M_BACKUP_CATALOG view. This view provides additional information about each database service that was involved in a backup. For example, with a data backup, each database service is listed with its specific backup information such as destination path and redo log position.</p>

Identifying Old Log Backups

The following system views provide information about the data and log backups:

- M_BACKUP_CATALOG
- M_BACKUP_CATALOG_FILES

Each service-specific data backup file in the complete data backup has its own REDO_LOG_POSITION. This service-specific redo log position identifies the log backup that belongs to the service-specific data backup file. The log position for each service specific data backup file is contained in M_BACKUP_CATALOG_FILES.REDO_LOG_POSITION.

Additionally, each service-specific log backup contains the range of log entries that are in the backup.

- The oldest log entry in a log backup is specified by
M_BACKUP_CATALOG_FILES.FIRST_REDO_LOG_POSITION
- The most recent entry in a log backup is specified by
M_BACKUP_CATALOG_FILES.LAST_REDO_LOG_POSITION

If you want to delete old log backups, you have to decide which complete data backup will be the oldest that you potentially want to recover. You can use the REDO_LOG_POSITION to find all the older service-specific log backups. Using M_BACKUP_CATALOG_FILES, search for all service specific log backups, for which the LAST_REDO_LOG_POSITION is smaller than the REDO_LOG_POSITION of the corresponding data backup.

Example

This is an example statement to find older service-specific log backups:

```
select destination_path from m_backup_catalog_files o where
o.last_redo_log_position is not null and exists(select 1 from
m_backup_catalog_files i where i.backup_id = <backup_id> and
i.source_id = o.source_id and i.redo_log_position >
o.last_redo_log_position) order by destination_path
```

Note

The backup catalog is not changed by this statement. For that reason, if the statement is executed multiple times, it still returns files that may already have been deleted.

Note

To find the <backup_id> of the requested complete data backup, query the monitoring view M_BACKUP_CATALOG .

Backing Up the SAP HANA Database

The sections that follow describe how to perform a data backup using different tools:

- SAP HANA studio
- SQL commands
- Batch mode

Note

You can also schedule periodic backups using DBA Cockpit. To work with DBA Cockpit, you need a compatible ABAP system.

More information: DBA Cockpit documentation in the SAP Help Portal

<http://help.sap.com/>

Performing a Data Backup Using SAP HANA Studio

To create a data backup, perform the following steps:

1. In the Navigator view, select the system for which you want to start a backup.
2. From the context menu, choose *Backup...*
3. Specify the location (directory) and the backup file prefix to use.
The default location shows the path specified in *global.ini* under the backup parameter *basepath_databackup*.

Note

A data backup performed with SAP HANA studio only saves the payload of the data volumes of the database. The database configuration files (*.ini files) are not backed up. Configuration files (.ini files) that contain customer-specific changes can be backed up manually in order to more easily identify and restore customer-specific changes in a recovery situation. .

More information: About the SAP HANA Configuration Files

4. When all the settings are correct, choose *Finish*.
The backup then starts.
The progress of the backup is shown for all types of services (for example, the statistics server, name server, and index servers).

When all the volumes have been backed up, a confirmation message is displayed.

Performing a Data Backup Using SQL Commands

You can enter SQL commands either by using the SQL editor in SAP HANA studio, or by using the **hdbsql** program on the command line.

Note

Backups using SQL commands are only recommended for batch mode (see section Using Batch Mode below).

The backup can be started with the following SQL command:

```
BACKUP DATA USING FILE ('<target>')
```

If no complete path is specified, the default backup location is prepended to the target.

In both cases, a suffix will be appended to the target.

Example

```
BACKUP DATA USING FILE  
( '/backupDir/data/MONDAY/COMPLETE_DATA_BACKUP' )
```

This would create the following files in the directory `/backupDir/data/MONDAY`:

```
COMPLETE_DATA_BACKUP_databackup_0_1 (name server topology)  
COMPLETE_DATA_BACKUP_databackup_1_1 (name server)  
COMPLETE_DATA_BACKUP_databackup_2_1 (for example, index server)  
...
```

Note

If the directory does not exist, we recommend that you create it.


More information: SAP Note [1651055](#) - Scheduling SAP HANA Database Backups in Linux


More information:

Backup File Names

Performing a Data Backup in Batch Mode

To perform data backups in batch mode, we recommend that you use the command line tool `hdbsql` on operating system level. `hdbsql` enables you to trigger backups through `crontab`.

 For batch processing, `hdbuserstore` is required. `hdbuserstore` is installed as part of the client package. We recommend that you set up a batch user with backup privileges (BACKUP ADMIN).

 **Caution**
We do not recommend that you start the backup command using `hdbsql` on operating system level without `hdbuserstore`. This is because the password used for `hdbsql connect` will be visible in the process list.

Setting Up a Data Backup Using Cron

1. Install the client software.

The client software enables access to the `hdbuserstore`.

Use the following command:

```
hdbinst -a client (default location: /usr/sap/hdbclient)
```

2. Create a user key.

Use the following command:

```
/usr/sap/hdbclient/hdbuserstore set <KEY>  
<host>:3<instance id>15 <user> <password>
```

Example:

```
/usr/sap/hdbclient/hdbuserstore set BACKUP  
vebwtests1:30015 user password
```

3. In `crontab`, call the following at the desired time:

```
/usr/sap/hdbclient/hdbsql -U<KEY> "BACKUP DATA USING FILE  
( '<target>' ) "
```

Example:

```
/usr/sap/hdbclient/hdbsql -U BACKUP "BACKUP DATA USING  
FILE ('MONDAY') "
```

A data backup is then created in the default location. In the above example, the prefix of all service-related backup files is "MONDAY".

Canceling a Running Data Backup

You can cancel a data backup that is still running. To cancel a running data backup, you can use either SAP HANA studio or an SQL command.

Prerequisites

A user needs the system privileges CATALOG READ and BACKUP ADMIN.

Canceling a Running Data Backup Using SAP HANA Studio

You can use SAP HANA studio to cancel running data backups that you started using the backup wizard. To cancel a running data backup, choose *Cancel*. The backup is then canceled and you are notified of this.

Canceling Other Running Data Backups

To cancel any data backup:

1. Find the Backup ID.

When a data backup is started, the system assigns a unique ID to the data backup.

To find the backup ID of the running data backup, use the monitoring view M_BACKUP_CATALOG, which provides an overview of information about backup and recovery activities.

To find the backup ID of the running data backup, use the following SQL command:

```
select BACKUP_ID from "SYS"."M_BACKUP_CATALOG" where  
entry_type_name = 'complete data backup' and state_name =  
'running' order by sys_start_time desc;
```

You can now use the backup ID to cancel the running data backup.

2. Cancel the data backup.

To cancel the running data backup use the following SQL command:

```
backup cancel <backup_id>
```

3. Check that the data backup was canceled.

The data backup is canceled asynchronously. To check that the running data backup was canceled, check the monitoring view M_BACKUP_CATALOG. This view shows the state of the canceled data backup.

If its state is "canceled", this confirms that the backup was canceled successfully. When the backup state is "canceled", you can start a new data backup.

 **Note: "cancel pending"**

If you cancel a running data backup that is waiting for other resources, the cancelation is postponed until the cause of the situation has been resolved. Until that time, the backup is flagged as "cancel pending".

To check the state of the data backup, use the following command:

```
select state_name from "SYS"."M_BACKUP_CATALOG" where  
backup_id = <backup_id>
```

About SAP HANA Database Recovery

It may be necessary to recover the SAP HANA database in the following situations:

- A disk in the data area is unusable.
- A disk in the log area is unusable.
- As a consequence of a logical error, the database needs to be reset to its state at a particular point in time.
- If you want to create a copy of the database

Recovery Scenarios

The steps to recover the database depend on the recovery scenario and the reason for the recovery. This section provides an overview of recovery scenarios.

Data Area is Unusable

If the data area is unusable, and all the data changes after the last complete data backup are still available in the log backups and log area, the data from committed transactions that was in memory at the time of failure can be recovered.

No committed data is lost.

For recovery, the data backups, the log backups, and the log area are used. When the data backup has been successfully restored, the log entries from the log backups and the log area are automatically replayed.

It is also possible to recover the database using an older data backup and log backups. All relevant log backups made after the data backup are needed for the recovery.

More information: SAP Note [1705945](#) (Determining the files needed for a recovery)



Note

In the Recovery Wizard dialog, ensure that the paths to the data and log backup files are correct.

More information:

Default Location of the Data Backup Files

Location of the Log Backup Files

Used for Recovery

- Data backups
- Log backups
- Log area

Steps for Recovery

Perform a database recovery to the most recent state.

More information: Recovering the Database

Log Area is Unusable

If the log area is unusable, it is only possible to replay the log backups. As a consequence, any changes that were made after the most recent log backup will be lost. In addition, all the transactions that were open during the log backup will be rolled back.

It is still possible to recover the database to a point in time within the existing log backups.

More information: `log_backup_timeout_s`

For recovery, the data backups and the log backups are used. When the data backup has been successfully restored, the log entries from the log backups are automatically replayed. In the Recovery Wizard, you must specify the option *Initialize log area* to prevent the recovery of entries from the unusable log area.

Initialize log area corresponds to the CLEAR LOG command.

Used for Recovery

- Data backups
- Log backups

Steps for Recovery

1. Perform a database recovery to the most recent state.
2. Specify the option *Initialize log area*.

Logical Error – Point in Time Recovery

To reset the database to a particular point in time, you need a data backup from before the point in time to recover to, the subsequent log backups, and the log area.

All changes made after the recovery time will be lost. If you need to perform this recovery, consider recovering the database to a different system.

Used for Recovery

- Data backups
- Log backups
- Log area

Steps for Recovery

Perform a database recovery to a point-in-time before the logical error.

Note

You need to specify a point in time to recover the database to. If you specify a point in time in the future, the effect is the same as recovering the database to the most recent state.

More information:

Copying a Database Using Backup and Recovery

Copying a Database

More information:

Copying a Database Using Backup and Recovery

Recovering the SAP HANA Database

The sections that follow describe the requirements and the steps to recover the SAP HANA database.

Prerequisites

To perform a SAP HANA database recovery, the following requirements must be met:

- The SAP HANA database software must be installed, so that an initial database exists. In a recovery situation, you can use SAP HANA studio to restore customer-specific changes to this initial database.

Note

If you want to restore customer-specific configuration settings, you can do this either before you restore the database and the log backups or at the end of the recovery.


- Ensure that the target system and the source system have identical configurations. The number of hosts and the number and types of services (for example, index server) on each host must be identical for both system landscapes.
- All the data and log backups must be available in the backup directories.
- At least one data backup must be available before the recovery is started.
- To restore the database to a particular point in time, a data backup and all the log backups up to the point in time for recovery are needed.

Constraints

- Recovery to a lower system release is not possible.
- Once it is running, a recovery cannot be canceled.
- If an error occurs during a recovery, the complete recovery must be repeated.

Recovering the Database

The recovery process must be started using SAP HANA studio.


-  To recover the SAP HANA database, the database needs to be shut down. For this reason, during recovery, the database cannot be accessed by end users or applications.



Procedure

To recover a SAP HANA database, perform the following steps:

1. From SAP HANA studio, open the context menu over the database.
2. Choose *Recovery...*
A dialog box is displayed, requesting that you confirm that the system can be shut down for the recovery.
3. Confirm and choose *Next*.
4. Specify the recovery type.

The following recovery types are available:

Recovery Type	Description
Recover the database to its most recent state	<p>This option recovers the database to as close as possible to the current time.</p> <p>This recovery option uses the following data:</p> <ul style="list-style-type: none"> • The most recent data backup • Log backups made since the most recent data backup • Log area
Recover the database to the following point in time	<p>This recovery option uses the following data:</p> <ul style="list-style-type: none"> • The last data backup available before the specified point in time • Log backups made since the data backup to be used • Log area <p> Note If you specify a point in time in the future, the effect will be the same as recovering the database to the most recent state.</p>

Recover the database to a specific data backup	<p>This recovery option uses the following data:</p> <ul style="list-style-type: none"> • The specified data backup <p> Note Log entries are not replayed, neither from the log backups nor from the log area. All log entries that still exist in the log area are deleted.</p>
Recover the database to the following log position	<p> Note This recovery type is an advanced option that can be used in exceptional cases if a previous recovery failed.</p> <p>This recovery option uses the following data:</p> <ul style="list-style-type: none"> • The most recent data backup available before the specified log position • Log backups made since the data backup to be used • Log area

5. Choose *Next*.
6. Specify the locations of the data and log backup files if they are not in the default location.

 **Note**

When you perform a recovery, you can specify the option *Initialize log area*. *Initialize log area* corresponds to the CLEAR LOG command.

With this option, no log entries from the log area are replayed, and the log area is initialized. The log entries from the log backups are replayed if they are needed.

In the following situations, you **must** specify the option *Initialize log area*:

- When the log area is unusable
- If you recover the database to a different system

7. Choose *Next*.
A summary of the selected recovery options is displayed.
8. If the settings are correct, choose *Finish*.
The recovery is then started.
The progress of the recovery is displayed in the dialog box.
When the recovery is complete, the system is online.

Copying a Database Using Backup and Recovery



You can create a homogenous copy of a database by recovering an existing source database backup to a different, but compatible target database. The source database backup consists of data backup files, and the log backup files.

A homogenous database copy is a quick way to set up cloned systems for training, testing, and development. For this reason, it can significantly reduce total cost of delivery (TCD).

Prerequisites


- A database backup (data backup and optionally log backups) of the source database at a dedicated backup location.
- The version of the target database system is the same or higher than the source database.
- The target system must be configured with sufficient disk and memory capacity.
- Ensure that the target system and the source system have identical configurations. The number of hosts and the number and types of services (for example, index server) on each host must be identical for both system landscapes.
- Customer-specific changes can be manually applied to the target system.

More information: About the SAP HANA Configuration Files

-  Do not change any entries in `nameserver.ini`, section `landscape`.
Do not change any entries in `sapprofile.ini`.
- Ensure that a license key file is available for the target database.
 -  The old license key will not work in the target system.

Procedure

1. Shut down the target database.
2. Move or delete available data and log backup files from the target database.
3. Move or delete the backup catalog files, if they exist:
`$(DIR_INSTANCE)/../SYS/global/hdb/metadata/BackupCatalog.xml`
`$(DIR_INSTANCE)/../SYS/global/hdb/metadata/BackupCatalog_Recover.xml`
4. Copy the data and optionally log backup files from the source database to the corresponding directories in the target database.

 **Note**

The directories for data and log backup are defined by the parameters `basepath_databackup` and `basepath_logbackup` in the `global.ini` file.

5. Start the recovery of the target database system.
In SAP HANA studio, select the target database system and choose *Recovery...* from the context menu.
6. Select the recovery type *Recover the database to its most recent state*.
7. Choose *Next*.
8. Check that the specified default locations of the data and log backup files for the target database matches the backup files that you have copied.
9. Select the option *Initialize log area*.
10. Confirm the warning.
11. Select the option *Install new license* and specify the license file.
12. Choose *Next*.
13. Choose *Finish* to start the recovery.
When the recovery is successfully completed, the database is started.

Result

The source database has now been recovered to the target database. The target database is online.

You can now work with the restored database.

Backup Checks


The following checks warn you of errors that are connected with the backup process:

Check	Check availability of volumes for backup
Alert ID	34
Description	Before a backup is performed, this check warns you if that backup cannot be performed because a volume is unavailable. This alert can be triggered in combination with the alerts NOT_ASSIGNED_VOLUMES and CHECK_INACTIVE_SERVICES.
Alert Text	A backup cannot be performed because volume <no> (<service>) is not available or because the service assigned is not in ACTIVE_STATE=YES
User Action	Find out why the volume is not available.
Default Interval	1 hour

Check	Check whether a data backup exists
Alert ID	35
Description	Checks whether at least one data backup exists and warns you if no successful data backup is available for the instance. You are warned before any actual data loss occurs.
Alert Text	No data backup exists.
User Action	To ensure that your database can be recovered, perform a data backup as soon as possible.
Default Interval	24 hours

Check	Check last data backup
Alert ID	36
Description	Checks whether the last data backup was successful and warns you if the last data backup failed. If a scheduled backup fails, this check can help you prevent a situation from arising where no current backups are available.
Alert Text	The last data backup was not successful
User Action	Check the reason why the last data backup was not successful, resolve the problem, and perform a new data backup as soon as possible.
Default Interval	24 hours

Check	Check the age of the last data backup
Alert ID	37
Description	Checks the age of the last successful data backup, and provides different levels of notification if the last successful data backup is too old.
Alert Text	The last data backup is already <days> days old.
User Action	To reduce your downtime in a recovery situation, perform a data backup as soon as possible.
Default Interval	24 hours

Check	Check last log backups
Alert ID	38
Description	<p>Checks whether the last log backups were successful, and provides information about a failed log backup for a service or volume.</p> <p>As log backups are performed automatically, this is the only way to notify users. For this reason, it is important that this check be performed frequently and with a high priority.</p> <p> Note Log backups are performed individually and independently for each volume. For this reason, log backups need to be checked for each volume.</p>
Alert Text	The last log backup was not successful for volume <no> (<service> at <host>:<port>)
User Action	Check the reason why a log backup was not successful and resolve the problem.
Default Interval	15 minutes

More information:

Monitoring Alerts in the SAP HANA Database – Administration Guide

SQL Syntax for Backup and Recovery

In addition to performing backup and recovery operations using the SAP HANA studio, you can also use SQL statements. This section details the syntax for these statements.

Prerequisites

To execute the following commands, the database system must be ONLINE, and the user requires the privilege BACKUP ADMIN:

- <execute_backup_data>
- <execute_backup_check>
- <execute_backup_cancel>

You can pass the following command to recoverSys.py.

- <execute_recover_database>

Commands for Data Backup

If no absolute path is specified for the data backup statement, the data backup will be created in the default location. The default location is specified by the database configuration parameter *basepath_databackup* in the *global.ini* configuration file.

Each complete data backup includes all the configured services (for example, name server, index server, and statistics server).

The name of a data backup file consists of an optional prefix, the string 'databackup', the volume ID of the affected service and an index.

<execute_backup_data>	::= BACKUP DATA <data_file_definition> BACKUP DATA <data_file_definition> <backup_option>
<data_file_definition>	::= USING <file_definition> ALL USING <file_definition>
<file_definition>	::= FILE (<file_name_prefix>)
<file_name_prefix>	::= <apostrophe> <string_const> <apostrophe>
<backup_option>	::= ASYNCHRONOUS

<apostrophe>	::= ' '
--------------	---------

Table: Syntax of Commands for Data Backup

Examples of Commands for Data Backup

- `BACKUP DATA ALL USING FILE ('COMPLETE')`

Perform a data backup for the entire database and create a data backup file for every relevant service in the default backup location.

Data backup files are prepended with the string `COMPLETE`.

The following files are created:

- `<default backup directory>/COMPLETE_databackup_0_1` (contains topology)
- `<default backup directory>/COMPLETE_databackup_1_1` (contains name server persistency)
- `<default backup directory>/COMPLETE_databackup_2_1` (normally contains index server persistency)
- `<default backup directory>/COMPLETE_databackup_3_1` (normally contains statistic server persistency)
- `BACKUP DATA ALL USING FILE ('/backup/MONDAY')`

Perform a data backup for the entire database and create a data backup file for every relevant service in the directory `/backup`.

Data backup files are prepended with the string `MONDAY`.

The following files are created:

- `/backup/MONDAY_databackup_0_1` (contains topology)
- `/backup/MONDAY_databackup_1_1` (contains name server persistency)
- `/backup/MONDAY_databackup_2_1` (normally contains index server persistency)
- `/backup/MONDAY_databackup_3_1` (normally contains statistic server persistency)
- `BACKUP DATA ALL USING FILE ('/backup/') ASYNCHRONOUS`

Perform a data backup for the entire database and create a data backup file for every relevant service in the directory `/backup`.

Data backup files are not labeled.

The `ASYNCHRONOUS` option returns to the client immediately after activating the backup operation, even though the data backup is still running.

The following files are created:

Example of Commands for Backup Cancel

```
BACKUP CANCEL 1331715084250
```

More information:

Canceling a Running Data Backup

Data Snapshot

SAP HANA database provides the capability to clone an entire database while online.

Cloning is done using the underlying storage system. For consistency, a database-wide snapshot of the data area is used, similar to the procedure for data backup.

After the database has been cloned, the snapshot is removed from the source database. The snapshot in the cloned database is restored during the first restart.

Then, the snapshot in the destination database is removed.

<execute_create_snapshot>	::= BACKUP DATA CREATE SNAPSHOT
<execute_drop_snapshot>	::= BACKUP DATA DROP SNAPSHOT

Table: Syntax of Commands for Data Snapshot

Examples of Commands for Data Snapshot

```
BACKUP DATA CREATE SNAPSHOT
```

Create a database-wide snapshot based on a transactional consistent savepoint similar to the data backup.

If a snapshot exists, no complete data backup is possible.



Every request of this kind is rejected with a notification that a data backup is still in process.

```
BACKUP DATA DROP SNAPSHOT
```

Drop an database-wide snapshot. From this point in time on, complete data backups are possible.

Database Recovery and Data Recovery

This section describes the syntax of SQL statements for database (data and log) and data recovery.

<execute_recover_database>	<pre>::= RECOVER DATABASE <until_definition> [<source_definition>] <execute_recover_data></pre>
<execute_recover_data>	<pre>:: = <execute_recover_normal> <execute_recover_legacy></pre> <p> This statement does not recover the redo log that was backed up before.</p>
<execute_recover_normal>	<pre>::= RECOVER DATA <data_file_definition> <recover_normal_option></pre>
<execute_recover_legacy>	<pre>::= RECOVER DATA <data_file_definition> [<recover_option>]</pre> <p> The <execute_recover_legacy> statement can be used only when the log mode is "legacy" or "overwrite": This is because both log modes simulate the log behavior in SAP HANA < SAP HANA 1.00.11.</p>
<until_definition>	<pre> UNTIL TIMESTAMP <timestamp> [<recover_option>] UNTIL LOG POSITION <int_const> AT VOLUME <int_const> [<recover_option>]</pre>
<timestamp>	<pre>::= 'YYYY-MM-DD HH:MM:SS'</pre>
<source_definition>	<pre>::= <data_source_definition> [<log_source_definition>] <log_source_definition> [<data_source_definition>]</pre>
<data_source_definition>	<pre>::= CHANGE ALL DATA USING PATH (<source_path_list>)</pre>
<log_source_definition>	<pre>::= CHANGE ALL LOG USING PATH (<source_path_list>)</pre>

<source_path_list>	::= <source_path> <source_path_list> <comma> <source_path>
<source_path>	::= <apostrophe> <string> <apostrophe>
<data_file_definition>	::= USING <file_definition> ALL USING <file_definition>
<file_definition>	::= FILE (<file_name_prefix>)
<file_name_prefix>	::= <apostrophe> <string_const> <apostrophe>
<recover_option>	::= <recover_option_clearlog> <recover_option_license>
<recover_normal_option>	::= <recover_option_clearlog> [<recover_option_license>] [<recover_option_license>] <recover_option_clearlog>
<recover_option_clearlog>	::= CLEAR LOG
<recover_option_license>	::= SET LICENSE <license>
<license>	::= <apostrophe> <string_const> <apostrophe>
<comma>	::= ,
<apostrophe>	::= '

Table: Syntax of Commands for Database Recovery and Data Recovery

Examples of Commands for Database Recovery and Data Recovery

Example Scenario 1

The last complete data backup was created on Monday 22nd at 8 am CEST (6am UTC). The database ran up to 6pm CEST (4 pm UTC) and must be recovered up to 5pm (3 pm UTC).


Use the following command:

```
RECOVER DATABASE UNTIL TIMESTAMP '2011-08-22 15:00:00'
```



Note

The timestamp given to the database is interpreted as UTC time. Therefore, 3pm means UTC and not 3pm CEST.

 Note that the specified timestamp is not included in the transactions to be redone. For that reason, you would need to redo all transactions with a commit timestamp that is before the target timestamp.

Example Scenario 2

Recover the database until a point in time with a timestamp using a dedicated directory for data backups and further directories containing automatically written log backups.

Use the following command:

```
RECOVER DATABASE UNTIL TIMESTAMP '2011-08-22 15:00:00' CHANGE ALL
DATA USING PATH ('/backup/MONDAY/') CHANGE ALL LOG USING PATH
('/backup/logs1/', '/backup/logs2/')
```

Example Scenario 3

Perform data recovery for the entire database without replaying the redo log. The data backup files will be taken from the specified directory `/backup`.

Data Backup files have to be labeled with prefix `MONDAY`.

```
RECOVER DATA ALL USING FILE ('/backup/MONDAY') CLEAR LOG
```