

2400005 - FAQ: SAP HANA Persistence

| | | | |
|-----------------------|--|------------------------|----------------------------|
| Version | 28 | Type | SAP Knowledge Base Article |
| Language | English | Master Language | English |
| Release Status | Released to Customer | Category | How To |
| Component | HAN-DB-PER (SAP HANA Database Persistence) | Released On | 17.07.2020 |

Please find the original document at <https://launchpad.support.sap.com/#/notes/2400005>

Symptom

You are interested in details about the SAP HANA persistence layer.

Environment

SAP HANA

Cause

1. [What is the SAP HANA persistence?](#)
2. [Which indications exist for issues on persistence layer?](#)
3. [How can disk I/O related issues be analyzed?](#)
4. [Where can persistence parameters be configured?](#)
5. [Which SAP HANA services have a persistence?](#)
6. [Why can the persistence become fragmented?](#)
7. [How can the persistence be defragmented?](#)
8. [How does garbage collection work on persistence level?](#)
9. [Is it possible to encrypt the persistence level?](#)
10. [What are shadow pages?](#)
11. [What are persistence containers?](#)
12. [How can corruptions on persistence level be analyzed?](#)
13. [What does "Disk information on the Overview screen may be inconsistent" mean in SAP HANA Studio?](#)
14. [What are typical reasons for significant size differences in memory vs. on disk?](#)
15. [How is the number and size of data files controlled?](#)
16. [Which commands exist for SAP HANA disk space analysis?](#)
17. [Can there be orphan files on persistence level?](#)
18. [Is it possible to re-initialize the persistence of a service?](#)
19. [Are all persistences considered in context of operations like backup, system replication or tenant move?](#)
20. [What is data volume partitioning?](#)

Resolution

1. What is the SAP HANA persistence?

The SAP HANA persistence is responsible for mapping data in memory to pages on disk. It performs disk I/O read and write operations and takes care for garbage collection on persistence level. It also makes sure that transaction consistency is maintained.

2. Which indications exist for issues on persistence layer?

The following SAP HANA alerts indicate problems in the persistence area:

| Alert | Name | Description |
|-------|--|---|
| 2 | Disk usage | Determines what percentage of each disk containing data, log, and trace files is used. This includes space used by non-SAP HANA files. |
| 28 | Most recent savepoint operation | Determines how long ago the last savepoint was defined, that is, how long ago a complete, consistent image of the database was persisted to disk. |
| 30 | Check internal disk full event | Determines whether or not the disks to which data and log files are written are full. A disk-full event causes your database to stop and must be resolved. |
| 77 | Database disk usage | Determines the total used disk space of the database. All data, logs, traces and backups are considered. |
| 92 | Root keys of persistent services are not properly synchronized | Not all services that persist data could be reached the last time the root key change of the data volume encryption service was changed. As a result, at least one service is running with an old root key. |

SQL: "HANA_Configuration_MiniChecks" (SAP Notes [1969700](#), [1999993](#)) returns a potentially critical issue (C = 'X') for one of the following individual checks:

| Check ID | Details |
|----------|--|
| 250 | Max. used disk size (%) |
| 370 | Unused space in data files (%) |
| 383 | Max. size of shadow pages (GB, last day) |

For I/O related issues see SAP Note [1999930](#).

3. How can disk I/O related issues be analyzed?

Disk I/O topics aren't covered in this SAP Note, instead you can check SAP Note [1999930](#).

4. Where can persistence parameters be configured?

Persistence related settings can be found in ini file sections [persistence] and [fileio]. With SAP HANA 1.0 also the hdbparam tool was used for configuring persistence related settings.

See SAP Note [2186744](#) for more information related to SAP HANA parameter settings.

5. Which SAP HANA services have a persistence?

Only a subset of SAP HANA services has the possibility to persist data to disk:

- computeserver
- dpserver
- indexserver
- nameserver
- standalone statistics server (SAP Note [2147247](#))

- scriptserver (SAP HANA <= 2.0 SPS 02, >= 2.0 SPS 04)
- xsengine

Other services only work in memory without persisting data to disk.

You can use SQL: "HANA_Topology" (PATH = '/volumes%', NAME = 'servicetype') available via SAP Note [1969700](#) to display services with a persistence.

6. Why can the persistence become fragmented?

Fragmentation in the context of persistence is defined as difference between used and allocated space in the data area on disk. It doesn't necessarily mean that a lot of small space fragments exist. Fragmentation in general shows up when the used space on disk is significantly smaller than before, because the allocated space won't automatically reduce. Typical scenarios are:

- Reduction of business data, e.g. due to archiving and deletion operations (SAP Note [2388483](#))
- Cleanup after blocked garbage collection (SAP Note [2169283](#))
- Table optimization (i.e. merge, compression) of a large table (SAP Note [2057046](#)), requiring temporarily double space
- Snapshots / savepoints preserved for a longer time (SAP Note [2100009](#)), resulting in an increased amount of shadow pages

You can use SQL: "HANA_Disks_Overview" (SAP Note [1969700](#)) to check the current and historic data fragmentation. The output column FRAGMENTATION_PCT indicates the percentage of fragmentation. Values below 30 % are usually acceptable, in case of larger values you can consider a defragmentation. If you know that the free space will be used in the future or the fragmentation will return in a similar way, also values of more than 30 % can be fine.

Fragmentation on persistence level is only a matter of disk space utilization, you don't have to expect other impact like performance regressions.

7. How can the persistence be defragmented?

The persistence can be defragmented with the following command (SAP Note [1870858](#)):

```
ALTER SYSTEM RECLAIM DATAVOLUME '<host>:<port>' 120 DEFRAGMENT
```

The "120" indicates the fragmentation overhead that can remain, in this case 20 % of fragmentation on top of the existing 100 % data is acceptable. 120 is a reasonable value because due to temporary space requirements for table optimizations and garbage it is quite normal that 20 % of space is allocated and deallocated. Smaller values can significantly increase the defragmentation runtime and only provide limited benefit.

The defragmentation happens online without blocking accesses to tables.

The SAP HANA database trace (SAP Note [2380176](#)) contains related messages like:

```
Shrink DataVolume to <pct>% of payload size
Reclaim[0, '<datafile>.dat']:<init>: sizes= (used= <used> mb / max=<max>mb | file= <file>mb),
curOverhead= <pct>%,
maxOverhead= <targetpct>%, moving up to <mb>mb pages per cycle
DVolPart::truncate(payload= <payload>, maxSuperblockIndex= <ind>size= <source> --> <target>)
Reclaim[0, '<datafile>.dat']:<truncate>: #<id>: truncated <mb>mb, sizes= (used= <mb>mb /
max=<mb>mb |
file= <mb>mb), curOverhead= <pct>%, maxOverhead= <target_pct>%, <seconds>s)
```

Reclaim[0, '<datafile>.dat']:<finished after <steps> steps>: moved <pages> pages (<mb>mb) in interval [0mb, 0mb[, sizes= (used=<used>mb / max=<max>mb | file= <mb>mb), curOverhead= <pct>%, maxOverhead= <target_pct>%, <mb>mb truncated (<seconds>s)

The progress of defragmentation can be monitored via SQL: "*HANA_Disks_Data_SuperBlockStatistics*" (SAP Note [1969700](#)).

Example: (output before, during and after defragmentation)

| HOST | PORT | VOLUME_ID | SB_SIZE_MB | USED_SB_COUNT | TOT_SB_COUNT | USED_GB | ALLOC_GB | FRAG_PCT |
|----------|-------|-----------|------------|---------------|--------------|---------|----------|----------|
| saphana1 | 30003 | 4 | 64.00 | 8945 | 11514 | 559 | 720 | 28.71 |
| saphana1 | 30003 | 4 | 64.00 | 8323 | 11514 | 520 | 720 | 38.33 |
| saphana1 | 30003 | 4 | 64.00 | 8146 | 8604 | 509 | 538 | 5.62 |

The RECLAIM task is resumable, so if it is terminated (e.g. due to "general error: Shrink canceled, probably because of snapshot pages", SAP Note [1999880](#)), it will continue next time at roughly the place where it stopped.

When RECLAIM is run in parallel to production load and modifications there is a certain risk of significant runtime overhead depending on the used SAP HANA Revision level:

- SAP HANA 1.0: Rather high risk of runtime overhead
- SAP HANA 2.0 <= SPS 03: Reduced risk of runtime overhead
- SAP HANA 2.0 >= SPS 04: Further optimizations to reduce risk of runtime overhead

8. How does garbage collection work on persistence level?

For proper transaction management undo and cleanup files exist on persistence level (also called history files). Once a transaction is finished and nobody else requires the information from a consistent view perspective, the persistence garbage collection kicks in and removes these history files. See SAP Note [2169283](#) for more information regarding SAP HANA garbage collection.

9. Is it possible to encrypt the persistence level?

It is possible to encrypt the SAP HANA data and log volume in order to make sure that data on disk and backup level isn't readable for unauthorized persons. See the [SAP HANA Security Guide](#) ("Data and Log Volume Encryption") for more information.

10. What are shadow pages?

Shadow pages are old copies of current pages that are required to make sure that you can restore the situation that existed at the time of the last savepoint (SAP Note [2100009](#)). A high number of shadow pages (e.g. due to a savepoint issues) can increase the persistence size unexpectedly. Check ID M0383 ("Max. size

of shadow pages (GB, last day)", SAP Note [1999993](#) issues a warning in this case.

Be aware that pages preserved by snapshots are treated as normal pages and not as shadow pages.

11. What are persistence containers?

Containers are collections of persisted pages, organized into container specific structures. Containers also provide an API that abstracts from the details how the structures are implemented internally.

A specific container type is the virtual file which provides a file-like API.

12. How can corruptions on persistence level be analyzed?

You can use the hdbpersdiag tool (SAP Note [2272121](#)) in order to analyze pages on disk level. See SAP Note [2116157](#) for more information related to SAP HANA consistency checks.

13. What does "Disk information on the Overview screen may be inconsistent" mean in SAP HANA Studio?

The warning

```
"Disk information on the Overview screen may be inconsistent; for more detailed information see the Volumes screen"
```

in the Administration screen of SAP HANA Studio (SAP Note [2073112](#)) indicates that there is a problem retrieving disk information from views like M_VOLUMES, M_VOLUME_SIZES or M_DISKS. This typically happens when one database doesn't use all nodes in a distributed multitenant environment. In this case M_DISKS contains information of all hosts while M_VOLUMES and M_VOLUME_FILES only contain the information of the hosts where the database is actually running. In this case you can ignore the warning.

In some cases the problem can be caused by problems accessing data in the underlying monitoring views. You may observe long times querying these views and size values of 0 being displayed. You can check the database trace files (SAP Note [2380176](#)) for further details and query the monitoring views M_VOLUMES, M_VOLUME_FILES and M_DISKS manually in order to understand the situation. Typically you will recognize that a referenced disk isn't accessible. If it is "only" a non-critical area like LOG_BACKUP, the system itself continues to run stable and only the queries to the persistence views have issues. Make sure that all defined disk areas are reachable, then the warning will disappear.

14. What are typical reasons for significant size differences in memory vs. on disk?

See SAP Note [1999997](#) ("What are typical reasons for significant size differences in memory vs. on disk?") for more information.

15. How is the number and size of data files controlled?

In most cases SAP HANA uses only one data file per service that can grow to significant sizes. In the following cases multiple files per service are created:

| Scenario | SAP Note | Details |
|---|----------|---|
| File system type with relevant file size limitation | 2154870 | If the used file system has a limited maximum file size (e.g. 2 TB for ext3), SAP HANA will automatically create a new data file once the limit is reached. |

| | | |
|---|---------|---|
| Explicit data volume striping configuration | 2631285 | It is possible to explicitly configure the maximum file size using the following parameters: global.ini -> [persistence] -> datavolume_striping = true global.ini -> [persistence] -> datavolume_striping_size_gb = <max_file_size_gb> A manual configuration is typically not required. |
| Explicit data volume partitioning commands | | Starting with SAP HANA 2.0 SPS 03 explicit commands exist to adjust the number of volume partitions respectively data files: ALTER SYSTEM ALTER DATAVOLUME ADD PARTITION ALTER SYSTEM ALTER DATAVOLUME DROP PARTITION <id> Starting with SAP HANA 2.0 SPS 04 you can optionally specify a system replication site ID (SAP Note 1999880): ALTER SYSTEM ALTER DATAVOLUME ADD PARTITION SYSTEM REPLICATION SITE <site_id> ALTER SYSTEM ALTER DATAVOLUME DROP PARTITION <id> SYSTEM REPLICATION SITE <site_id> See What is data volume partitioning? below for more information regarding data volume partitioning. |

16. Which commands exist for SAP HANA disk space analysis?

The following analysis commands are available via SAP Note [1969700](#):

| SQL statement | Details |
|---|--|
| SQL: "HANA_Disks_Data_Overview" | Rough breakdown of disk space allocation in data area (column store tables, row store tables, disk LOBs, history files, shadow pages, ...) |
| SQL: "HANA_Disks_Data_Pages" | Information related to pages in data area |
| SQL: "HANA_Disks_Data_SuperblockStatistics" | Superblock information for data area, can e.g. be used to monitor process of RECLAIM DATAVOLUME |
| SQL: "HANA_Disks_Overview" | Disk space allocation and fragmentation information |

17. Can there be orphan files on persistence level?

Orphan virtual files in the DATA area are files that have no longer any link to an accessible object. The following scenarios can be responsible for orphan files:

- SAP Note [2007021](#) ("Orphaned persistence files after lob column type conversion")
- SAP Note [2684367](#) ("ALTER TABLE XXX MOVE TO LOCATION XXX" Fails With Error: "ContainerName XXX - RuntimeData already assigned in session")
- SAP Note [2910004](#) ("Increase Disk Usage After Performing DDL Operation on Table With Packed LOB Column")

If you face orphan files in other contexts it is recommended to open a SAP incident on component HAN-DB-PER in order to understand the root cause and the best cleanup approach.

Cleanup options are provided via SAP Note [2007021](#) (virtual files) and [2910004](#) (packed LOBs).

18. Is it possible to re-initialize the persistence of a service?

Re-initializing the persistence of a service would mean to erase all contents and start from scratch. This is in general no option and in case of persistence issues it is usually better to clean up manually or restore a good backup (SAP Note [1642148](#)). In some cases it can be better to re-initialize a less important persistence rather than going for a more difficult option like restore and incomplete recovery.

Persistences that can be re-initialized in some circumstances are:

| Area | SAP Note |
|------------------------------|----------|
| Standalone statistics server | 1950221 |
| XS engine (classic) | 1697613 |
| System database (MDC) | 2588284 |

Problems that can be repaired with a re-initialization (if related to a non-critical persistence as described above) are:

| Problem / error | Details |
|--|---|
| Any kind of logical or physical persistence corruption | All kinds of inconsistencies and corruptions on persistence layer can be repaired with a re-initialization (SAP Note 2116157, 2572224). |
| log position inconsistency detected - please contact HANA dev support before clearing or throwing away your redo log files | In different scenarios SAP HANA can reach a state of inconsistent log positions. If the inconsistency is linked to a service with a non-critical persistence, a re-initialization can be used to resolve the problem. |
| Missing first log segment, cannot recover starting from <id>, first known segment <id> | If missing redo logs result in recovery terminations, it may be possible to temporarily remove and then recreate a non-critical service. |

19. Are all persistences considered in context of operations like backup, system replication or tenant move?

Usually persistence are considered by operations like backup, system replication or tenant move. These persistences are called transactional persistences because they are required to guarantee transactional consistency. In contrast to this standard behavior sometimes local persistences exist that aren't required for transactional consistency. A local persistence can be identified by a virtual volume ID of 1024 or higher. As an example the scriptserver with SAP HANA >= 2.00.040 has a local persistence. These local persistences aren't backed up, they aren't replicated and they aren't copied when doing a tenant move.

20. What is data volume partitioning?

Since SAP HANA 2.0 SPS 03 administrators have the option to partition the data volume for the indexserver so that data files can be segmented and stored in different locations and can then be accessed in parallel threads. See [How is the number and size of data files controlled?](#) -> "Explicit data volume partitioning commands" for related commands.

In the case of Network File Systems data can also be written in parallel across multiple connections. Partitioning data volumes in this way will therefore speed up all read/write operations on the volume including savepoints, table optimizations, restarts, table loading operations and backups.

SAP HANA data volume partitioning can be beneficial primarily for customers using Network File System (NFS) for SAP HANA data and log volumes as this feature enables the filesystem to have more parallel channels for processing I/O. To truly benefit from this feature additional mountpoints need to be configured for the additional locations of data volume partitions in order to leverage additional TCP connections in case of NFS. No further network configuration is required as long as the network infrastructure can sustain the additional workload.

For non-NFS type filesystems the benefits of adopting the feature depends on setup provided by the hardware vendor / TDI design. Please discuss the feasibility with your hardware / storage vendor or team responsible for designing the storage layout.

Consider the following for SAP HANA data volume partitioning:

- Data is not immediately distributed when a data volume partition is added to an existing system. Fresh I/O writes are distributed to the new data volume partition and eventually the database achieves even distribution from a size point of view.
- If immediate even distribution of data is required, consider using SAP HANA backup and recovery (only file and backint based backup and recovery).
- Dropping a SAP HANA data volume partition involves reading data from the dropped partition and writing it into existing partitions. Since this could involve significant I/O activity depending on the quantity of reads/writes, such an activity should be performed during low business workload timeframes.
- In a running system replication setup, you may not be able to drop an active data volume partition as system replication uses data volume snapshot technology. You may see the error "Cannot move page inside/out of DataVolume". In this case it may be necessary to disable system replication functionality, drop the partition, and then setup system replication again.

Don't mix up SAP HANA data volume partitioning with SAP HANA data volume striping:

- SAP HANA data volume partitioning distributes fresh incoming pages across data volume partitions to achieve parallelization of I/O operations wherever possible.
- SAP HANA data volume striping provides the possibility to limit the size of the existing data volume files and create a new data volume file and redirect incoming pages to the new file if no space exists in the older file. There is no even distribution of I/O writes as achieved with data volume partitioning.

See [How is the number and size of data files controlled?](#) for more details.

Starting with SAP HANA 2.0 SPS 04 you can use SQL: "*HANA_Disks_Data_Partitions*" (SAP Note [1969700](#)) to display an overview of existing data volume partitions.

Keywords

SAP HANA persistence persistency reclaim disk I/O

Products

SAP HANA, platform edition all versions

Other Components

| Component | Description |
|-----------|-------------------|
| HAN-DB | SAP HANA Database |

This document refers to

| SAP Note/KBA | Title |
|--------------|--|
| 2572224 | How-To: Repairing SAP HANA Tables |
| 2388483 | How-To: Data Management for Technical Tables |
| 2380176 | FAQ: SAP HANA Database Trace |
| 2272121 | How-To: Analyzing Physical Corruptions with the SAP HANA Persistence Diagnosis Tool (hdbpersdiag) |
| 2186744 | FAQ: SAP HANA Parameters |
| 2169283 | FAQ: SAP HANA Garbage Collection |
| 2154870 | How-To: Understanding and defining SAP HANA Limitations |
| 2116157 | FAQ: SAP HANA Consistency Checks and Corruptions |
| 2100009 | FAQ: SAP HANA Savepoints |
| 2073112 | FAQ: SAP HANA Studio |
| 2057046 | FAQ: SAP HANA Delta Merges |
| 1999997 | FAQ: SAP HANA Memory |
| 1999993 | How-To: Interpreting SAP HANA Mini Check Results |
| 1999930 | FAQ: SAP HANA I/O Analysis |
| 1999880 | FAQ: SAP HANA System Replication |
| 1870858 | HANA Alerts related to file system utilization |
| 2910004 | Increase Disk Usage After Performing DDL Operation on Table With Packed LOB Column |
| 2684367 | "ALTER TABLE XXX MOVE TO LOCATION XXX" Fails With Error: "ContainerName XXX - RuntimeData already assigned in session" |
| 2631285 | Indexserver Fails to Start up With "Cannot open file" After Upgrade to SAP HANA 2.0 Rev. 30 |
| 2588284 | Re-initialize a not recoverable system database |
| 2007021 | Orphaned Persistence Files After LOB Column Type Conversion |
| 1969700 | SQL Statement Collection for SAP HANA |
| 1950221 | Reinitializing the Statisticsserver |
| 1697613 | Removing the SAP HANA XS Classic Engine service from the topology |

| | |
|---------|--|
| 1642148 | FAQ: SAP HANA Database Backup & Recovery |
| | SAP HANA Security Guide |

[Terms of use](#) | [Copyright](#) | [Trademark](#) | [Legal Disclosure](#) | [Privacy](#)