

Header Data

Released On 15.09.2015 17:54:48

Release Status Released for Customer

Component HAN-DB-BAC SAP HANA Backup & Recovery

Priority Recommendations / Additional Info

Category FAQ

Symptom

To ensure optimal performance, 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. If there is 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.

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.

This SAP Note provides answers to frequently asked questions about backup and recovery of SAP HANA database. It does not describe how to perform a backup and recovery of all the components that can be part of a SAP HANA installation.

1. Where can I find detailed information about SAP HANA backup and recovery?
2. Which users are allowed to back up and recover a SAP HANA database?
3. Do I have to stop the SAP HANA database to perform a backup?
4. Which SAP HANA backup types are there?
5. Are there any preconditions to be checked before I can start a data backup?
6. How do I perform a SAP HANA data backup?
7. What happens if a scheduled data backup starts before the previous backup is completed?
8. Is there a difference between the backup procedure in single-host and multi-host (scale-out) SAP HANA databases?
9. Is the log area also saved during a data backup?
10. Can data and log backups run in parallel?
11. What data is saved in data and log backups?
12. Can I back up or recover individual database objects such as tables or schemata?
13. Are third-party backup tools supported?
14. With a file-based backup, can I configure the directories where the backups are written?
15. Does SAP HANA database support file-based backups to NFS-mounted file systems?
16. What are the naming conventions for data backup files?
17. How can I estimate the size of a data backup?
18. Where can I see a history of my data and log backups?
19. Can I cancel a running data backup?
20. Are backups compressed?
21. Are incremental and differential backups available?
22. How can I check the integrity of data and log backups?

23. How often should I execute a SAP HANA data backup?
24. Are the configuration files also saved during a backup?
25. Are the configuration files needed for a recovery?
26. What general configuration options do I have for saving the log area?
27. How can I delete log segments that are no longer needed from the log volumes?
28. Where can I find the backup catalog?
29. How can I identify data and log backup files needed for a recovery?
30. How can I identify old log backups?
31. How do I recover my SAP HANA database?
32. How do I recover my SAP HANA database if the data area is unusable?
33. How do I recover my SAP HANA database if the log area is unusable?
34. How do I recover my SAP HANA database if a logical error occurred (for example, some data was deleted by mistake)?
35. Is online recovery possible?
36. Where can I see that my recovery is finished?
37. Can I cancel a running recovery?
38. What do I have to do if the recovery fails?
39. What information is needed to analyze backup or recovery problems?
40. Does the recovery tolerate corrupt or missing backups?
41. Can backup and recovery be resumed?
42. Can I copy a SAP HANA database using backup and recovery?
43. Are backups encrypted?
44. Can backups be recovered to a SAP HANA database with a higher software version?
45. Can I recover a HANA backup made on INTEL architecture on IBM Power architecture?

Other Terms

HANA; Backup; Recovery; FAQ;

Reason and Prerequisites

You use SAP HANA database revision 46 or higher.

Solution

1. Where can I find detailed information about SAP HANA backup and recovery?

You can find further details about backup and recovery in the SAP HANA Administration Guide on http://help.sap.com/hana_appliance.

2. Which users are allowed to back up and recover a SAP HANA database?

SQL users require the system privilege BACKUP ADMIN or BACKUP OPERATOR to perform a data backup. If you execute data backups using SAP HANA studio, the privilege CATALOG READ is also needed. It is not recommended to use the SYSTEM user 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 can only be started by the operating system user <sid>adm.

3. Do I have to stop the SAP HANA database for backup?

No.

Data and log areas can only be backed up when the SAP HANA database is online. During data and log backup, the system is available as usual.

4. Which SAP HANA backup types are there?

SAP HANA database support full data backups as well as delta data backups. A full data backup is a complete data backup or a data snapshot of the SAP HANA data area. A delta data backup covers incremental and differential data backups. In addition to this log backups are supported too.

5. Are there any preconditions to be checked before I can start a data backup?

All the configured services (see SAP Note 1697613 and SAP Note 1649519) such as indexserver, statisticsserver, etc. must be running. You can check this in the Overview of SAP HANA studio -> Operational State: All Services are started.

6. How do I perform a SAP HANA data backup?

Basically each data backup can be executed using an arbitrary SQL interface like hdbsql. In addition SAP HANA database provides with the SAP HANA studio and the web-based SAP HANA cockpit two powerful UIs to fulfil that task. In addition DBA cockpit (ABAP stack only) can be used too.

You will find more detailed information in the SAP HANA Administration Guide and the DBA cockpit for SAP HANA Guide.

7. What happens if a scheduled data backup starts before the previous backup is completed?

The first data backup will continue normally; the second data backup will not start and an error message will be displayed.

8. Is there a difference between the backup procedure in single host and multi-host (scale-out) SAP HANA databases?

No.

There is no difference. All hosts will be backed up simultaneous.

9. Is the log area also saved during a data backup?

No.

A SAP HANA data backup saves the data area only.

After an initial data backup, the log area is automatically saved to log backups at regular intervals.

10. Can data and log backups run in parallel?

Yes.

If the auto log backup parameter is set, log backups can be started automatically and independently of data backups.

11. What data is saved in data and log backups?

When the data area or log area is backed up, the entire payload data from all server nodes of the SAP HANA database instance is backed up. This applies to both single-host and multi-host environments.

12. Can I back up or recover individual database objects such as tables or schemata?

No.

Backup and recovery always applies to the whole database. It is not possible to back up or recover individual database objects.

13. Are third-party backup tools supported?

Yes.

In addition to backing up the SAP HANA database to files in the file system, you can back up and recover the SAP HANA database using a third-party backup tool that supports the Backint for SAP HANA interface. Third-party backup tools communicate with the SAP HANA database through the Backint for SAP HANA interface. Backint for SAP HANA uses pipes to back up the content of the database and performs all the actions needed to write the backup data to external storage.

Note that backup tools using Backint for SAP HANA must be certified by SAP. For more information, see SAP Note 1730932 and 2031547.

14. With a file-based backup, can I configure the directories where the backups are written?

Yes.

The location of file-based data and log backups can be changed.

For improved data safety, it is recommended that you specify a path to an external backup destination. The backup destination should never be on the same file system as the data or log areas. It is recommended that you create the destination directory structures before the backup is started.

a) Data backup

By default, data backups are written to the following destination: \$DIR_INSTANCE/backup/data. You can specify a different destination when you perform the backup. Alternatively, you can change the default backup destination in the Backup section in SAP HANA studio.

To change the default destination for data backups:

1. In SAP HANA studio, go to Backup.

2. Go to the Configuration tab and choose Data Backup Settings.
3. Specify the new default destination in the Destination field.
4. Save.

Different data backups can be written to different destinations, but all the parts of one particular data backup are written to the same destination.

b) Log backup

By default, log backups are written to the following destination: \$DIR_INSTANCE/backup/log. You can change the backup destination in the Backup editor in SAP HANA studio.

To change the default destination for log backups:

1. In SAP HANA studio, go to Backup.
2. Go to the Configuration tab and choose Log Backup Settings.
3. Specify the new default destination in the Destination field.
4. Save.

In general changes to the data or log backup destinations take effect immediately.

For more information, see the SAP HANA Administration Guide.

15. Does SAP HANA database support file-based backups to NFS-mounted file systems?

Yes.

Backup to a NFS mounted file system is supported.

16. What are the naming conventions for data backup files?

Each backup name comprises the following elements: <path><prefix>

If no complete **path** is specified, the default backup destination is prepended to the backup name. For file-based backups, the backups are written to a destination in the file system. For backups made using third-party backup tools, a pipe is created in the file system. The pipe is always created in the directory /usr/sap/<SID>/SYS/global/hdb/backupint. The backup tool decides where the backups are actually written.

You can use the **prefix** proposed by the system or you can specify a different prefix for the backup name. During backup, each service backs up its data to the specified backup destination.

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

- /backup/data/COMPLETE_DATA_BACKUP_databackup_0_1 (name server topology)
- /backup/data/COMPLETE_DATA_BACKUP_databackup_1_1 (name server)
- /backup/data/COMPLETE_DATA_BACKUP_databackup_2_1 (for example, index server)...

In the above example, the <path> is /backup/data/, the <prefix> is COMPLETE_DATA_BACKUP. databackup_0_1 is the suffix, which is automatically added by the system.

17. How can I estimate the size of a data backup?

To estimate the size of a data backup, you can query the system table M_CONVERTER_STATISTICS. This system table contains the number of currently used blocks. To estimate the size of the next complete data backup, you can use either of the following SQL 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 slightly between SELECT statement and DATA BACKUP execution.

18. Where can I see a history of my data and log backups?

The backup catalog provides information about the execution of backups, recoveries and their history. It enables the system to determine whether a recovery is possible and to choose which data and log backup to use to recover the database.

The monitoring views M_BACKUP_CATALOG and M_BACKUP_CATALOG_FILES can be used to display information about the backup catalog. The system privilege CATALOG READ is needed for this.

For more information, see the SAP HANA Administration Guide.

19. Can I cancel a running data backup?

Yes.

Starting with SAP HANA database revision 28 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 - provided the user has the system privileges CATALOG READ and BACKUP ADMIN. This feature is available with revision 28 onwards.

For more information, see the SAP HANA Administration Guide.

20. Are backups compressed?

Currently, SAP HANA database does not offer backup compression. However, third-party backup tools normally offer backup compression. If you are using a third-party backup tool, contact your tool vendor about what compression options are offered.

21. Are incremental and differential backups available?

Yes.

Starting with HANA SPS 10 incremental and differential data backups are available. It's also possible to combine both kinds of backups in one backup strategy, e.g. an incremental backup can be created based on a complete data backup, a differential or an incremental data backup.

22. How can I check the integrity of data and log backups?

When a SAP HANA data or log backup is performed, the integrity of the data to be backed up is checked automatically.

In addition, when a recovery is performed, the integrity of the data to be recovered is checked automatically.

SAP HANA also allows you to manually check the integrity of data and log backups without performing a recovery. More information: SAP Note 1869119

23. How often should I execute a SAP HANA data backup?

- After the initial load has finished successfully, a data backup is strongly recommended.
- It is recommended to perform a backup (including configuration files) before the database software is upgraded to a new version. If the software upgrade fails, the database can then be recovered with no data loss.
- After adding or removing a host or a service, a data backup is recommended.
- To ensure that you can recover your database to the most recent point in time without needing to replay a large number of log backups, you should perform data backups at regular intervals.

Archiving of log backups should be part of your overall backup strategy.

24. Are the configuration files also saved during a backup?

No.

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.

Configuration files that contain customer-specific changes are located by default in the following directories:

For global configuration settings:

`$DIR_INSTANCE/.. /SYS/global/hdb/custom/config`

For host-specific configuration settings:

`$SAP_RETRIEVAL_PATH`

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.

25. Are the configuration files needed for a recovery?

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

26. What general configuration options do I have for saving the log area?

You can configure the log backup behaviour using the Backup editor in SAP HANA studio. Note that you cannot configure the the log mode using the Backup editor. To configure the log mode, you have to change the configuration parameter in the global.ini configuration file.

Automatic Log Backup

Automatic log backup can be enabled or disabled using parameter `enable_auto_log_backup`. By default, automatic log backup is enabled.

During normal system operation (log mode normal), it is recommended that automatic log backup be kept activated. If automatic log backup is disabled, the log area grows until the file system is full. If the file system is full, the database will freeze until the situation is resolved.

Log Modes

After changing the log mode parameter `log_mode`, you must restart the SAP HANA database 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 (default)**

Log segments are automatically backed up if automatic log backup is enabled.
`log_mode = normal` is recommended to provide support for point-in-time recovery.

- **log_mode = overwrite**

Log segments are freed by savepoints and no log backup is performed. 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 Backup Timeout

You can specify a timeout for log backups. The timeout determines the intervals at which the logs are backed up. If the log segments become full before the timeout, the logs will be backed up anyway. If the timeout exceeds but the log segment contains no commit entry, the log segment will not be backed up at this point in time. The log backup timeout is controlled by the parameter `log_backup_timeout_s`.

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.

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

If you specify a timeout 0, the log backups are only made when a log segment is full and when services are restarted.

Note: The backup timeout only takes effect if automatic log backups are enabled.

27. How can I delete log segments that are no longer needed from the log volumes?

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

For more information, see SAP HANA Administration Guide.

28. Where can I find the backup catalog?

The backup catalog is stored in the SAP HANA database and is part of the log backups. This allows access to the backup catalog during recovery. Even in situations such as when log mode overwrite is set, where logs backups are not written, the backup catalog is still backed up. The backup catalog is backed up and versioned after every completed backup operation. If the backup catalog is saved using a third-party backup tool, the versioning of the backup catalog is handled by the backup tool.

SAP Note 1812980 provides more information about the backup catalog.

29. How can I identify data and log backup files needed for a recovery

SAP Note 1705945 describes which data and log backups are needed for a recovery.

30. How can I identify old log backups?

The SAP HANA Administration Guide describe in detail how to identify old data and log backups.

31. How do I recover my SAP HANA database?

SAP HANA studio provides point-in-time recovery with full GUI support.

In a recovery situation, you can use SAP HANA studio to restore customer-specific changes (configuration settings) either before or at the end of the recovery. To recover your database from SAP HANA studio -> navigator view -> open context menu of the database -> choose "Recover...". Confirm that the database can be shut down - recovery is only done when the database is offline. In the recovery wizard choose one of the recovery types:

- Recover the database to its most recent state.
- Recover the database to a certain point in time.
- Recover the database to a specific data backup. (Just a specific data backup is used for recovery; Log entries are not replayed neither from the log backups nor from the log area.)

Caution: The SAP HANA database will be initialized with the specified data backup. This backup becomes the first backup of a new database life cycle. Older backups are not compatible with the log written after the recovery.

- Recover the database to a certain log position. (Advanced option to be used in exceptional cases if a previous recovery failed.)

The SAP HANA database uses the backup catalog to identify the destination of the backups. You do not need to specify whether File or Backint was used for the backup.

32. How do I recover my SAP HANA database, if the 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, all relevant log backups made after the data backup are needed.

Perform a database recovery to the **most recent state**.

33. How do I recover my SAP HANA database if the 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.

For recovery, the data backups and the log backups are needed.

Perform a database recovery to the **most recent state** and in the Recovery Wizard choose the option **Initialize log area** to prevent the recovery of entries from the unusable log area.

34. How do I recover my SAP HANA database if a logical error occurred (for example, some data were deleted by mistake)?

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.

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

35. Is online recovery possible?

No.

Database recovery is always a full recovery of the whole database and therefore always requires a downtime.

36. Where can I see that my recovery is finished?

SAP HANA studio shows the progress of the recovery and reports when it is finished.

37. Can I cancel a running recovery?

Yes.

You can cancel a recovery while it is running. However, canceling a recovery makes the database inconsistent. To be able to work with the database, a further recovery would then be necessary.

38. What do I have to do if the recovery fails?

That depends on the situation. Often, not all necessary backups are available. Therefore you will have to provide the missing backups and repeat the recovery. Otherwise you will have to use another recovery strategy corresponding to your available backups.

39. Which information is needed to analyze backup or recovery problems?

The following information has to be provided to SAP support in case of backup or recovery problems:

- content of directory `$SAP_RETRIEVAL_PATH/trace` as minimum `backup.log`, `backint.log` (if you are using a third-party backup tool) and all `*.trc` files
- list of the content (recursive) of your backup directories, typically `$DIR_INSTANCE/backup`

40. Does the recovery tolerate corrupt or missing backups?

The SAP HANA database recovery always leaves the database in a consistent state. If the latest data backup is corrupt, you can use an older data backup (plus the log backups between the two data backups) to reach the desired point in time. If log backups are missing or corrupt, the reachable point in time might be earlier than specified by the database administrator for the recovery.

41. Can backup and recovery be resumed?

You cannot pause and resume a data backup with SAP HANA. This is not necessary, as SAP HANA database backup has no impact on normal database operation.

Database recovery cannot be paused/resumed either; it is always a full recovery and carried out as fast as possible in offline mode.

42. Can I copy a SAP HANA database using backup and recovery?

Yes.

For details read the SAP HANA Administration Guide.

43. Are backups encrypted?

Currently, SAP HANA database does not offer backup encryption. However, third-party backup tools normally offer backup encryption. If you are using a third-party backup tool, contact your tool vendor about what encryption options are offered.

44. Can backups be recovered to a SAP HANA database with a higher software version?

Yes.

You can recover data and log backups created with a SAP HANA database software version lower or equal to the target database. Data and log backups cannot be recovered to a SAP HANA database with a lower software version.

If you are updating your SAP HANA database from a maintenance revision to a support package, certain restrictions apply.

More information: SAP Note [1948334](#) (SAP HANA Database Update Paths for Maintenance Revisions)

45. Can I recover a HANA backup made on INTEL architecture on IBM Power architecture?

No.

It's not possible to recover a HANA data or log backup made on INTEL on POWER or vice versa.

Validity

This document is not restricted to a software component or software component version

References

This document refers to:

SAP Notes

- 1910535 [Recovery of SAP HANA database not possible](#)
- 1867324 [SAP HANA Database: Deactivating the Script Server](#)
- 1852242 [Confusion about file-based backup catalog naming](#)
- 1835075 [Analyze backup and recovery performance issues](#)
- 1820553 [Incompatibility of data backup and DB configuration](#)
- 1820529 [Unsuitable network file systems for backup and recovery](#)
- 1817703 [Archiving file system-based backups](#)
- 1812980 [Changes to the backup catalog as of Revision 45](#)
- 1749467 [Copying SAP HANA From a Multiple- to a Single-Host System](#)
- 1730932 [Using backup tools with Backint for HANA](#)
- 1730928 [Using external software in a HANA appliance](#)
- 1716418 [Data backup hangs when writing to NFS server](#)
- 1705945 [Determining required restoration files](#)
- 1703435 [Generating a copy of a running instance with database snapshots](#)
- 1697613 [Removing SAP HANA XS from topology](#)
- 1651055 [Scheduling SAP HANA Database Backups in Linux](#)
- 1645183 [Changing log mode for SAP HANA DB after SPS2 to SPS3 upgrade](#)

This document is referenced by:

SAP Notes (22)

- 1817703 [Archiving file system-based backups](#)
- 1812980 [Changes to the backup catalog as of Revision 45](#)
- 1730928 [Using external software in a HANA appliance](#)
- 1730932 [Using backup tools with Backint for HANA](#)
- 1749467 [Copying SAP HANA From a Multiple- to a Single-Host System](#)
- 2039883 [FAQ: SAP HANA database and storage snapshots](#)
- 2051052 [GPFS "No space left on device" when df shows free space - Lenovo x86_64](#)
- 1910535 [Recovery of SAP HANA database not possible](#)
- 2165547 [FAQ: SAP HANA Database Backup & Recovery in an SAP HANA System Replication Landscape](#)
- 1867324 [SAP HANA Database: Deactivating the Script Server](#)
- 1820529 [Unsuitable network file systems for backup and recovery](#)
- 1651055 [Scheduling SAP HANA Database Backups in Linux](#)
- 1650957 [SAP HANA Database: Starting the Script Server](#)
- 1645183 [Changing log mode for SAP HANA DB after SPS2 to SPS3 upgrade](#)
- 2009486 [Disable SAP HANA log backups during upgrade of third-party backup tool that supports the Backint for SAP HANA interface](#)
- 1820553 [Incompatibility of data backup and DB configuration](#)
- 1852242 [Confusion about file-based backup catalog naming](#)
- 1716418 [Data backup hangs when writing to NFS server](#)
- 1705945 [Determining required restoration files](#)
- 1703435 [Generating a copy of a running instance with database snapshots](#)
- 1835075 [Analyze backup and recovery performance issues](#)
- 1697613 [Removing SAP HANA XS from topology](#)