

1999880 - FAQ: SAP HANA System Replication

Version	173	Type	SAP Knowledge Base Article
Language	English	Master Language	English
Release Status	Released to Customer	Category	How To
Component	HAN-DB (SAP HANA Database)	Released On	22.02.2019

Please find the original document at <https://launchpad.support.sap.com/#/notes/1999880>

Symptom

You are interested in further details related to SAP HANA system replication.

Environment

SAP HANA

Cause

1. [What is SAP HANA system replication?](#)
2. [Where do I find further information related to SAP HANA system replication?](#)
3. [Which indications exist for problems with the SAP HANA system replication?](#)
4. [Which modes of SAP HANA system replication are available?](#)
5. [Can SAP HANA system replication be used for system copies?](#)
6. [How can I determine details related to SAP HANA system replication?](#)
7. [Is it possible to initialize a replicated system based on backup / restore?](#)
8. [What are the basic requirements for the SAP HANA system replication network connection regarding throughput and latency?](#)
9. [Can problems with system replication impact the performance on the primary system?](#)
10. [Is it possible to take backups on the secondary site?](#)
11. [Is it possible to set up a cascaded replication environment?](#)
12. [In which situations does a SAP HANA system replication takeover make sense?](#)
13. [Can multiple SAP HANA databases be replicated to the same target?](#)
14. [Is it supported to set up a replication scenario between systems with different hardware?](#)
15. [Is it supported to set up a replication scenario between systems with a different SAP HANA patch level?](#)
16. [What has to be considered when upgrading SAP HANA in system replication environments?](#)
17. [What has to be considered in terms of SAP HANA parameter settings in replicated scenarios?](#)
18. [Is it possible to perform a takeover even if the secondary database is down?](#)
19. [How can RECLAIM DATAVOLUME be executed when system replication is active?](#)
20. [What can I do if too old system replication snapshots exist?](#)
21. [Does local log buffer writing and log buffer shipping happen in parallel?](#)
22. [Is it possible to compress data before sending it to the secondary site?](#)
23. [Which factors influence the takeover time?](#)
24. [What has to be done for a failback?](#)
25. [How long does it typically take to initialize a synchronization of primary and secondary from scratch?](#)
26. [Can a takeover or failback happen without terminating active application processes?](#)
27. [Which log sizes are typically shipped?](#)
28. [Why is the secondary system consistent although undo logs aren't shipped?](#)
29. [Are there any instance number / port restrictions when using system replication?](#)
30. [How can I check the memory allocation details on the secondary system?](#)
31. [Are SAP HANA histories still available after a takeover?](#)

32. [What has to be considered in terms of licensing in SAP HANA system replication environments?](#)
33. [Can I set up a SAP HANA system replication based on a storage snapshot \(instead of a full data backup\)?](#)
34. [Is it possible to monitor remote system replication sites on the primary system?](#)
35. [What is the logreplay operation mode?](#)
36. [How can I determine the current logreplay delay?](#)
37. [What happens with the content when the secondary site of a system replication environment is upgraded?](#)
38. [Does the secondary SAP HANA system replication site need to be up and running before activating system replication?](#)
39. [What has to be considered for the row store reorganization in system replication scenarios?](#)
40. [What can be done if I am not able to disable secondary in the normal way?](#)
41. [Are there specific SAP HANA users and schemas related to system replication?](#)
42. [What are indications for log shipping timeouts?](#)
43. [Can disk I/O on secondary site have an impact on SYNCMEM system replication scenarios?](#)
44. [What is the state of a secondary, if services were not in sync at takeover time?](#)
45. [What has to be considered when changing the number of SAP HANA nodes?](#)
46. [What has to be considered for Active/Active \(read enabled\)?](#)
47. [Why do I see deviating values in the system replication mode parameter?](#)
48. [How can I check the number of parallel channels used for data shipping?](#)
49. [How can tracing for system replication be activated?](#)
50. [Is it recommended to disable system replication during larger upgrade or migration activities?](#)
51. [How can pending savepoints and large page cache sizes on secondary site be explained?](#)
52. [Can I set up system replication between systems with different topologies?](#)
53. [Can a primary site be replicated to more than one secondary site?](#)
54. [What happens with active transactions and statements during takeover?](#)
55. [What is secondary time travel?](#)

Resolution

1. What is SAP HANA system replication?

SAP HANA system replication provides the possibility to copy and continuously synchronize a SAP HANA database to a secondary location in the same or another data center. Usually system replication is used to support high availability and disaster recovery.

It must not be mixed up with the SAP HANA LT Replication Server (SLT) described in SAP Note [2014562](#).

2. Where do I find further information related to SAP HANA system replication?

Among others SAP Note [2407186](#) references the "How To Perform System Replication with SAP HANA" guide that contains a lot of information and references to further documents that can help to implement and use system replication.

SAP Note [2165547](#) provides further information about backup and recovery in SAP HANA system replication environments.

3. Which indications exist for problems with the SAP HANA system replication?

The following SAP HANA alerts indicate problems in the SAP HANA system replication area:

Alert	Name	Description
-------	------	-------------

21	internal event	Identifies internal database events.
78	Connection between systems in system replication setup	Identifies closed connections between the primary system and a secondary system. If connections are closed, the primary system is no longer being replicated.
79	Configuration consistency of systems in system replication setup	Identifies configuration parameters that do not have the same value on the primary system and a secondary system. Most configuration parameters should have the same value on both systems because the secondary system has to take over in the event of a disaster.
94	Log replay backlog for system replication secondary	System Replication secondary site has a higher log replay backlog than expected.
104	Log shipping backlog for system replication secondary	System Replication secondary site has a higher log shipping backlog than expected.
106	ASYNCR replication in-memory buffer overflow	Checks if local in-memory buffer in ASYNCR replication mode runs full

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
M1160	Average COMMIT time (ms)
M1810	Services with replication error
M1811	Services with unknown replication state
M1813	Replication connection closed (last day)
M1814	Log position gap (MB)
M1815	Current log shipping delay (s)
M1816	Filling level of async shipping buffer (%)
M1818	Async log shipping buffer full (last day)
M1820	Parameter deviations primary vs. secondary site
M1830	Age of oldest replication snapshot (h)
M1840	Avg. sync log shipping time (ms/req, last hour)
M1843	Avg. sync log shipping time (ms/req, last day)
M1846	Max. sync log shipping time (ms/req, history)
M1852	Current log replay backlog (GB)
M1853	Current log replay delay (s)
M1854	Max. log replay backlog (GB, last week)
M1855	Max. log replay delay (s, last week)
M1860	Datashipping combined with log retention
M1865	System replication savepoint delay (h)

An increased number of SAP HANA threads with one of the following details exist (SAP Note [2114710](#)):

Thread State	Lock Name
--------------	-----------

Semaphore Wait	LoggerBufferSwitch
Semaphore Wait	POSTCOMMIT_FINISH_SMP

4. Which modes of SAP HANA system replication are available?

The following SAP HANA system replication modes exist:

Mode	Details	Behaviour when secondary system is not available
Synchronous	Primary system waits until secondary system has received data and persisted it to disk.	Primary system waits until global.ini -> [system_replication] -> logshipping_timeout is exceeded (default: 30 s) and then proceeds without replicating data.
Synchronous in-memory	Primary system waits until secondary system has received data.	Primary system waits until global.ini -> [system_replication] -> logshipping_timeout is exceeded (default: 30 s) and then proceeds without replicating data.
Synchronous full sync (SPS 08 and higher)	Primary system waits until secondary system has received data and persisted it to disk.	Primary system is blocked until secondary system becomes available.
Asynchronous	Primary system doesn't have to wait for secondary system (unless the asynchronous log buffer is full and waiting is configured, see "Can problems with system replication impact the performance on the primary system?" below for details)	Primary system proceeds without replicating data.

5. Can SAP HANA system replication be used for system copies?

The normal reason for SAP HANA system replication is high availability and disaster recovery, but it can also be used for system copies. Once the initial synchronization has happened the replication side can be used as copy system by executing a takeover. It is much quicker to establish the system copy based on an existing replication than doing the same with traditional approaches like backup and restore. If the SAP HANA database has to be renamed for the system copy, the hdbrename tool can be used.

Be aware that the number of working hosts on the copied system needs to be identical to the original number of working hosts (standby hosts can vary).

6. How can I determine details related to SAP HANA system replication?

System replication information can be found in monitoring view M_SERVICE_REPLICATION.

The following SQL statements are available in SAP Note [1969700](#) in order to evaluate system replication related details:

SQL statement	Details
SQL: "HANA_Replication_SystemReplication_Bandwidth"	Calculates the required network bandwidth between primary and replication side
SQL:	Can be used to show current key figures in

"HANA_Replication_SystemReplication_KeyFigures_Current_CommandGenerator"	terms of bandwidth and throughput, manual maintenance of reference values required
SQL: "HANA_Replication_SystemReplication_KeyFigures"	Possibility to analyze the system replication key figures historically (SAP HANA >= 1.0 SPS 09)
SQL: "HANA_Replication_SystemReplication_LogShipping_RetentionTime"	Maximum time a disconnection of primary and secondary site is possible in "logreplay" mode (SAP HANA >= 1.0 SPS 11) before the retained log segments for delta resync are lost
SQL: "HANA_Replication_SystemReplication_Overview"	Provides a general overview of system replication information (configuration, activity details)
SQL: "HANA_Replication_SystemReplication_ParameterDeviations"	Checks for SAP HANA parameters which deviate between primary and replication side
SQL: "HANA_Replication_SystemReplication_Status"	System replication status of individual services
SQL: "HANA_Replication_SystemReplication_Takeovers"	Information about historic system replication takeover activities (SAP HANA >= 2.00.030)

7. Is it possible to initialize a replicated system based on backup / restore?

Up to SAP HANA SPS 11 it is only possible to initialize the replication side based on an snapshot that is created via full data shipping from primary site. An initialization via backup and restore is not possible.

Starting with SAP HANA SPS 12 you can initialize the secondary with a binary copy (via storage snapshot or while primary site is stopped) from the primary site.

8. What are the basic requirements for the SAP HANA system replication network connection regarding throughput and latency?

For system replication it is important that the network throughput (or bandwidth) and the latency (or roundtrip time) are sufficient to fulfill the requirements. SAP Note [1100926](#) provides further details about these KPIs.

Further details are discussed in the whitepaper for Network Recommendations for SAP HANA System Replication available at [Network Recommendations for SAP HANA System Replication](#).

As of SAP HANA 1.0 SPS 09 you can compress data before sending it across the network (see "Is it possible to compress data before sending it to the secondary site?" below). This will reduce the network bandwidth requirements.

Due to the need of data transfer in addition to the log transfer the bandwidth requirements of SAP HANA can be higher than expected. Starting with SAP HANA 1.0 SPS 11 the continuous log replay feature (operation_mode = logreplay) is available that significantly reduces the bandwidth requirements.

For asynchronous system replication in long-distance environments and a related high latency you need to make sure that a reasonable TCP windows size is used:

- Check the network utilization profile for the network link to see if the maximum capacity of the network has been reached
- If the network is not fully utilized, check the linux kernel TCP configuration with:
`sysctl -a | egrep "net.core|net.ipv4.tcp"`
- Check that windows scaling is set to the default value of 1:
`net.ipv4.tcp_window_scaling = 1`
- Check whether the max size can be increased for:
`net.ipv4.tcp_wmem` and `net.ipv4.tcp_rmem`

- Calculate the Bandwidth Delay Product (BDP): Bandwidth * Latency (for example, BDP = 50ms * 3 Gbps = 19.2 MB). The BDP tells you what TCP window size is needed to use the network link fully.

9. Can problems with system replication impact the performance on the primary system?

In the following scenarios the performance of the primary system is impacted by system replication:

Scenario	Replication mode	Details
Connection to secondary system not working	synchronous full sync	Change operations and COMMITs on the primary system have to wait permanently until a connection to the secondary system is established again.
Connection to secondary system not working	synchronous synchronous in-memory	Change operations and COMMITs on the primary system have to wait until a connection to the secondary system is established again or the timeout configured with the following parameter is reached: global.ini -> [system_replication] -> logshipping_timeout = <seconds> Be aware that this timeout only applies to situations where no error is sent to the primary site (e.g. slow network connection). In case of explicit errors (e.g. when the secondary site is terminated or when an explicit network error is returned), the primary site continues processing immediately.
Inadequate network connection in terms of latency and throughput	synchronous synchronous full sync synchronous in-memory	Change operations and COMMITs are slowed down because the network communication is within the critical path and they can only proceed after the changes were successfully processed on secondary site.
Slow disk I/O while processing logs on secondary system	synchronous synchronous full sync	Change operations and COMMITs are slowed down because persisting the log information on secondary site is within the critical path.
Asynchronous log shipping buffer full	asynchronous	Per default change operations and COMMITs on the primary system wait until space in the asynchronous log shipping buffer is again available. If you don't accept these waits, you can set the following parameter to 'false' on primary site: global.ini -> [system_replication] -> logshipping_async_wait_on_buffer_full = false In this case SAP HANA will temporarily close the system replication to the secondary system and try to re-establish the connection after the time defined with the following parameter: global.ini -> [system_replication] -> reconnect_time_interval = <seconds> Attention: If a full data shipment takes place and system replication is disconnected, it has to start from scratch. So it is useful to keep logshipping_async_wait_on_buffer_full activated until initial data shipment is finished. Be aware that a restarted full data shipment from primary to secondary site also results in a completely new full data shipment to a tertiary site (if configured). In order to reduce the risk of buffer full situations, you can adjust the size of the asynchronous log shipping buffer on primary site using the following parameter: <service>.ini -> [system_replication] ->

		<p>logshipping_async_buffer_size = <size_in_byte> It is recommended that you set this parameter only for the services with a high log generation, so typically the indexserver (<service>.ini = indexserver.ini), because a setting in global.ini would increase the log shipping buffer for all services and so memory is wasted unnecessarily. See SAP Note 1999997 (heap allocator Pool/PersistenceManager/DisasterRecoveryPrimary) for more details. Starting with SAP HANA 1.00.122.17 and 2.00.024.01 the default of this parameter is increased from 64 MB to 256 MB for the indexserver (SAP Note 2678164). You can use SQL: "HANA_Replication_SystemReplication_KeyFigures" (SAP Note 1969700) to check if and when the asynchronous log shipping buffer was already full (column "BF"). Column BUFF_MB indicates the filling level of the asynchronous log shipping buffer. If system replication is currently inactive, BUFF_MB shows the accumulated redo log size that needs to be shipped once system replication comes back online. In order to activate the new buffer size the secondary site needs to re-connect to the primary site. A reconnect can e.g. be enforced by restarting the secondary site. Be aware that in a cascaded system replication environment (see "Is it possible to set up a cascaded replication environment?" below) a full asynchronous log shipping buffer on the secondary system has a direct impact on the performance of the primary system. Therefore the recommendations above also apply to secondary systems serving tertiary systems asynchronously.</p>
Connection to secondary not working	all	Due to the SAP HANA bug described in SAP Note 2544075 a race condition can result in permanent "LoggerBufferSwitch" waits (SAP Note 1999998).
File system full because log shipping retention size is larger than file system size	asynchronous synchronous synchronous in-memory	In context of logreplay operations modes the following SAP HANA parameter defines the maximum amount of redo logs that are kept on primary site for synchronization with the secondary site (default: 1 TB): global.ini -> [system_replication] -> logshipping_max_retention_size = <size_mb> If the underlying file system isn't large enough to hold the complete configured retention size, it can happen in the worst case that the file system runs full and the primary site comes to a standstill. Therefore SAP recommends to configure this parameter with not more than 80 % of the log volume size (SAP Note 2600030).

In case of issues you can use SQL:

"HANA_Replication_SystemReplication_KeyFigures_Current_CommandGenerator" and double-check the key figures. LOG_AVG_SHIP_TIME_MS should typically not be higher than a few milli seconds and the LOG_SHIP_MB_PER_S should be significantly above 10 MB / s. If you see worse values, you have to check both the network and the secondary system and eliminate issues and bottlenecks.

10. Is it possible to take backups on the secondary site?

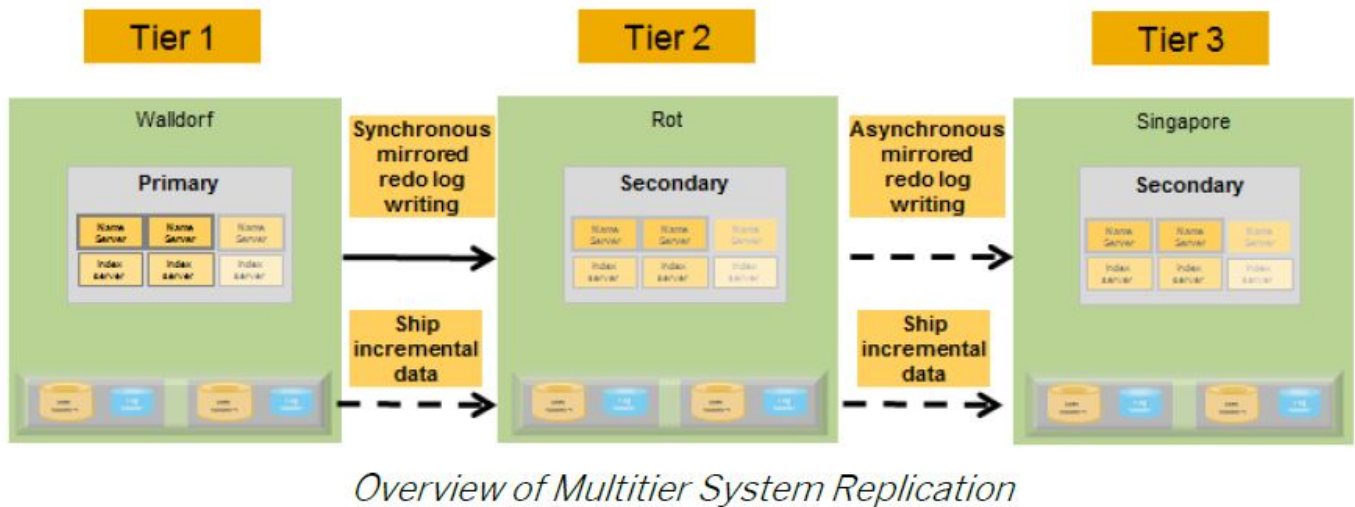
No, currently you can only take data and log backups on the primary site.

If automatic log backup is enabled, after takeover to the secondary site the log backups will automatically be written there.

11. Is it possible to set up a cascaded replication environment?

Yes, it is possible to set up a sequence of system replications, a so-called multitier system replication.

Example:



See SAP Note [2303243](#) for more details on supported replication modes.

The operation mode must be identical for all replications, a mixture (e.g. delta_datashipping from tier 1 to tier 2 and logreplay from tier 2 to tier 3) isn't allowed. Only exception: logreplay_readaccess from tier 1 to tier 2 and logreplay from tier 2 to tier 3.

12. In which situations does a SAP HANA system replication takeover make sense?

See SAP Note [2063657](#) that describes what aspects should be considered before performing a takeover. It is important to follow these steps, because otherwise a takeover can even make everything worse (e.g. data loss).

Starting with SAP HANA 2.0 SPS 03 the Python script getTakeoverRecommendation.py is delivered together with SAP HANA that can be used to check if a takeover is useful or possible in a specific scenario.

13. Can multiple SAP HANA databases be replicated to the same target?

It is not possible to replicate different source SAP HANA databases into the same target SAP HANA database.

You can replicate different source SAP HANA databases into different target SAP HANA databases on the same target host. Make sure that the restrictions of SAP Note [1681092](#) are considered. This includes that all other SAP HANA databases have to be stopped on replication side as soon as one SAP HANA database takes over the productive role for a primary system.

14. Is it supported to set up a replication scenario between systems with different hardware?

Yes, it is possible to set up replication between systems that differ in terms of hardware (e.g. different manufacturer, different memory sizes, different number of CPUs). From a system replication perspective you only need to make sure that the number of hosts is identical on all sites and that every service has a partner on all sites.

It is also possible to run different hardware configurations between the primary and the secondary site in terms of virtualization configurations. In other words, the secondary site can be a virtual configuration

compared to the primary site being physically configured and vice versa.

Be aware that it is not possible to set up system replication between platforms with different endianness, because the logs aren't binary compatible. In general, SAP HANA systems are little-endian, with the exception of Power in combination with SLES 11.x.

SAP Note [1984882](#) describes how system replication can be used for a hardware exchange with minimum downtime.

15. Is it supported to set up a replication scenario between systems with a different SAP HANA patch level?

It is allowed to use different SAP HANA patch levels in a replication scenario as long as the patch level of the replicated system is not lower than the patch level of the primary system. It is also allowed that the replicated system is on a higher SPS level than the primary system. This possibility can be used for near-zero downtime upgrades (NZDU) upgrades in replicated environments. For more information see the near zero downtime upgrade section in [How to Perform System Replication for SAP HANA](#). SAP Note [2386973](#) describes NZDU in 3-tier system replication scenarios.

16. What has to be considered when upgrading SAP HANA in system replication environments?

You can stop both the primary and secondary database, then upgrade both databases and then restart both systems. If you want to minimize downtime, you can proceed according to SAP Note [1984882](#).

17. What has to be considered in terms of SAP HANA parameter settings in replicated scenarios?

In general you should make sure that SAP HANA parameters are set identical on primary and secondary system replication side whenever possible. Deviations can result in unexpected behavior during or after a takeover. You can check SAP HANA alert 21 (≤ 1.0 SPS 08) or 79 (≥ 1.0 SPS 09) or use SQL: *"HANA_Replication_SystemReplication_ParameterDeviations"* (SAP Note [1969700](#)) in order to check for parameters that are set differently on primary and replication side.

Sometimes different parameter settings are necessary on the primary and the secondary site, e.g. if additional non-productive systems are running on the secondary site. Then the table reload must be turned off and the global allocation limit must be set to a minimum value to make sure that the memory allocation of the replicated database remains on a low level. See SAP Note [2127458](#) for more information regarding table reload.

In the SAP HANA Admin guide the "parameter checker" is documented, which generates the alerts. It can be manually adjusted to exclude ranges of parameters from these checks. See the section "Monitoring INI File Parameter Changes" in [How to Perform System Replication for SAP HANA](#) for more details.

Starting with SAP HANA 1.0 SPS 12 it is possible to automatically replicate parameters from the primary to the secondary site by activating the following parameter on primary site:

```
global.ini -> [inifile_checker] -> replicate = 'true'
```

When a parameter is replicated, the following type of message will be written to the database trace (SAP Note [2380176](#)):

```
Replicate following ini file entry to site '<site_id>':  
<file>.ini/<layer>//[<section>]/<parameter> = <value>
```

Be aware that even with activated automatic parameter replication there can still be mismatches reported in case a parameter is explicitly set to its default value (SAP Note [2737227](#)).

18. Is it possible to perform a takeover even if the secondary database is down?

Yes, a takeover can be performed even if the secondary database is down. In order to make it available, it of course has to be started.

19. How can RECLAIM DATAVOLUME be executed when system replication is active?

With ALTER SYSTEM RECLAIM DATAVOLUME DEFRAGMENT you can reduce fragmentation on disk level and reclaim disk space (SAP Note [1870858](#)). When system replication is used, there are typically snapshots preserved and as a consequence the RECLAIM DATAVOLUME command fails with:

```
general error: Shrink canceled, probably because of snapshot pages
```

See SAP Note [2332284](#) for required steps to perform a successful RECLAIM DATAVOLUME in context of an active system replication.

Be aware that the error can also be caused by a snapshot-based backup rather than system replication. So in order to make sure that a RECLAIM DATAVOLUME succeeds, you have to make sure that no backup based on snapshots is scheduled at the same time.

20. What can I do if too old system replication snapshots exist?

System replication takes advantage of SAP HANA database snapshots. The retention time of these snapshots is defined with the following parameter:

Parameter	Default	Unit	Details
global.ini -> [system_replication] -> datashipping_snapshot_max _retention_time	120 (SAP HANA 1.0) 300 (SAP HANA >= 2.00)	minutes	This parameter defines how many minutes a system replication related database snapshot is preserved, that has been completely shipped to the secondary system. Shipped snapshots older than datashipping_snapshot_max _retention_time will be dropped automatically. Snapshots currently used in data shipping are not affected and are not dropped, if data shipping takes longer than datashipping_snapshot_max _retention_time. They can be dropped if data shipping has been finished. If the parameter is set to 0, snapshots are immediately dropped after data replication has been finished. Its default value corresponds to 2 hours. There is normally no need to change this parameter (exception: See "How can RECLAIM DATAVOLUME be executed when system replication is active?").

In rare cases it can happen that a system replication related database snapshot is not deleted and its age significantly exceeds the configured retention time. You can identify old snapshots (in this example: older than 5 hours) by running SQL: "HANA_IO_Snapshots" (SNAPSHOT_repli = 'REPLICATION', MIN_SNAPSHOT_AGE_H = 5) available via SAP Note [1969700](#).

Example:

PORT	SERVICE_NAME	VOLUME_ID	SNAPSHOT_TIME	SNAPSHOT_AGE_H	FOR_BACKUP
30005	statisticsserver	2	2015/02/27 21:37:16	35.61	FALSE
30007	xsengine	3	2015/02/27 21:36:44	35.61	FALSE

Now you can check the current replication status using SQL: *"HANA_Replication_SystemReplication_Status"* (SAP Note [1969700](#)). If the STATUS_DETAILS are empty for the services, you can proceed. Otherwise you should check if the replication of the service permanently remains in the status and analyze what is going on.

You can clean up no longer needed system replication related database snapshots using the hdbcons command line tool (SAP Note [2222218](#)). At first you have to identify the related snapshot ID for the impacted service:

```
hdbcons -e hdb<service_name> 'snapshot l'
```

It will return an information like the following including the snapshot ID:

```
SnapshotID: 290360 / ForBackup: 0 / ForReplication: 1 / RestartPage: 0x700000020004P
[0x20004,16k,0] / TimeStamp: ...
```

Using the snapshot ID the space allocated by a database snapshot can be checked via:

```
hdbcons -e hdb<service_name> 'snapshot a <snapshot_id>'
```

If you have confirmed that the database snapshot is no longer required, the snapshot ID can be dropped with the following hdbcons command:

```
hdbcons -e hdb<service_name> 'snapshot d <snapshot_id>'
```

21. Does local log buffer writing and log buffer shipping happen in parallel?

Both activities happen in parallel. If SQL: *"HANA_Replication_SystemReplication_Overview"* (SAP Note [1969700](#)) returns the following numbers

```
Avg. local log buffer write time (ms): 2.81
Avg. log buffer shipping time (ms): 2.46
```

you can assume that the total duration of both activities doesn't exceed 3 ms.

22. Is it possible to compress data before sending it to the secondary site?

Starting with SPS 09 you can activate compression of log and / or data using the following parameters:

Parameter	Availability	Default	Details
global.ini -> [system_replication] -> enable_log_compression	SPS 09 and higher	false	Enables compression of log when it is sent to the secondary site Needs to be set on target site (secondary and / or tertiary), afterwards target needs to be reconnected to source site
global.ini -> [system_replication] -> enable_data_compression	SPS 09 and higher	false	Enables compression of data when it is sent to the secondary site Needs to be set on target site (secondary and / or tertiary), afterwards target needs to be reconnected to source site

The activation of the compression reduces the required network bandwidth, but at the same time there is some CPU overhead for compressing and decompressing the information. Using compression is particularly

useful in case of long distances between primary and secondary site or in case of bandwidth limitations.

23. Which factors influence the takeover time?

The main steps during takeover are:

- Open data persistence based on last savepoint
- Load row store
- Replay redo log
- Rebuild row store indexes

As with normal startups, the majority of the time is caused by the row store activities. As a rough rule of thumb it takes 1 minute to load 10 to 20 GB of row store data.

Starting with SPS 09 the row store is already preloaded on the secondary system, so that normally no time needs to be spent during the takeover to load the row store. Exception: When delta data shipping was active at the time the takeover happens, the row store has to be re-loaded from scratch.

Loading of column store data generally happens after the actual takeover. So it doesn't increase the takeover time, but it can impact the performance of the system for some more time after the takeover. Per default columns used on the primary system are regularly loaded into the memory of the secondary system in order to minimize the reload overhead after takeover. See SAP Note [2127458](#) for more information.

24. What has to be done for a failback?

If a takeover has happened and you want to revert back to the former primary system, you basically have to execute the same steps like during the takeover, but with exchanged system roles. Further details can be found in the [SAP HANA Administration Guide](#).

25. How long does it typically take to initialize a synchronization of primary and secondary from scratch?

The initial synchronization is dominated by the data shipping time across the network. As a rough rule of thumb the following formula can be used:

```
"initial synchronization time" >= "backup size" / "available network bandwidth"
```

If data compression is used, the compression factor (compressed data size compared to uncompressed data size) also needs to be considered. The initial synchronization time for a SAP HANA host can then be calculated via:

```
"initial synchronization time" >= "backup size" * "compression factor" / "available network bandwidth"
```

The overall synchronization time of system replication in a scale-out environment is the maximum of the host specific synchronization times calculated above.

Example: If the backup size is 2 TB, the system replication data compression factor is 0.3 (so 2 TB are compressed to 600 GB) and the available network bandwidth is 300 Mbit / s, the minimum initial synchronization time is around 4.5 hours (2 TB * 0.3 / 300 Mbit/s).

26. Can a takeover or failback happen without terminating active application processes?

The influence of a takeover or failback on client processes depends both on the implementation of the switch process and the client reaction.

In case of SAP ABAP servers the standard reconnect mechanism makes sure that a temporary unavailability of the SAP HANA database is recognized and can take place without resulting in many terminations (see SAP Note [24806](#)).

In SAP HANA environments you can additionally consider to suspend database connections as described in SAP Note [1913302](#). With this approach a takeover or failback can happen without any termination on SAP ABAP side.

27. Which log sizes are typically shipped?

Log shipping happens on a log buffer level. The size of a log buffer typically varies between 4 KB and global.ini -> [persistence] -> log_buffer_size_kb (default: 1 MB), depending on the amount of COMMITs. The "Avg. shipped log buffer size (KB)" provided via SQL: "HANA_Replication_SystemReplication_Overview" (SAP Note [1969700](#)) contains the average size of a shipped log buffer in the system.

The shipped log buffers are in line with the local log buffers, so mechanisms for SAP HANA disk accesses described in SAP Note [1900823](#) also apply to the shipped log buffer sizes.

28. Why is the secondary system consistent although undo logs aren't shipped?

Undo logs are part of the data area and so they are not transferred via log shipping. This is no problem, because:

- Undo logs are part of the data shipping
- When redo logs are applied on the secondary system, undo logs for these changes are automatically created.

29. Are there any instance number / port restrictions when using system replication?

The system replication communication internally communicates via the subsequent port number of the installed system. So if SAP HANA is installed with instance number 00, the communication will also use the 01 related ports (e.g. 30103 for the indexserver). As a consequence it is not possible to install another SAP HANA instance with this increased port number. See SAP Note [2176528](#) for more information. An error like the following will appear when you try to start an instance that wants to use a port that is already occupied by the system replication:

```
Internal Error. Basis/IO/Stream/impl/NetworkChannel.cpp: 1119: Error address in use: $msg$,  
rc=98:  
Address already in use. NetworkChannelBase::bindLocal. bind failed
```

30. How can I check the memory allocation details on the secondary system?

It is not possible to use SQL requests on the secondary system, so you have to use the operating system hdbcons as a fallback (see SAP Note [2222218](#)) and execute "mm list -S" to display the top heap allocators. See SAP Note [1999997](#) for more information.

31. Are SAP HANA histories still available after a takeover?

During a takeover the whole database is opened on the secondary site, so all statistics server histories are available, even if the content is linked to the primary site. Nevertheless there are the following monitoring views with history information that are based on trace files and so the primary site information is no longer

available after a takeover:

- M_CS_LOADS
- M_CS_UNLOADS
- M_EXECUTED_STATEMENTS
- M_EXPENSIVE_STATEMENTS
- M_LOAD_HISTORY_HOST
- M_LOAD_HISTORY_SERVICE

32. What has to be considered in terms of licensing in SAP HANA system replication environments?

For details about licensing see section "License Validity" in the how-to guide [How to perform System Replication for SAP HANA](#) and SAP Note [2685661](#).

33. Can I set up a SAP HANA system replication based on a storage snapshot (instead of a full data backup)?

Yes, this is possible. Prior to setting up SAP HANA system replication a full data backup is required, which can also be provided as a storage snapshot.

However, enabling the system replication based on a storage snapshot (no full data backup was created) is currently (i. e. <= SAP HANA SPS 10) only possible from command line using "hdbnsutil -sr_enable ..." and not from SAP HANA Studio.

34. Is it possible to monitor remote system replication sites on the primary system?

Starting with SAP HANA SPS 11 several monitoring views of the secondary site are available on the primary system via schema `_SYS_SR_SITE_<site_name>`. You can manually query these views using:

```
SELECT ... FROM "_SYS_SR_SITE_<site_name>".<view_name> ...
```

Example: (secondary site: hana02)

```
SELECT * FROM "_SYS_SR_SITE_hana02".M_LOG_BUFFERS
```

It is also possible to use the following SQL statements available via SAP Note [1969700](#) to monitor the secondary site:

```
SQL: "HANA_ColumnStore_Merges"
SQL: "HANA_GarbageCollection_Persistence"
SQL: "HANA_Heap_MemoryObjects"
SQL: "HANA_IO_Commits"
SQL: "HANA_IO_ConverterStatistics"
SQL: "HANA_IO_KeyFigures_Details"
SQL: "HANA_IO_KeyFigures_Total"
SQL: "HANA_IO_Savepoints"
SQL: "HANA_IO_Snapshots"
SQL: "HANA_Jobs_Executors"
SQL: "HANA_liveCache_Containers"
SQL: "HANA_liveCache_Procedures"
SQL: "HANA_Logs_LogBuffers"
SQL: "HANA_Locks_Internal_LockWaits_Overview"
SQL: "HANA_Memory_TopConsumers"
```

SQL: "HANA_Transactions_UndoCleanupFiles"

In order to make sure that the data from the secondary site is read, you can execute the following command in advance:

```
SET SCHEMA "_SYS_SR_SITE_<site_name>"
```

If you use DBACOCKPIT (SAP Note [2222220](#)), you should lock the used work process beforehand (report RSTRC000), because otherwise the adjusted schema is lost in case a change of the work process happens.

Alternatively you can prefix the monitoring view names in the SQL statement manually with "_SYS_SR_SITE_<site_name>".

Be aware that statistics server history tables can't be evaluated for the secondary site, only a subset of M_* monitoring views.

35. What is the logreplay operation mode?

The following table compares the original operation mode delta_datashipping with the operation mode logreplay available as of SAP HANA >= SPS 11:

Operation mode	SAP HANA Version	Description	Advantages
delta_datashipping	no restriction	Synchronization between primary and secondary site is implemented via regular delta data shippings from primary to secondary site.	Memory footprint of secondary site can be smaller as column store tables don't have to be loaded during system replication As a consequence it is possible to set the global_allocation_limit (SAP Note 1999997) and run an additional non-production system on the same machine. History tables are generally supported (for logreplay modes the restrictions described in SAP Note 2480889 apply)
logreplay	>= 1.00.110	Synchronization between primary and secondary site is implemented purely based on redo logs Column store generally has to be loaded into memory to provide reasonable delta merge speed	Delta data shipping no longer required Reduced network traffic between primary and secondary site Reduced network bandwidth requirements between primary and secondary site Reduced takeover times (shorter log replay, no persistence garbage collection, no file ID mapping and disk LOB initialization, no adjustment of resource container dispositions; see SAP Note 2222217) No propagation of logical persistence corruptions on disk level (e.g. introduced by the bug described in SAP Note 2375691)
logreplay_readaccess	>= 2.00	Synchronization between primary and secondary site is implemented purely based on redo logs Column store generally has to be loaded into	Same advantages like "logreplay" above Query load can be moved away from primary site

		memory to provide reasonable delta merge speed See SAP Note 2391079 for restrictions.	
--	--	---	--

The operation mode can be defined with the following parameter:

```
global.ini -> [system_replication] -> operation_mode
```

In 3-tier environments the following details can be considered for the logreplay mode:

- If the tertiary system is down, the logs are only saved on the secondary system, not on the primary system.
- If the secondary system is down and you want to use the tertiary system as secondary system, usually a full sync is required, because the logs retained on primary site are meant for the secondary site and so they may not be compatible with the tertiary system.

SAP Note [2661878](#) provides settings that can be used to optimize the performance of logreplay on secondary system replication site. These settings are a trade-off between logreplay performance and memory utilization and secondary site, so tests are required in order to identify an optimal compromise.

36. How can I determine the current log replay delay?

The following options exist do determine the current log replay delay (i.e. the amount of logs that is already shipped, but not replayed, yet):

Tool	Availability	SAP Note	Details
Statistics server	>= SPS 12	2147247	Starting with SAP HANA SPS 12 alert 94 ("Log replay backlog for system replication secondary") reports increased log replay backlogs.
Mini Checks	>= SPS 12	1999993	Mini check 1852 ("Current log replay backlog (GB)") reports significant log replay backlogs as potentially critical.
SQL	>= SPS 12	1969700	SQL: "HANA_Replication_SystemReplication_KeyFigures" displays among others the log replay backlog (REPLAY_BACKLOG_MB). As a fallback option you can use the contents of M_SERVICE_REPLICATION to determine the log replay delay on the secondary site: SELECT SHIPPED_LOG_POSITION, REPLAYED_LOG_POSITION FROM M_SERVICE_REPLICATION Now you can calculate the difference and multiply it with the log position size of 64 byte: (SHIPPED_LOG_POSITION - REPLAYED_LOG_POSITION) * 64 = <replay_backlog_byte>
hdbcons	>= SPS 11	2222218	As a fallback option you can use the hdbcons option 'replication info' in order to determine the last shipped log position (shippedLogPos) and the last replayed log position (replayFinishLogPos): > hdbcons 'replication info' - shippedLogPos : 0x1098b6940 ... - replayFinishLogPos : 0x10987c1a7 Now you can convert these hexadecimal numbers into decimal numbers, calculate the difference and multiply it with the log position size of 64 byte. The result is the replay backlog: (0x1098b6940 - 0x10987c1a7) * 64 = 239513 * 64 = (4455098688 - 4454859175) * 64 = 239513 * 64 = 15328832 = 14 GB

If you face a high log replay backlog and there are no obvious issues (e.g. in the network between primary

and secondary site) you can check SAP Note [2409671](#) for optimizations.

37. What happens with the content when the secondary site of a system replication environment is upgraded?

During the SAP HANA upgrade of the secondary site new content is installed. This new content is overwritten at the time the system replication from the primary site is established again, so new objects like views will disappear.

This behavior works as designed because during a takeover from primary to secondary site a catalog upgrade will take place on secondary site and the missing content is automatically recreated.

38. Does the secondary SAP HANA system replication site need to be up and running before activating system replication?

The secondary site needs to be stopped before registering it on primary site. If the secondary site can't start (e.g. due to a page corruption), you can re-initialize it using:

```
hdbnsutil -sr_register --force_full_replica
```

39. What has to be considered for the row store reorganization in system replication scenarios?

See SAP Note [2222277](#) ("What has to be considered for the row store reorganization in system replication scenarios?") for more details.

40. What can be done if I do not manage to disable secondary in the normal way?

If a normal stop of the secondary site via "hdbnsutil -sr_unregister --name=<secondary site name>" doesn't work, you can proceed according to SAP Note [2386231](#) and try to remove the system replication configuration on secondary site by stopping secondary and executing "hdbnsutil -sr_cleanup --force".

41. Are there specific SAP HANA users and schemas related to system replication?

The following specific schemas and users exist in the context of SAP HANA system replication:

Name	Available with	Type	Details
__SYS_SR	>= 1.00.110	schema, user	Required for create __SYS_SR_SITE_<site_name> schemas and own their data (SAP Note 2499182)
__SYS_SR_SITE_<site_name>	>= 1.00.110	schema	Contains monitoring views holding data from remote system replication sites, see "Is it possible to monitor remote system replication sites on the primary system?" above for more details

42. What are indications for log shipping timeouts?

The following indications exist for log shipping timeouts:

Indication	SAP Note	Details
Database trace entry	2380176	Closing connection to siteID <site_id>. LogShipping

		was waiting for <seconds> seconds (logshipping_timeout = <timeout_s>)
M_SERVICE_REPLICATION columns SECONDARY_CONNECT_TIME, SECONDARY_RECONNECT_COUNT, REPLICATION_STATUS_DETAILS	1969700	When a log shipping timeout happens, the column SECONDARY_CONNECT_TIME is updated with the reconnect timestamp and SECONDARY_RECONNECT_COUNT is incremented. REPLICATION_STATUS_DETAILS contains 'Log shipping timeout occurred'. SQL statements like SQL: "HANA_Replication_SystemReplication_Overview" and SQL: "HANA_Replication_SystemReplication_KeyFigures" can be used to display these details (among others).

43. Can disk I/O on secondary site have an impact on SYNCMEM system replication scenarios?

SYNCMEM doesn't wait until all data is written to disk on secondary site, so there is no direct impact on the primary site when there are disk I/O issues. If, however, disk I/O is so slow that it can't keep up with the amount of data being received, memory areas like the log buffer can run full. In this case system replication has to wait until space is available in the memory buffer and so an indirect impact of slow disk I/O is possible on the primary site.

44. What is the state of a secondary, if services were not in sync at takeover time?

If the secondary site was in sync at least once and a takeover is performed, the most recent shipped log will determine the timestamp the secondary is running on after takeover. Transactional consistency is given, but newer data, that were created on the primary but could not be shipped, are lost. If the primary and secondary systems have never been in sync or if the secondary site is only started after the primary went down (without the takeover), the secondary system cannot start.

45. What has to be considered when changing the number of SAP HANA nodes?

When nodes are added to or removed from the SAP HANA database and you are anyway in downtime mode, you can proceed as follows:

- Stop system replication
- Adjust node layout first on primary, then on secondary site
- Register secondary site again

Alternatively the change to the host layout can also be done without stopping system replication:

- If a node is added: Add it at first on secondary site, then on primary site
- If a node is removed: Remove it at first on primary site, then on secondary site

46. What has to be considered for Active/Active (read enabled)?

Active/Active (read enabled) is available starting with SAP HANA 2.0 and allows to execute read requests on the secondary system replication site. It is activated during registration of the secondary system via option "--operationMode=logreplay_readaccess" which adds the following entry to global.ini:

```
global.ini -> [system_replication] -> operation_mode =logreplay_readaccess
```

For access restrictions see SAP Notes [2737255](#) and [2391079](#).

Accesses to the secondary system replication site can be executed by using the following hint:

```
RESULT_LAG('hana_sr' [, <seconds>])
```

For hint details see SAP Note [2142945](#).

A default lag threshold can be configured with the following parameter:

```
indexserver.ini -> [hint_result_lag_hana_sr] -> sr_default_lag_time = <seconds>
```

Also a direct connection to the read enabled secondary system is possible.

Be aware that even with synchronous system replication you can't assume that a query to the secondary site always sees the most up-to-date information. Synchronous system replication only guarantees that log shipping happens synchronously, but the actual log replay happens asynchronously. So in rare cases there can be a delay of several minutes or more before a current change is visible on secondary site.

47. Why do I see deviating values in the system replication mode parameter?

The value in parameter global.ini -> [system_replication] -> mode depends on the original role of the site:

- If site was originally configured as primary site: mode = 'primary'
- If site was originally configured as secondary / tertiary site: mode = 'sync', 'async', ... (dependent on the system replication mode)

As a consequence the mode value can be different in two identically configured systems if a takeover happened in one system, but not in the other.

48. How can I check the number of parallel channels used for data shipping?

The number of parallel data shipping channels can be configured with the following SAP HANA parameter:

```
global.ini -> [system_replication] -> datashipping_parallel_channels
```

Per default 4 streams are used for a persistence size of ≥ 4 GB. To monitor the number of used streams, you can activate the sr_dataaccess trace on level "info" (SAP Note [2380176](#)) on secondary system replication site:

```
global.ini -> [trace] -> sr_dataaccess = info
```

In the database trace files you can then grep for "multistream".

Example:

```
sr_dataaccess DisasterRecoverySecondaryImpl.cpp(02207) : multistream, getStreamCount:
connections=4 minStreamSize=1073741824 statusbits=000
sr_dataaccess DisasterRecoverySecondaryImpl.cpp(02235) : multistream,
splitTransferSizeToStreams: maxParallelityPossible=2 is less than that yielded from
configuration=4 streams with maxParallelityPossible
sr_dataaccess DisasterRecoverySecondaryImpl.cpp(02286) : multistream,
splitTransferSizeToStreams: connections=2 restSize=0 (maxParallelityPossible=2
transferSize=2969567232 minStreamSize=1073741824
```

Here one can see that the default setting (4 connections) is configured, but – because the data size (transferSize) is smaller - only 2 parallel channels (connections) are used.

49. How can tracing for system replication be activated?

The following database trace options (SAP Note [2380176](#)) can be useful in context of system replication:

Trace setting	SAP Note	Details
global.ini -> [trace] -> sr_dataaccess = info		System replication data access details (including channel information)
global.ini -> [trace] -> sr_nameserver = debug	2694960	System replication nameserver trace
global.ini -> [trace] -> trexnet = info	2694960	System replication network trace

50. Is it recommended to disable system replication during larger upgrade or migration activities?

In general SAP recommends to keep system replication enabled:

- During regular release / SP upgrades within S/4HANA, SAP recommends to keep high availability mechanisms in regular / production mode. This includes ZDO.
- During Suite on HANA -> S/4HANA system conversion the following options are feasible:
 - Keep system replication active during the entire process. Optionally switch into ASYNC mode to remove pressure from system replication during SUM downtime. This is the recommended option if using Downtime-optimized Conversion (DoC), and it should already be part of the dress rehearsal testing to ensure sufficient hardware resources are available (network bandwidth, latency, ...)
 - Disable system replication just before entering SUM downtime. Re-establish system replication right after SUM downtime has ended.
- During ERP -> S4H system conversion: Establish system replication after SUM downtime has ended.

If (re-)establishing system replication after SUM downtime has ended, take explicit decision whether hand-over to business can happen before or only after system replication has caught up into production mode.

Beyond these best practices the decision to disable system replication during larger upgrade or migration activities has to be taken individually and depends on the following factors:

- Amount of database changes performed
- Capabilities of system replication infrastructure in terms of throughput and latency
- Impact of system replication on primary site performance
- Necessity and overhead of initial data shipping after the activities are finished

Another reason for stopping system replication is to preserve a system state from before the upgrade or migration activity as a fallback in case of problems. On SAP HANA >= 2.0 SPS 03 you can alternatively think about keeping the secondary site up and running and using timetravel in order to get back to the previous state in case of problems.

51. How can pending savepoints and large page cache sizes on secondary site be explained?

When a full data shipment or a delta data shipment takes place, the pages are transferred to the memory page cache (Pool/PersistenceManager/PersistentSpace/DefaultLPA/Page, SAP Note [1999997](#)) and subsequently they are written down to disk on continuous basis. In cases where the shipment across the

network is faster than the I/O write throughput on secondary site, the following symptoms can be observed:

- Secondary site is shown in status "Savepoint Pending" for a longer time and actual synchronization can't start
- Heap allocator Pool/PersistenceManager/PersistentSpace/DefaultLPA/Page on secondary site shows large size and can even be responsible for OOMs
- Many page cache entries are shown with disposition nonSwappable (can be checked with "pageaccess a" option of hdbcons, SAP Note [2222218](#))

All these symptoms are normal in case the I/O on secondary site can't keep up with the data transferred via network from primary to secondary. The OOM situations will disappear as soon as enough page cache entries are written down to disk.

In order to improve the situation you should check if the performance of the I/O stack on secondary system replication site can be improved. See SAP Note [1999930](#) for more information related to I/O analysis in SAP HANA environments.

52. Can I set up system replication between systems with different topologies?

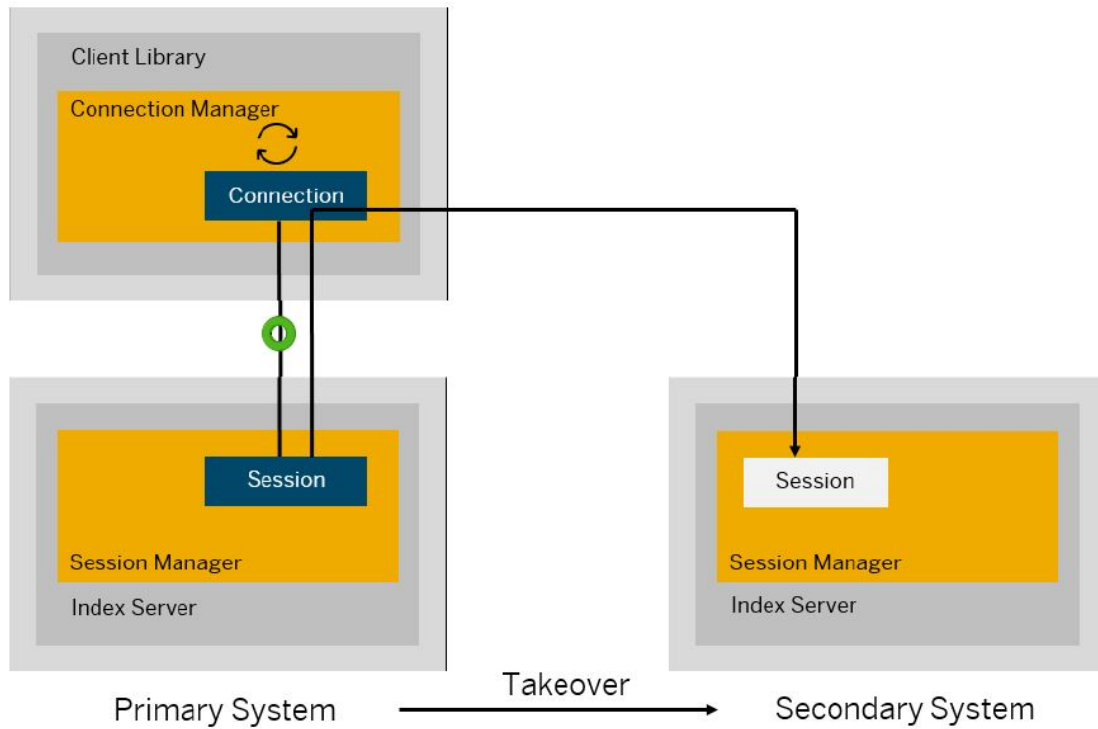
The topology of primary and secondary site of a system replication scenario must be identical. As a consequence it isn't possible to replicate from non-MDC to MDC (SAP Note [2101244](#)) and vice versa, and it is also not possible to replicate from single-node to scale-out and vice versa.

53. Can a primary site be replicated to more than one secondary site?

Up to SAP HANA 2.0 SPS 02 a primary system can only be replicated to a single secondary system. Starting with SAP HANA 2.0 SPS 03 multi-target replication is available and so you can connect several secondary sites to the same primary system.

54. What happens with active transactions and statements during takeover?

Active transactions and statements are usually terminated and rolled back when a takeover happens. Starting with SAP HANA 2.0 SPS 03 it is possible to take advantage of the invisible takeover feature. When invisible takeover is activated, the SAP HANA client (SAP Note [2393013](#)) keeps track of the required connection information from before the switch and continues to use it after the takeover. This so-called transparent session recovery is currently only available for SQLDBC clients. The following diagram visualizes an invisible takeover:



Per default the invisible takeover is deactivated. It can be activated with the following SAP HANA parameter setting:

```
<service>.ini -> [session] -> enable_session_recovery = true
```

Be aware that invisible takeover isn't possible in case of update transactions, i.e. transactions that have already performed modifications to table data.

55. What is secondary time travel?

Secondary time travel is available with SAP HANA >= 2.0 SPS 03 in context of system replication with Active/Active (read enabled) and allows to access data on secondary system replication sites from earlier points in time.

The following parameters can be used to control secondary time travel:

Parameter	Default	Unit	Details
global.ini -> [system_replication] -> timetravel_call_takeover_hooks	false		Indicates if takeover hook should be called during secondary time travel
global.ini -> [system_replication] -> timetravel_max_retention_time	0 (deactivated)	minutes	Controls for how long historic snapshots are preserved for time travel purposes
global.ini -> [system_replication] -> timetravel_snapshot_creation_interval	1440	minutes	Defines how frequently snapshots are created for secondary time travel

Be aware that time travel increases the disk space requirements both in the data area (for snapshots) and the log area (for log replay based on snapshots). Increasing timetravel_max_retention_time and / or reducing timetravel_snapshot_creation_interval increases the disk space requirements.

In order to reset the secondary site to different point in time you can stop the SAP HANA database and execute the following command:

```
hdbnsutil -sr_timetravel --startTime=<start_time>
```

The <start_time> has to be specified as UTC timestamp in the format of "dd.mm.yyyy-hh24.mm.ss".

Example:

```
hdbnsutil -sr_timetravel --startTime=2019/02/22-17:49:30
```

Afterwards the SAP HANA database can be started again.

Keywords

SAP HANA system replication M_SERVICE_REPLICATION high availability disaster recovery

Products

SAP HANA, platform edition all versions

Other Components

Component	Description
HAN-DB-HA	SAP HANA High Availability (System Replication, etc.)

This document refers to

SAP Note/KBA	Title
2737227	Parameter mismatch error is returned even though inifile_checker is configured
2694960	How to Activate Traces and Reproduce HANA System Replication Issues for Troubleshooting and Analysis
2685661	Licensing Required for HANA System Replication
2600030	Parameter Recommendations in SAP HANA Environments
2499182	_SYS_SR system account
2480889	Error: History tables are not allowed in operation mode 'logreplay'
2407186	How-To Guides & Whitepapers For SAP HANA High Availability

2393013	FAQ: SAP HANA Clients
2386231	Recovery issue after disable HSR on HANA
2380176	FAQ: SAP HANA Database Trace
2332284	Data volume reclaim failed because of snapshot pages
2222277	FAQ: SAP HANA Column Store and Row Store
2222220	FAQ: SAP HANA DBACOCKPIT
2222218	FAQ: SAP HANA Database Server Management Console (hdbcons)
2222217	How-To: Troubleshooting SAP HANA Startup Times
2176528	Cannot start two HANA system on one host at the same time
2147247	FAQ: SAP HANA Statistics Server
2142945	FAQ: SAP HANA Hints
2127458	FAQ: SAP HANA Loads and Unloads
2116157	FAQ: SAP HANA Consistency Checks and Corruptions
2101244	FAQ: SAP HANA Multitenant Database Containers (MDC)
2100009	FAQ: SAP HANA Savepoints
1999998	FAQ: SAP HANA Lock Analysis
1999997	FAQ: SAP HANA Memory
1999993	How-To: Interpreting SAP HANA Mini Check Results
1999930	FAQ: SAP HANA I/O Analysis
1977252	How to handle HANA Alert 21: 'Internal event'
1870858	HANA Alerts related to file system utilization
2737255	Limitations on Active/Active(Read Enabled) Feature of SAP HANA in System Replication
2678164	Default Configuration of Parameter [system_replication] logshipping_async_buffer_size Increased for Indexserver
2661878	HANA System Replication log replay setting recommendations for large systems
2544075	Primary Site is Becoming Unresponsive When Connection to Secondary Site is Lost
2391079	Access restrictions in Active/Active (read enabled) system setup
1969700	SQL Statement Collection for SAP HANA
	How to perform System Replication for SAP HANA
	Network Recommendations for SAP HANA System Replication