SAP How-to Guide
Business Analytics

SAP HANA™ Appliance

How To Perform System Replication for SAP HANA

Applicable Release

SAP HANA 2.0 SPS02

Version 5.4

January 2018

For additional information contact:
mechthild.bore-wuesthof@sap.com

The Best-Run Businesses Run SAP™

© Copyright 2017 SAP SE. All rights reserved.

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

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

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

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

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

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

Oracle is a registered trademark of Oracle Corporation.

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

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

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

Java is a registered trademark of Sun Microsystems, Inc.

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

SAP, R/3, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP BusinessObjects Explorer, StreamWork, and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and other countries.

Business Objects and the Business Objects logo, BusinessObjects, Crystal Reports, Crystal Decisions, Web Intelligence, Xcelsius, and other Business Objects products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Business Objects Software Ltd. Business Objects is an SAP company.

Sybase and Adaptive Server, iAnywhere, Sybase 365, SQL Anywhere, and other Sybase products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Sybase, Inc. Sybase is an SAP company.

All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

The information in this document is proprietary to SAP. No part of this document may be reproduced, copied, or transmitted in any form or for any purpose without the express prior written permission of SAP AG.

This document is a preliminary version and not subject to your license agreement or any other agreement with SAP. This document contains only intended strategies, developments, and functionalities of the SAP® product and is not intended to be binding upon SAP to any particular course of business, product strategy, and/or development. Please note that this document is subject to change and may be changed by SAP at any time without notice.

SAP assumes no responsibility for errors or omissions in this document. SAP does not warrant the accuracy or completeness of the information, text, graphics, links, or other items contained within this material. This document is provided without a warranty of any kind, either express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or non-infringement.

SAP shall have no liability for damages of any kind including without limitation direct, special, indirect, or consequential damages that may result from the use of these materials. This limitation shall not apply in cases of intent or gross negligence.

The statutory liability for personal injury and defective products is not affected. SAP has no control over the information that you may access through the use of hot links contained in these materials and does not endorse your use of third-party Web pages nor provide any warranty whatsoever relating to third-party Web pages.

SAP "How-to" Guides are intended to simplify the product implementation. While specific product features and procedures typically are explained in a practical business context, it is not implied that those features and procedures are the only approach in solving a specific business problem using SAP NetWeaver. Should you wish to receive additional information, clarification or support, please refer to SAP Consulting.

Any software coding and/or code lines / strings ("Code") included in this documentation are only examples and are not intended to be used in a productive system environment. The Code is only intended better explain and visualize the syntax and phrasing rules of certain coding. SAP does not warrant the correctness and completeness of the Code given herein, and SAP shall not be liable for errors or damages caused by the usage of the Code, except if such damages were caused by SAP intentionally or grossly negligent.

Disclaimer

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

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

The Best-Run Businesses Run SAP™

Document History

Document Version	Description
1.0	First official release of this guide
1.1	Small updates and language edit
1.2	Some descriptions were corrected
2.0	Update with HANA 1.0 SPS07 and SPS08 features
2.3	Minor corrections
2.4	Secondary site usage for QA/DEV – small update
2.6	Reference to docu on e-mail notification for alerts
3.0	Update with HANA 1.0 SPS09 features
3.1	Minor corrections
3.2	Fixed return codes of systemReplicationStatus.py
3.3	Updated network settings & sr_state output
3.4	Corrected system replication hostname resolution
3.5	Update with HANA 1.0 SPS10 features
4.0	Update with HANA 1.0 SPS11 features & minor corrections
4.1	Update with HANA 1.0 SPS12 features & minor additions
4.2	Minor corrections
4.3	Minor corrections
5.0	HANA 2.0 SPS00 – including Active/Active (read enabled)
5.1	Update with HANA 2.0 SPS02 features
5.2	Added some network relevant information
5.3	Data retention clarification
5.4	Minor corrections

Typographic Conventions

Type Style	Description
Example Text	Words or characters quoted from the screen. These include field names, screen titles, pushbuttons labels, menu names, menu paths, and menu options. Cross-references to other documentation
Example text	Emphasized words or phrases in body text, graphic titles, and table titles
`Example text`	File and directory names and their paths, messages, names of variables and parameters, source text, and names of installation, upgrade and database tools.
`Example text`	User entry texts. These are words or characters that you enter in the system exactly as they appear in the documentation.
`<Example text>`	Variable user entry. Angle brackets indicate that you replace these words and characters with appropriate entries to make entries in the system.
`EXAMPLE TEXT`	Keys on the keyboard, for example, `F2` or `ENTER`.

Icons

Icon	Description
	Caution
	Note or Important
	Example
	Recommendation or Tip

The Best-Run Businesses Run SAP™

Table of Contents

1. Business Scenario 1
2. Before you start 1
3. Background Information 2
   - 3.1 High Availability 2
   - 3.2 System Replication 3
     - 3.2.1 Replication modes 4
     - 3.2.2 Operation modes 4
     - 3.2.3 Data transferred to the secondary 6
   - 3.3 Takeover 8
4. Planning 9
   - 4.1 Prerequisites 9
   - 4.2 Network recommendations 10
     - 4.2.1 Distance between data centers 10
     - 4.2.2 Throughput 11
   - 4.3 Use secondary site for other systems (DEV/QA/...) 13
     - 4.3.1 Resources required on the secondary site 15
   - 4.4 License validity 16
5. Configuration steps 16
   - 5.1 Configure system replication 17
     - 5.1.1 Using SAP HANA cockpit 2.0 17
     - 5.1.2 Using SAP HANA studio 23
     - 5.1.3 Using command line tool `hdbnsutil` 26
     - 5.1.4 Creating a tenant DB in a running system replication 27
   - 5.2 Disable system replication 28
     - 5.2.1 Using SAP HANA cockpit 2.0 28
     - 5.2.2 Using SAP HANA studio 30
     - 5.2.3 Using command line tool `hdbnsutil` 31
   - 5.3 Setting up Multitier System Replication 31
     - 5.3.1 Using SAP HANA cockpit 2.0 33
     - 5.3.2 Using SAP HANA studio 36
     - 5.3.3 Using command line tool `hdbnsutil` 38
   - 5.4 Enabling full sync replication 38
   - 5.5 Changing the replication mode 39
   - 5.6 Changing the operation mode 40
6. Takeover 41
   - 6.1 Perform Takeover 41
     - 6.1.1 Using SAP HANA cockpit 2.0 41

6.1.2 Using SAP HANA studio 43
6.1.3 Using command line tool `hdbnsutil` 44
6.2 Client connection recovery 44
6.2.1 Virtual IP addresses in Active/Active (read enabled) system replication 45

7. Resync optimization 46
7.1 Data Retention 46
7.2 Log Retention 47
7.2.1 Log Retention for Secondary Disconnect (on primary) 47
7.2.2 Log Retention for Failback (on secondary) 48
7.2.3 Log Retention Parameters 49
7.2.4 Maximum Retention Time Estimation 50

8. Active/Active (read enabled) secondary 53
8.1 Active/Active read-only specifics 53
8.2 Check Active/Active configuration 54
8.2.1 Using the command line 54
8.2.2 Using SAP HANA cockpit 2.0 54
8.2.3 Using HANA studio 56
8.3 Connection types 56
8.3.1 Explicit read-only connection to secondary 56
8.3.2 Statement routing from primary to secondary 57
8.3.3 Virtual IP address handling 58

9. Testing 58

10. Operation / Maintenance 59
10.1 Alerts 59
10.2 Verification 61
10.2.1 Using SAP HANA cockpit 2.0 61
10.2.2 Using SAP HANA studio 68
10.2.3 Check via SQL query 68
10.2.4 Using command line tool `hdbnsutil` 69
10.3 System Replication status checks 71
10.3.1 Using `landscapeHostConfiguration.py` 71
10.3.2 Using `systemReplicationStatus.py` 73
10.3.3 Using console 73
10.3.4 Using predefined SQL statement 76
10.4 Monitoring and replicating `ini` parameter changes 79
10.5 Monitoring the secondary site 81
10.6 System replication network connection 84
10.6.1 Secure configuration of the network connection 84

10.6.2 Allowed senders for the connection 86
10.6.3 Encryption of the connection 87
10.6.4 Network security settings in HANA cockpit 88
10.6.5 Data and Log compression for transfer 88
10.6.6 Monitoring the network connection.....89
10.6.7 Monitoring the network latency 90
10.6.8 Network Speed Checks for System Replication 93
10.6.9 Network Speed Checks for Scale-out System.....93
10.7 Tenant Copy/Move within System Replication 94
10.8 Data and Log Volume Encryption 95
10.9 Upgrade and Maintenance 96
10.9.1 "Normal" Upgrade 97
10.9.2 Near Zero Downtime Upgrade 97
10.9.3 Hardware Exchange.....99

11. Further documentation 99
11.1 Whitepapers and HowTo guides 99
11.2 Official Guides 100

1. Business Scenario

Business Continuity requires that the operation of business-critical systems remain highly available at all time, even in the presence of failures. High availability and disaster recovery are the building blocks to support this.

Among other features, SAP HANA provides the possibility to replicate your SAP HANA system within the same or over two data centers. This paper briefly describes SAP HANA System Replication in a step-by-step manner to support High Availability and Disaster Recovery and references the needed guides for details.

2. Before you start

It is recommended to have studied the corresponding parts of the following documents and SAP Notes, which are frequently referred to in this document:

- *SAP HANA High Availability Whitepaper*:
  Available over this SAP Note 2407186 - How-To Guides & Whitepapers For SAP HANA High Availability
- *SAP HANA Administration Guide*, *SAP HANA Master Guide*, *SAP HANA Server Installation Guide*, *SAP HANA Security Guide*; which are all accessible from here:
  https://help.sap.com/viewer/p/SAP_HANA_PLATFORM
- SAP Note 2423367 - Multitenant database containers will become the standard and only operation mode

The first document will give a broad overview and basic knowledge to understand what this paper discusses.

Additionally, there is a set of SAP HANA Academy videos available, which are worth watching:

- *SAP HANA Academy on system replication*:
  http://scn.sap.com/community/hana-in-memory/blog/2015/05/19/sap-hana-system-replication
- *SAP HANA Academy on "What's New with HANA 2.0 SPS00"*:
  https://blogs.sap.com/2016/12/21/sap-hana-2.0-sps-00-whats-new-high-availability-by-the-sap-hana-academy/

You should also be aware of these SAP notes containing valuable information on SAP HANA system replication:

- SAP Note 1999880 - FAQ: SAP HANA System Replication

- SAP Note 2057595 - FAQ: SAP HANA High Availability
- SAP Note 2165547 - FAQ: SAP HANA Database Backup & Recovery in an SAP HANA System Replication Landscape
- SAP Note 1945676 - Correct usage of hdbnsutil -sr_unregister
- SAP Note 1984882 – Using HANA system replication for Hardware Exchange with Minimum Downtime
- SAP Note 2063657 - HANA System Replication takeover decision guideline
- SAP Note 1913302 - Connectivity suspend of Appserver during takeover.

The following blogs also discuss the topic; please feel free to comment:

- *HANA System Replication - Takeover process:* http://scn.sap.com/docs/DOC-52345
- *HANA System Replication – Backup:* http://scn.sap.com/docs/DOC-53608
- *HANA System Replication – Switching back and forth:* http://scn.sap.com/community/hana-in-memory/blog/2013/12/16/sap-hana-system-replication--using-hdbnsutil-sr

For information about SAP HANA in general, see:
http://help.sap.com/hana_appliance

3. Background Information

3.1 High Availability

SAP HANA offers different kinds of high availability mechanisms, supporting a broad range of recovery scenarios from various faults.

There are three basic scenarios:

- *Host Auto-Failover*: One (or more) standby nodes are added to an SAP HANA system and configured to work in *standby mode*. (SAP HANA scale-out)[1].
- *Storage Replication*: The storage itself replicates all data to another location (this solution is provided by hardware partners). Disks are mirrored without a control process from the SAP HANA system.
- *System Replication*: SAP HANA replicates all data to a secondary SAP HANA system (standard SAP HANA feature). Data is constantly pre-

[1] For technical details of this solution, please refer to the Host Auto-Failover technical document: http://scn.sap.com/docs/DOC-62494

loaded on the secondary system to minimize the recovery time objective (RTO).

This paper focuses on supporting decision making on SAP HANA system replication including setting up, testing and maintaining such a system landscape. Of course, a comprehensive high availability solution offers more design choices and requires the discussion of more details than can be covered in a short paper; thus, additional consultation may be required.

3.2 System Replication

SAP HANA system replication ships all data to a secondary system located at another site:

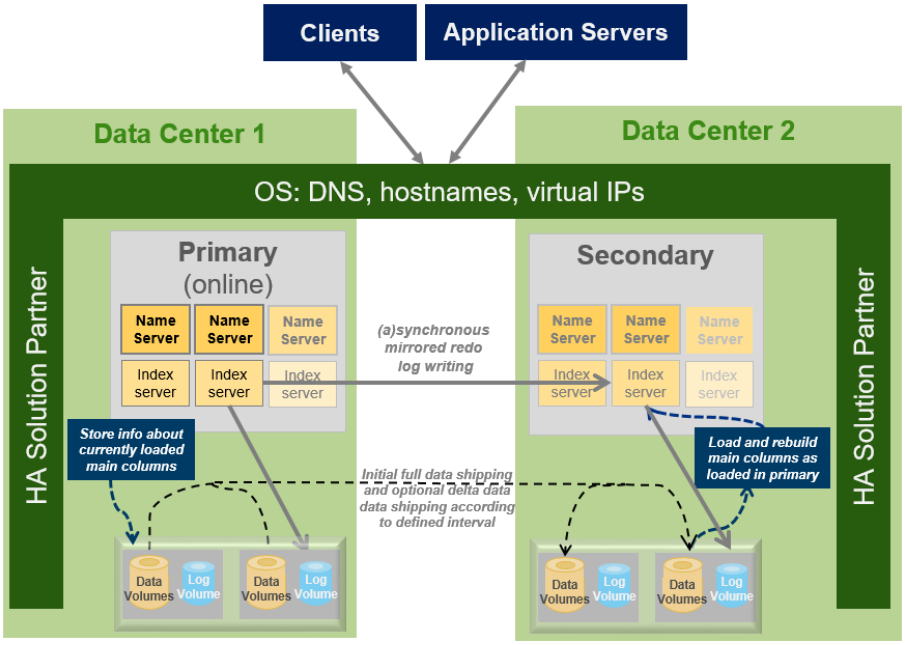

Overview of System Replication with single node HANA databases

Once SAP HANA system replication is enabled, each server process on the secondary system establishes a connection with its primary counterpart and requests a snapshot of the data. Now all logged changes in the primary system are replicated continuously. Each persisted redo log in the primary system is sent to the secondary system. A transaction in the primary system is not committed before the redo logs are replicated.

SAP *HANA Tenant Database Systems*[2] (introduced with HANA 1.0 SPS09 and

[2] Formerly called "SAP HANA Multitenant Database Containers".

default as of HANA 2.0 SPS01[3]) can also run in an SAP HANA system replication configuration. The system as a whole is replicated[4], i. e. the System DB and all tenant DBs. Just like in the HANA single container system each service with a persistency (i. e. data and log volume) of the primary site replicates to its counterpart on the secondary site.

While the system replication is running, the secondary system, which is configured identically to the primary system, will be on standby until a takeover takes place. As of HANA 2.0 SPS00 the secondary system can also be used for read access, if configured as Active/Active (read enabled) system.

3.2.1 Replication modes

Depending on customer requirements, SAP HANA offers different modes for replication of the redo log:

- *Synchronous*: Secondary system sends acknowledgement back to primary as soon as data are received and persisted to the log volumes on disk.
- *Synchronous in-memory*: Secondary system sends acknowledgement back to primary as soon as data is received in memory; the disk I/O speed on the secondary does not influence the primary's performance.
- *Asynchronous*: As per design of asynchronous replication, the primary does not wait until the secondary sends an acknowledgement.

Additionally, (as of HANA 1.0 SPS08) the *synchronous* replication mode (SYNC) can run with "*full sync*" enabled. In *full sync* operation, transaction processing on the primary site blocks, when the secondary is currently not connected and newly created redo log buffers cannot be shipped to the secondary site. This behavior ensures that no transaction can be committed locally without shipping the redo log buffers to the secondary site.

Changing the replication mode is done as described in the section "Changing the replication mode" of this document.

3.2.2 Operation modes

Since HANA 2.0 SPS00 SAP HANA system replication can be run in three different operation modes:

[3] Please refer to SAP Note 2423367 - *Multitenant database containers will become the standard and only operation mode*

[4] Note that SAP HANA system replication on tenant database level is not supported.

- *delta_datashipping*: In addition to the continuous redo log shipping taking place the secondary system requests a delta data shipping from time to time (per default every 10 minutes). During takeover, the redo log needs to be replayed up to the last arrived delta data shipment. (This is the "classical" operation mode of SAP HANA system replication.)

- *logreplay*: In this operation mode (since HANA 1.0 SPS11) pure redo log shipping is done after the system replication was initially set up with one full data shipping. The redo log is replayed on the secondary immediately after arrival making this step superfluous during a takeover, which shortens the RTO by factors. Additionally, the amount of data which needs to be transferred to the secondary site is reduced dramatically, because no delta data shipping is required anymore.

- *logreplay_readaccess*: Regarding the continuous log shipping, the redo log replay on the secondary site as well as the required initial full data shipping and the takeover, this operation mode (since HANA 2.0 SPS00) behaves just like the "`logreplay`" operation mode. The difference is that here the secondary system is read enabled; i. e. SQL SELECT queries are possible[5]. By establishing a direct connection to the secondary database or by providing a SELECT statement from the primary with a HINT, read access is possible on the Active/Active (read enabled) secondary system.

Using the operation mode `logreplay` makes your secondary site in the SAP HANA system replication a HotStandby system; using operation mode `logreplay_readaccess` allows for read-only access on the read enabled secondary system making an Active/Active (read enabled) system out of your system replication landscape.

Note

In a multitier system replication, only one operation mode is allowed for the whole landscape, with one exception:

If operation mode `logreplay_readaccess` is configured between primary (tier-1) and the tier-2 secondary, then operation mode `logreplay` is allowed between tier-2 and tier-3 secondary.

[5] For read access details on the secondary system please refer to 2391079 - Access restrictions in Active/Active system setup.

Switching the operation mode in a configured and running system replication requires to shut down the secondary and re-register it to the primary again. For details on the required steps please read section "Changing the operation mode" in this document.

3.2.3 Data transferred to the secondary

The SAP HANA database system sends two resp. three types of data "packages" over the network to the secondary side (depending on the configured operation mode), when system replication is configured:

- *Full data shipping*: A full set of the data created as HANA in-place snapshot on the disk of the primary is initially sent when system replication is set up.[6]

- *Delta data shipping*: Only in **`delta_datashipping`** operation mode the increment of the data (i.e. every data that has changed since the last full or the last delta data shipping), is transported from time to time (default every 10 minutes) from the data area of the primary to the data area of the secondary.

- *Redo Log shipping*: Every committing write transaction on the primary generates redo log buffers that are continuously sent to the secondary site.

The following picture visualizes this traffic on the transportation channel between primary and secondary for the **`delta_datashipping`** operation mode.

[6] If the connection between primary and secondary is weak and it would take very long to get this initial full data shipping through the network channel, as of HANA 1.0 SPS12 the secondary system can be initialized with a consistent storage snapshot. Please see the SAP HANA Administration guide – section: *Initialize the Secondary with Storage Copy from Primary*

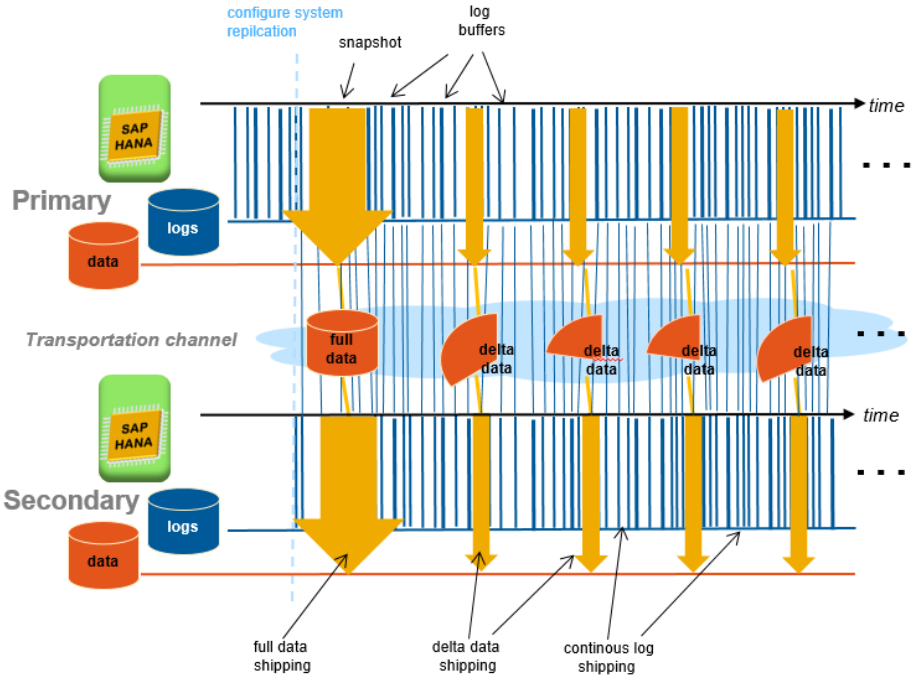

Operation mode ***delta_datashipping****:*
Initial full data shipping, frequent delta data shipping
and continuous redo log shipped to secondary

In **logreplay** or **logreplay_readaccess** operation mode the delta data shipping is not required.

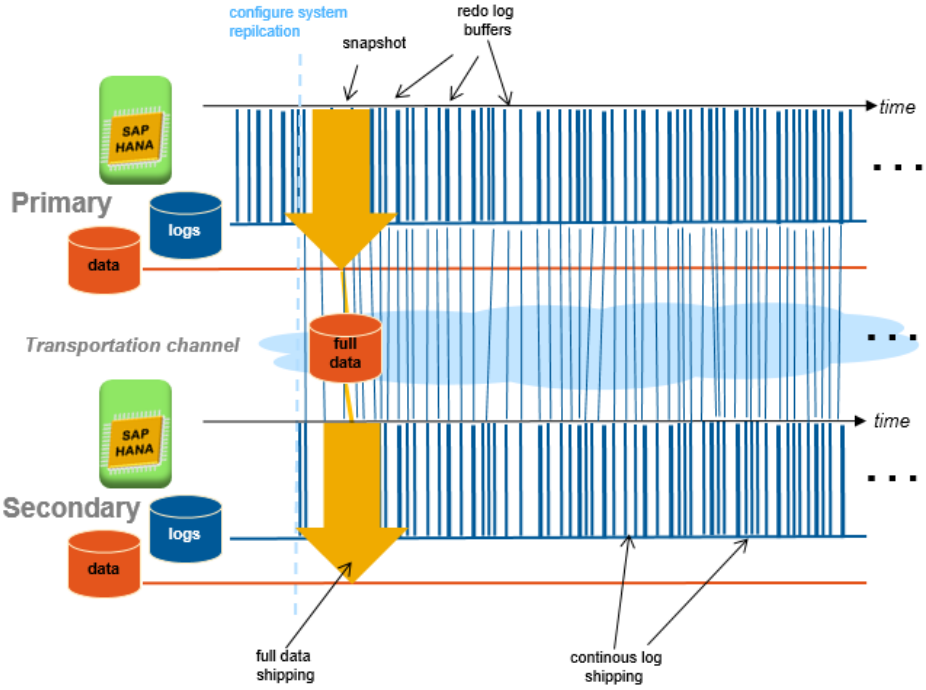

Operation mode ***logreplay*** *or* ***logreplay_readaccess***
Initial full data shipping and continuous redo log shipping to secondary

Tipp

With the `ini` file parameter `datashipping_parallel_channels` (default: 4) the full and the delta data shipping is done using parallel network channels. It can be changed on the secondary system in `global.ini` section `[system_replication]`.

3.3 Takeover

The takeover process is the name for the task of switching your active system from the current primary system onto the secondary system[7]. Once the takeover command runs, the former secondary system becomes the new "primary" system – more correctly, it becomes your new actively running system.

The takeover automatically performs some tasks before the system is fully available:

- Until HANA 1.0 SPS08 the row store tables were loaded into memory during takeover; since HANA 1.0 SPS09 the row store is kept in shared memory and thus is "pre-loaded".
- Until HANA 1.0 SPS09 the row store indexes were rebuilt during takeover; with HANA 1.0 SPS10 rebuilding the secondary indexes during reactivation of the row store is done in a decoupled way, so it does not influence the takeover performance.
- Until HANA 1.0 SPS10 the redo log buffers shipped to the secondary site since the last delta data transport could first be replayed during takeover; with the two **`logreplay*`** operation modes (`logreply` since HANA 1.0 SPS11 and `logreplay_readaccess` since HANA 2.0 SPS00) the log is continuously replayed on the secondary site, improving the takeover performance.
- If preload is used, the main parts of the column tables are already loaded into memory, as they were loaded in the primary. The first access to a table that was previously used in the primary loads the delta part only. The delta part is typically much smaller than the main part and can be loaded within seconds in most cases.

[7] Note that a takeover does not include stopping the previous primary, if it is still running after takeover, if you did not take measures to stop it!

4. Planning

Let us discuss some facts, which need to be considered or decided during the planning phase.

4.1 Prerequisites

Before you start setting up SAP HANA system replication, your HANA databases need to fulfill the following prerequisites:

- The primary and secondary systems are both installed and configured. You have verified that both are independently up and running.
- In HANA 2.0 the System PKI SSFS key and data files need to be copied from the primary to the secondary systems as described in this SAP Note 2369981. The files can be found here:

  ```
  /usr/sap/<SID>/SYS/global/security/rsecssfs/data/SSFS_<SID>.DAT
  /usr/sap/<SID>/SYS/global/security/rsecssfs/key/SSFS_<SID>.KEY
  ```
- The number of nodes in the secondary system must be equal to the number of active nodes in the primary system. (As of HANA 1.0 SPS06 the secondary system does not need to have standby nodes.)
- All configuration steps have to be executed on the master name server node; for SAP HANA Tenant Database Systems this means on the System DB (and not on the tenant DBs).
- The SAP HANA software version of the secondary must be equal to or newer than the one on the primary; however, if you want to make use of a read enabled secondary system in an Active/Active (read enabled) configuration, the SAP HANA software versions and the CPU architectures (same endianness) must be identical.
- The secondary system must have the same SAP system ID, `<SID>`, and instance number as the primary system. The primary replicates all relevant license information to the secondary.
- System replication between two systems on the same host is not supported.
- Configuring system replication between SAP HANA single container systems (default until HANA 2.0 SPS00) and SAP HANA Tenant Database System (default as of HANA 2.0 SPS01) is not supported.
- Changes to the `ini` file configuration parameters made on one system should be duplicated on the other system; except for some very host-specific parameters, like host names.

  As of HANA 1.0 SPS06 the `configuration parameter checker` reports differences between primary and secondary parameter settings (generating alerts in the SAP HANA studio). As of HANA 1.0 SPS12 `INI` parameters can be replicated to the secondary system.

- The required ports must be available. The same `<instance number>` is used for primary and secondary systems. The `<instance number>+1` must be free on both systems, because this port range is used for system replication communication.[8]
- An initial data backup or snapshot must be performed on the primary before the system replication can be set up. In an SAP HANA Tenant Database System all databases must have been backed up, i. e. the system DB as well as all tenant DBs[9].

4.2 Network recommendations

The network connection between the primary and the secondary systems is of importance for the overall performance of SAP HANA databases configured in a SAP HANA system replication. Please find some basic information on this topic in this section.

Network throughput requirements of the communication channel used in SAP HANA system replication differs for the used operation mode. The network latency requirement is only important for synchronous replication modes, because the primary's performance is influenced by the time it takes until the acknowledgement from the secondary site arrives.

For further information on monitoring and configuring the network connection between the primary and the secondary systems for SAP HANA system replication please refer to the section "*System Replication Network Connection*" in this document.

4.2.1 Distance between data centers

System replication offers synchronous and asynchronous replication modes to accommodate network latency. If the distance between your sites is less than 100 km you can use a synchronous replication mode: `SYNC` or `SYNCMEM`. For all data centers that are more than 100 km apart, the asynchronous replication mode `ASYNC` is recommended.

Note

Depending on latency, data volume, volume of changed data records, this could lead to loss of changes because of missing redo logs. Please also consider monitoring requirements for asynchronous mode.

[8] For additional port specific information in Multitenant Database Containers running in System Replication please refer to http://help.sap.com/hana/SAP_HANA_Administration_Guide_en.pdf - section: *SAP HANA System Replication with Multitenant Database Containers*

[9] In an already running SAP HANA system replication for a Tenant Database System, every newly created tenant DB has to be backed up for the replication to start for this tenant.

Tipp

For further information on `ASYNC` replication see section "*Check Network Configuration (Long Distance)*" in the SAP HANA Troubleshooting and Performance Analysis Guide.

4.2.2 Throughput

To estimate the required throughput, you need to know the size of data and log that are generated during your daily workload. To gain this information you can use one of the SQL statements contained in the zip file, which is attached to this SAP Note 1969700. It provides a set of complex SQL statements – including some that are relevant for SAP HANA system replication – that can be executed for example by importing them to the SAP HANA studio as follows:

For the primary system go to the *System Information* tab and right-click in the "Name" column → *Import SQL Statements*.

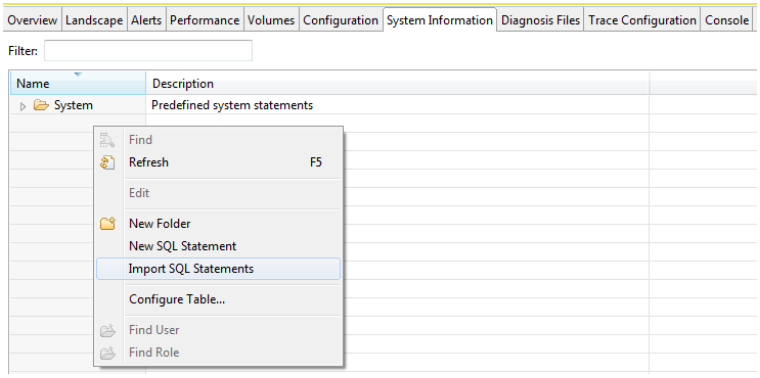

Select the "`SQL Statements.zip`" file you downloaded from the SAP note:

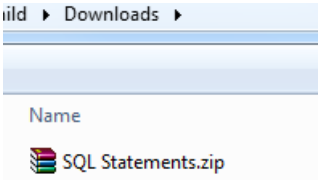

A folder with the corresponding SQL statements will be imported.

Right-click on the "Bandwidth" statement under *Replication* → *Bandwidth* and select *Open in SQL console*.

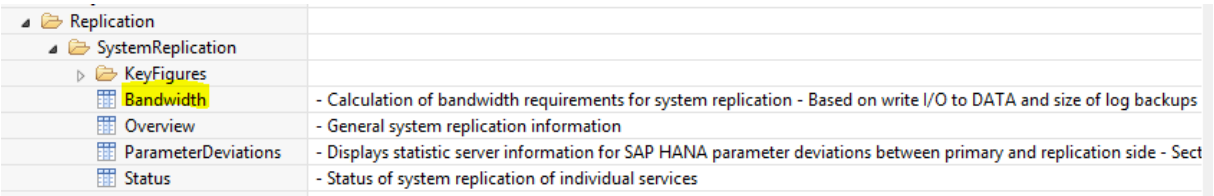

Then you will get the complete SQL statement (which is also contained in the mentioned ZIP file as `HANA_Replication_SystemReplication_Bandwidth.txt`).
In this SQL statement, you can change the **`/* Modification section */`** so that the results are displayed per `'DAY'`; for example like this:

SQL | Result

```
FROM
( SELECT                                    /* Modification section */
    TO_TIMESTAMP('1000/01/01 18:00:00', 'YYYY/MM/DD HH24:MI:SS') BEGIN_TIME,
    TO_TIMESTAMP('9999/12/31 18:10:00', 'YYYY/MM/DD HH24:MI:SS') END_TIME,
    '%' HOST,
    'TIME' AGGREGATE_BY,          /* TIME, HOST and comma separated combinations or NONE for no aggregation */
    'DAY' TIME_AGGREGATE_BY       /* HOUR, DAY, HOUR_OF_DAY or HANA time pattern, NONE for no aggregation */
  FROM
```

	SNAPSHOT_TIME	HOST	PERSISTENCE_GB	DATA_SIZE_GB	LOG_SIZE_GB	TOTAL_SIZE_GB	LOG_PCT	AVG_BANDWIDTH_MBIT	SIMPLE_BANDWIDTH_MBIT
1	2016/01/08 (FRI)	any	18.24	28.60	31.38	59.99	52.31	5.68	1.73
2	2016/01/07 (THU)	any	18.24	8.02	0.61	8.63	7.13	0.81	1.73
3	2016/01/06 (WED)	any	18.24	11.41	3.04	14.45	21.04	1.37	1.73
4	2016/01/05 (TUE)	any	18.24	2.96	0.32	3.29	9.91	0.31	1.73

When executing this statement, you receive the above shown result which provides the following information:

`SNAPSHOT_TIME`	Time slot for which the results are valid
`HOST`	Host name
`PERSISTENCE_GB`	(Current) persistence data size (GB)
`DATA_SIZE_GB`	Total size of data written to disk (GB)
`LOG_SIZE_GB`	Total size of logs generated (GB)
`TOTAL_SIZE_GB`	Total amount of data and logs generated (GB)
`LOG_PCT`	Percentage of log compared to total size (%)
`AVG_BANDWIDTH_MBIT`	Average required network bandwidth to replication side (Mbit), only available for certain `TIME_AGGREGATE_BY` values
`SIMPLE_BANDWITH_MBIT`	Simple network bandwidth calculation (Mbit) based on the formula that it should be possible to ship the persistence once per day

The requirements regarding the network throughput depend on the operation mode selected, because (as described above) different sizes of data are shipped over the network. **`PERSISTENCE_GB`** and **`LOG_SIZE_GB`** are the most important values in this context. The following overview distinguishes between the operation modes making use of these values:

Operation mode	Throughput requirements
`delta_datashipping`	*Throughput* It must be possible to transport the size of the persistently stored data within one day from the primary to the secondary. The following example further explains this requirement:

	Given: 4,3 TB of persistently stored data Throughput: 4,3 TB/day ⇔ ~ 50 MByte/s => ~0,5 GBit/s connection required The size of the persistently stored data can be obtained from the above-mentioned bandwidth SQL from column `PERSISTENCE_GB`.
`logreplay & logreplay_readaccess`	With this operation mode, no delta data shippings are necessary; pure log shipping is done. The network throughput requirements mainly depend on the transactional workload on your primary system. *Throughput* It must be possible to transport the size of the log backups of one day from the primary to the secondary within one day. When executing the mentioned *bandwidth SQL statement*, you receive the shown result – among others the sizes of log that was created per day (= `LOG_SIZE_GB`). During normal system replication operation, your network channel must be capable of handling this amount of data per day – focusing on peak values like 31.38 GB in the simple example above.

4.3 Use secondary site for other systems (DEV/QA/…)

In system replication landscapes, it is possible to run other systems (like DEV/QA systems, or even productive systems) on the secondary site's hardware while the primary system is in production.

Hint

Please refer to SAP Note 1681092 - Multiple SAP HANA DBMSs (SIDs) on one SAP HANA system - describing that as of HANA 1.0 SPS09 multiple SAP HANA systems (identified by SIDs) can run on one host productively. For SAP HANA system replication, this SAP note includes that for example setups like the following – cross replication of two systems with two hosts – are supported (provided that the hosts have enough resources to handle the load):

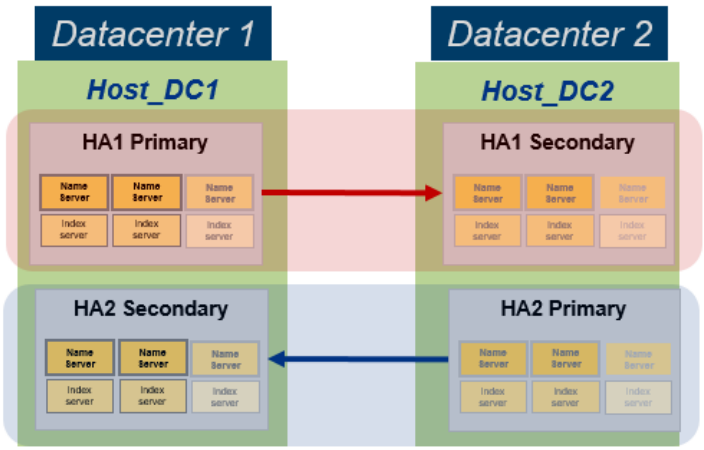

Additionally, the following prerequisites must be considered:

- Additional independent disk volume is needed for the other systems; since the secondary system requires the same I/O capacity as the primary, the additional systems must not have a negative impact on the secondary's I/O – thus it is recommended to have a separated storage infrastructure for each system.
- The SIDs and instance numbers used for the additional systems running on the secondary hardware, must be different from the system replication SID.
- The `<instance number>+1` of the productive system must not be used but must be free on both sites, because this port range is used for system replication communication.
- To save memory resources preload of tables could be switched off on the secondary system:

  ```
  global.ini/[system_replication]->
  preload_column_tables=false
  ```

  - The takeover process will take longer as no data is preloaded to memory on the secondary site (could still meet SLAs for disaster recovery)
- The available memory on the secondary site's hardware must be shared among the systems running there – taking care that the secondary system is assigned the amount of memory it needs. Only the remaining memory can be shared among the DEV/QA systems.
  - Memory allocation for every system running on the secondary's hardware is done by setting the global allocation limit for each system:

    ```
    global.ini/[memorymanager]->
    global_allocation_limit
    ```
- The additional systems running on the secondary site need to be shut down in case of a takeover, if there are not enough resources to handle the load of all systems. Have a look at the following section for details.

- As of HANA 1.0 SPS07 it is possible to change the `global.ini` on the secondary accordingly and then activate the change with "`hdbnsutil -reconfig`" (because no SQL is possible in this state).

4.3.1 Resources required on the secondary site

Running other systems on the secondary site depends on the available hardware resources, the table preload option (`preload_column_tables`) set for and the memory (`global_allocation_limit`) needed by the secondary system of the system replication.

The following table lists these requirements of the secondary system for each combination of table preload and operation mode:

Load option	Operation mode	Memory needed for the secondary system (`global_allocation_limit`)
Preload on	`delta_datashipping`	Set `global_allocation_limit` to the same value as memory is available on the primary
Preload on	`logreplay` (HANA 1.0 SPS11)	Set `global_allocation_limit` to the same value as memory is available on the primary
Preload on	`logreplay_readaccess` (HANA 2.0 SPS00)	Set `global_allocation_limit` to the same value as memory is available on the primary ⚠ Caution Consider that read access on the Active/Active (read enabled) secondary system requires additional CPU capacity
Preload off	`delta_datashipping`	minimum 64 GB or row store size[10] + 20 GB (if this sum is higher)[11] ⚠ Caution If the row store size grows during operation of the primary, it might become necessary to increase the `global_allocation_limit` on the secondary site.
Preload off	`logreplay` (HANA 1.0 SPS11)	size of loaded column tables (in-memory)[12] + row store size[10] + 50 GB ⚠ Caution

[10] The row store size can be determined with this SQL statement:

```
select host, round(sum(page_size*USED_BLOCK_COUNT)/1024/1024/1024,2) as "RowStore Size GB"
from m_data_volume_page_statistics where page_sizeclass = '16k-RowStore' group by host;
```

In SAP HANA Tenant Database Systems the sum of the row store size of all databases (SystemDB and all tenant DBs) is to be used here.

[11] If this limit is not set, the HANA database on the secondary site uses as much memory as it can get and possibly takes it away from the DEV/QA systems, which could run into out-of-memory.

[12] The size of loaded column tables (in-memory) can be found with this SQL statement:

```
select round(sum(memory_size_in_total)/1024/1024/1024) size_GB from m_cs_tables;
```

Load option	Operation mode	Memory needed for the secondary system (`global_allocation_limit`)
		If the row store size grows during operation of the primary, it might become necessary to increase the `global_allocation_limit` on the secondary site.
Preload off	`logreplay_readaccess` (HANA 2.0 SPS00)	Set `global_allocation_limit` to the same value as memory is available on the primary. Note Preload off reduces the required memory size in the Active/Active (read enabled) secondary by tables which are used read-only in the primary system. Caution Consider that read access on the Active/Active (read enabled) secondary system requires additional CPU capacity

All resources that are not needed by the secondary system can be used to run further systems on the secondary site. These "left over" resources must be granted to them by setting the global allocation limit explicitly.

4.4 License validity

The primary system automatically replicates relevant license information to the secondary. No additional license needs to be installed, since the primary and secondary system have the same SID. Further information on licensing in SAP HANA system replication can be found in SAP note 2211663.

In Active/Active (read enabled) system replication configurations also the secondary system is subject to licensing. For further details please refer to the SAP HANA Feature Scope Description.

5. Configuration steps

This section describes the following steps:

- Perform an initial data backup or a storage snapshot using native HANA options. An initial data backup or snapshot is mandatory but an up-to-date backup is highly recommended anyway.
- Enable the primary system for system replication
- Establish a connection between secondary and primary system
- Initiate a full data replication by configuring system replication on the secondary and starting it – thereafter incremental data replication (only in `delta_datashipping` operation mode) and continuous redo log replication (in both `logreplay` operation modes) starts automatically
- Disable system replication on secondary system
- Disable system replication on primary system

- Monitor status of system replication to ensure that both systems are active and in sync

System replication can be set up in three ways: on the console via command line, using the SAP HANA cockpit 2.0 or the SAP HANA studio. The primary system stays online during this procedure.

SAP HANA cockpit 2.0 and SAP HANA studio provide an easy way to set up and maintain system replication, whereas during run time the command line will probably be used, because it can be a part of a script, which executes further steps beyond system replication.

5.1 Configure system replication

To configure system replication the primary system must have been backed up (at least) once. Thus, in the following sections the steps to create a full data backup or database snapshot are described as well although they do not have to be executed, if the system was backed up before in its lifetime (also see the section *Backup and Recovery* in the SAP HANA Administration Guide).

Note
To configure SAP HANA system replication for SAP HANA Tenant Database Systems all configuration steps must be done on the System DB.

5.1.1 Using SAP HANA cockpit 2.0

Use the SAP HANA cockpit 2.0 to set up system replication between two identically configured systems. You have registered these systems in the HANA cockpit 2.0 and they are accessible there.

In the following picture the System DBs as well as the tenant DBs of the single tenant primary and secondary system are registered to the SAP HANA cockpit and the credentials are stored:

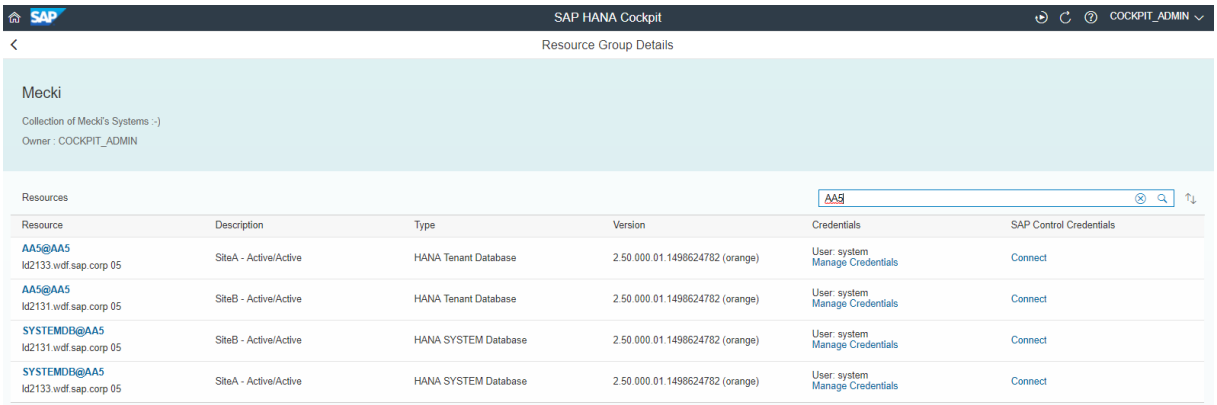

Note

In SAP HANA Tenant Database Systems, the data backups must be created for the System DB as well as for all tenant DBs. However, the SAP HANA system replication setup steps described in the following sections must be executed on the System DB only.

Please look at this example of creating a full data backup of a tenant DB, which works the same way as for the System DB.

On the system overview page of your tenant DB in the primary system click on “Manage database backups”:

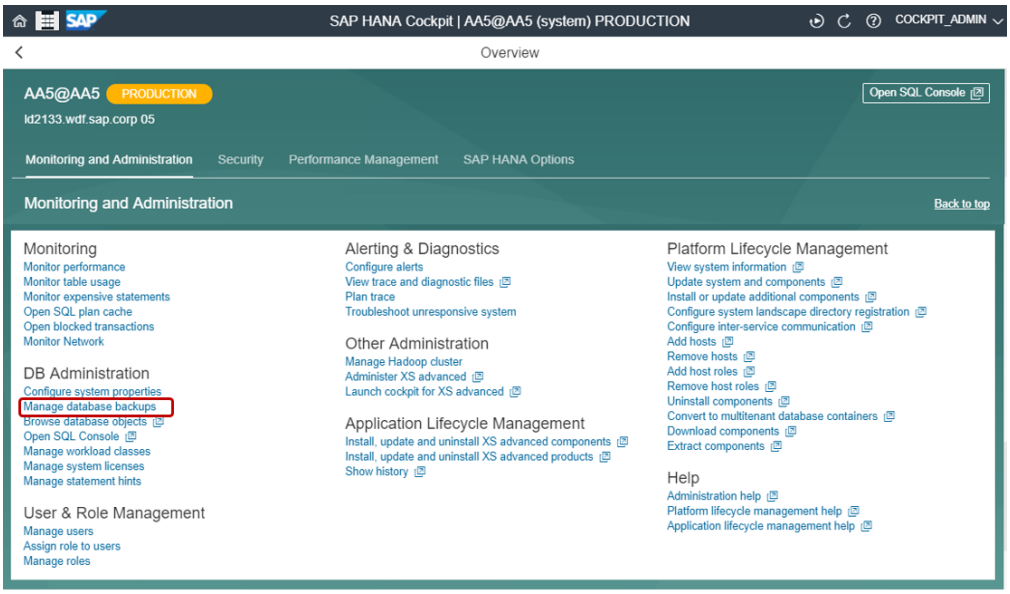

Click on “Create Backup”:

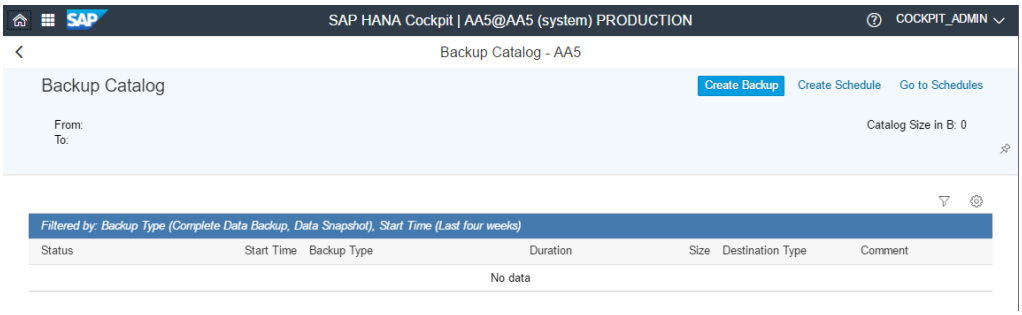

Create a complete data backup and chose a backup location – click "Back Up":

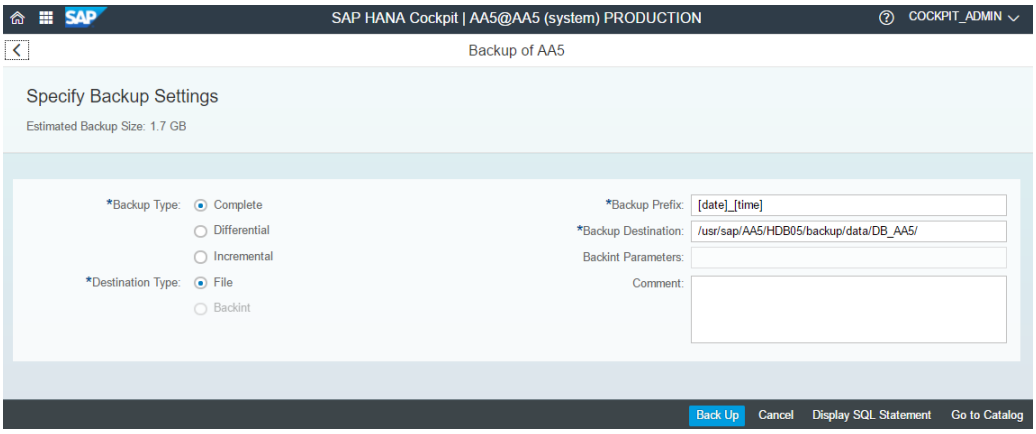

The backup progress running per service is shown:

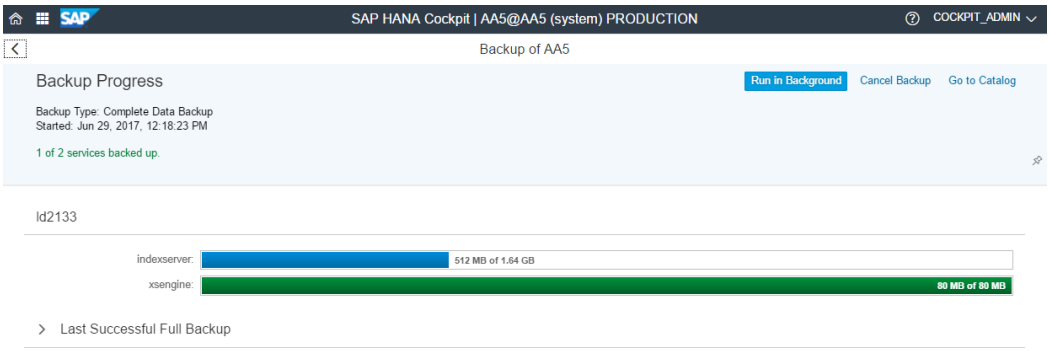

The backup success is then reported with the following screen – showing the backup details:

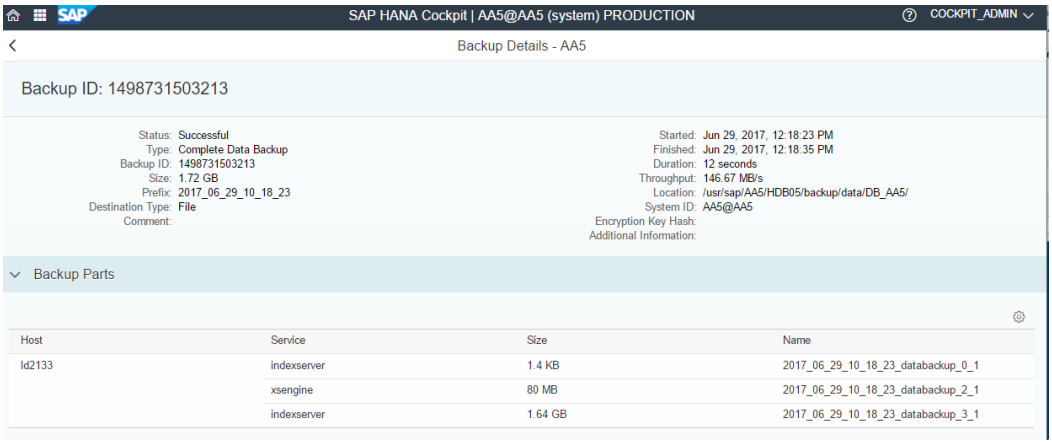

Go back to the previous screen (by clicking on "<") where you get a view of the data backups contained in the backup catalog:

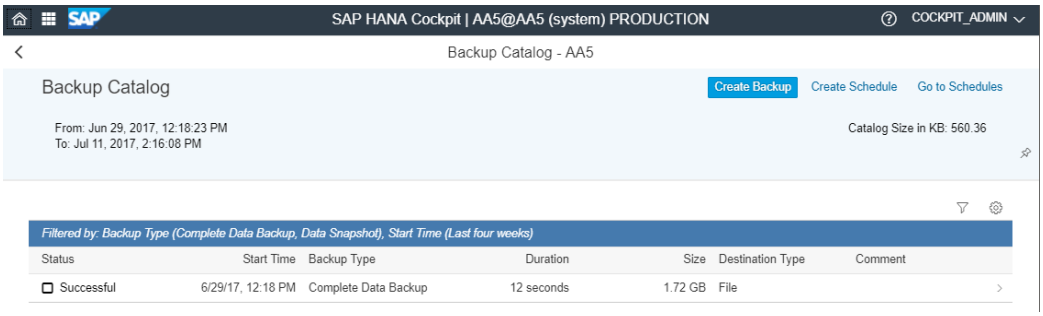

To configure system replication, proceed as follows:
Go to the system overview page of the System DB of your primary system. There you will see that "System replication is not yet enabled for this system" on your freshly installed system. Click on the System Replication tile to configure this primary site:

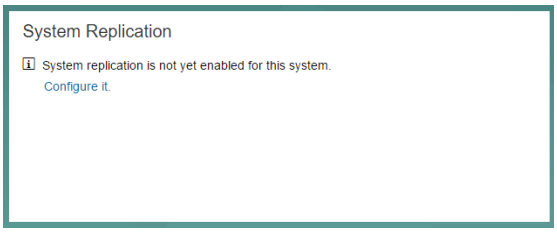

In the System Replication overview you will see "Not Configured". Click on "*Configure System Replication*" to configure system replication in one dialog – without having to switch to the to-be-secondary system:

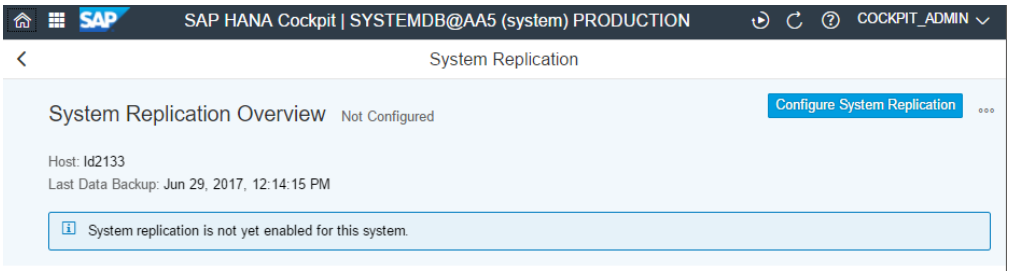

Enter the following information into the corresponding fields – before you click on "*configure*":

- Site Name for the primary site
- Site Name for the secondary site
- Secondary System Host can be selected when clicking on - all hostnames that run systems with the same SID and instance number and are registered to the cockpit will be listed
- *Optional* Checkbox if you did not stop the secondary system prior to the configuration
- Replication Mode
- Operation Mode

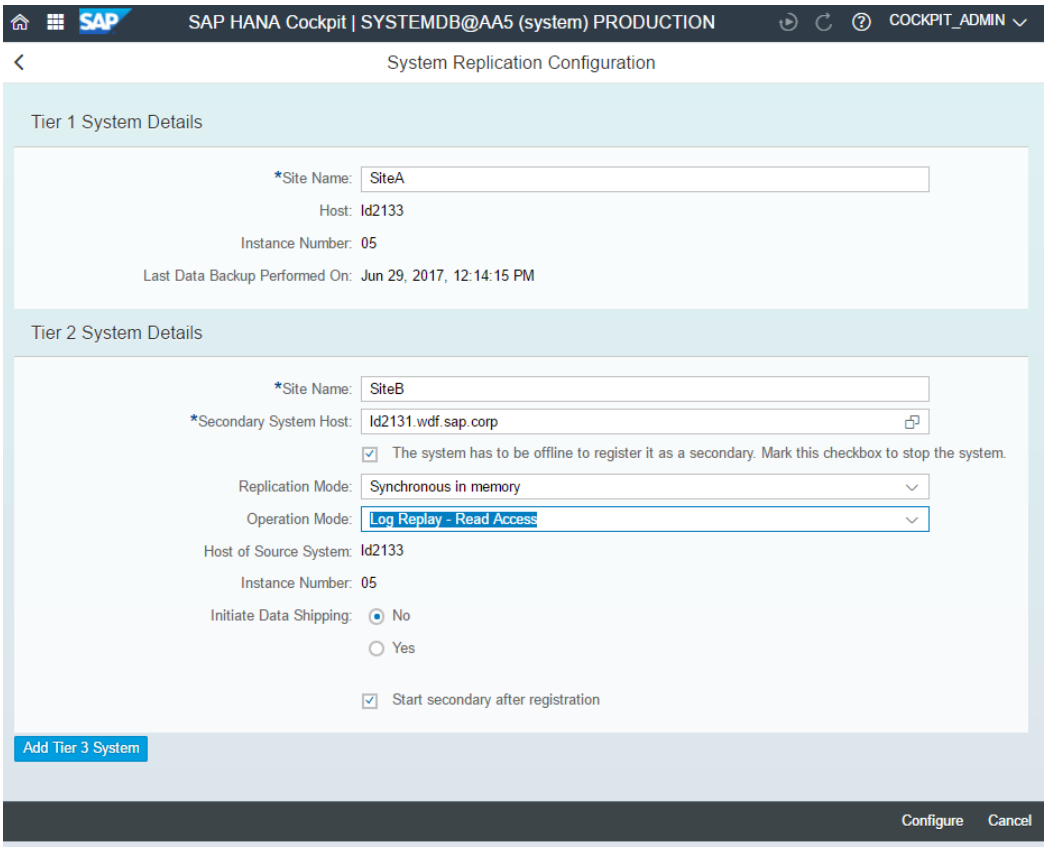

Note

If you want to configure a third tier making it a Multitier System Replication (see details below), you can simply click on "*Add Tier 3 System*" in the above-mentioned screen and follow the instructions.

The overview now shows that this HANA system is configured in a SAP HANA system replication:

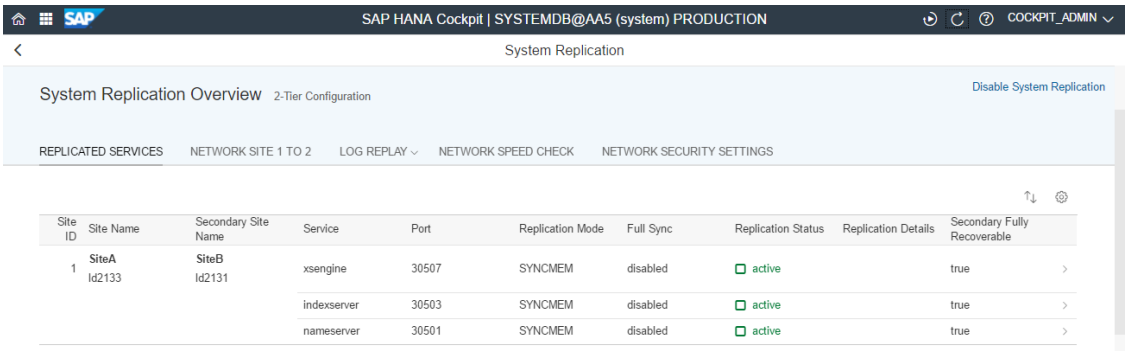

On the HANA cockpit 2.0 system overview pages for the primary system and the secondary system you will see these tiles – in this case telling you that everything worked well and that the replication is active and in sync for all services:

	System DB	Tenant DB
Primary	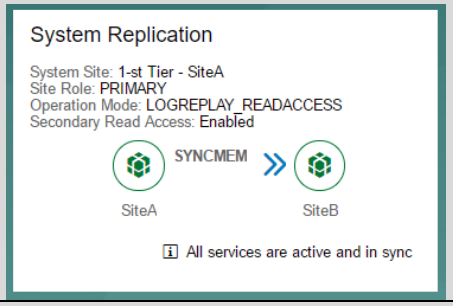	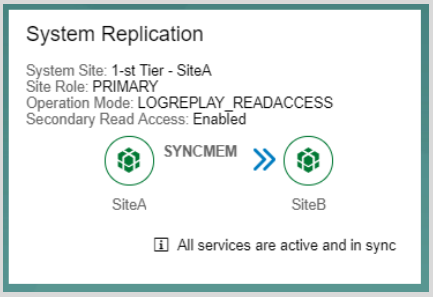
Secondary Active/Active (read-enabled)	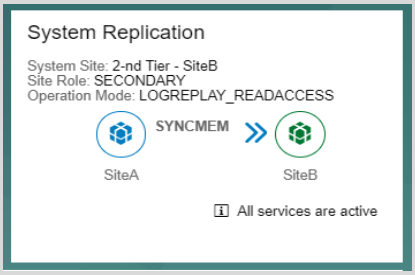	

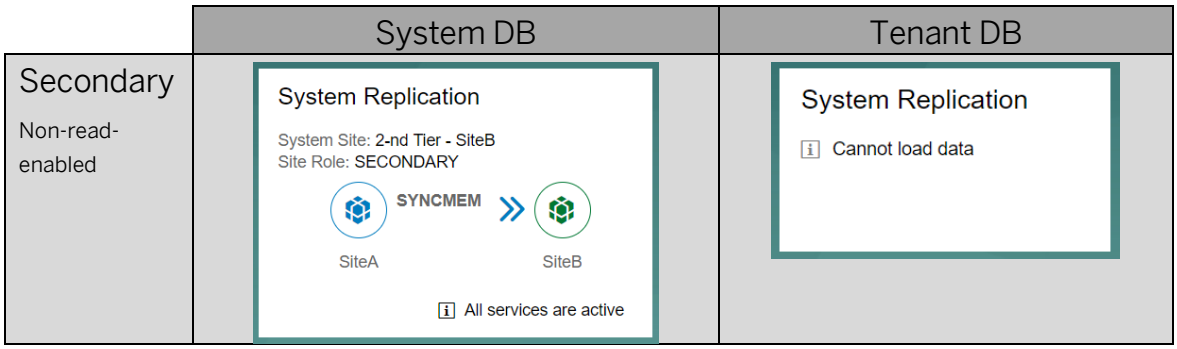

5.1.2 Using SAP HANA studio

Use the SAP HANA studio to set up system replication between two identically configured systems:

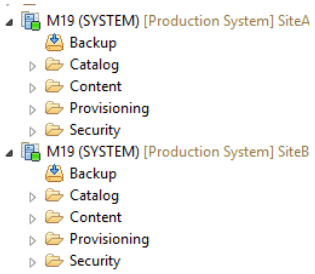

Create a data backup of the primary system[13]. Right mouse-click on dedicated primary → *Backup and Recovery* → *Back Up System*:

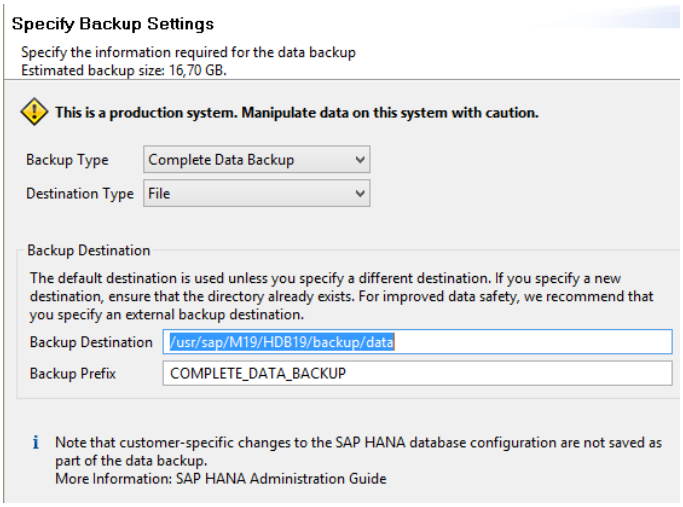

[13] In SAP HANA Tenant Database Systems, the System DB as well as all tenant DBs have to be backed up.

Note

Alternatively, you could create a storage snapshot. Right mouse-click on dedicated primary → *Backup and Recovery* → *Manage Storage Snapshot:*

Prepare:

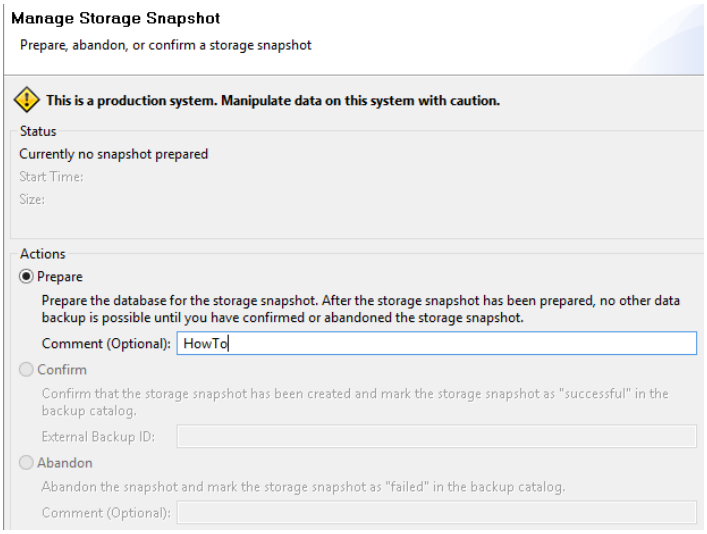

Confirm:

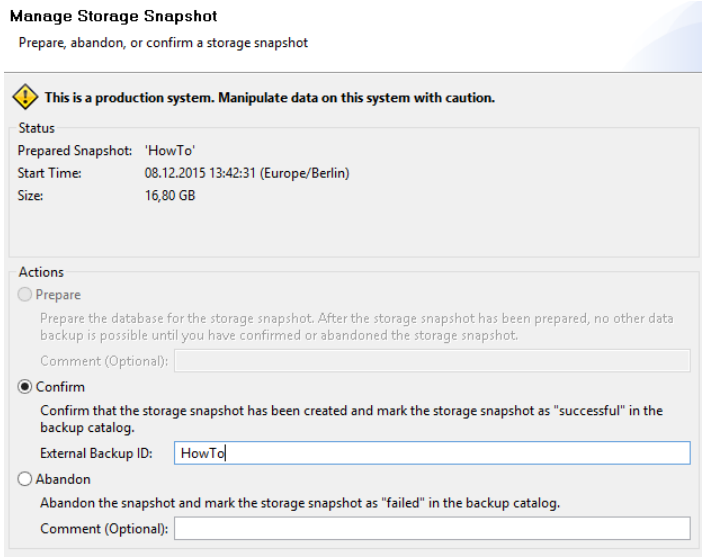

To configure system replication, proceed as follows:

Right mouse-click on Primary System → *Configuration and Monitoring* → *Configure System Replication*[14] … . Check the radio button to enable system replication:

[14] Only the actions that are possible in the current system state will be offered to you. In this case only the "enable" is possible.

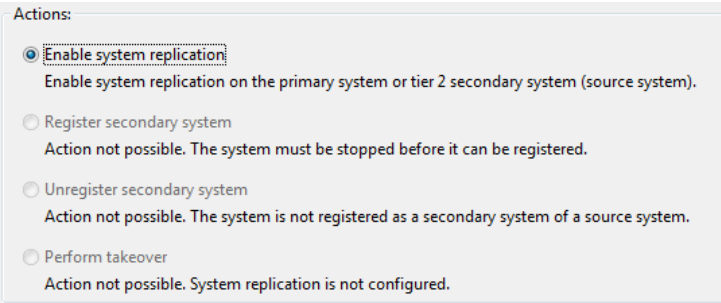

Give the primary a logical site name, for example SITEA:

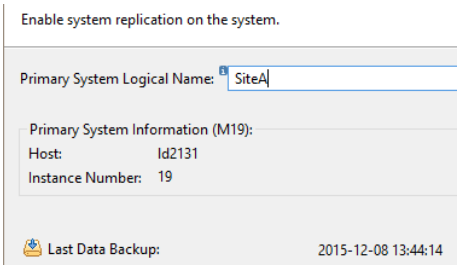

Stop the secondary system with right mouse-click on Secondary System → *Configuration and Monitoring* → *Stop System*

> **Caution**
>
> If you are running with HANA 2.0 you will need to copy the systemPKI SSFS key and data file from the primary to the secondary before registering the secondary site. The corresponding files can be found on the primary:
>
> ```
> /usr/sap/<SID>/SYS/global/security/rsecssfs/data/SSFS_<SID>.DAT
> /usr/sap/<SID>/SYS/global/security/rsecssfs/key/SSFS_<SID>.KEY
> ```

Register[15] the secondary: Right mouse-click on Secondary System → *Configuration and Monitoring* → *Configure System Replication* … Check radio button "Register secondary system":

[15] If problems occur indicating an error like "unable to contact primary site" and "bad certificate", please follow the procedure described in the SAP HANA Security guide – section: *Secure Internal Communication between Sites in System Replication Scenarios.*

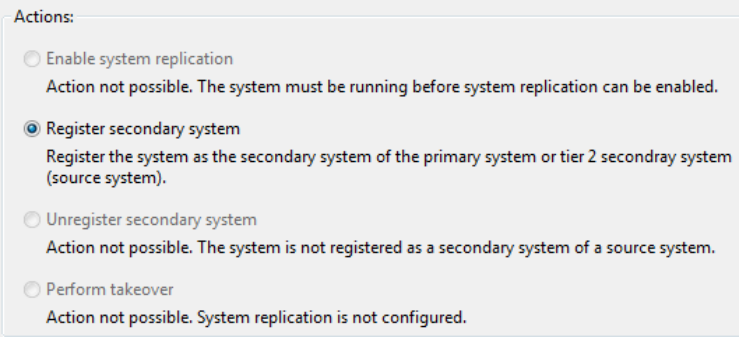

Type a logical name (= site name) for the secondary, choose a replication mode, an operation mode[16], and enter the primary site's host name:

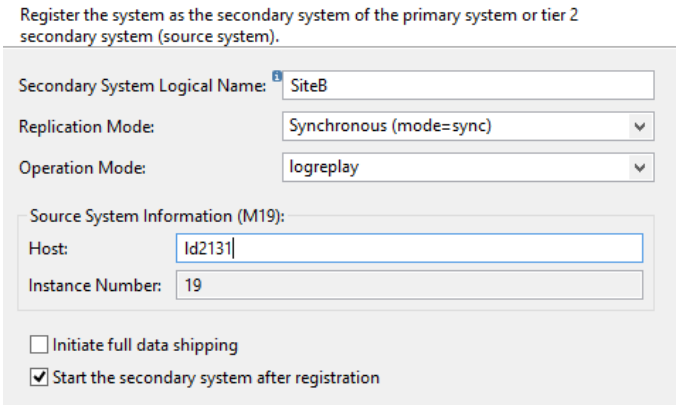

Once the secondary system is automatically started, the replication process will also start automatically.

5.1.3 Using command line tool `hdbnsutil`

Alternatively use the command line tool `hdbnsutil` as `<sid>adm` on OS level.

Create a data backup of the primary system:

```
hdbsql BACKUP DATA USING FILE ('<path><prefix>')
```

> **Note**
> In *SAP HANA Tenant Database Systems* all databases must be backed up using the "`hdbsql`" tool via the database name option:
> - for the system DB "`-d SystemDB`" resp.
> - for the tenant DBs "`-d <tenantDBName>`".

[16] **Caution**: Since SAP HANA studio will not be supported anymore and will be replaced by SAP HANA cockpit, the operation mode "logreplay_readaccess" to configure Active/Active (read enabled) system replication is not available here.

Enable the primary system and give the primary a logical name, for example SITEA:

```
hdbnsutil -sr_enable --name=SITEA
```

> **Caution**
> If you are running with HANA 2.0 you will need to copy the systemPKI SSFS key and data file from the primary to the secondary before registering the secondary site – if you not already have done that. The corresponding files can be found on the primary:
>
> ```
> /usr/sap/<SID>/SYS/global/security/rsecssfs/data/SSFS_<SID>.DAT
> /usr/sap/<SID>/SYS/global/security/rsecssfs/key/SSFS_<SID>.KEY
> ```

Stop the secondary system:

```
sapcontrol -nr <instance_number> -function StopSystem HDB
```

Register the secondary system, provide a logical name (for example SITEB), and choose a replication mode and the operation mode:

```
hdbnsutil -sr_register
--remoteHost=<primary hostname>
--remoteInstance=<instance number>
--replicationMode=<sync|syncmem|async>
--
operationMode=<delta_datashipping|logreplay|logreplay_readaccess>
--name=SITEB
```

Start the secondary system to start replication:

```
sapcontrol -nr <instance_number> -function StartSystem HDB
```

Once the secondary system is started, the replication process will start automatically.

5.1.4 Creating a tenant DB in a running system replication

After a new tenant DB was created in an SAP HANA Tenant Database System running with SAP HANA system replication, a backup of this new tenant DB is necessary to automatically start its replication. Otherwise the replication for this tenant DB will not start.

5.2 Disable system replication

To remove a system replication configuration, because you want to run the two systems separately or simply do not need this high availability and disaster recovery feature anymore you must unregister the secondary and disable the primary system accordingly.

> Note
>
> Please refer to SAP Note 1945676 - Correct usage of hdbnsutil - sr_unregister to understand the only scenarios where unregistering a secondary system makes sense. You do not need to unregister a secondary, if you will re-register it right afterwards again; this is not necessary!

5.2.1 Using SAP HANA cockpit 2.0

On the stopped secondary system (see above) click on the System Replication tile in the system overview page:

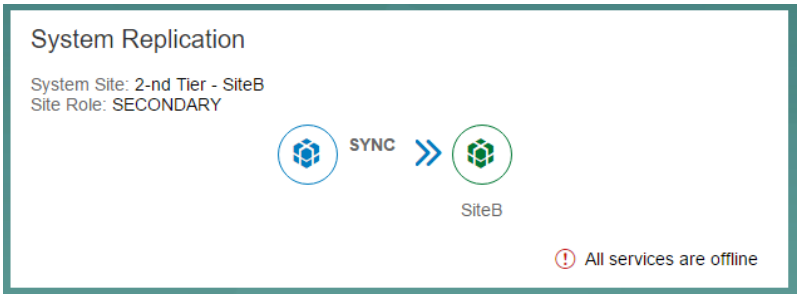

In the System Replication Overview – now showing that "All services are offline" – click on "Unregister Secondary System":

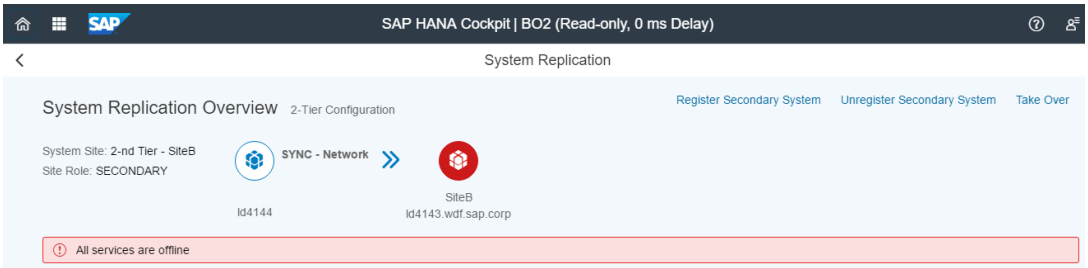

Depending whether your system should be online or offline after the unregister command, check or uncheck the checkbox "Start system after unregistration" and hit "OK":

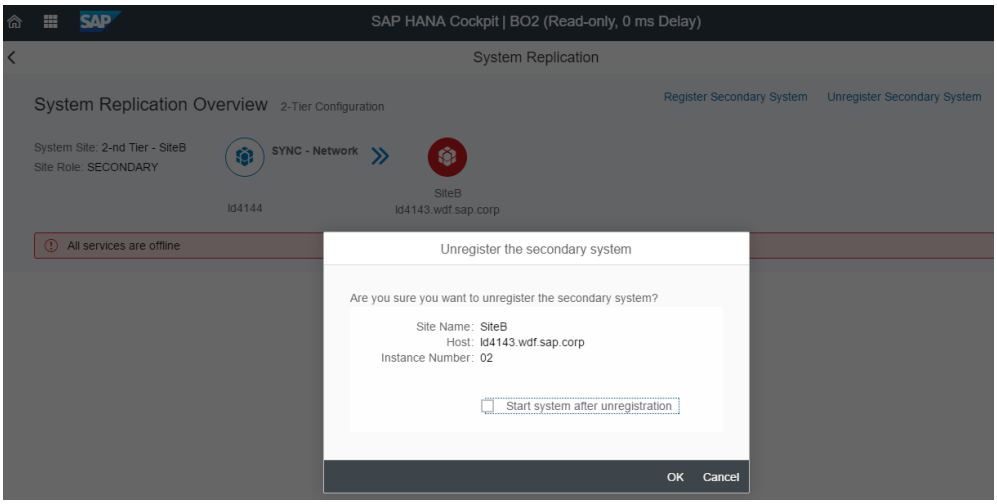

On primary system: Disable system replication on the primary system by clicking on the system replication tile:

On the System Replication Overview click on “Disable System Replication”:

Verify that you want to disable system replication by clicking “OK”:

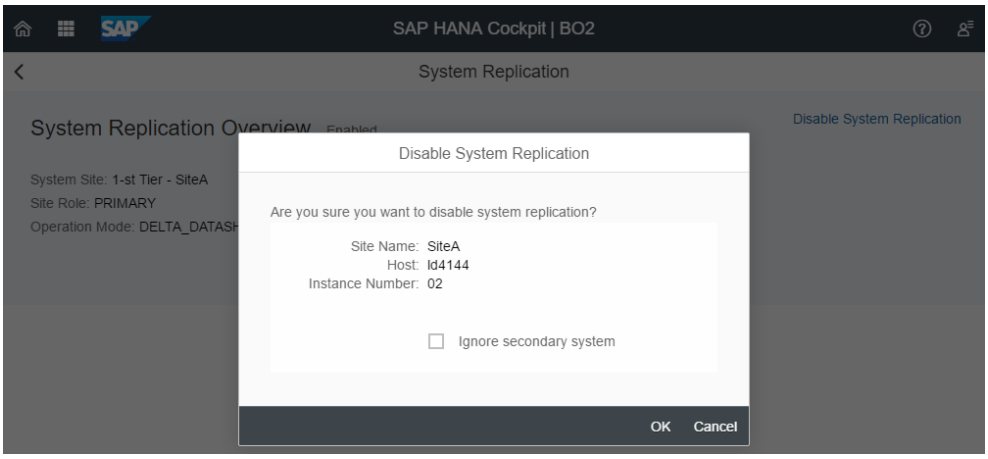

Afterwards the former primary system reports, that System Replication is not configured:

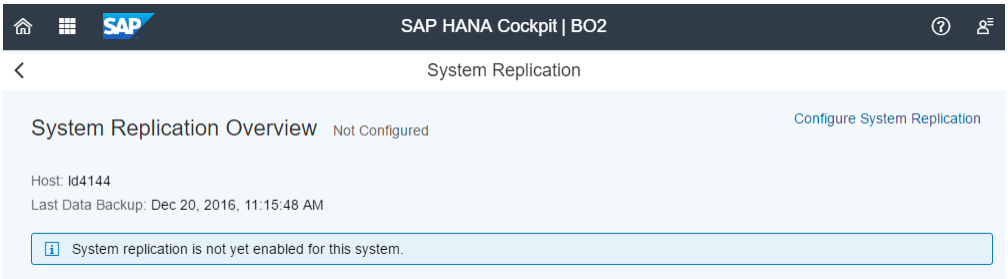

5.2.2 Using SAP HANA studio

Stop the secondary system with right mouse-click on Secondary System → *Configuration and Monitoring* → *Stop System*

On secondary system: Unregister system replication for the secondary system with right mouse-click on Secondary System → *Configuration and Monitoring* → *Configure System Replication* … :

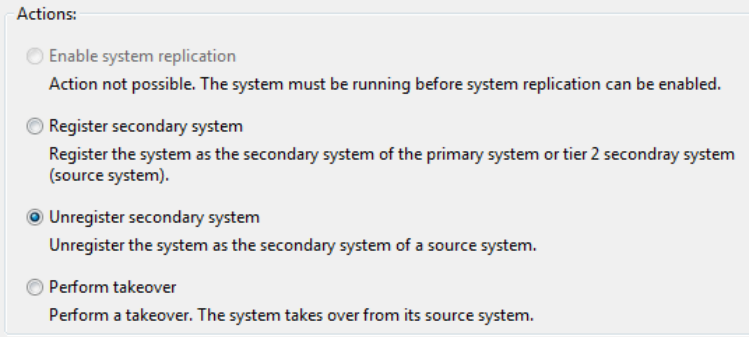

On primary system: Disable system replication on the primary system with right mouse-click on Primary System → *Configuration and Monitoring* → *Configure System Replication* ...:

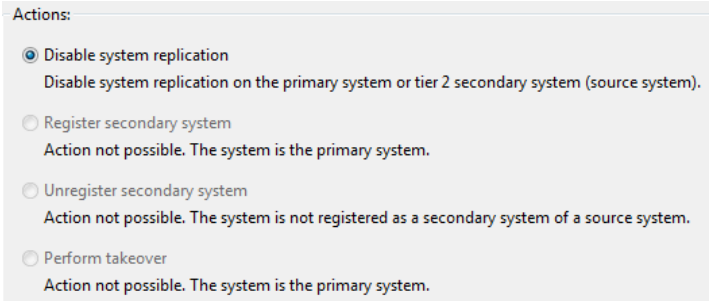

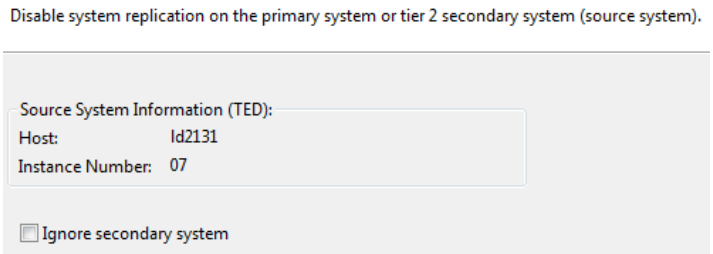

5.2.3 Using command line tool `hdbnsutil`

Stop the secondary system:

```
sapcontrol -nr <instance_number> -function StopSystem HDB
```

On secondary system unregister the secondary system:

```
hdbnsutil -sr_unregister
```

If you want to use this secondary as a normal SAP HANA installation from now on, you have to start it to complete the unregistration. On the secondary execute:

```
sapcontrol -nr <instance_number> -function StartSystem HDB
```

On primary system disable system replication:

```
hdbnsutil -sr_disable
```

5.3 Setting up Multitier System Replication

As of HANA 1.0 SPS07 with the Multitier System Replication, a synchronous system replication can be used as the source for asynchronous replication in a chained setup of primary site, tier-2 secondary site and tier-3 secondary site.

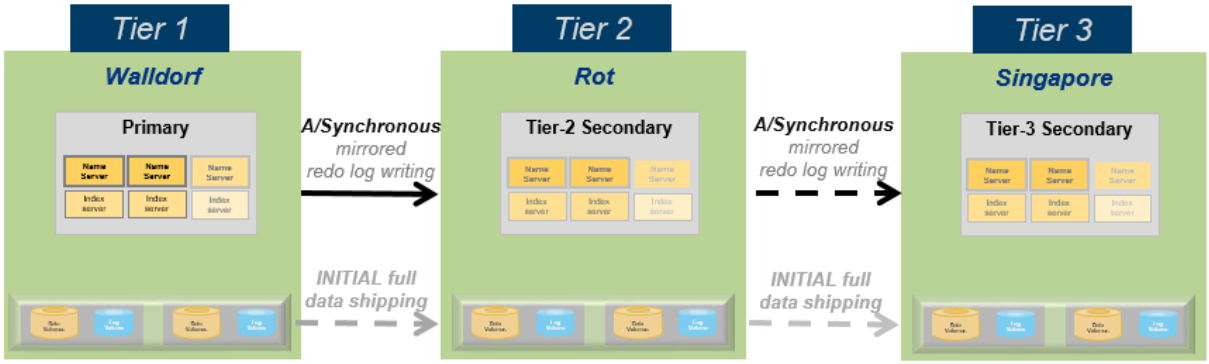

Overview of Multitier System Replication

Until HANA 1.0 SPS11 the primary system had to *synchronously* replicate to the tier-2 secondary system and the tier-2 secondary had to *asynchronously*[17] replicate to the tier-3 secondary system.

As of HANA 1.0 SPS11 more combinations of replication modes (`SYNC`, `SYNCMEM`, and `ASYNC`) in a multitier landscape are possible. For details please have a look at SAP Note 2303243 where the supported combinations are listed.

Tier 1 -> Tier 2	Tier 2 -> Tier 3	Supported since
SYNC	SYNC	SPS12
SYNC	SYNCMEM	SPS12
SYNC	ASYNC	SPS07
SYNCMEM	SYNC	SPS11
SYNCMEM	SYNCMEM	SPS12
SYNCMEM	ASYNC	SPS07
ASYNC	ASYNC	SPS11

Supported replication mode combinations in Multitier System Replication

Caution

In *Multitier System Replication* the operation mode must be the same for all sites. However, if the Active/Active (read enabled) operation mode `logreplay_readaccess` is used between tier-1 and tier-2, as an exception from this rule only operation mode `logreplay` can be used between tier-2 and tier-3!

Given a running 2-tier system replication (as described above) the following steps are to be executed to add the tier-3 secondary; this third site must fulfill the same prerequisites as described in 4.1.

[17] Currently only asynchronous replication is supported for the connection between the tier 2 and the tier 3 secondary site.

5.3.1 Using SAP HANA cockpit 2.0

Use the SAP HANA cockpit to add a tier-3 secondary to a system replication landscape.

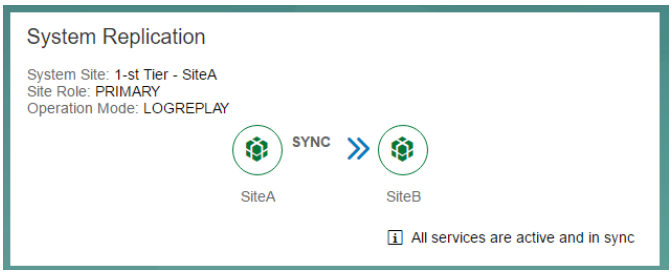

All three systems are registered in the SAP HANA cockpit.

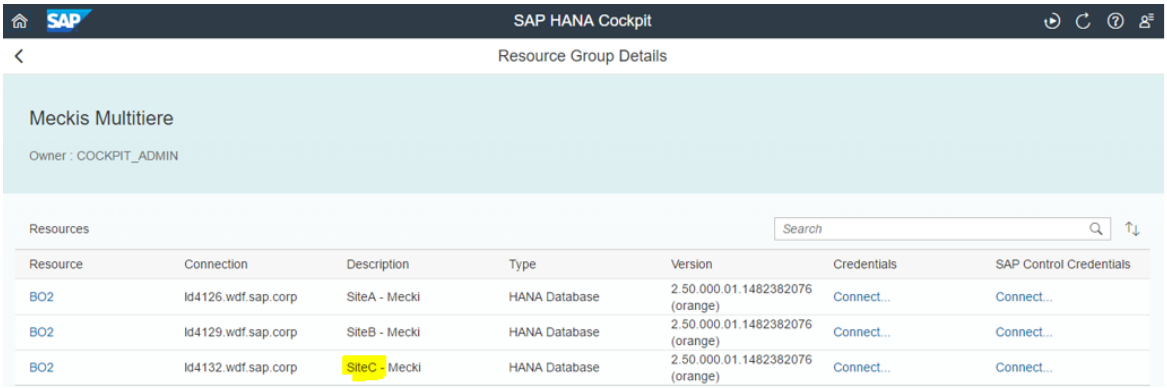

Click on the HANA database that is currently configures as tier-2 secondary of the existing 2-tier system replication landscape.

On the system overview page click on the “System Replication” tile to access the system replication application:

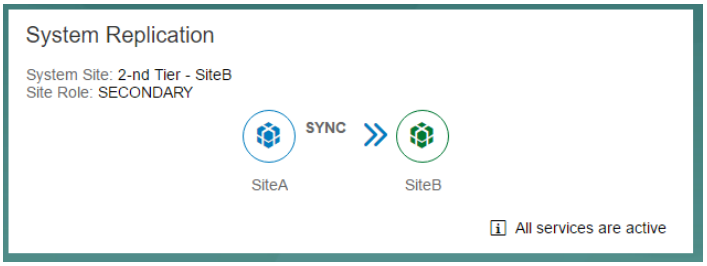

Enable this tier-2 secondary to function as the source for the added tier-3 secondary by clicking on “Configure System Replication”:

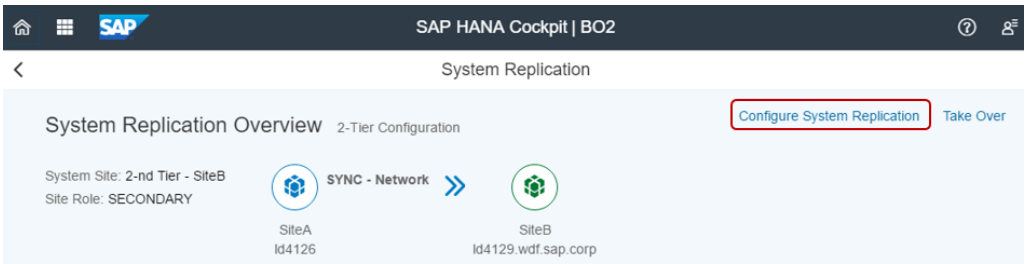

The Site Name is already known from topology. Simply click on “Configure”:

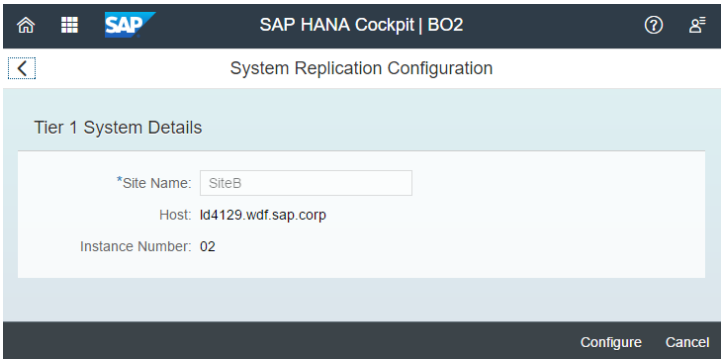

To stop the to-be tier-3 secondary from the system overview page of this HANA database go via the “Overall Database Status” tile to the “Manage Services” application:

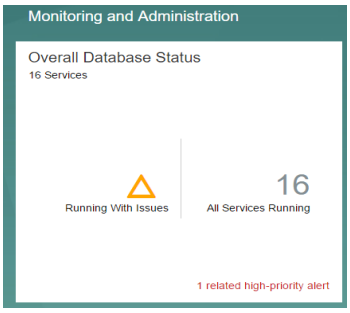

Stop the tier-3 secondary-to-be system:

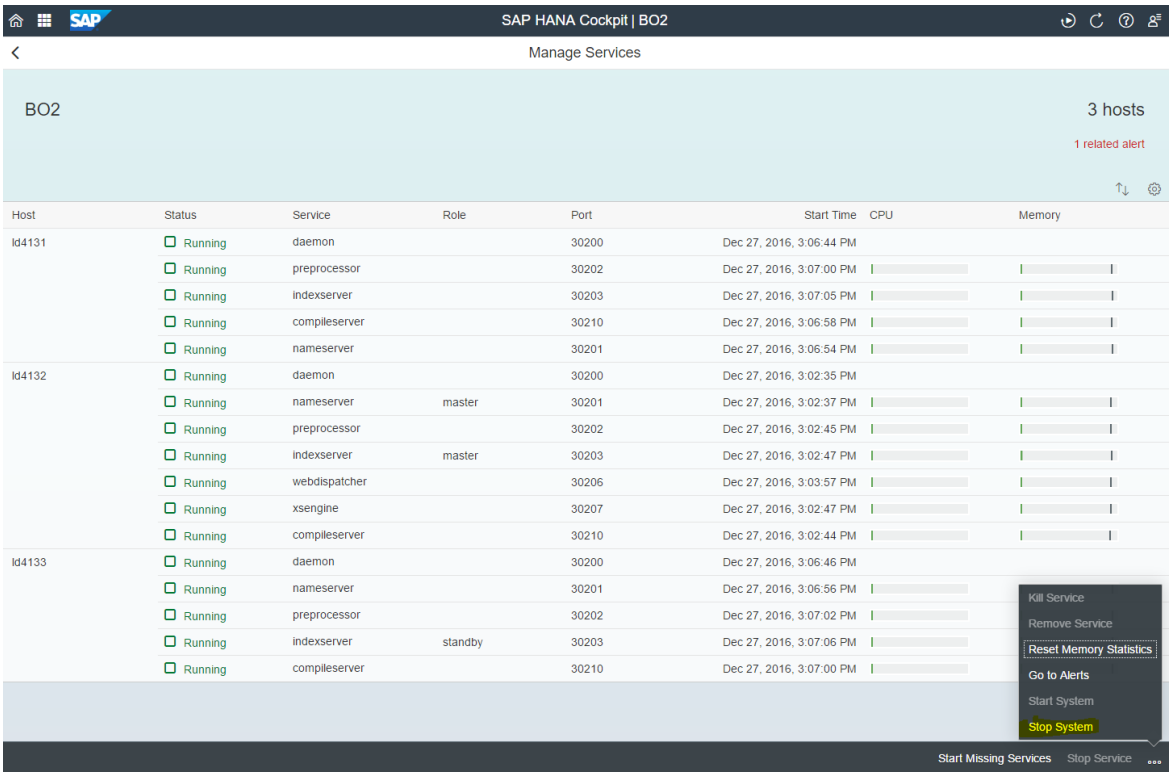

Register the tier-3 secondary by clicking on the "System Replication" tile and then hitting "*Register Secondary System*":

Type a site name for the tier-3 secondary, choose replication mode (in this example `async`), choose operation mode (in this example: `logreplay`) and enter the tier-2 secondary system's master host name:

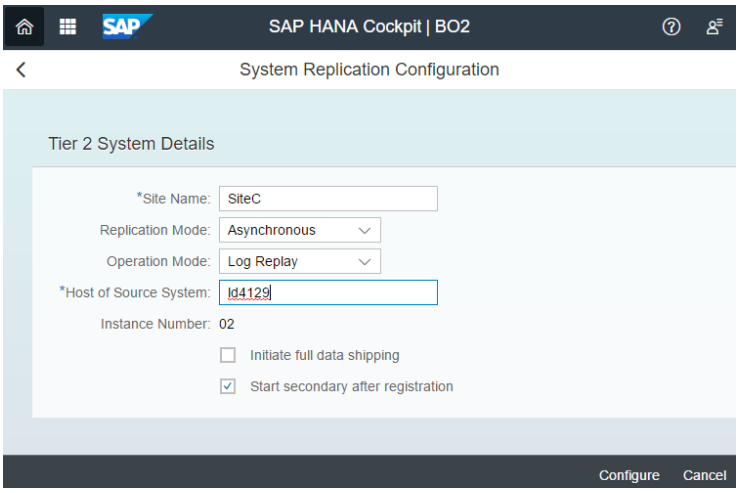

Once the secondary system is automatically started the replication process to the tier-3 secondary will also start automatically.

5.3.2 Using SAP HANA studio

Use the SAP HANA studio to add a tier-3 secondary to a system replication landscape:

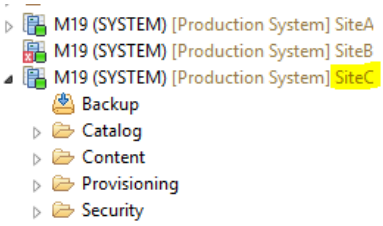

Right mouse-click on the tier-2 secondary → *Configuration and Monitoring* → *Configure System Replication* … .

Check the radio button to enable system replication – site name is already known from topology:

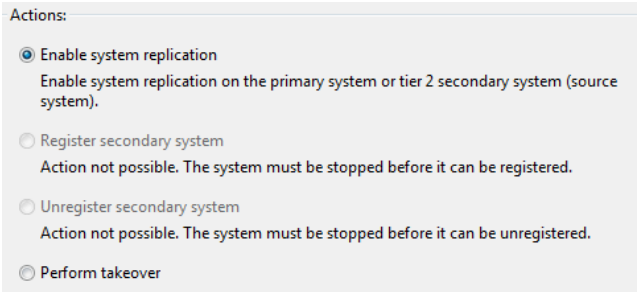

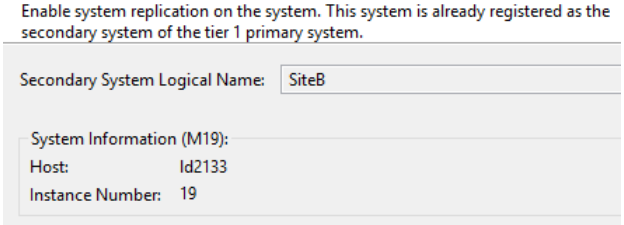

Stop the tier-3 secondary system with right mouse-click → *Configuration and Monitoring* → *Stop System*:

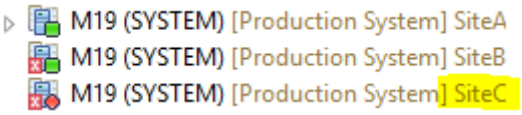

Register the tier-3 secondary: Right mouse-click on tier-3 secondary system → *Configuration and Monitoring* → *Configure System Replication* ...

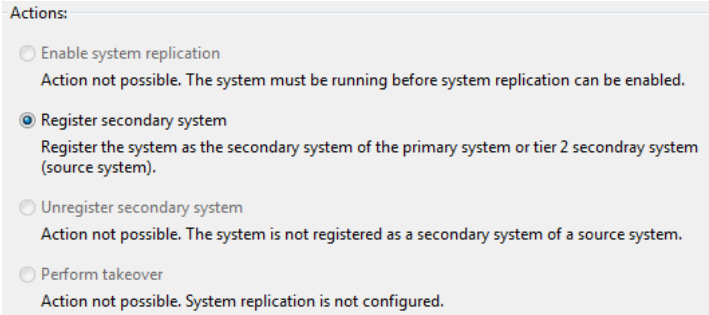

Type a logical name for the tier-3 secondary, choose replication mode ASYNC, the same operation mode as for the primary and tier-2 secondary (in this example: `logreplay`) and enter the tier-2 secondary site's host name:

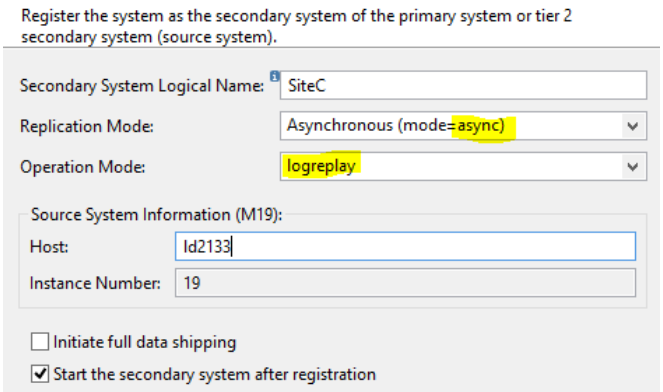

Once the secondary system is automatically started the replication process to the tier-3 secondary will also start automatically.

5.3.3 Using command line tool `hdbnsutil`

1. Tier-2 secondary: `hdbnsutil -sr_enable`
2. Tier-3 secondary: `sapcontrol -nr <instance_number> -function StopSystem HDB`
3. Tier-3 secondary:
   ```
   hdbnsutil -sr_register --remoteHost=<tier_2_host>
   --remoteInstance=<instance number>
   --replicationMode=<sync|syncmem|async>
   --operationMode=<delta_datashipping|logreplay>
   --name=<siteName>
   ```
4. Tier-3 secondary: `sapcontrol -nr <instance_number> -function StartSystem HDB`

5.4 Enabling full sync replication

As of HANA 1.0 SPS08 to reach a true RPO=0 for synchronous system replication, the *full sync* option can be enabled for **`SYNC`** replication mode (i.e. not for `SYNCMEM`). With the activated *full sync* option, transaction processing on the primary system blocks when the secondary is currently not connected and newly created log buffers cannot be shipped to the secondary system. This behavior ensures that no transaction can be committed on the primary without shipping the log buffers to the secondary site.

The full sync option can be switched on and off using the command

```
hdbnsutil -sr_fullsync --enable|--disable
```

This changes the setting of the `global.ini` file accordingly:

```
global.ini/[system_replication]/enable_full_sync
```

Recommendation

Proceed in the following order, when configuring the *full sync* option:

1. Configure your system replication with the replication mode **`sync`** , which is the prerequisite for enabling the *full sync* option
2. Check that the system replication status is "*active and in sync*" for all services.
3. Enable the *full sync* option with `hdbnsutil -sr_fullsync --enable`

In the `M_SERVICE_REPLICATION` system view the setting of the *full sync* option can be viewed in the column "`FULL_SYNC`". It can have the following values:

- `DISABLED`: full sync is not configured at all `global.ini/[system_replication]/enable_full_sync = false`
- `ENABLED`: full sync is configured, but it is not yet active, so transactions do not block in this state. To become active the secondary must connect and `REPLICATION_STATUS` has to be `ACTIVE`.
- `ACTIVE`: full sync mode is configured and active. If a connection of a connected secondary is getting closed, transactions on the primary side will block in this state.[18]

If full sync is enabled when an active secondary is currently connected, the `FULL_SYNC` will be immediately set to ACTIVE.

Note
Resolving a blocking situation of the primary caused by the enabled full sync option must be done with the `hdbnsutil` command, since also a configuration changing command could block in this state.

Hint
The Cluster Manager that is used to operate SAP HANA system replication landscapes could provide a timeout after which the blocking situation is resolved automatically using the `hdbnsutil` command and deactivating this full-sync option. However, after the reason for the blocking situation disappeared, it also must be taken care to re-activate the *full sync* option again (manually or automatically with the help of the cluster manager tool).

5.5 Changing the replication mode

The replication mode can be changed without having to go through a full data shipping from the primary to the secondary afterwards.

Command on online or offline Secondary:

```
hdbnsutil -sr_changemode --mode=sync|syncmem|async
```

If the mode was changed correctly can be checked in the `M_SERVICE_REPLICATION` view or with this command:

[18] The *full sync* option might not become active immediately, but waits until the initial data transfer for the different services running in the HANA instance is completed. This prevents the system from blocking transactions right after setting the parameter to true. This means, when *full sync* option is first enabled by the administrator, it is not immediately active; only after the status of the system replication for the first service going into active also *full sync* is internally activated for this service.

```
hdbnsutil -sr_state --sapcontrol=1
```

5.6 Changing the operation mode

The operation mode can only be changed by stopping and re-registering the secondary system with the desired operation mode.

> **Note**
> As already mentioned above in this document: It is NOT necessary to unregister the secondary in this situation; just execute the "–sr_register" command, which will overwrite the previous register configuration.
> Please refer to SAP Note 1945676 - Correct usage of hdbnsutil -sr_unregister to understand the only scenarios where unregistering a secondary system makes sense.

On the command line of the stopped secondary system as `<sid>adm` you need to do the following to newly register the secondary system:

```
hdbnsutil -sr_register
--remoteHost=<primary hostname>
--remoteInstance=<instance number>
--replicationMode=<sync|syncmem|async>
--operationMode=<delta_datashipping|logreplay|logreplay_readaccess>
--name=<siteName>
```

Start the secondary system to start replication with the new operation mode:

```
sapcontrol –nr <instance_number> -function StartSystem HDB
```

Once the secondary system is started, the replication process will start automatically.

> **Note**
> When changing the operation mode from `delta_datashipping` to `logreplay` *or* `logreplay_readaccess` no full data shipping is necessary during re-registration.
> When changing from `logreplay` *or* `logreplay_readaccess` operation modes back to `delta_datashipping`, however, a full data shipping is necessary.

6. Takeover

6.1 Perform Takeover

The following steps are performed:

- Trigger a takeover to the secondary system in the event of a disaster.
- Register the former primary system as new secondary when it becomes available again.

The takeover can be triggered from SAP HANA cockpit, from the command line or from the SAP HANA Studio.

6.1.1 Using SAP HANA cockpit 2.0

On secondary system overview page click on the "System Replication" tile:

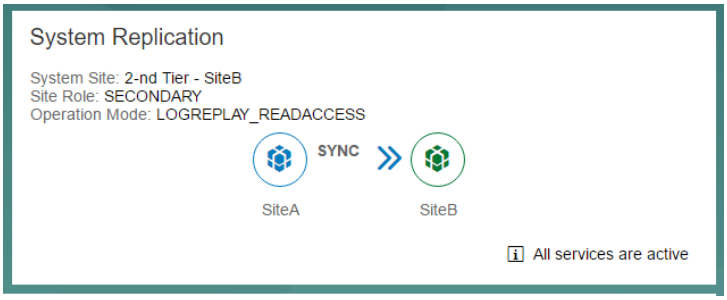

Perform a takeover by selecting "Take Over":

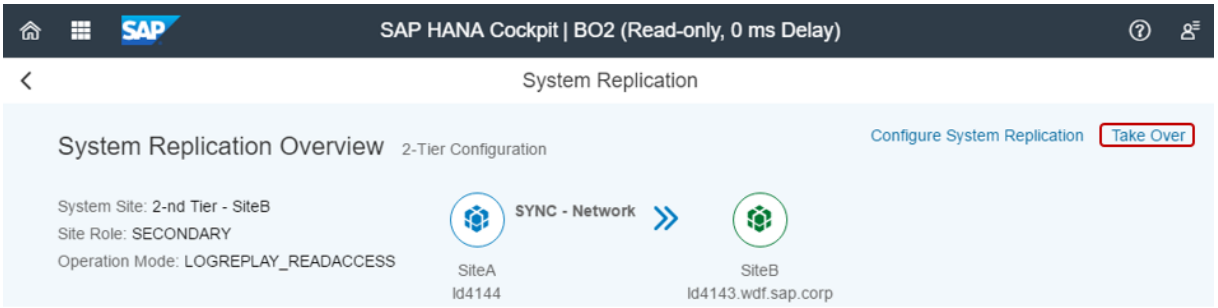

Afterwards this system is running as the active HANA database not functioning as a secondary anymore.

When the former primary is available again it can be registered as secondary.

Stop original (former) primary system.

On original (former) primary system: Register system as new secondary by clicking the System Replication tile on this new to-be secondary and then clicking on "Register Secondary System":

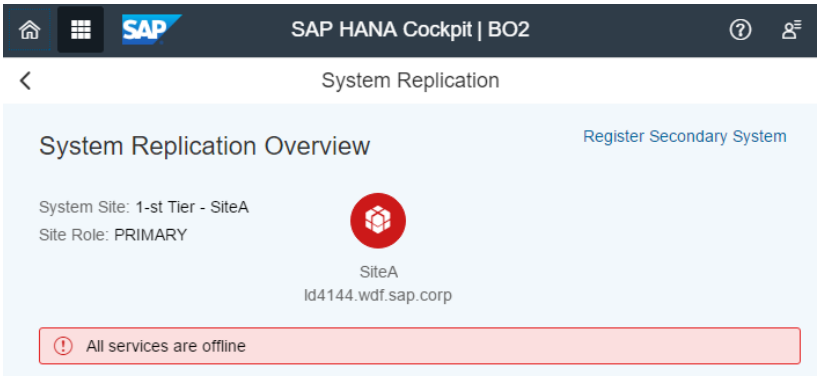

Add the site name of the original (former) primary's site, the replication and operation mode and the master host name of the former secondary (now primary) and click on "Configure":

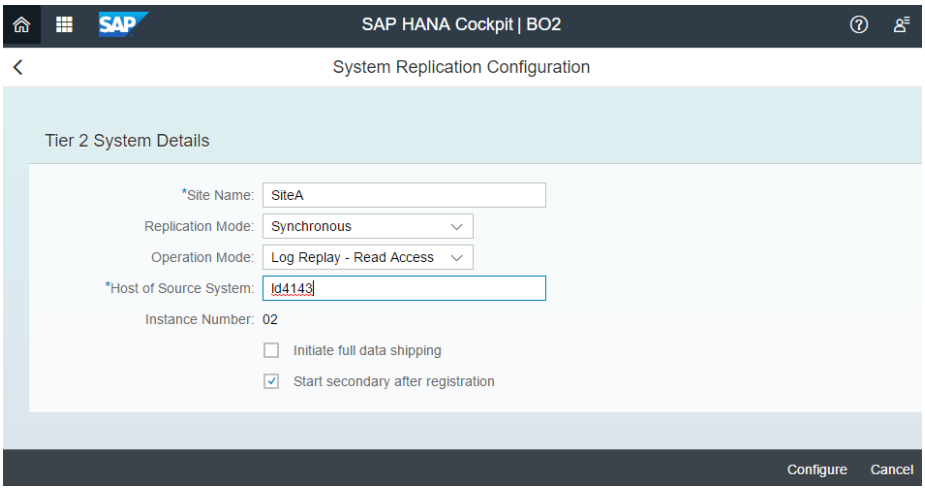

Afterwards the roles are switched: The former primary now runs as secondary to the new primary, which is the former secondary after takeover.

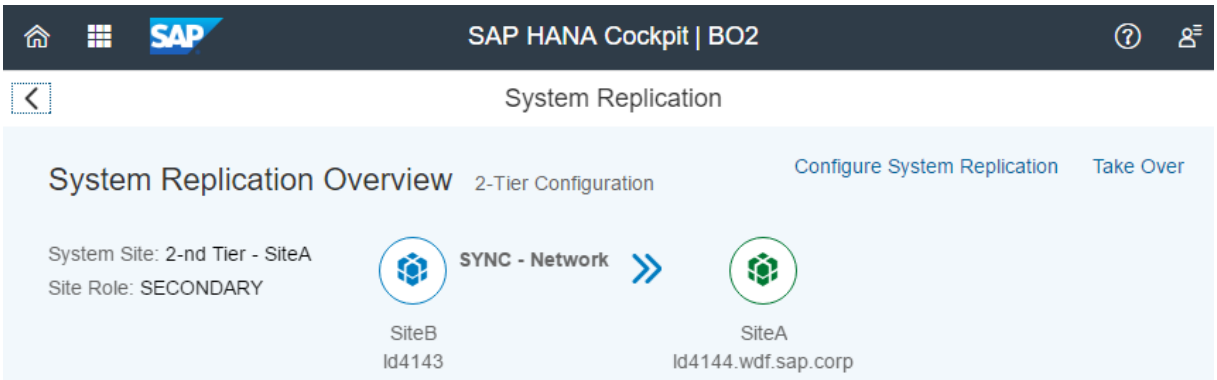

6.1.2 Using SAP HANA studio

On secondary system: Perform a takeover with right mouse-click on Secondary System → *Configuration and Monitoring* → *Configure System Replication* ...:

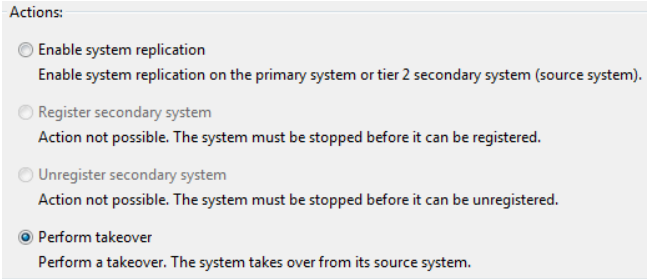

When the former primary is available again it can be registered as secondary.

Stop original (former) primary system.

On original (former) primary system: Register system as secondary with right mouse-click on former Primary System → *Configuration and Monitoring* → *Configure System Replication* ... :

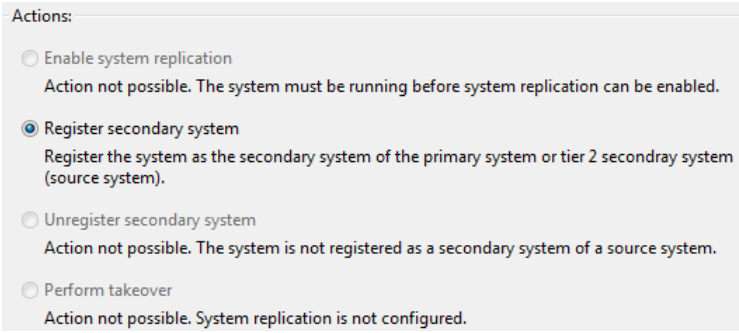

You will be informed that this system used to be the primary system before.

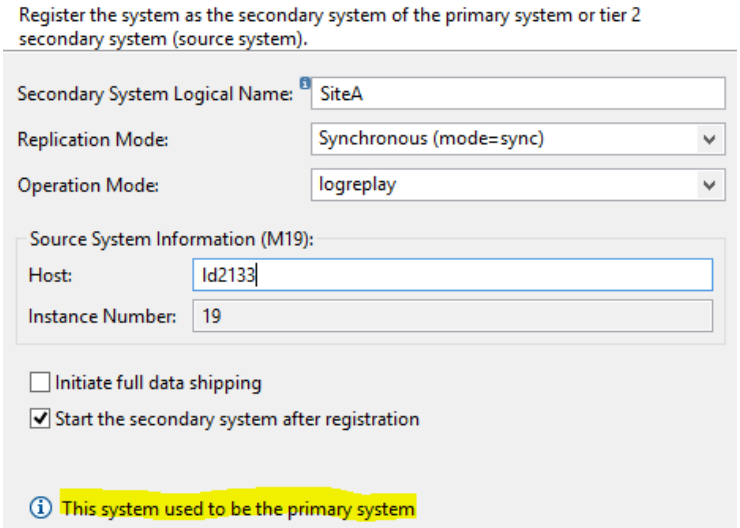

6.1.3 Using command line tool `hdbnsutil`

1. Perform a takeover on the secondary:

```
hdbnsutil -sr_takeover
```

2. When the former primary is available again it can be registered as the new secondary:

```
hdbnsutil -sr_register --remoteHost=<new primary hostname>
--remoteInstance=<instance number>
--replicationMode=<sync/syncmem/async>
--
operationMode=<delta_datashipping|logreplay|logreplay_readaccess>
--name=<siteName>
```

6.2 Client connection recovery

To perform the takeover only on the SAP HANA system will, in most cases, not be enough. Somehow the client or application server needs to be able to continuously reach the SAP HANA system, no matter which site is currently the primary resp. the active system after takeover.

There are several methods:

- *IP redirection*: A virtual IP address is assigned to the virtual host name. In case of a takeover, the virtual IP will unbind from the network adapter of the primary system and bind to the adapter on the secondary system.

- *DNS redirection*: In this scenario the IP for the host name in the DNS will be changed from the address of the primary to the address of the secondary system.

Both methods have their advantages but it will be mostly decided by the IT policies and existing configuration. If there are no existing constraints, IP redirection has the clear benefit of being faster to process in a script rather than synchronizing changes of DNS entries over a global network.

Since HANA 1.0 SPS09 SAP HANA offers the so-called HA/DR providers which can inform external entities about activities inside SAP HANA scale-out (such as Host Auto-Failover) and SAP HANA system replication setups. In a Python script actions can be defined which should be executed before or after certain HANA activities (like startup, shutdown, failover, takeover, connection changed, service state changed, ...). One example for these so-called hooks is moving virtual IP addresses after takeover in SAP HANA system replication.

Additionally, external cluster management software can be used to perform the client reconnect after takeover. For further details on the section "Implementing HA/DR Providers" in the SAP HANA Administration Guide.

Please refer to the section "*Virtual IP address handling*" in the "*Active/Active (read enabled)*" chapter of this document for the specifics of virtual IP handling in such a system replication configuration.

6.2.1 Virtual IP addresses in Active/Active (read enabled) system replication

In an Active/Active (read enabled) configuration a second virtual IP address for the read access on the secondary system is needed – besides the virtual IP address handling the access to the primary system.

During takeover you can keep the virtual IP address of the secondary system, which will be used for read access until a reconnect occurs. The former virtual IP address of the primary is also rebound to access the former secondary which is the now active system. In this situation two virtual IP addresses are available for accessing the former secondary system after takeover.

> **Note**
> Already prior to a takeover you must take care that the secondary read enabled systems – after takeover – can handle the workload of the former primary and the read enabled secondary!

During failback, the system replication sites switch roles and the virtual IP addresses switch their location binding accordingly. For further information please check the corresponding section in the Active/Active chapter of this document.

7. Resync optimization

Whenever the primary and the secondary sites are disconnected (e. g. due to network problems, a temporarily stopped primary or secondary, or after a takeover and prior to a failback where the former primary is registered as new secondary), the replication is out of sync. To get in sync again after reconnect the SAP HANA system replication tries to achieve this by initiating a delta shipping of the missing data (instead of a full data shipping).

Depending on the chosen operation mode (`delta_datashipping` or one of the `logreplay*` modes) two different techniques are in place to achieve this: *Data Retention* and *Log Retention*.

7.1 Data Retention

In the SAP HANA system replication operation mode "`delta_datashipping`" the primary creates snapshots that are sent as delta data shipping to the secondary system. It depends on the value of the parameter `datashipping_snapshot_max_retention_time` (default: 300 minutes[19]) how long the primary system keeps this snapshot. If it is not available anymore, a full data shipping is necessary to get the two systems in sync again.

In the following two scenarios, this parameter setting determines, if a full data shipping is necessary or a delta data shipping will be sufficient:

- The failback scenario, where after takeover of the secondary system the former primary system is registered as new secondary.
- The re-register scenario, where after takeover of the secondary system this very system is re-registered as secondary to the original primary system.

[19] As of HANA 2.0 SPS00 this default was increased from 120 minutes in lower versions.

Note

If the connection to your secondary system is lost temporarily (e. g. secondary offline, network issues) – even if it was down longer than `datashipping_snapshot_max_retention_time` – always a delta data shipping is possible in this operation mode.

7.2 Log Retention

In the SAP HANA system replication operation modes "`logreplay`" and "`logreplay_readaccess`" the secondary system only uses the log of the online log area of the primary for re-syncing. After the reconnect or a failback the primary sends the incremental log. Thus, the log must be retained for a longer time (i. e. longer than in `delta_datashipping` operation mode); log segments will not be freed, while the secondary is disconnected.

7.2.1 Log Retention for Secondary Disconnect (on primary)

The primary will not reuse log segments in the online log area that are required to sync the secondary via delta log shipping.

If the secondary is disconnected – but still registered

- the log segments are retained on the primary and marked as `RetainedFree` until secondary has successfully synced again
- the log volume will grow on the primary site, until it has filled up with log segments

In HANA studio this can be monitored on the primary system using tab *Volumes* by selecting the corresponding log volume:

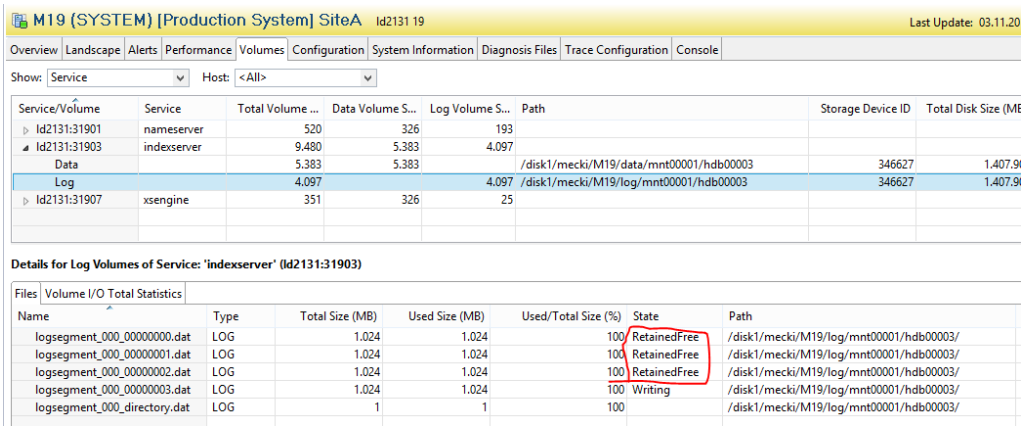

Once the secondary system reconnects and has synced the missing log, these log segments are set to `Free` and can be reused[20] afterwards.

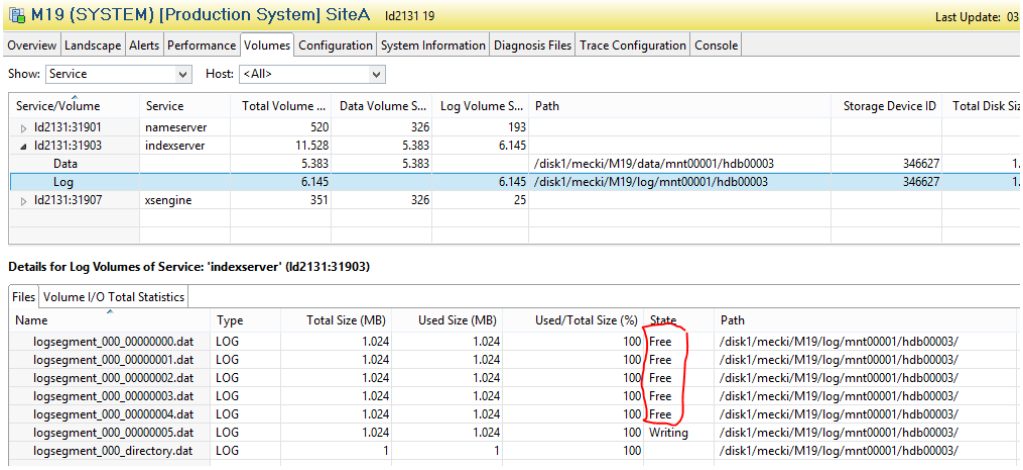

Depending on the setting of the `logshipping_max_retention_size` parameter a full log volume can be prevented at the price of a possibly necessary full data shipping when the system reconnects. Additionally, take care that the setting of this parameter (default is 1 TB per log volume) does not exceed the size of the file system reserved for all log volumes.

This behavior is automatically turned on, if a secondary system with operation mode `logreplay` or `logreplay_readaccess` is registered.

> **Caution**
>
> If a secondary system is shut down and not used for a longer period of time unregister (`hdbnsutil -sr_unregister`[21]) it to prevent log volumes from filling up on the primary site! You can unregister it using the HANA cockpit System Replication app, the HANA studio, or the command line.

7.2.2 Log Retention for Failback (on secondary)

On the secondary site, log retention is required to do a failback with optimized data synchronization. The primary periodically creates persistence snapshots during replication (every 20 min resp. 5 GB) and provides the log position

[20] Log segments marked as "Free" can be reclaimed to free the disk space of these currently unused log segments using this console command: `hdbcons -e <service> "log release"`.

[21] Please also check SAP Note 1945676 - Correct usage of hdbnsutil -sr_unregister

information to the secondary. After takeover, when the old primary is started as secondary, the most recent snapshot is opened on the old primary and the missing log – up to this snapshot – is requested from the new primary.

Log retention can occur in two situations:

- While replication is active
  - The secondary keeps all log starting from the last snapshot position provided by the primary site
  - The old log is automatically released after a new snapshot has been created on the primary site
  - This is active by default and ensures that during replication only a few `RetainedFree` segments are kept online needed to fill the gap between the primary snapshot and the current potential takeover log position
- After a takeover
  - The new primary has to keep the log until a new secondary site is registered and has synced the missing log
  - Because syncing can take some time this behavior has to be explicitly turned on by setting this parameter on the new primary `global.ini/[system_replication]/enable_log_retention = on`

Caution

If the old primary will not be reused as new secondary (failback), it should be disabled after the takeover (`hdbnsutil -sr_disable`) to prevent log volumes from filling up on the new primary site. You can disable it with SAP HANA cockpit, SAP HANA studio, or via command line.

Note

If you have a setup in which there will be frequent failbacks between two sites, we recommend that you set the following parameter on both sites to simplify configuration:

`global.ini/[system_replication]/enable_log_retention = on`

7.2.3 Log Retention Parameters

There are two `ini` parameters to be mentioned in the context of log retention with operation mode `logreplay` in `global.ini/[system_replication]` section:

- `enable_log_retention = auto|on|off`

- `auto`
  - Enable log retention on primary for re-connect
  - Enable log retention on secondary during replication (consider last primary snapshot position)
  - Disable log retention on secondary after takeover
- `On`
  - Enable log retention always
    → required after takeover for failback with delta log shipping

- `logshipping_max_retention_size = 1048576` (MB), default: 1 TB
  - Specifies how the system behaves when many log segments of type `RetainedFree` are created
  - Maximum amount of log that will be kept on primary side for syncing a system replication secondary system:
  - Soft limit if set <> `0`

    → If the limit is reached, segment in state `RetainedFree` are reused in disk full; then a full data shipping is required, i. e. the sedondary needs to be newly registered (`sr_register`)
  - Hard limit if set to `0`

    → Primary standstill, in case disk on primary runs full

7.2.4 Maximum Retention Time Estimation

How long will your SAP HANA system replication landscape "survive" the above described disconnect situations (secondary disconnect, failback) before running into a log full situation, when operation mode `logreplay` or `logreplay_readaccess` is configured? In its system replication application HANA 2.0 SPS01 cockpit provides information about the "Log Retention".

On the System DB overview page click on the system replication tile:

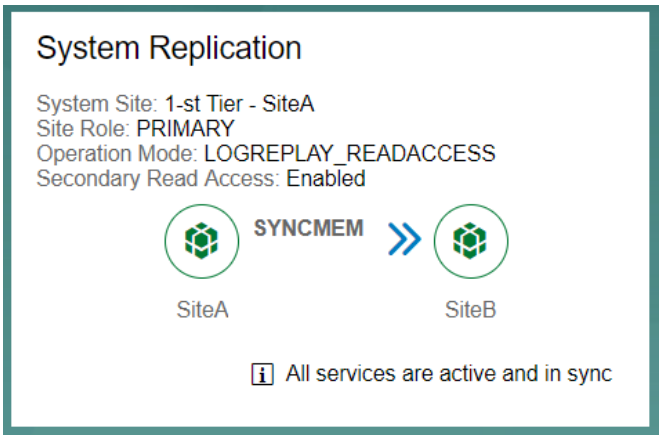

In the System Replication overview page in the upper left area, there is a value shown under a label starting with "*Estimated …*" – in the following example it is "*Estimated log full time*: 210 days".

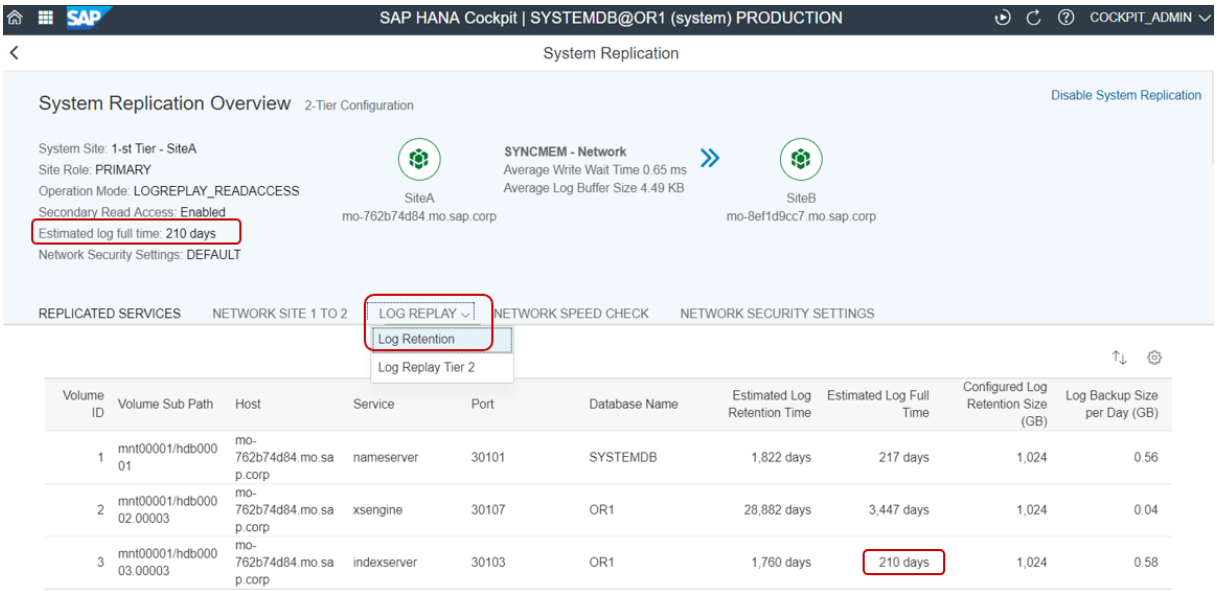

The label could be either one of these two – depending on which type delivers the smallest calculated value:

- *Estimated log retention time*
- *Estimated log full time*

The given value says, how much time is left before this primary HANA database either

- will start to overwrite the `RetainedFree` marked log segments, which then will be no longer available for an optimized resync after the primary could not replicate the redo logs (= *Estimated log retention time*) – resulting in a full data shipping to get the systems in sync again or
- will run into a log full (= *Estimated log full time*)

This value is taken from the list shown in the lower part of the above picture. In this services list one row is shown for each service that is relevant for replication, because it has its own persistence consisting of data and log volume. The calculation is based on the setting of the parameter `LOGSHIPPING_MAX_RETENTION_SIZE`, the write load on the primary system (taken from the log backup size), and the available disk space.

The columns provide the log retention details for the log volumes only:

- Volume ID: of the service
- Volume Sub Path: The *subpath* refers to the last two directories of the file system path where the log volume of this service is located.
- Host: where the service is running
- Service: system replication relevant service (indexserver, nameserver, xsengine)
- Port
- Database Name: SID resp. System DB and tenant DB names
- Estimated Log Retention Time: Maximum number of days/hours logs can be retained to resync primary and secondary after replication was interrupted without requiring a full data shipping.
- Estimated Log Full Time: Maximum number of days/hours until the log volume runs full when created redo logs can no longer be reused.
- Configured Log Retention Size (GB): Determined by the setting of `ini` parameter `LOGSHIPPING_MAX_RETENTION_SIZE` that is valid for this service. (For example, the setting of this parameter in the `indexserver.ini` on the host on which the service is running overwrites the setting of this parameter in `global.ini`.)
- Log Backup Size per Day (GB): The maximum log backup size per day, indicating which replication relevant write load occurred for this service.
- Disk ID: Database Internal Device ID (can also be shown optionally)
- Free Disk Space (GB): (can also be shown optionally)

The value in the upper left part of the overview is taken from the columns "Estimated Log Retention Time" or "Estimated Log Full Time" for the corresponding service – because once the first service starts overwriting its `RetainedFree` log segments a full data shipping becomes necessary, or this service is in "danger" of running into a log full. In the example the indexserver of the tenant DB OR1 would run into a log full in 210 days.

Note
Log full means that no more log segments can be created in the log volume, because the log segment directory is full; currently the number of log segments is roughly 10000.
Disk full means that the disk is full, which is not necessarily the case in log full.
Log retention usually deals with disk full situations, because with 10000 log segments each 1 GB of size – you can create 10 TB of log segments.

8. Active/Active (read enabled) secondary

Since HANA 2.0 SPS00 the SAP HANA system replication can be configured as Active/Active (read enabled) landscape where the SQL ports on the tier-2 secondary are open for read access.

This allows to actively make use of the secondary system, which in earlier HANA revisions could not be accessed via SQL, but purely functioned as passive replication site prepared for a fast takeover. Reporting load can be taken from the primary system and executed on the secondary system in such an Active/Active (read enabled) configuration.

8.1 Active/Active read-only specifics

The following applies for an Active/Active (read enabled) secondary system:

- Secondary allows read-only access, no write accesses resulting in redo logging is possible
- All replication modes (sync, syncmen, async) are supported
- The secondary must have the same HANA version as the primary (i. e. read-only access to secondary is not supported during Near Zero Downtime Upgrade [see below])
- Exporting tables is supported with CSV file as target (but no binary export is available).
- The Redo Log Replay on the secondary system runs asynchronously to the primary operations. Thus, the secondary system provides statement level snapshot isolation[22] with a delayed view on the data and no minimum delay guarantee.
- The secondary system allows read accesses via SQL. The XS service port remains closed until a takeover took place.
  - In this context please read SAP note 2391079 - Access restrictions in Active/Active system setup
  - The table consistency checker can also run on the read-enabled secondary; for details refer to SAP Note 2116157 - FAQ: SAP HANA Consistency Checks and Corruptions
- The secondary will not accept new connections if the primary is down in HANA 2.0 SPS00. Existing connections can continue.
- Single Host and scale-out HANA systems are supported
- Since HANA 2.0 SPS02 an Active/Active (read enabled) configuration can coexist with Dynamic Tiering; however, no read-access on DT data is possible on the secondary system. For a list of System Replication features

[22] This is the isolation level *Commited Read* where different statements in a transaction may see different snapshots of the database, i.e. each statement sees changes that were committed when execution of the statement started.

that Dynamic Tiering supports, see SAP Note 2447994 – SAP HANA Dynamic Tiering Support for SAP HANA System Replication.

8.2 Check Active/Active configuration

You can check, if your system replication is configured as an Active/Active (read enabled) system using command line tools, SAP HANA cockpit and SAP HANA studio.

8.2.1 Using the command line

As `<sid>adm` on command line, run one of the following commands and look for the operation mode `logreplay_readaccess`.

systemReplicationStatus.py:

```
python
$DIR_INSTANCE/exe/python_support/systemReplicationStatus.py --
sapcontrol=1 | grep OPERATION_MODE
   service/ld4144/30207/OPERATION_MODE=logreplay_readaccess
   service/ld4144/30201/OPERATION_MODE=logreplay_readaccess
   service/ld4144/30203/OPERATION_MODE=logreplay_readaccess
```

hdbnsutil:

```
hdbnsutil -sr_state | grep "operation mode"
   operation mode: logreplay_readaccess
```

8.2.2 Using SAP HANA cockpit 2.0

On the primary system check the System Replication tile of your system DB in the system overview page:

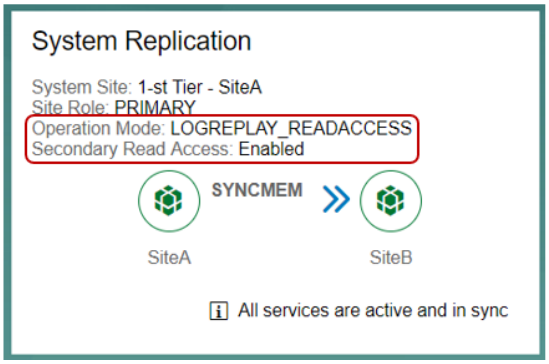

It shows as configured "*Operation Mode*" `logreplay_readaccess` indicating, that this is an Active/Active (read enabled) configured system replication landscape. Additionally, the "*Secondary Read Access*" informs you, that the

SQL ports are open for reading on the Active/Active (read enabled) secondary system.

If you go to the system overview page of this read enabled secondary system, you are informed about this fact that this is a "Read-only" system. Additionally, the delay (ms) is shown on top, indicating how far the consistent view on the data on this secondary system is behind the current data of the primary system.

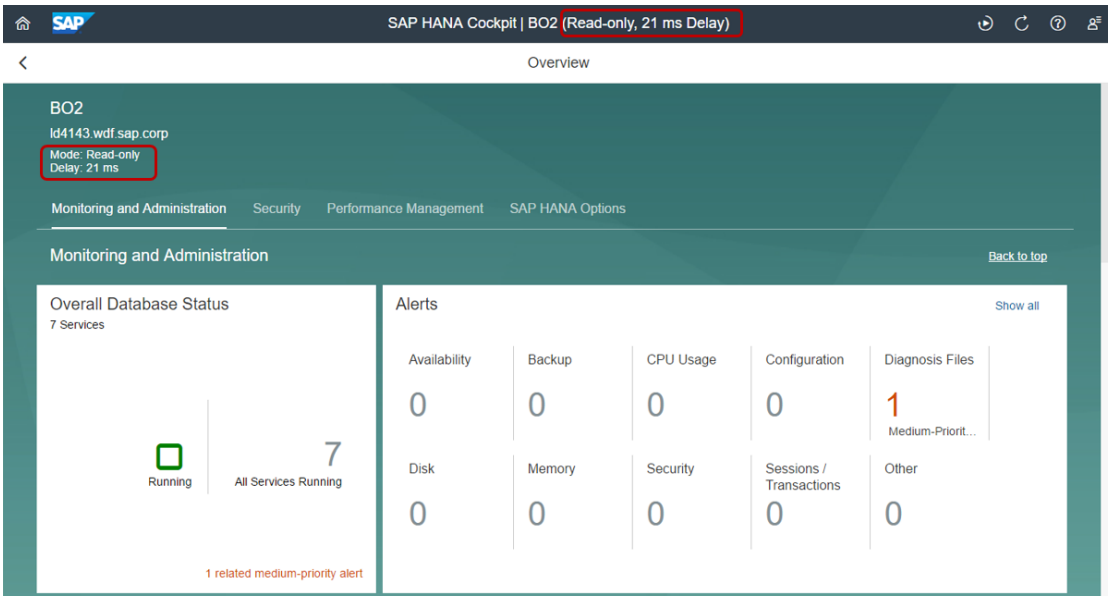

Some of the monitoring tiles deliver data about the state of this secondary system, like Memory Usage, CPU Usage, Disk Usage, active Threads, ... because SQL access is possible. Also, the System Replication tile shows information.

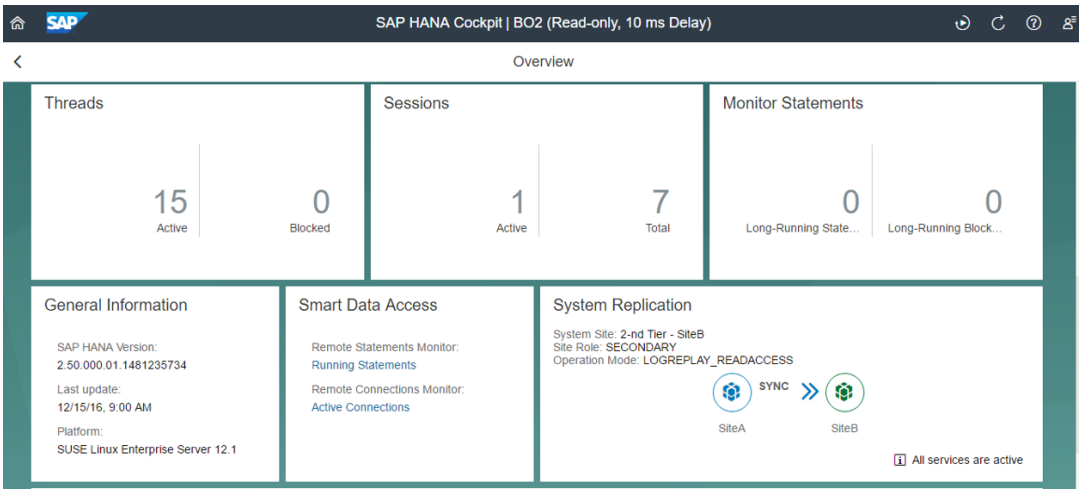

8.2.3 Using HANA studio

On the primary system select from the monitoring view `M_SYSTEM_REPLICATION` and check columns "`OPERATION_MODE`" and "`SECONDARY_READ_ACCESS_STATUS`" to find out, if this is an Active/Active (read enabled) system, where reading is possible on the secondary system:

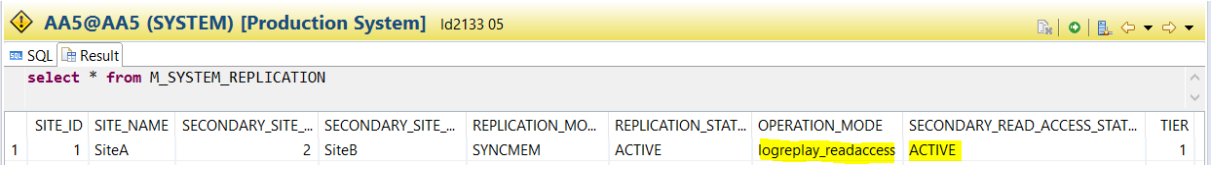

	SITE_ID	SITE_NAME	SECONDARY_SITE_...	SECONDARY_SITE_...	REPLICATION_MO...	REPLICATION_STAT...	OPERATION_MODE	SECONDARY_READ_ACCESS_STAT...	TIER
1	1	SiteA	2	SiteB	SYNCMEM	ACTIVE	logreplay_readaccess	ACTIVE	1

8.3 Connection types

There are two ways to execute SQL queries on the read enabled secondary system: opening an explicit connection to the secondary or executing an SQL statement on the primary which is redirected to the secondary according to a hint. Applications making use of this feature need to go the one way or the other.

8.3.1 Explicit read-only connection to secondary

The application can directly connect to the read enabled secondary system and execute the SQL queries against it.

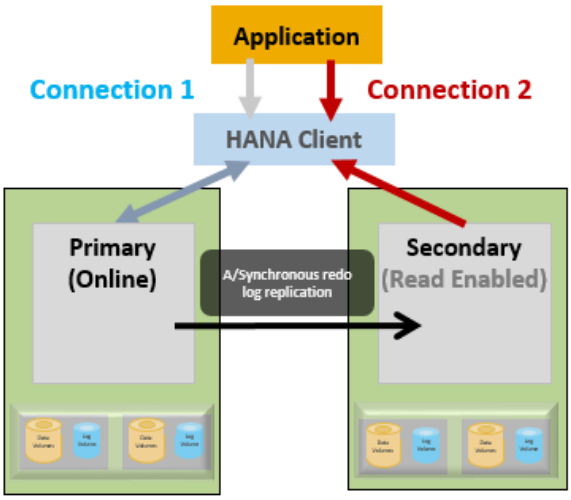

Direct connection to read enabled secondary

You can do that for example using the SAP HANA studio simply by opening an SQL console on the read enabled secondary system. Alternatively, you can use

`hdbsql` when logged on as `<sid>adm` to select on the read enabled secondary.[23]

8.3.2 Statement routing from primary to secondary

You can also pass the SQL query to the primary and add a hint saying that this statement should be preferably executed on the read enabled secondary system.

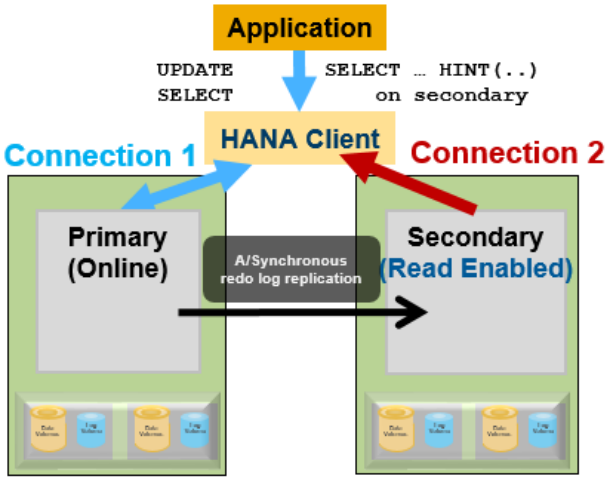

HINT-based statement routing to read enabled secondary

For such a hinted statement the HANA client opens a second connection to the read enabled secondary system according to host information returned by primary. The client then sends the statement to the secondary. If the secondary cannot execute a statement for any reason it returns an error and the client sends the statement to the primary (*automatic failback*).

With the so-called `RESULT_LAG` hint `'hana_sr'` you can also provide a maximum acceptable snapshot delay in `<seconds>`:

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

Examples of such hinted `SELECT` statements are:

```
SELECT * FROM T1 WITH HINT( RESULT_LAG ('hana_sr') );
SELECT * FROM T1 WITH HINT( RESULT_LAG ('hana_sr', 60) );
```

Hint

[23] Currently the HANA cockpit is not yet capable of accessing the read enabled secondary.

If the routed statement shows up in the `M_SQL_PLAN_CACHE` monitoring view on the secondary system, your statement was executed there.

For further details, please have a look at the SAP HANA SQL and System Views Reference guide – section: *Hints for Active/Active (Read-Enabled)*.

8.3.3 Virtual IP address handling

Since in an Active/Active (read enabled) configuration both sites of the system replication setup are open for SQL access, a *second virtual IP address* for the read access on the secondary system is needed[24].

During the lifetime of such a system the following scenarios are relevant for handling these virtual IP addresses in an Active/Active (read enabled) system:

- *Read access*: As said for read access on the secondary system a second virtual IP address is required – besides the one handling the access to the primary system.
- *Takeover*: During takeover you can keep this secondary's virtual IP address. It will be used for read-only access until a reconnect occurs. The former primary's virtual IP address is also rebound to access the former secondary which is the now active system (see section "Client Connection Recovery" in this document). In this situation two virtual IP addresses are available for accessing the former secondary after takeover.

  Caution

  Make sure that the now active system can handle this "double" workload of the original primary and the read-only secondary.

- *Failback*: Just like the system replication sites switched roles, the virtual IP addresses switched the locations (hosts) they are bound to.

9. Testing

The test phase is a very important phase to verify if KPIs are met and the landscape performs the way it was configured. Therefore, a few test cases are suggested below as guideline, which should be enhanced by your specific requirements. The tests should be performed with realistic data load and size.

[24] In scale-out HANA databases every master candidate needs its own virtual IP address.

Test case	Description
Full Replication	Measure how long the initial synchronization takes, from when replication is started until primary and secondary are in sync.
Lost Connection	Measure how long it takes until primary and secondary are back in sync after the connection is re-established.
Takeover	Measure how long it takes for the secondary system to be fully available after a takeover command.
Data Consistency	Create or change data, then perform a takeover and check if the data is still available.
Client Reconnect	Test client access after a take-over, to check if the DNS/Virtual IP switch worked.
Primary becomes secondary	Measure how long it takes until both systems are in sync, when the former primary becomes the secondary after a takeover.

10. Operation / Maintenance

There are multiple ways to monitor SAP HANA, which are described in the SAP HANA Administration Guide and SAP Solution Manager[25].

And there are various ways to verify if the primary and secondary systems are in sync and are running correctly.

10.1 Alerts

With HANA 1.0 SPS09 system replication specific alerts were introduced (they are no longer hidden behind "Internal Events"):

- System Replication *Connection Closed* (Alert ID 78)
- System Replication *Configuration Parameter Mismatch* (Alert ID 79)

As of HANA 1.0 SPS12 a new system replication specific alert was introduced for systems running in operation mode "logreplay":

- System Replication *Logreplay Backlog* (Alert ID 94)
  The alert is thrown when logreplay is delayed on the secondary site causing a longer takeover time.
  The alert has a different priority based on the size of the redo log that was not yet replayed:

Low	10 GB < logreplay backlog < 50 GB
Medium	50 GB <= logreplay backlog < 500 GB

[25] For SAP Solution Manager please consider Note 1747682 - SolMan 7.1: Managed System Setup for HANA

High	logreplay backlog >= 500 GB

- To identify the reason for the increased system replication logreplay backlog, check the state of the services on the secondary system. To get more information, monitor the secondary site. Possible causes for the increased system replication logreplay backlog can be, for example, a slow or not functioning log replay, or non-running service on the secondary system.

Another new alert (as of HANA 2.0 SPS02) monitors the log shipping:

- System Replication *Log Shipping Backlog* (Alert ID 104)
  This alert is raised when the system replication log shipping is delayed or will not work properly causing data loss on the secondary system, if a takeover is done.
  The alert has a different priority based on the threshold reached:

Low	1 GB < log shipping backlog < 10 GB
Medium	10 GB <= log shipping backlog < 50 GB
High	log shipping backlog >= 50 GB

- To identify the reason for the increased system replication log shipping backlog, check the status of the secondary system. Possible causes for the increased system replication log shipping backlog can be a slow network performance, connection problems, or other internal issues (for example, in `sync` or `syncmem` replication modes).

These alerts are only visible with the Embedded Statistics Server (ESS)[26]; however, old style alerts are still generated in order not to invalidate any reporting infrastructure after migration. Old alerting can be disabled by setting the following configuration parameter in `global.ini`:

```
[system_replication]
keep_old_style_alert = false (default=false)
```

How to configure e-mail notifications of alerts, please follow the instructions in the SAP HANA Admin guide (see *Configure E-Mail Notifications for Alerts*).

As the monitoring of the secondary site was improved by the introduction of the so-called proxy views (see below in section "*Monitoring the secondary site*") with HANA 1.0 SPS12 an alerting for the secondary site(s) was established. On the primary site alerts occurring on the secondary hosts are shown as alerts and associated with the host where they occurred by providing this `host:port` information.

[26] Migration of the classic statistics server to the Embedded Statistics Server is described in SAP note 1917938.

In this example a secondary host[27] reported a DiskFull event:

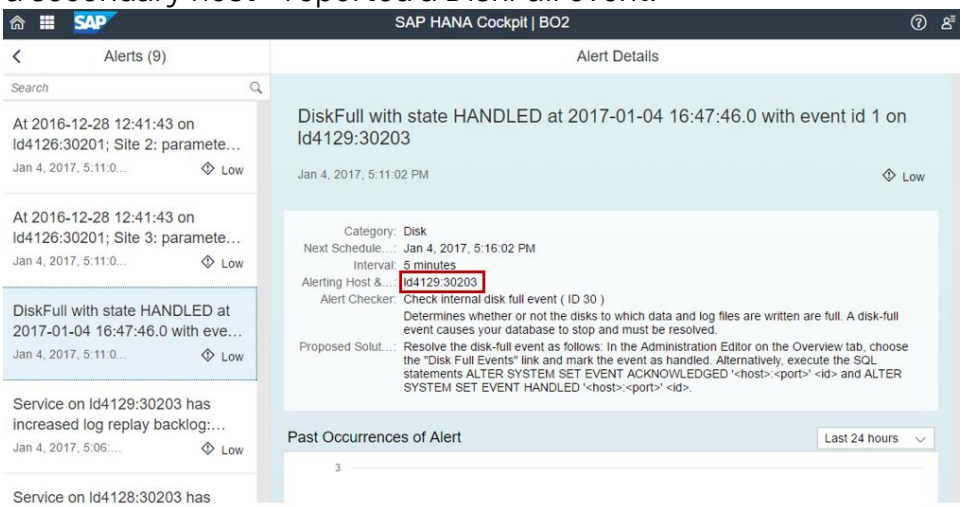

HANA Cockpit:

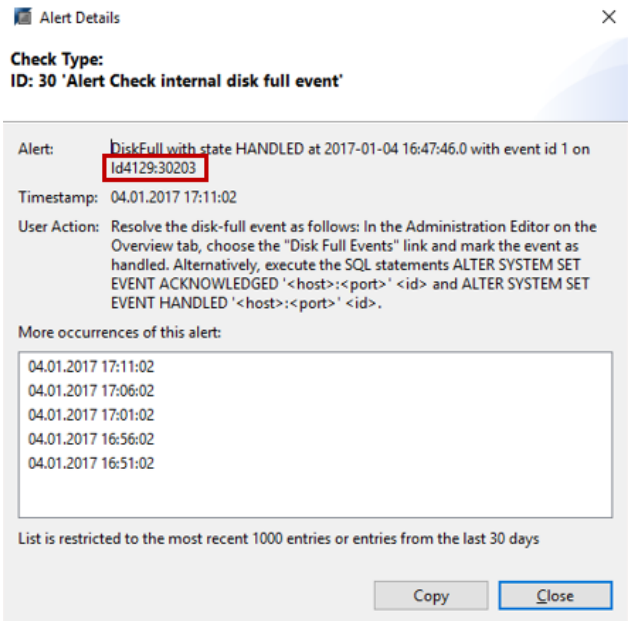

HANA studio:

10.2 Verification

10.2.1 Using SAP HANA cockpit 2.0

Since HANA 2.0 SPS00 with the XSA based SAP HANA cockpit[28] you can monitor multiple HANA systems within one cockpit from a browser. When opening the HANA cockpit for one HANA database an extended system replication monitoring application is offered.

[27] A DB admin currently must know which host belongs to which site.

[28] More information on SAP HANA Cockpit can be found here: http://help.sap.com/hana/SAP_HANA_Administration_Guide_en.pdf

If system replication is configured, the corresponding tile appears on the main screen of the system overview page providing information about the type of landscape (*2-tier* or *3-tier*), the replication mode(s) between the primary and the secondary(s), the operation mode as well as an overall replication status:

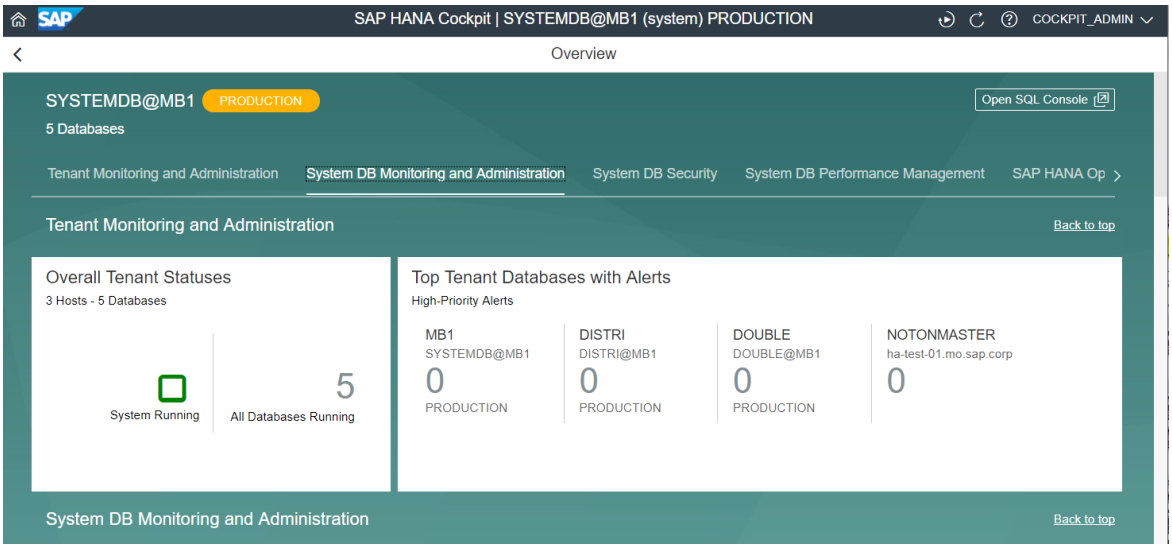

...

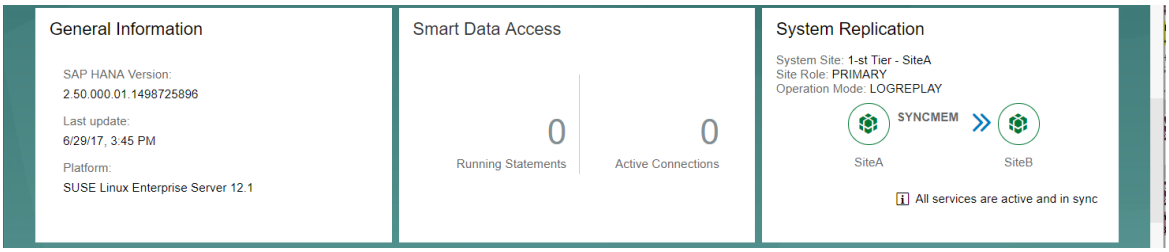

If all tiers are shown in green and the *System Replication* tile[29] tells you "All services are active and in sync" your system is doing well. Red tiers would indicate a problem with the replication. And if you are accessing one of the secondary HANAs with the cockpit system overview, this tile shows the other tiers in blue and the one you are on with the HANA cockpit in green.

To check the status of replication in detail, click on the *System Replication tile*. The application provides an overview on the system replication configuration and status. On top, the "chain" of systems with their replication modes is shown containing further information about the sites and the network connections between them.

[29] If the tile does not show up, you have to grant the system replication role to the corresponding user, e. g. in the HANA studio right mouse-click on the corresponding user in the landscape overview under *Security* → on the "*Granted Roles tab*" click on *"+"* → filter for "sysrep" and select corresponding role.

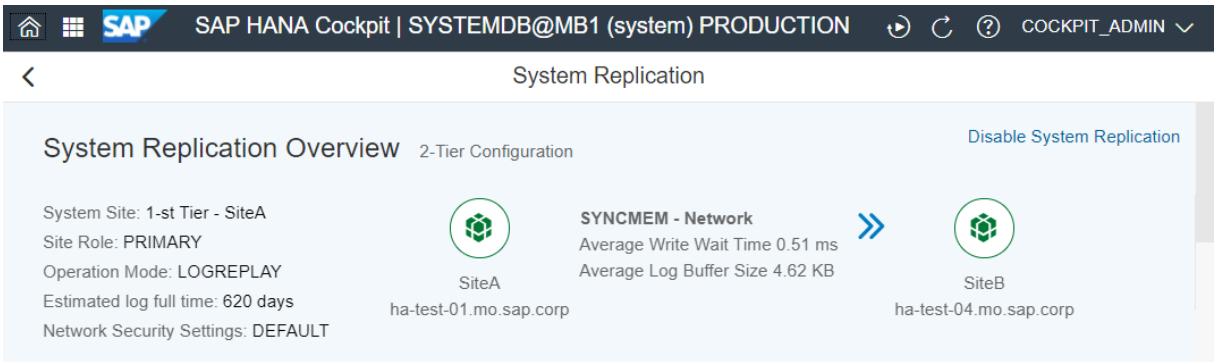

A graphical representation of your system replication landscape is given. It tells you the chosen site names, the replication mode used between the sites and even provides a snapshot view on the current average redo log shipping time ("*Log Buffer Write Wait Time in ms*") and the average size of shipped redo log buffers. It describes how long it took on average to send redo log buffers of "*Avg. Log Buffer Size*" to the secondary site based on measurements of the last 24 hours.[30]

Below there is a selection of tabs providing additional information for different system replication relevant topics.

Related Alerts tab

If a system replication relevant alert occurred, the first tab is the "*Related Alerts*" tab (if no alert exists which is relevant for system replication, this tab is not shown):

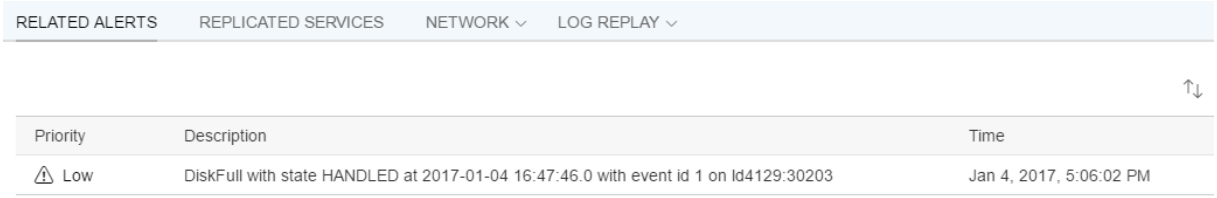

Replicated Services tab

When activating the "Replicated Services" tab an excerpt of the monitoring view `M_SERVICE_REPLICATION` is shown. The displayed table shows at a glance per site and service the replication state per service.

[30] For synchronous replication, this is the round trip time for sending the redo log buffer and receiving the acknowledgement; for asynchronous replication it refers to the time, it takes, until the log buffer was sent after its creation.

REPLICATED SERVICES | NETWORK SITE 1 TO 2 | LOG REPLAY | NETWORK SPEED CHECK | NETWORK SECURITY SETTINGS

Site ID	Site Name	Secondary Site Name	Service	Port	Replication Mode	Full Sync	Replication Status	Replication Details	Secondary Fully Recoverable
1	SiteA ha-test-01.mo.sap.corp	SiteB ha-test-04.mo.sap.corp	xsengine	30107	SYNCMEM	disabled	active		true
	SiteA ha-test-02.mo.sap.corp	SiteB ha-test-05.mo.sap.corp	indexserver	30143	SYNCMEM	disabled	active		true
	SiteA ha-test-01.mo.sap.corp	SiteB ha-test-04.mo.sap.corp	indexserver	30146	SYNCMEM	disabled	active		true
			indexserver	30140	SYNCMEM	disabled	active		true
	SiteA ha-test-02.mo.sap.corp	SiteB ha-test-05.mo.sap.corp	indexserver	30140	SYNCMEM	disabled	active		true
	SiteA ha-test-01.mo.sap.corp	SiteB ha-test-04.mo.sap.corp	indexserver	30143	SYNCMEM	disabled	active		true
			indexserver	30103	SYNCMEM	disabled	active		true
	SiteA ha-test-02.mo.sap.corp	SiteB ha-test-05.mo.sap.corp	indexserver	30103	SYNCMEM	disabled	active		true
	SiteA ha-test-01.mo.sap.corp	SiteB ha-test-04.mo.sap.corp	nameserver	30101	SYNCMEM	disabled	active		true
	SiteA ha-test-03.mo.sap.corp	SiteB ha-test-06.mo.sap.corp	nameserver	30101	STANDBY	disabled	active		false
	SiteA ha-test-02.mo.sap.corp	SiteB ha-test-05.mo.sap.corp	nameserver	30101	STANDBY	disabled	active		false

Tip

If you want to have a look at the trace files of all sites of your system replication landscape, just click on the link "*View trace and diagnostic files*" on the primary's system overview page in HANA cockpit.

Alerting & Diagnostics
Configure alerts
View trace and diagnostic files
Plan trace
Troubleshoot unresponsive system

This makes all diagnosis files from the trace directories of all sites visible in the browser:

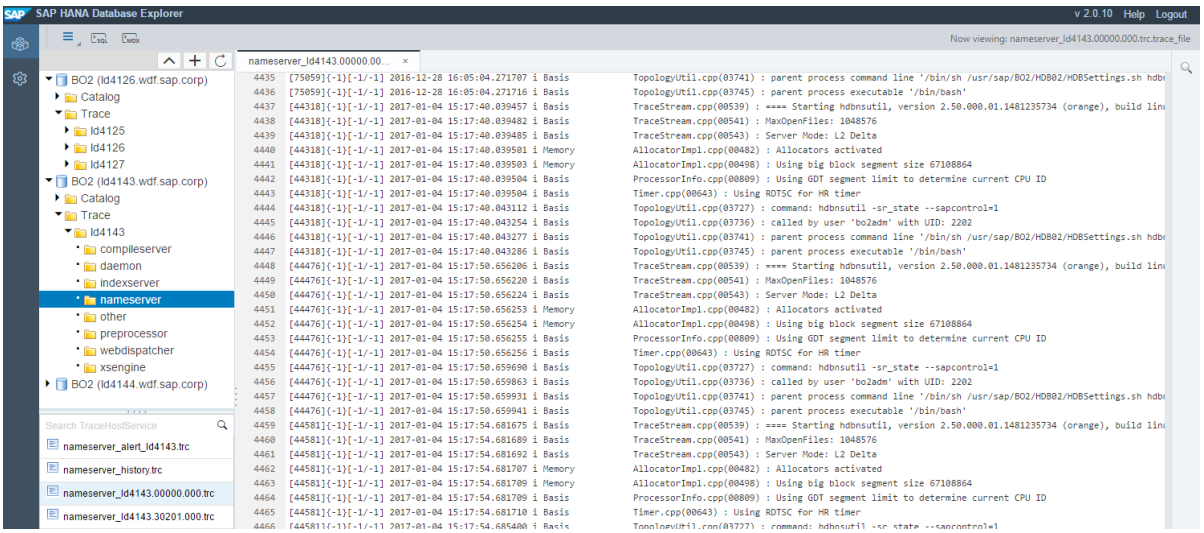

Trace file viewer in HANA cockpit

If you click on one row, you can see the details for the corresponding service grouped thematically, like in the below example for one *indexserver*. Since this

information is “context aware”, you only get the information required for this system. Thus, since this example system is running in operation mode “logreplay” no information on delta data shipping is shown here. But the context sensitive information about the “log replay delay” is displayed. The delta between “Last Log Position” and “Replayed Log Position” indicates how far the log replay is behind on the secondary.

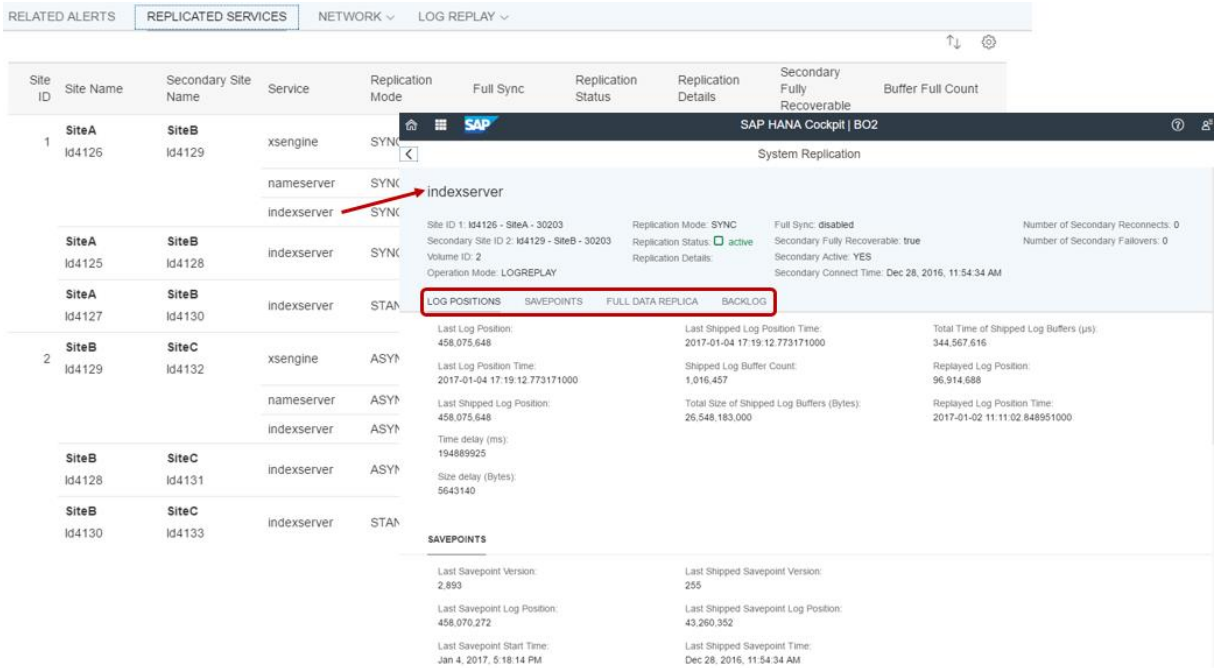

Network tab

If the “network” tab is activated – you can select the network connection you want to analyze (pull down menu for *Network Site 1 to 2*, *Network Site 2 to 3*). A graph appears comparing the *local write wait time* (i.e. writing redo log buffer into the local log volume) with the *remote write wait time* (i.e. shipping the redo log and receiving the acknowledgement) monitored over the last 24 hours. At a glance, one can see if peak times occurred and how the network connection reacted.

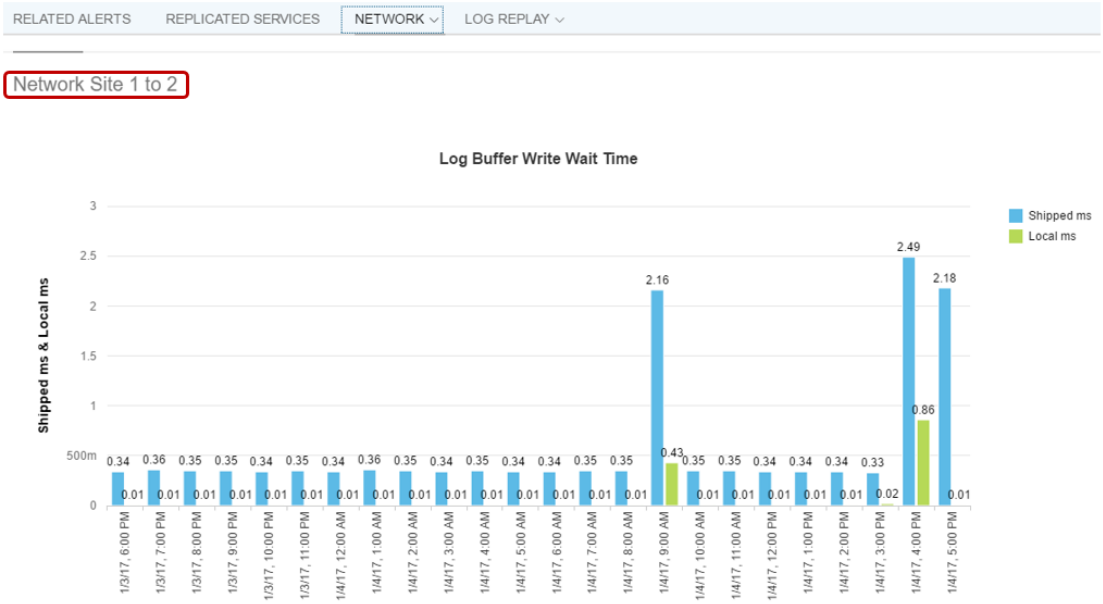

If the ASYNC replication mode is configured between two sites, like in this example between tier-2 and tier-3, you also receive information about the network performance by selecting the corresponding connection between tier-2 and tier-3:

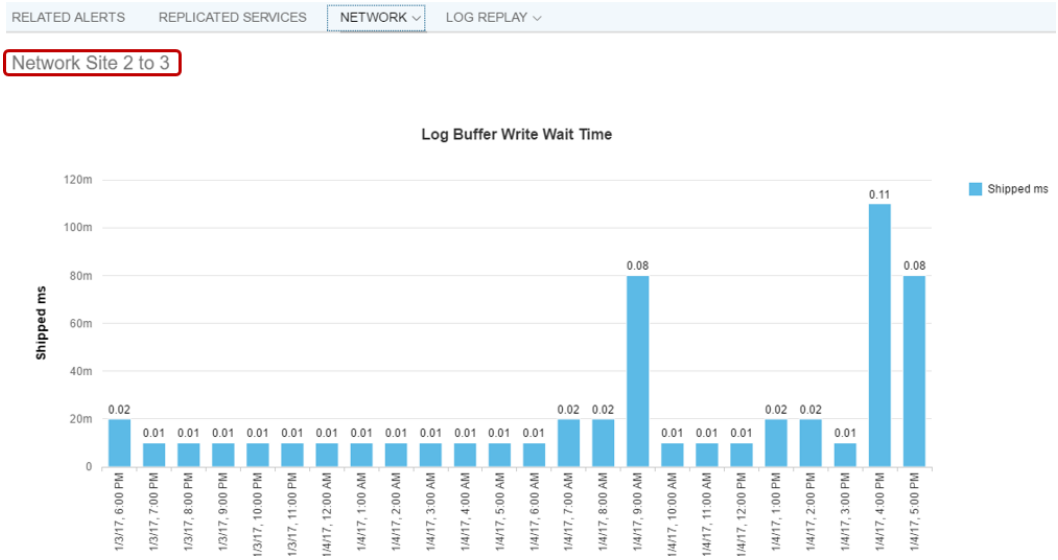

Here the "*Avg. Write Wait Time ms*" describes the time it took from the creation of the redo log buffer (i.e. committing a write transaction) until the redo log buffer in fact was sent out to the network. This value is an indicator for peak load phases

and could point to network or I/O problems[31] on the secondary site, which can influence the primary's performance as well.

Log Replay tab

This "Log Replay" tab is only visible, if the operation mode `logreplay` or `logreplay_readaccess` is configured for the system replication landscape. When this tab is activated for a secondary (selectable via pull-down menu) the log replay delay on the secondary system is show as size in GB for the last 24 hours. And if at some point in time the threshold of the corresponding alert (ID 94 - *System Replication Logreplay Backlog*) was exceeded, this is indicated accordingly (like in the below example):

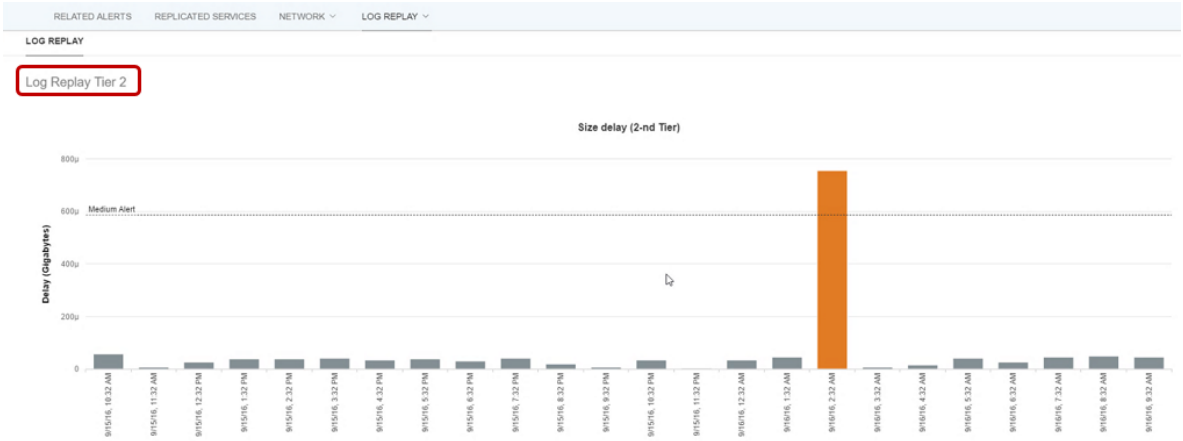

How to analyze a high replay backlog is described in this SAP Note 2409671 - *High replay backlog on HANA System Replication Secondary Site*.

Additionally, (with HANA 2.0 SPS01) under this tab you can select the "Log Retention" information of your system (please see section "*Maximum Retention Time Estimation*" in this document for details).

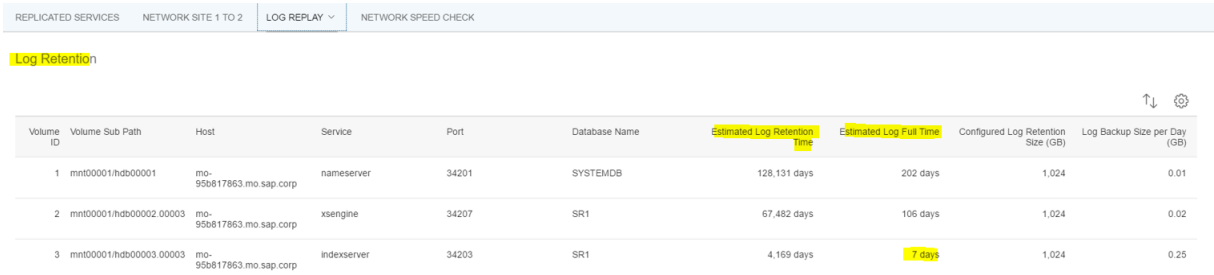

Volume ID	Volume Sub Path	Host	Service	Port	Database Name	Estimated Log Retention Time	Estimated Log Full Time	Configured Log Retention Size (GB)	Log Backup Size per Day (GB)
1	mnt00001/hdb00001	mo-95b817863.mo.sap.corp	nameserver	34201	SYSTEMDB	128,131 days	202 days	1,024	0.01
2	mnt00001/hdb00002.00003	mo-95b817863.mo.sap.corp	xsengine	34207	SR1	67,482 days	106 days	1,024	0.02
3	mnt00001/hdb00003.00003	mo-95b817863.mo.sap.corp	indexserver	34203	SR1	4,169 days	7 days	1,024	0.25

[31] If the receiving OS buffer on the secondary cannot write down the incoming redo log buffers to disk due to I/O problems, this buffer can run full and is not able to accept newly shipped buffer fast enough.

10.2.2 Using SAP HANA studio

Check the overall status on the primary's *Overview* tab. This should state "All services are active and in sync":

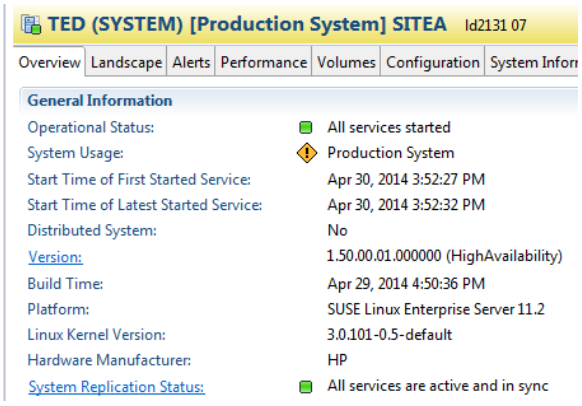

To check the status of the replication in detail: Select *Landscape tab → system replication* (shows information from system view `M_SERVICE_REPLICATION` with a lot of columns):

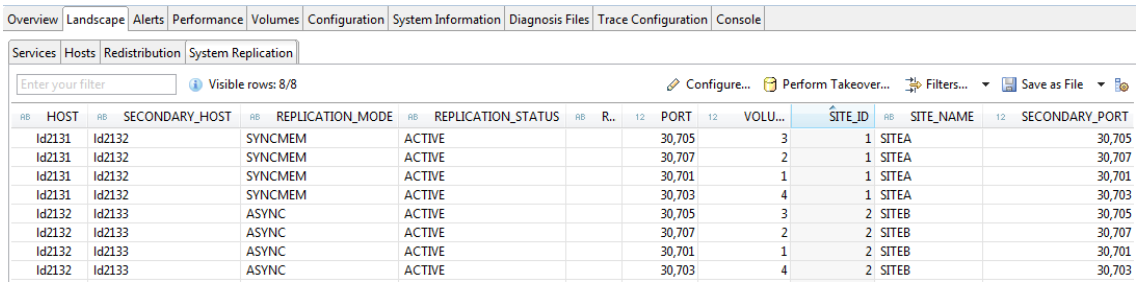

HOST	SECONDARY_HOST	REPLICATION_MODE	REPLICATION_STATUS	R..	PORT	VOLU...	SITE_ID	SITE_NAME	SECONDARY_PORT
Id2131	Id2132	SYNCMEM	ACTIVE		30,705	3	1	SITEA	30,705
Id2131	Id2132	SYNCMEM	ACTIVE		30,707	2	1	SITEA	30,707
Id2131	Id2132	SYNCMEM	ACTIVE		30,701	1	1	SITEA	30,701
Id2131	Id2132	SYNCMEM	ACTIVE		30,703	4	1	SITEA	30,703
Id2132	Id2133	ASYNC	ACTIVE		30,705	3	2	SITEB	30,705
Id2132	Id2133	ASYNC	ACTIVE		30,707	2	2	SITEB	30,707
Id2132	Id2133	ASYNC	ACTIVE		30,701	1	2	SITEB	30,701
Id2132	Id2133	ASYNC	ACTIVE		30,703	4	2	SITEB	30,703

For all services, the `REPLICATION_STATUS` should be "`ACTIVE`". Detailed information about shipped sizes and shipping times are available.

10.2.3 Check via SQL query

Or directly get system replication specific information from the system view **`M_SERVICE_REPLICATION`**. On the primary execute:

```
select * from "SYS"."M_SERVICE_REPLICATION";
```

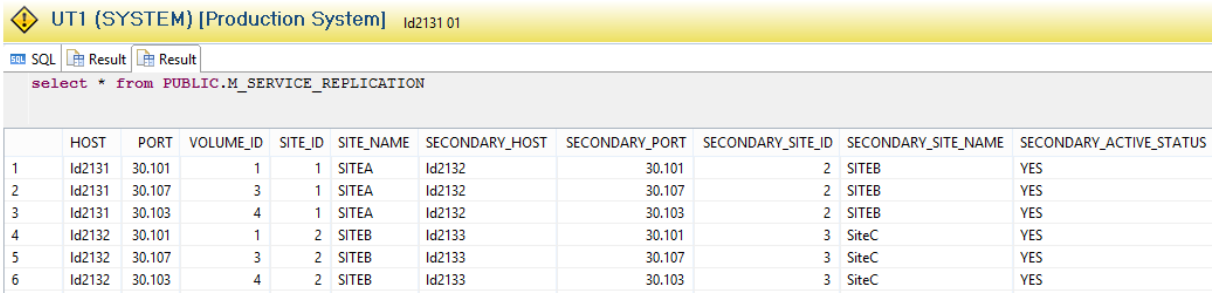

	HOST	PORT	VOLUME_ID	SITE_ID	SITE_NAME	SECONDARY_HOST	SECONDARY_PORT	SECONDARY_SITE_ID	SECONDARY_SITE_NAME	SECONDARY_ACTIVE_STATUS
1	Id2131	30.101	1	1	SITEA	Id2132	30.101	2	SITEB	YES
2	Id2131	30.107	3	1	SITEA	Id2132	30.107	2	SITEB	YES
3	Id2131	30.103	4	1	SITEA	Id2132	30.103	2	SITEB	YES
4	Id2132	30.101	1	2	SITEB	Id2133	30.101	3	SiteC	YES
5	Id2132	30.107	3	2	SITEB	Id2133	30.107	3	SiteC	YES
6	Id2132	30.103	4	2	SITEB	Id2133	30.103	3	SiteC	YES

Since HANA 1.0 SPS09 the contents of the view `M_SERVICE_REPLICATION` is collected by the statistics server every hour. Thus, the history of data and log replication can be viewed in the table. On the primary execute the following command to view the data replicated by the indexservers (volume 4 in this example) from the primary to the tier-2 secondary:

```
select * from "_SYS_STATISTICS"."HOST_SERVICE_REPLICATION"
where volume_id=4 and site_id=1;
```

UT1 (SYSTEM) [Production System] ld2131 01

SQL | Result | Result

```
select * from "_SYS_STATISTICS"."HOST_SERVICE_REPLICATION" where volume_id=4 and site_id=1
```

	SNAPSHOT_ID	SERVER_TIMESTAMP	INDEX	HOST	PORT	VOLUME_ID	SITE_ID	SITE_NAME	SECONDARY_HOST	SECONDARY_PORT	SECONDARY_SITE_ID
1	12.06.2015 01:09:52.0	12.06.2015 03:09:55.52	ld2131:30103:4:1:2	ld2131	30.103	4	1	SITEA	ld2132	30.103	2
2	12.06.2015 02:09:52.0	12.06.2015 04:10:08.082	ld2131:30103:4:1:2	ld2131	30.103	4	1	SITEA	ld2132	30.103	2
3	12.06.2015 03:09:52.0	12.06.2015 05:10:08.491	ld2131:30103:4:1:2	ld2131	30.103	4	1	SITEA	ld2132	30.103	2
4	12.06.2015 04:09:52.0	12.06.2015 06:10:08.512	ld2131:30103:4:1:2	ld2131	30.103	4	1	SITEA	ld2132	30.103	2
5	12.06.2015 05:09:52.0	12.06.2015 07:10:08.157	ld2131:30103:4:1:2	ld2131	30.103	4	1	SITEA	ld2132	30.103	2
6	12.06.2015 06:09:52.0	12.06.2015 08:09:56.789	ld2131:30103:4:1:2	ld2131	30.103	4	1	SITEA	ld2132	30.103	2
7	12.06.2015 07:09:52.0	12.06.2015 09:10:08.584	ld2131:30103:4:1:2	ld2131	30.103	4	1	SITEA	ld2132	30.103	2
8	12.06.2015 08:09:52.0	12.06.2015 10:10:09.229	ld2131:30103:4:1:2	ld2131	30.103	4	1	SITEA	ld2132	30.103	2
9	12.06.2015 09:09:52.0	12.06.2015 11:10:08.822	ld2131:30103:4:1:2	ld2131	30.103	4	1	SITEA	ld2132	30.103	2
10	12.06.2015 10:09:52.0	12.06.2015 12:10:08.568	ld2131:30103:4:1:2	ld2131	30.103	4	1	SITEA	ld2132	30.103	2

Since HANA 1.0 SPS11 there is the new system view **`M_SYSTEM_REPLICATION`** providing general system replication relevant information about the whole system, e. g. which replication mode is used, which operation mode, and as which "`TIER`" a site is configured. In the below example `SiteA` with `SITE_ID=1` is currently configured as `TIER=1` (i. e. as primary). On the primary execute:

```
select * from "SYS"."M_SYSTEM_REPLICATION";
```

UT1 (SYSTEM) [Production System] ld2131 01

SQL | Result

```
select * from "SYS"."M_SYSTEM_REPLICATION"
```

	SITE_ID	SITE_NAME	SECONDARY_SITE_ID	SECONDARY_SITE_NAME	REPLICATION_MODE	REPLICATION_STATUS	OPERATION_MODE	TIER
1	1	SiteA	2	SiteB	SYNC	ACTIVE	logreplay	1
2	2	SiteB	3	SiteC	ASYNC	ACTIVE	logreplay	2

10.2.4 Using command line tool `hdbnsutil`

To view the system replication topology configuration status on both systems, execute `hdbnsutil -sr_state` on the primary and the secondary:

```
tedadm@ld2131:/usr/sap/TED/HDB07> hdbnsutil -sr_state
checking for active or inactive nameserver ...
System Replication State
~~~~~~~~~~~~~~~~~~~~~~~~
mode: primary
site id: 1
site name: SITEA
Host Mappings:
~~~~~~~~~~~~~~
ld2131 -> [SITEA] ld2131
ld2131 -> [SITEB] ld2132
```

```
done.
```

For a Multitier System Replication the mappings of all three sites are displayed:

```
ut1adm@ld2131:/usr/sap/UT1/HDB01> hdbnsutil -sr_state
checking for active or inactive nameserver ...
System Replication State
~~~~~~~~~~~~~~~~~~~~~~~~
mode: primary
site id: 1
site name: SITEA
Host Mappings:
~~~~~~~~~~~~~~
ld2131 -> [SITEA] ld2131
ld2131 -> [SITEC] ld2133
ld2131 -> [SITEB] ld2132
done.
```

When using the additional option `--sapcontrol=1` the key-value-pair output can be parsed by a script line by line.

Here is the output where the `-sr_state` command was executed on a primary site of a Multitier System Replication:

```
ut1adm@ld2131:/usr/sap/UT1/HDB01> hdbnsutil -sr_state --sapcontrol=1
checking for active or inactive nameserver ...
SAPCONTROL-OK: <begin>
mode=primary
site id=1
site name=SITEA
mapping/ld2131=SITEA/ld2131
mapping/ld2131=SITEC/ld2133
mapping/ld2131=SITEB/ld2132
SAPCONTROL-OK: <end>
Done
```

Here is the output where the `-sr_state` command was executed on a tier-2 secondary site of a Multitier System Replication:

```
ut1adm@ld2132:/usr/sap/UT1/HDB01> hdbnsutil -sr_state --sapcontrol=1
checking for active or inactive nameserver ...
SAPCONTROL-OK: <begin>
mode=sync
site id=2
site name=SITEB
active primary site=1
mapping/ld2132=SITEA/ld2131
mapping/ld2132=SiteC/ld2133
mapping/ld2132=SITEB/ld2132
primary masters=ld2131
SAPCONTROL-OK: <end>
done.
```

Further explanation of the output:

- `mode` – can have the values `primary`, `sync`, `async`, and `syncmem` to represent the mode relevant on the site where the command is executed (e. g. in a Multitier System Replication on the primary the mode would be `primary`, on the tier-2 secondary it could be either `sync` or `syncmem`, and on the tier-3 secondary it is `async`).
- `site id` – is a unique identifier of a site which is incremented for each site attached to a SAP HANA system replication. It is first removed, when the system replication is disabled.
- `site name` – is the name you give your sites during the *enable* and *register* steps of system replication configuration.
- `mapping/<currentHost>` – shows which hosts are involved in this SAP HANA system replication together with their `site name`; if the HANA database is offline, this host mapping cannot be shown on the secondaries.
- `active primary site` – shows the `site id` of the currently active site.
- `primary masters` – shows the hostname(s) of the currently active master candidates of the primary[32].

Starting with HANA 1.0 SPS12 and Rev112.03 when running "`hdbnsutil -sr_state`" on an offline HANA, no host mapping will be available anymore. Please refer to SAP note 2315257 for more details.

10.3 System Replication status checks

There are some ways to gather information about the overall status of the sites and of the system replication.

10.3.1 Using `landscapeHostConfiguration.py`

Check the overall status of the primary system using as `<sid>adm` OS user the script `landscapeHostConfiguration.py` (located in `$DIR_INSTANCE/ /exe/python_support`).

```
<sid>adm># python $DIR_INSTANCE/exe/python_support/landscapeHostConfiguration.py

| Host  | Host   | Host   |    ...    | NameServer | NameServer  |   ...
|       | Active | Status |           | Config Role| Actual Role |
| ----- | ------ | ------ | --------- | ---------- | ----------- | ------
| host1 | yes    | ok     |    ...    | master 1   | master      |   ...
| host2 | yes    | ok     |    ...    | master 2   | slave       |   ...

overall host status: ok
```

[32] IMPORTANT: In a Multitier System Replication on tier-3 the given "primary" is the tier-2 secondary – which from this perspective is the primary for this tier-3.

The following host states are possible:

- `OK`: System is OK.
- `WARNING`: A host auto-failover to a standby host is taking place.
- `INFORMATION`: The landscape is completely functional, but the current (actual) role of the host differs from the configured role.
- `ERROR`: There are not enough active hosts.

Please use the parameter "`--sapcontrol=1`", if you require a reliable and parsable output:

```
<sid>adm># python $DIR_INSTANCE/exe/python_support/landscapeHostConfiguration.py --
sapcontrol=1

   SAPCONTROL-OK: <begin>
   host/ld2131/hostActualRoles=worker
   host/ld2131/removeStatus=
   host/ld2131/nameServerConfigRole=master 1
   host/ld2131/failoverStatus=
   host/ld2131/hostConfigRoles=worker
   host/ld2131/failoverActualGroup=default
   host/ld2131/storageConfigPartition=1
   host/ld2131/host=ld2131
   host/ld2131/indexServerConfigRole=worker
   host/ld2131/failoverConfigGroup=default
   host/ld2131/storageActualPartition=1
   host/ld2131/indexServerActualRole=master
   host/ld2131/nameServerActualRole=master
   host/ld2131/hostActive=yes
   host/ld2131/hostStatus=ok
   host/ld2131/storagePartition=0
   host/ld2132/hostActualRoles=worker
   host/ld2132/removeStatus=
   host/ld2132/nameServerConfigRole=master 3
   host/ld2132/failoverStatus=
   host/ld2132/hostConfigRoles=worker
   host/ld2132/failoverActualGroup=default
   host/ld2132/storageConfigPartition=2
   host/ld2132/host=ld2132
   host/ld2132/indexServerConfigRole=worker
   host/ld2132/failoverConfigGroup=default
   host/ld2132/storageActualPartition=2
   host/ld2132/indexServerActualRole=slave
   host/ld2132/nameServerActualRole=slave
   host/ld2132/hostActive=yes
   host/ld2132/hostStatus=ok
   host/ld2133/hostActualRoles=standby
   host/ld2133/removeStatus=
   host/ld2133/nameServerConfigRole=master 2
   host/ld2133/failoverStatus=
   host/ld2133/hostConfigRoles=standby
   host/ld2133/failoverActualGroup=default
   host/ld2133/storageConfigPartition=0
   host/ld2133/host=ld2133
   host/ld2133/indexServerConfigRole=standby
   host/ld2133/failoverConfigGroup=default
   host/ld2133/storageActualPartition=0
   host/ld2133/indexServerActualRole=standby
   host/ld2133/nameServerActualRole=slave
```

```
host/ld2133/hostActive=yes
host/ld2133/hostStatus=ignore
overall_status=ok
SAPCONTROL-OK: <end>
```

10.3.2 Using systemReplicationStatus.py

Check the overall status of the system replication using as <sid>adm OS user the script systemReplicationStatus.py (located in $DIR_INSTANCE/ /exe/python_support).

```
<sid>adm># python $DIR_INSTANCE/exe/python_support/systemReplicationStatus.py
| Host   | Service Name     | Site Name | Secondary  |  ...  | Replication |...
|        |                  |           | Host       |       | Status      |
| ------ | ---------------- | --------- |  --------- | ----- | ----------- |---
| ld7805 | indexserver      | WALLDORF  | ld8475     |  ...  | ACTIVE      |
| ld8513 | statisticsserver | WALLDORF  | ld8476     |       | ACTIVE      |
| ld8513 | xsengine         | WALLDORF  | ld8476     |       | ACTIVE      |
| ld8513 | nameserver       | WALLDORF  | ld8476     |       | ACTIVE      |
| ld8513 | indexserver      | WALLDORF  | ld8476     |       | ACTIVE      |
| ld8559 | indexserver      | WALLDORF  | NOT MAPPED |       |             |

status system replication site "2": ACTIVE
status system replication site "3": ACTIVE
overall system replication status: ACTIVE

Local System Replication State
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

mode: PRIMARY
site id: 1
site name: WALLDORF
```

The script provides the following return codes.

```
10: No System Replication
11: Error
12: Unknown
13: Initializing
14: Syncing
15: Active
```

10.3.3 Using console

To check the replication status on all hosts and for all services the HDB console can be used. Especially in case of ASYNC replication this will provide some additional information currently not shown by the system view, because in this mode the primary does not await the acknowledgement upon arrival of the shipped redo log buffer.

In this case an option is to check the current log position on the secondary using hdbcons on the secondary side – where this information is not available via SQL – on each node and for each persistency relevant service[33]:

[33] The service can be passed to "hdbcons" by providing the service name with the option "-e <service>", like hdbnameserver, hdbindexserver, etc. In MultiDB HANAs, however, the service is to be passed via its PID using the option "-p

```
<sid>adm># hdbcons -e hdbindexserver "replication info"
   SAP HANA DB Management Client Console (type '\?' to get help for client
   commands)
   Try to open connection to server process 'hdbindexserver' on system 'M19',
   instance '19'
   SAP HANA DB Management Server Console (type 'help' to get help for server
   commands)
   Executable: hdbindexserver (PID: 66110)
   [OK]
   --
   listing default statistics for volume 3
   System Replication Secondary Information
   ========================================
   System Replication Secondary Configuration
    [system_replication] site_id                                    = 2
    [system_replication] site_name                                  = SiteA
    [system_replication] mode                                       = sync
    [system_replication] operation_mode                             = logreplay
    [system_replication] datashipping_min_logsize_threshold         = 5368709120
    [system_replication] datashipping_min_time_interval             = 600
    [system_replication] reconnect_time_interval                    = 30
    [system_replication] enable_log_compression                     = false
    [system_replication] preload_column_tables                      = true
    [system_replication] ensure_backup_history                      = true
    [system_replication] enable_ssl                                 = off
    [system_replication] keep_old_style_alert                       = false
    [system_replication] enable_log_retention                       = 1
    [system_replication] logshipping_max_retention_size             = 1048576

   Last Primary Host: ld2133
   Last Primary Port: 32003

   Log Connection
    - ptr                 : 0x00007fd58931a400
    - channel             : NetworkChannel FD 25 [0x00007fd5ad064a98]  {refCnt=3,
   idx=2} 10.96.4.20/65117_tcp->10.96.4.22/32003_tcp Connected,[r---]

    - mode                : ReplicationMode_Synchronous
    - logSinceLastBackup  : 663552 bytes
    - timeSinceLastBackup : 67431655 microseconds

   Data Connection
    - ptr                 : 0x00007fd589315000
    - channel             : NetworkChannel FD 31 [0x00007fd5ad064c58]  {refCnt=2,
   idx=3} 10.96.4.20/65118_tcp->10.96.4.22/32003_tcp Connected,[----]

   Secondary Statistics
    - Creation Timestamp              : 08.12.2015-14.25.27 (1449584727282603)
    - Last Reset Timestamp            : 08.12.2015-14.25.27 (1449584727282603)
    - Statistic Reset Count           : 0

    - ReplicationMode                 : sync
    - OperationMode                   : logreplay

    - ReplicationStatus               : ReplicationStatus_Active
    - ReplicationStatusDetails        :
    - ReplicationFullSync             : DISABLED
    - shippedLogPos                   : 0x641cbb00
    - shippedLogPosTimestamp          : 08.12.2015-14.59.17 (1449586757965706)
    - sentLogPos                      : 0x0
```

<PID_of_service>". You get the PID of the corresponding service for example running "HDB proc" as user <sid>adm on OS commandline.

```
- sentLogPosTimestamp              : 01.01.1970-00.00.00 (0)
- shippedLogBuffersCount           : 11241
- shippedLogBuffersSize            : 8335585280 bytes
- shippedLogBuffersSizeUsed        : 8309875456 bytes (99.69%)
- shippedLogBuffersSizeNet         : 8309875456 bytes (99.69%)
- shippedLogBufferDuration         : 0 microseconds
- shippedLogBufferDurationMin      : 0 microseconds
- shippedLogBufferDurationMax      : 0 microseconds
- shippedLogBufferDurationSend     : 0 microseconds
- shippedLogBufferDurationComp     : 0 microseconds
- shippedLogBufferThroughput       : 0.00 bytes/s

- replayFinishLogPos               : 0x641cbb00
- replayFinishLogPosTimestamp      : 08.12.2015-14.59.17 (1449586757965706)
- replayStartLogPos                : 0x641cbb00
- replayPushLogPos                 : 0x641cbb00
- replayRetentionLogPos            : 0x62a66fcb
- replayStepCount                  : 61709
- replayLogSize                    : 8335581056 bytes
- replayDuration                   : 111608005 microseconds

- shippedSavepointVersion          : 2252
- shippedSavepointLogPos           : 0x5c595f82
- shippedSavepointTimestamp        : 08.12.2015-14.25.28 (1449584728678668)

- shippedFullBackupCount           : 1
- shippedFullBackupSize            : 17884512256 bytes
- shippedFullBackupSizeNet         : 17884512256 bytes (100.00%)
- shippedFullBackupDuration        : 81098893 microseconds
- shippedFullBackupDurationComp    : 0 microseconds
- shippedFullBackupThroughput      : 220527205.67 bytes/s

- shippedLastFullBackupSize        : 17884512256 bytes
- shippedLastFullBackupSizeNet     : 17884512256 bytes (100.00%)
- shippedLastFullBackupStart       : 08.12.2015-14.25.28 (1449584728678668)
- shippedLastFullBackupEnd         : 08.12.2015-14.26.49 (1449584809777561)
- shippedLastFullBackupDuration    : 81098893 microseconds

- shippedDeltaBackupCount          : 0
- shippedDeltaBackupSize           : 0 bytes
- shippedDeltaBackupSizeNet        : 0 bytes (-nan%)
- shippedDeltaBackupDuration       : 0 microseconds
- shippedDeltaBackupDurationComp   : 0 microseconds
- shippedDeltaBackupThroughput     : 0.00 bytes/s

- shippedLastDeltaBackupSize       : 0 bytes
- shippedLastDeltaBackupSizeNet    : 0 bytes (-nan%)
- shippedLastDeltaBackupStart      : not set
- shippedLastDeltaBackupEnd        : not set
- shippedLastDeltaBackupDuration   : 0 microseconds

- Secondary sync'ed via Log Count : 0
- syncLogCount                     : 0
- syncLogSize                      : 0 bytes
- Secondary Backup History         : complete
- shippedMissingLogCount           : 0
- shippedMissingLogSize            : 0 bytes
- backlogSize                      : 0 bytes
- backlogTime                      : 0 microseconds
- backlogSizeMax                   : 0 bytes
- backlogTimeMax                   : 0 microseconds

- Secondary Log Connect time       : 08.12.2015-14.25.27 (1449584727296916)
- Secondary Data Connect time      : 08.12.2015-14.25.27 (1449584727491743)
- Secondary Log Close time         : not set
```

```
 - Secondary Data Close time        : not set
 - Secondary Log Reconnect Count    : 0
 - Secondary Log Failover Count     : 0
 - Secondary Data Reconnect Count   : 0
 - Secondary Data Failover Count    : 0
[OK]
--
[EXIT]
--
[BYE]
```

Here you get information about the used replication and operation modes (`mode`, `operation_mode`). You see which IP address is used for data and log transfer (`Log connection` and `Data connection`) and – since this system replication example is running with operation mode `logreplay` – you can see how far the log replay is hanging behind the shipped log on this secondary (the delta between `shippedLogPos` and `replayFinishLogPos`[34]). For all services the `ReplicationStatus` should be `ReplicationStatus_Active`.

10.3.4 Using predefined SQL statement

Attached to SAP Note 1969700 is a set of complex SQL statements – including some system replication relevant statements.

The statements contained in the text files can simply be copied and executed from some client.

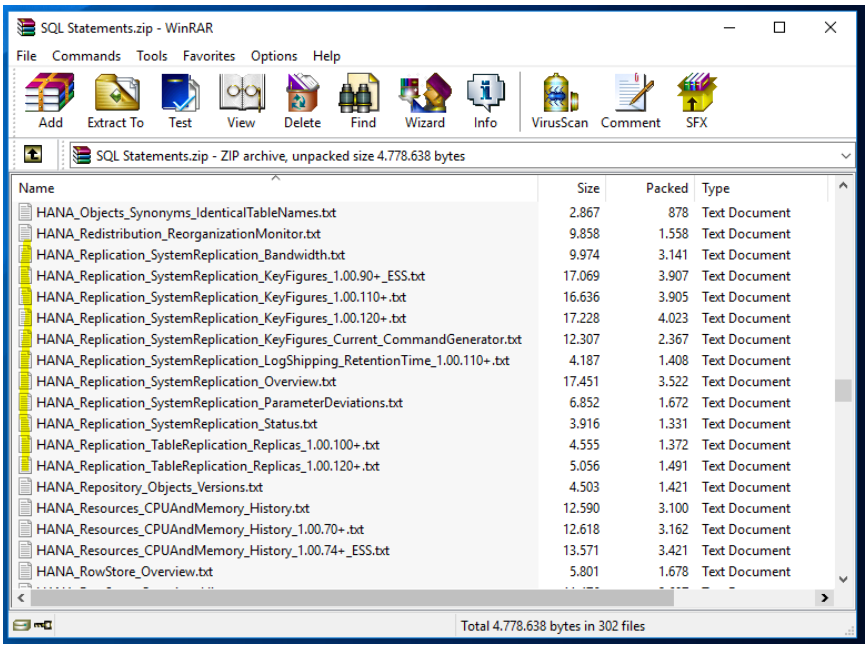

[34] For this specific topic please also refer to the SAP HANA system replication FAQ SAP Note 1999880 – question: *How can I determine the current log replay delay?*

In this example the "System Replication Bandwidth" SQL copied to and executed in the HANA cockpit "Execute SQL" application:

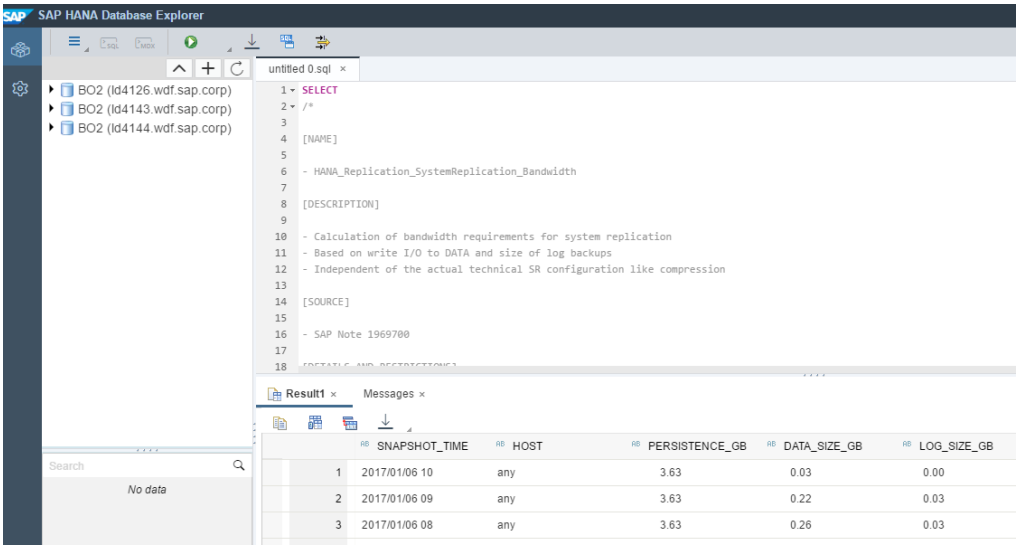

Additionally, the zip file can be imported to and executed in the SAP HANA studio as follows:

For the primary system go to the *System Information* tab and right-click on the "Name" column → *Import SQL Statements*.

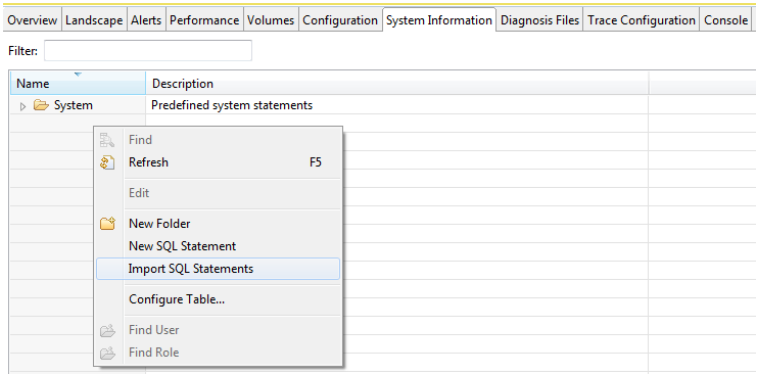

Select the "`SQL Statements.zip`" you downloaded from the SAP note:

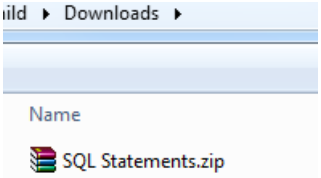

A folder with the SQL statements will be imported. Right-click on the statements under *Replication* → *Overview* and Executed a statement – for example the Overview:

Replication	
SystemReplication	
KeyFigures	
Current	
CommandGenerator	- Collects load and throughput figures for SAP HANA system replication
Rev90+	
ESS	- Current and historic system replication key figures
Rev110+	- Current and historic system replication key figures
LogShipping	
RetentionTime	
Rev110+	- Continuous log shipping retention time calculation
Bandwidth	- Calculation of bandwidth requirements for system replication - Based on write I/O to DATA and size of log backups - Independent of the actual technical SR configu
Overview	- General system replication information
ParameterDeviations	- Displays statistic server information for SAP HANA parameter deviations between primary and replication side - Section "ALERT CONFIGURATION" contains configu
Status	- Status of system replication of individual services

You will receive a lot of information about the system replication landscape and the per service replication state:

Raw Data | Distinct values

Filter pattern | 103 rows retrieved - 116 ms

REPLICATION_PATH	HOSTS	SERV...	KEY	VALUE
SITEA -> SITEB	ld2131 -> ld2132	indexserver	Replication mode	SYNCMEM
			Secondary connect time	2014/05/02 10:50:39
			Days since secondary connect time	0.02
			Used persistence size (GB)	4.98
			Log backup size / day (GB)	1.50
			Local log buffer write size (MB)	4623.18
			Shipped log buffer size (MB)	3593.53
			Avg. local log buffer write size (KB)	803.62
			Avg. shipped log buffer size (KB)	819.55
			Avg. local log buffer write time (ms)	6.60
			Avg. log buffer shipping time (ms)	6.71
			Local log buffer write throughput (MB/s)	118.86
			Log buffer shipping throughput (MB/s)	119.13
			Initial data shipping size (MB)	1122.84
			Initial data shipping time (s)	4.28
			Last delta data shipping size (MB)	1.01
			Last delta data shipping time (s)	0.00
			Delta data shipping size (MB)	3.24
			Delta data shipping time (s)	0.36
			Delta data shipping throughput (MB/s)	8.81
			Delta data shipping size / day (MB)	129.42
SITEA -> SITEB	ld2131 -> ld2132	nameserver	Replication mode	SYNCMEM
			Secondary connect time	2014/05/02 10:50:34
			Days since secondary connect time	0.02
			Used persistence size (GB)	0.00
			Log backup size / day (GB)	0.00

Of interest are for example the "Local log buffer write throughput (MB/s)" compared to the "Log buffer shipping throughput (MB/s)" in synchronous replication. For synchronous replication, this could for example be an indication for network problems or a problem with the I/O on the secondary side (for SYNC), if these two values differ too much.[35]

[35] For details on this please refer to the HowTo https://scn.sap.com/docs/DOC-58553 on system replication network configuration.

Tipp
SAP Note 1999993 - How-To: Interpreting SAP HANA Mini Check Results was extended with the SAP HANAChecker Python script that can be used to automatically execute the mini checks and send e-mails in case of potentially critical issues.

10.4 Monitoring and replicating `ini` parameter changes

`ini` parameters basically should be the same on the primary and secondary system and are checked automatically. The *configuration parameter checker* reports differences between primary and secondary if parameters differ.

In the "Alerts" tile on the system overview page of HANA cockpit you can click on "Show All" to see all alerts – including the ones created for "parameter mismatch":

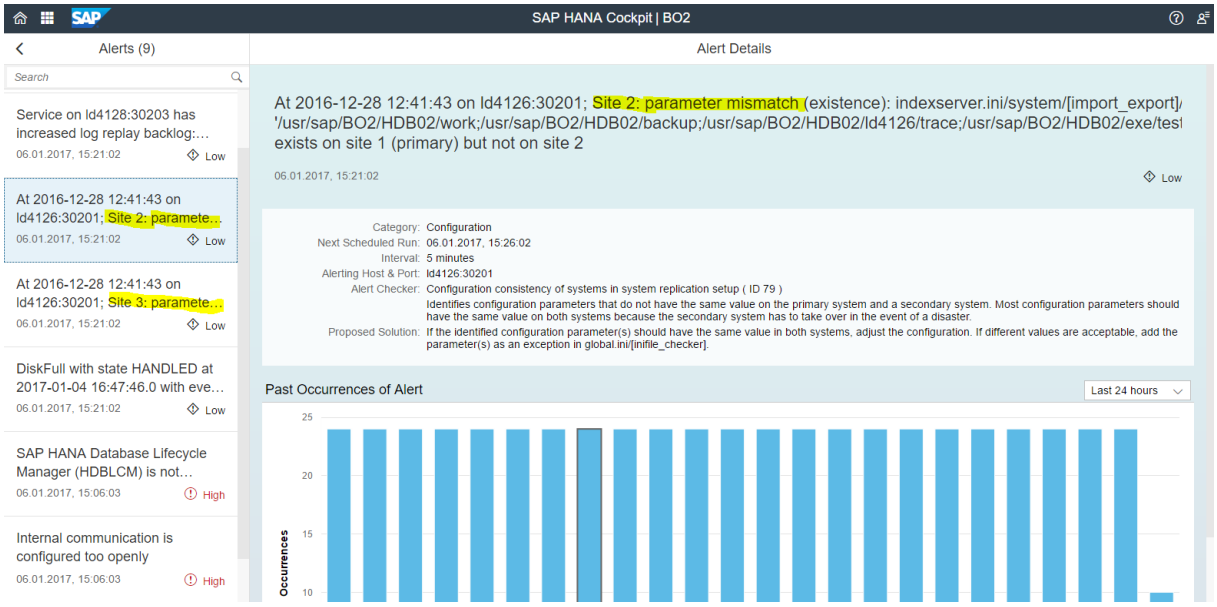

Also in the "Alerts" tab in HANA studio you can see the generated alerts:

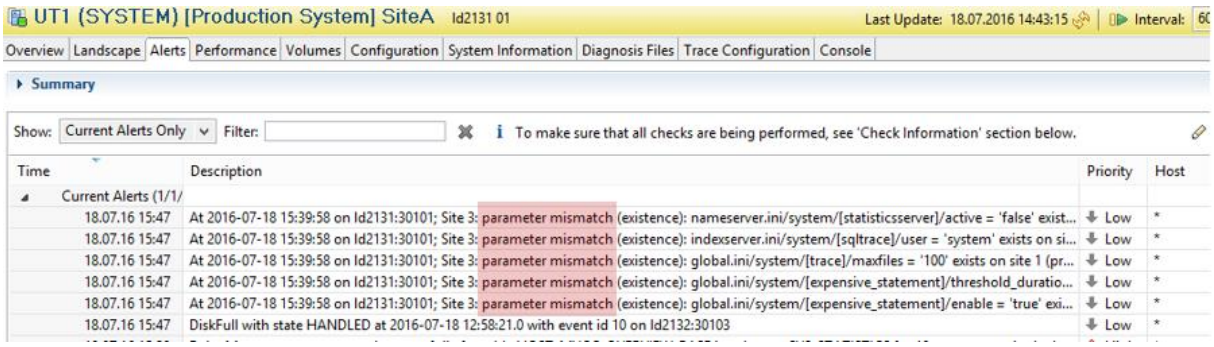

Since HANA 1.0 SPS12 you can activate `ini` parameter replication, where changes made on the primary are automatically replicated to the secondary sites.

`ini` parameter checks:

- Are done every hour by default.
- Generate alerts (visible in SAP HANA studio as internal event and the system view M_EVENTS)[36].
- Are optimized for the most recently changed parameters.

Enable and disable the `ini` parameter check on the primary site with
`[inifile_checker]/enable = true|false` (default: `true`)

Enable and disable the `ini` parameter replication on the primary site with
`[inifile_checker]/replicate = true|false` (default: `false`)

The `ini` parameter replication follows these rules:

Parameter set on		Action
Primary	Secondary	
Set	Not set	Copy parameter to secondary
Not set	Set	Delete parameter on the secondary
Set to value x	Set to value y	Copy value x to secondary

In the `global.ini` an exclusion list can be maintained to exclude parameters to be checked or replicated. Just follow the syntax to define exclusion rules for certain parameters:

```
exclusion_[inifile name|*][/] = [section with
wildcards|*][/parameter with wildcards|*], ... :=
SYSTEM\|HOST\|TENANT\|\*
```

If for example you intend to use your secondary system for DEV/QA systems and set the global allocation limit to its minimal value (as described above), you may exclude this parameter `global_allocation_limit` from these checks like this:

```
[inifile_checker]
enable = true
interval = 3600
exclusion_global.ini/SYSTEM =
memorymanager/global_allocation_limit
```

[36] How you can setup an e-mail notification is described in the SAP HANA Admin guide in section: *Configure E-Mail Notifications for Alerts*.

10.5 Monitoring the secondary site

Since HANA 1.0 SPS11 there is a possibility to monitor the secondary site using so-called proxy views. They provide remote SQL access on the primary – through *proxy schemas and views* – allowing for monitoring and reporting of secondary site statistics (for any replication mode). During registration of a secondary system, the new proxy schema on the primary site is created for each registered secondary site. The schema follows the naming convention `_SYS_SR_SITE_<siteName>` and contains a selected subset of monitoring views, which proxies the statistics from the secondary site. Proxy views have the same column definitions as the equivalently named public synonyms already available for the primary. When a secondary site is unregistered the corresponding schema will be dropped.

In the HANA cockpit on the system overview page for the primary click on "*Execute SQL*" to get to the SAP HANA Database Explorer. Then open the primary's "*Catalog*" and go to the corresponding schema:

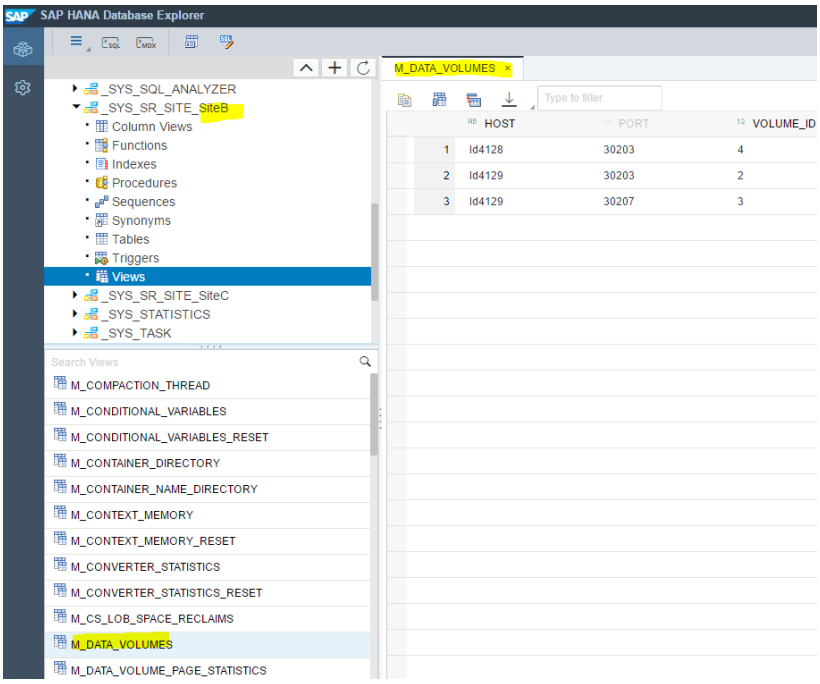

Proxy views of the secondary site's monitoring views (HANA cockpit)

In the HANA studio in the landscape overview just open the *Catalog* and the corresponding schema:

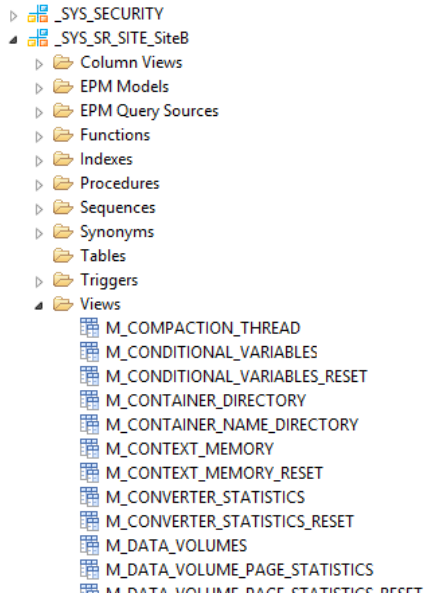

Proxy views of the secondary site's monitoring views (HANA studio)

There are some limitations of proxy views which need to be considered:

- Monitoring view access is only possible if primary and secondary sites run with the same software version.
- When such a proxy view is queried against and the secondary site is not started, no results are shown without the report of an SQL error.
- Querying against *SAP HANA multitenant database containers* landscapes is limited to single tenant databases or the system database, meaning there are no views unifying all tenants on the system database similar to the `SYS_DATABASES` schema.

As of HANA 2.0 SPS00 and only if the system replication is configured as Active/Active (read enabled) system using the operation mode "`logreplay_readaccess`" even more data of the secondary systems are available in the proxy schema `_SYS_SR_SITE_<siteName>`. Making use of the read access to the secondary system many of the other monitoring views of the secondary can be made accessible via virtual tables.

In the SAP HANA Database Explorer for the primary system (after clicking on "Execute SQL" on the system overview page) open the proxy schema and click on the "*tables*"item. A long list of accessible monitoring views from the secondary is shown:

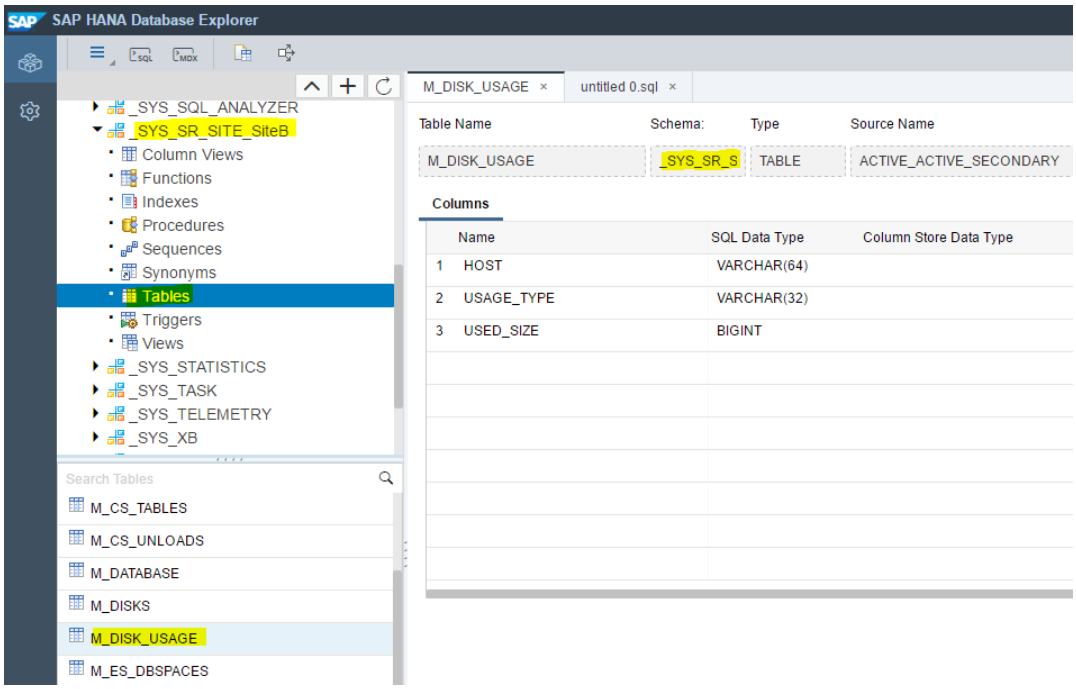

Virtual tables available in the proxy schema of the secondary system (HANA cockpit)

Equally you can access these virtual tables with the SAP HANA studio:

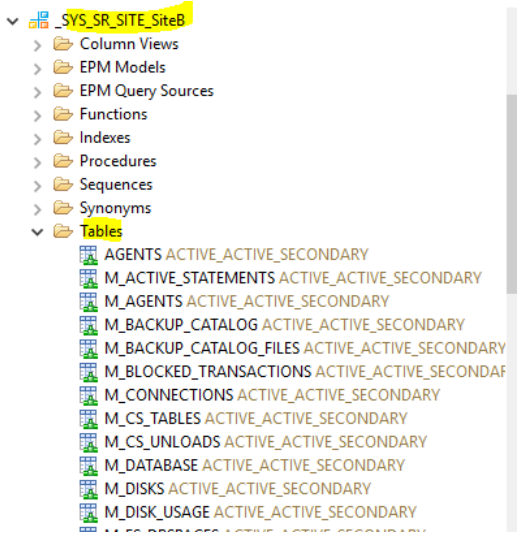

Virtual tables available in the proxy schema of the secondary system (HANA studio)

Any of these proxy views or virtual tables can be accessed via SQL from the primary simply by providing the correct secondary's schema name, for example:

```
select * from "_SYS_SR_SITE_SiteB"."M_HOST_INFORMATION";
```

Note
Based on these views and tables available in the proxy schema the statistics server is able to generate alerts on the secondary sites (identified by `host:port`) of a system replication landscape.

10.6 System replication network connection

The replication in a configured system replication uses either a public or a separate network channel between the involved data centers.[37]

10.6.1 Secure configuration of the network connection

By default, the primary and secondary systems establish communication using the internal host names[38].

With an `IPaddress-virtualHostname` mapping on the involved sites the system replication hostname resolution can be set configuring a separate network for system replication data traffic between primary and secondary[39].

This is done in the section `[system_replication_hostname_resolution]` in `global.ini`, where all hosts of the primary and the secondary sites must be defined on each site:

```
global.ini/[system_replication_hostname_resolution]
   <ip-address_same_site>=<internal_host_same_site>
   <ip-address_other_site>=<internal_host_other_site>
```

This also holds valid for a multitier system replication consisting of three sites (primary, tier-2 secondary and tier-3 secondary) because roles can switch after takeovers and failbacks.

Note
The parameters in the `global.ini` file must be set prior to registering the secondary system, because the `hdbnsutil -sr_register` command uses this mapping. Registration is one step in the process of configuring the secondary system.

The entries in the `[system_replication_hostname_resolution]` section are used in combination with the `listeninterface` parameter in the `[system_replication_communication]` section. The following combinations are possible:

[37] Please also check the HowTo guide on system replication network configuration https://scn.sap.com/docs/DOC-58553
[38] All SAP HANA system views containing a `HOST` column show these internal host names, e. g. `M_DATABASE`.
[39] As mentioned, in Multitier System Replication the tier 2 secondary serves as primary for the replication to the tier 3 secondary.

[system_replication_communication] listeninterface	[system_replication_hostname_resolution]	Additional Information
.global	No mappings specified	Default if nothing is specified. The default network route is used for system replication communication. This is normally the public network. Caution If you use a public network instead of a separate network, you must secure this connection with additional measures such as a firewall or a virtual private network and/or SSL.
.global	Entries for the primary and secondary hosts (for all hosts in multitier setups)	A separate network is used for system replication communication. Tip This way you can use a separate network for multitier system replication.
.internal	Entries for the primary and secondary hosts	A separate network is used for system replication communication. The primary hosts listen on the dedicated ports of the separate network only and incoming requests on the public interfaces are rejected. Caution In SAP HANA 1.0 SPS 11, network communication for system replication with `listeninterface=.internal` is supported for two-tier replication but not for three-tier setups!

There are two ways to activate `[system_replication_hostname_resolution]` in your system:

1. Restart all sites after setting the parameter
2. Temporarily resolve the system replication configuration – here no restart of the primary is necessary:
    a. Stop secondary
    b. Unregister secondary
    c. Disable primary
    d. Enable primary
    e. Register secondary
    f. Start secondary

Here is an example of the settings for a 2-tier system replication (3 node system) using a separate internal network per site and a separate connection for the system replication.

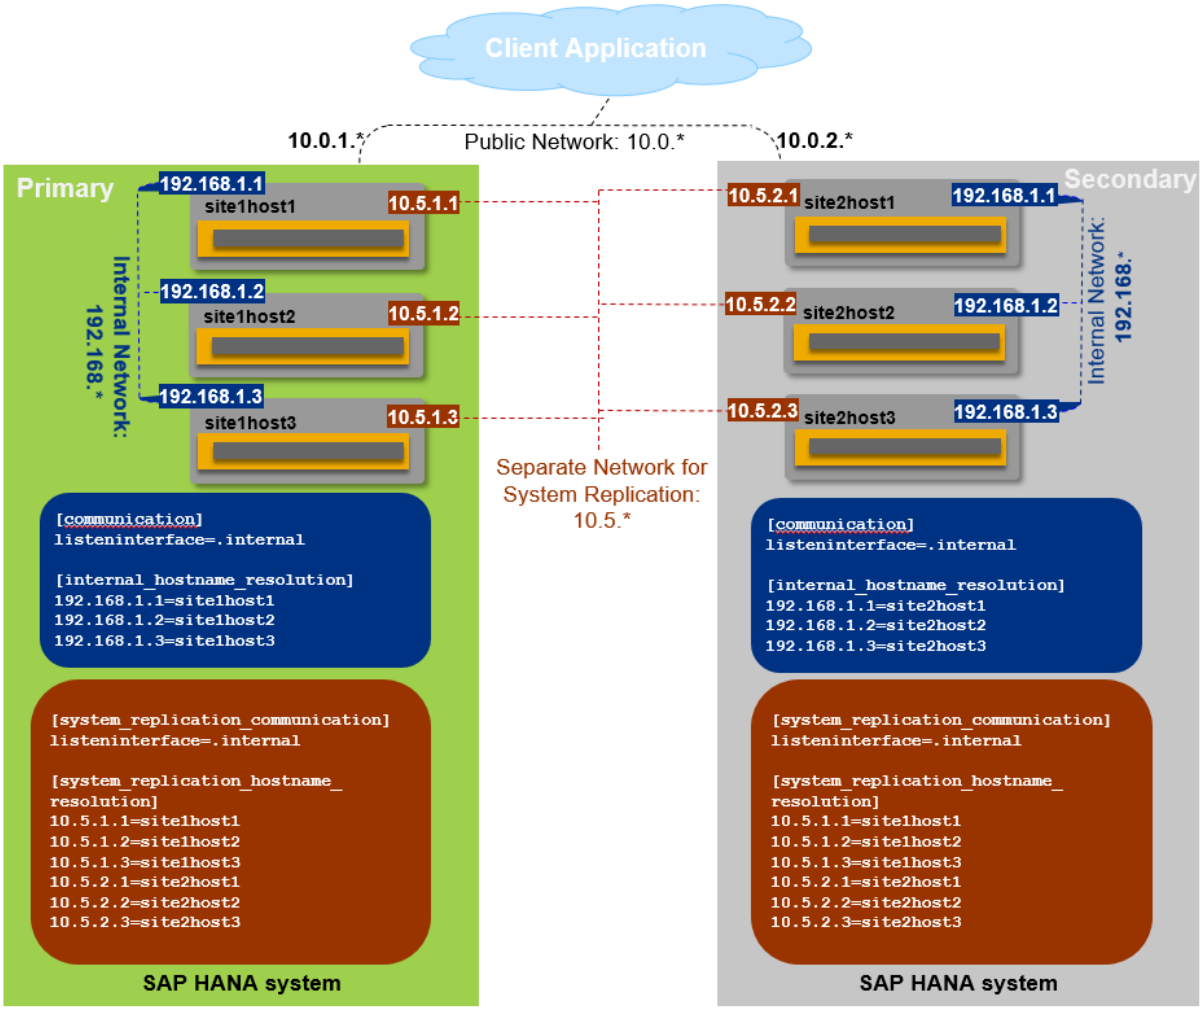

Multi-node SAP HANA System Replication over separate network with separate internal network

10.6.2 Allowed senders for the connection

If for some reason no separate network channel was configured for the SAP HANA system replication communication between the involved sites, the parameter `allowed_sender` could be used to restrict communication between primary and secondary to certain hosts. For this, the following settings can be configured in the `global.ini` file on the primary site:

`global.ini/[system_replication_communication]`	
Parameter:	`allowed_sender`
Value:	`<list of IP-addresses of secondary or CIDR-netmasks>`
Example:	`10.0.1.0/30`

The default is no restriction.

10.6.3 Encryption of the connection

With HANA 2.0 SPS02 the system replication connections (for data, redo log and metadata of the nameserver) are secured with `SSL` by `systemPKI` without the need to switch on `SSL` for the internal communication, as required in previous versions.[40] Now all system replication connections are specified only by setting `[system_replication_communication]/enable_ssl`. This simplifies security setting for system replication.

SAP HANA System Replication supports secure network communication (SSL) for data and log shipping to the secondary site. The following settings can be configured in the `global.ini` file:

`global.ini/[system_replication_communication]`
Parameter: `enable_ssl`
Values:
- `off`: ssl is disabled for source and target replication channels (default)
- `on`: ssl is enabled for source and target replication channels
- `source`: ssl is enabled as source replication channel only
- `target`: ssl is enabled as target replication channel only

In this context, a new port `4xx06` in system replication communication was added – covering encrypted metadata communication of the nameserver; now the following ports are used:

TCP Port	Service	Used For
4xx01	nameserver	Log and data shipping (System DB)
4xx02	nameserver	Metadata communication (System DB) – unencrypted
4xx06	nameserver	Metadata communication (System DB) – encrypted via SSL
4xx40 - 4xx97	indexserver	Log and data shipping (tenant DBs)
4xx40 - 4xx97	scriptserver	Log and data shipping

[40] Before HANA 2.0 SPS02, the `[system replication communication]/enable ssl` parameter only was valid for the data and the log connection. The connection for the nameserver-metadata ignored this setting, and was only affected by the internal communication SSL setting of `global.ini/[communication]/ssl=on` or `systempki`.

TCP Port	Service	Used For
		(optional, tenant DBs)
`4xx40 - 4xx97`	`docstore`	Log and data shipping (optional, tenant DBs)

Caution
Before you plan to upgrade HANA from a previous version to HANA 2.0 SPS02 or higher using Near Zero Downtime Upgrade (NZDU), please check SAP Note 2494079 - Near-Zero-Downtime-Upgrade to HANA 2 SPS02 or above when internal communication SSL is used.

10.6.4 Network security settings in HANA cockpit

With HANA 2.0 SPS02 the current settings for the system replication network security can be viewed from within the system replication overview application by clicking the "*Network Security Settings*" tab:

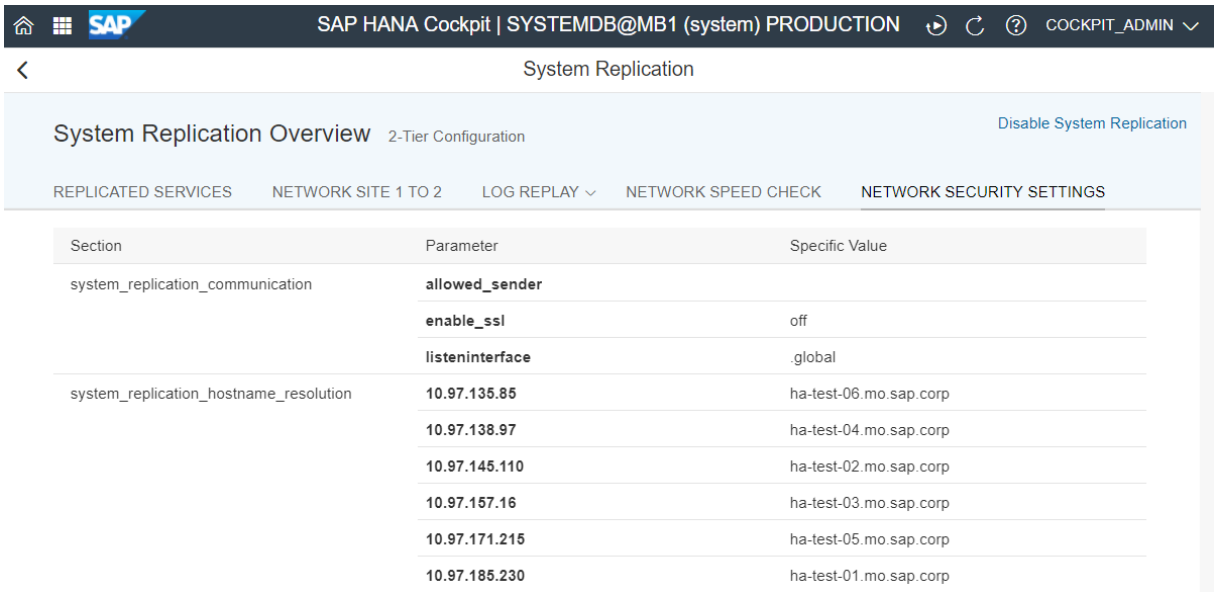

The security settings – described above in this document – are shown for your system.

10.6.5 Data and Log compression for transfer

Since HANA 1.0 SPS09 compression can be configured to reduce the amount of traffic between sites especially over long distances. It will be used for the initial full data shipping, the sub sequential delta data shipping as well as for the continuous log shipping.
Configuration is done in `global.ini` on the secondary site.

```
[system_replication]
enable_log_compression = true (default = false)
enable_data_compression = true (default = false)
```

By default, content compression is turned off; log buffer tail compression (default = true) and log buffer content compression can be combined.

10.6.6 Monitoring the network connection

The connection between the primary and the secondary system should be available for replication. If this is not the case for a certain time, the redo log cannot be shipped to the secondary system, the log segments start piling up on the primary (see Log Retention section), the secondary system is not takeover ready.

For the primary to stay operational at all times, even if the connection is lost occasionally, an internal event is generated which is visible as an internal event alert in the SAP HANA cockpit, the SAP HANA studio and in the system view `M_EVENTS` (if the old – not embedded – statistics server is used). With HANA 1.0 SPS09 system replication a specific alert was introduced (as described above in the "*Alerts*" section):

- System Replication Connection Closed (Alert ID 78)

Additionally, the replication connection can be checked using the HDB console:

```
hdbcons -e hdbindexserver "replication info"
```

The output delivers "Log Connection" information for the connection used by the provided service. It also shows errors if the connection cannot be resolved properly:

```
…
Log Connection
- ptr                  : 0x00007fdb6e8e3410
- channel              : NetworkChannel FD 158
[0x00007fdb6f1bbc90]  {refCnt=3, idx=1} 10.68.91.226/3
0103_tcp->10.68.92.13/49537_tcp Connected,[r---]
...
```

Use the OS command

```
lsof -n -p <indexserver-pid>
```

to check, if the configured connection is actually used. The output delivers "Log Connection" information for the connection used by the provided service.

For a more detailed analysis of the network connection used for system replication refer to the "*Troubleshoot System Replication*" chapter in the SAP_HANA_Troubleshooting_and_Performance_Analysis_Guide_en.pdf; also SAP Note 2081065 - Troubleshooting SAP HANA Network provides some information on SAP HANA network topics.

10.6.7 Monitoring the network latency

Only the latency for the *redo log shipping* is of interest in a running system replication and thus it is independent of the operation mode used; additionally, only for synchronous replication modes `SYNC` or `SYNCMEM` it plays a role. The following requirement should be met:

> *Latency*
> The redo log shipping wait time for 4 KB log buffers must be less than a millisecond or in a low single-digit millisecond range – depending on the application requirements (relevant for synchronous replication modes only).

All changes to data are captured in the redo log. The SAP HANA database asynchronously persists the redo log with I/O orders of 4 KB to 1 MB size into log segment files in the log volume (i. e. on disk) of the primary system. A transaction writing a commit into the redo log waits until the buffer containing the commit has been written to the log volume. This wait time for 4 KB log buffers should be less than a millisecond or in a low single-digit millisecond range.

Once you have set up SAP HANA system replication, you can check the local and the remote *log write wait times* in HANA cockpit or collect it with the SQL statement *HANA_Replication_Overview* (attached to SAP Note 1969700).

SAP HANA Cockpit 2.0

Click on the *System Replication tile* on the SystemDB of your primary system. As mentioned in this document in the section "*Verification – Using SAP HANA Cockpit*" you get an overview of your running system replication. Between the primary and the secondary the average log buffer size as well as the average redo log shipping wait time is shown.

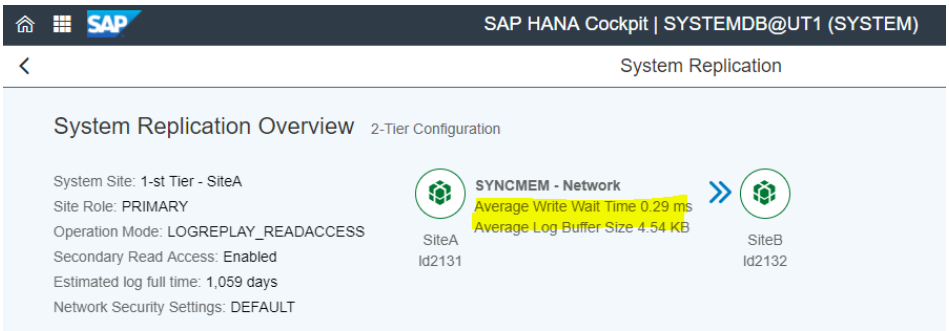

In this system replication cockpit app you also can have a look at the 24 hour history of these local log write wait times and the remote log shipping wait times:

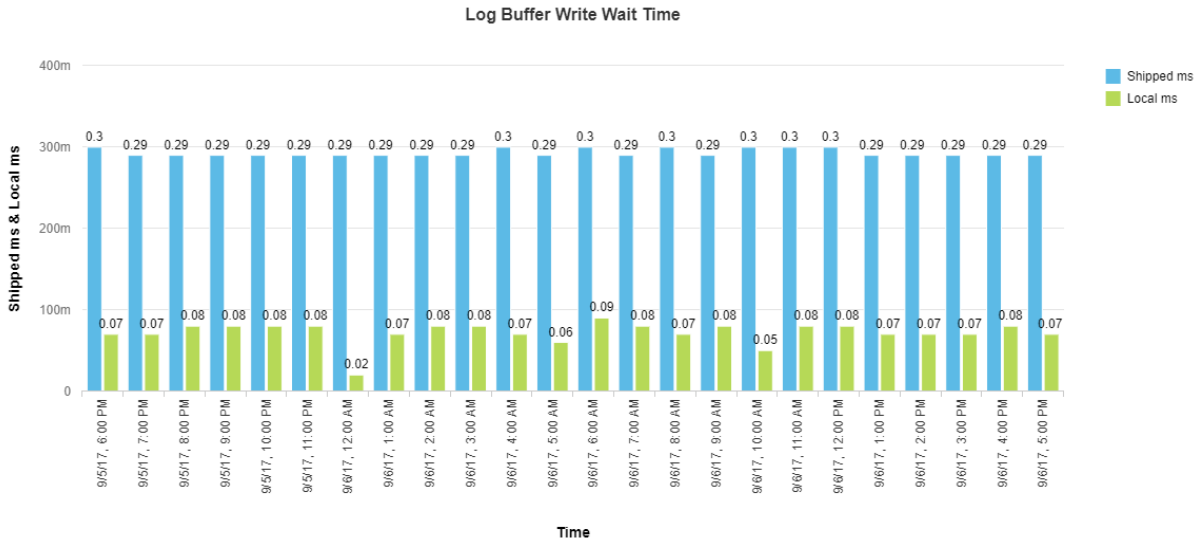

SAP HANA studio

You can gather these values also from SAP HANA studio.

Right-click on the statements under *Replication* → *Overview* and select *Execute* from the context menu:

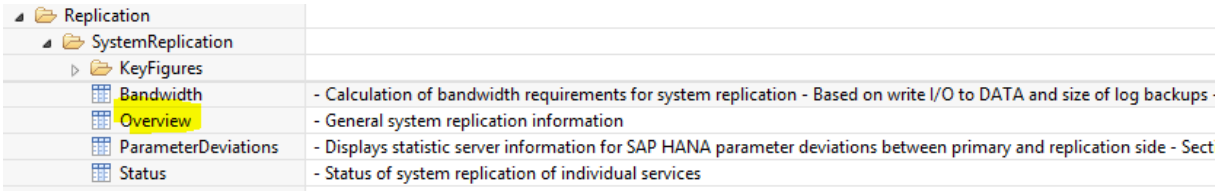

You will receive a lot of information about the system replication landscape and the per service replication state:

REPLICATION_PATH	HOSTS	PORT	KEY	VALUE
SiteA -> SiteB	Id2131 -> Id2132	30103	Replication mode	SYNCMEM
			Secondary connect time	2017/08/24 09:31:15
			Days since secondary connect time	13.32
			Used persistence size (GB)	1.99
			Log backup size / day (GB)	0.75
			Local log buffer write size (MB)	9850.19
			Shipped log buffer size (MB)	7758.82
			Avg. local log buffer write size (KB)	5.72
			Avg. shipped log buffer size (KB)	4.52
			Avg. local log buffer write time (ms)	0.60
			Avg. log buffer shipping time (ms)	0.29
			Local log buffer write throughput (MB/s)	9.28
			Log buffer shipping throughput (MB/s)	14.80
			Initial data shipping size (MB)	1584.00
			Initial data shipping time (s)	10.11
			Last delta data shipping size (MB)	0.00
			Last delta data shipping time (s)	0.00
			Delta data shipping size (MB)	0.00
			Delta data shipping time (s)	0.00
			Delta data shipping throughput (MB/s)	0.00
			Delta data shipping size / day (MB)	0.00
			Replication delay (s)	0.00

Local redo log write wait times

The log write wait times on the "local" system (i.e. on the persistence of the primary site) is returned by the mentioned SQL statement as "*Avg. local log buffer write time (ms)*" per service, where the index server is the one of interest.

System replication redo log shipping wait time

Additionally, the mentioned SQL statement returns the redo log write latency for the shipping to the secondary which can be slightly higher than the locally measured log write wait time. The returned value "*Avg. log buffer shipping time (ms)*" represents the time between enqueueing and finishing a request.

Also of interest are the "*Local log buffer write throughput (MB/s)*" compared to the "*Log buffer shipping throughput (MB/s)*" in synchronous replication. For synchronous replication, this could be an indication for network problems or a problem with the I/O on the secondary site (for `SYNC`), if these two values differ too much.

> **Tipp**
>
> For further information on `ASYNC` replication see section "*Check Network Configuration (Long Distance)*" in the SAP HANA Troubleshooting and Performance Analysis Guide.

10.6.8 Network Speed Checks for System Replication

To measure the network speed for the system replication host-to-host-mappings from the primary system to the secondary system (and even to a third tier) can be done from the HANA Cockpit 2.0 system replication application as of HANA 2.0 SPS02.

In the system replication application of your primary system click on the “Network Speed Check” tab. To measure the network speed (MB/s) choose a Packet Size from the pull-down menu and afterwards click on “Measure Network Speed”:

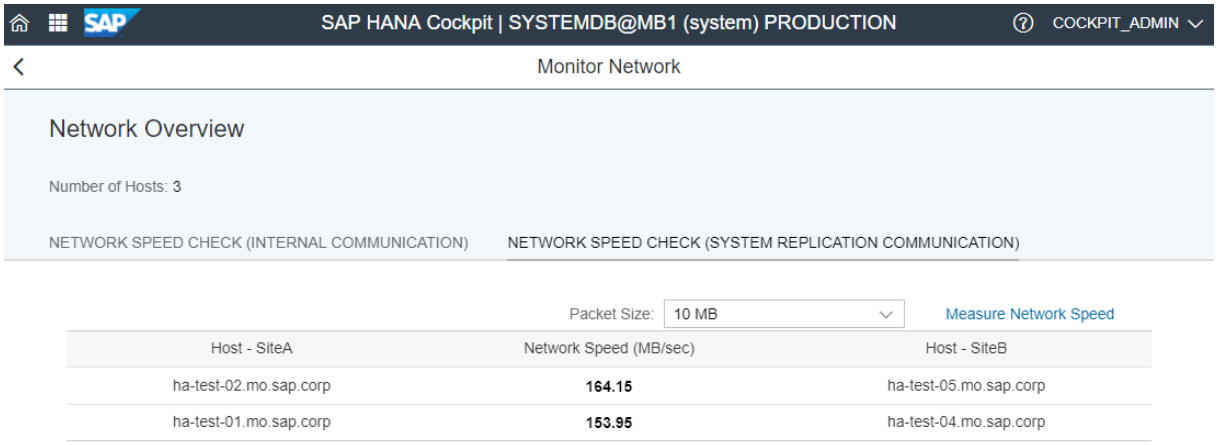

In this example, there are standby hosts configured on the primary (ha-test-03) and the secondary (ha-test-06) system. The network speed measurement cannot be done for standby hosts, because they do not have access to the data and log volumes of the HANA database and thus are not relevant for replication in this state.

> **Caution**
> Measuring the network speed of your SAP HANA system replication communication channels will have an impact on your network performance and thus on your running systems for the time of the measurement.

10.6.9 Network Speed Checks for Scale-out System

Although this is not a HANA system replication feature, it should be mentioned at this point. There is also a network speed check for scale-out HANA databases. It measures the network speed in both directions for each host to each other host of the scale-out system.

To run this network speed check, click on the *Monitor Network* link in the SAP HANA cockpit system overview page. This takes you to the *Monitor Network* page. Select a Packet Size from the scroll-down menu and click on *Measure Network Speed*:

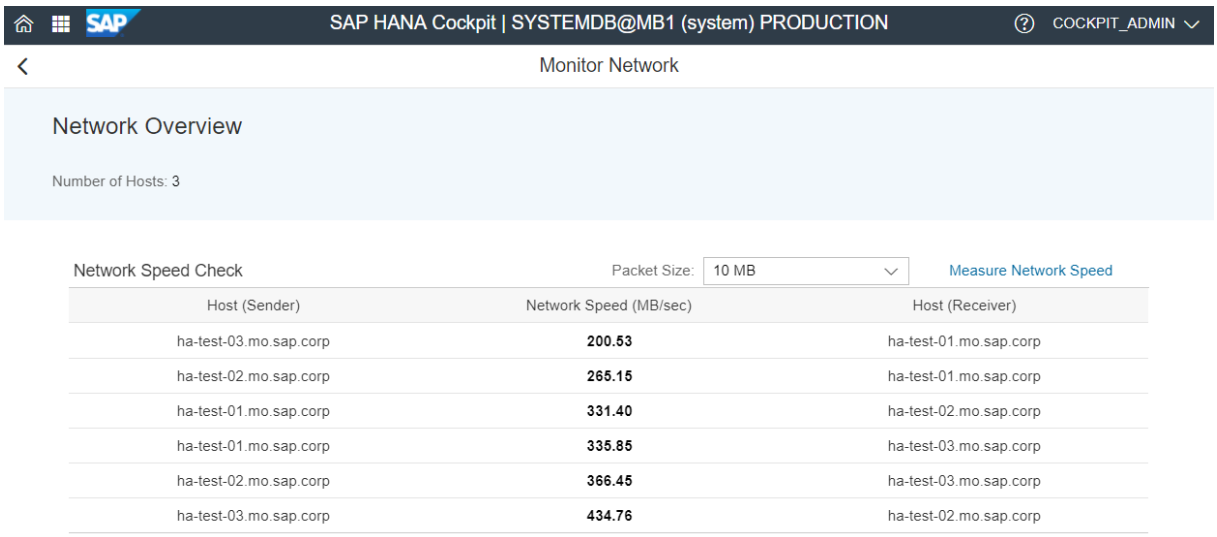

Host (Sender)	Network Speed (MB/sec)	Host (Receiver)
ha-test-03.mo.sap.corp	200.53	ha-test-01.mo.sap.corp
ha-test-02.mo.sap.corp	265.15	ha-test-01.mo.sap.corp
ha-test-01.mo.sap.corp	331.40	ha-test-02.mo.sap.corp
ha-test-01.mo.sap.corp	335.85	ha-test-03.mo.sap.corp
ha-test-02.mo.sap.corp	366.45	ha-test-03.mo.sap.corp
ha-test-03.mo.sap.corp	434.76	ha-test-02.mo.sap.corp

The result is a list representing the network speed for every host in the scale-out system to every other host of the system. The measurement is done in both directions; thus, every host-mapping shows up twice in this list – with switched roles of "Sender" and "Receiver". On top of the list, the fastest connection is shown.

> **Caution**
> Measuring the network speed of your internode communication channels in your SAP HANA scale-out system will have an impact on your network performance and thus on your running systems for the time of the measurement.

10.7 Tenant Copy/Move within System Replication

As of HANA 2.0 SPS02 it is possible to do a tenant copy or move operation in a SAP HANA tenant database system into or out of a primary system running in a SAP HANA system replication configuration.[41]

[41] Before either your system was configured in system replication OR you could do a tenant copy / move – not both at the same time.

There are three tenant copy/move scenarios that are applicable in a primary system. The following table describes the scenarios and briefly explains, what happens to the copied / moved tenant with regards to replication

Scenario	Tenant copy / move behavior
Copy / Move Tenant **Into Primary** 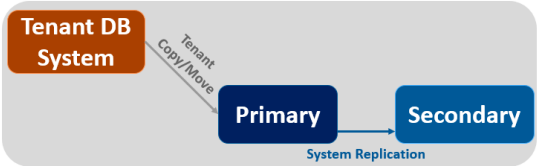 	A tenant database is copied/moved into the primary system of a system replication configuration. The copied / moved tenant starts replicating after the FINALIZE command was executed.
Copy/Move tenant **from primary to a Tenant Database System** (≠ secondary) 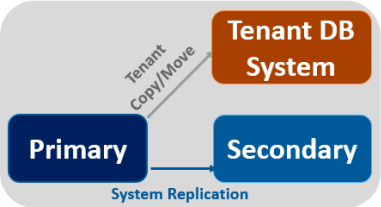 	A tenant database is copied/moved out of the primary system into another target Tenant Database than the Secondary. The copied / moved tenant arrives in the target Tenant Database System and is runnable after the FINALIZE command was executed.
Copy/Move tenant **within primary** 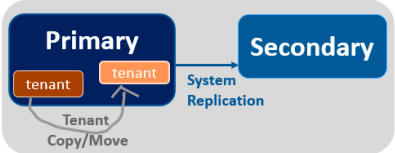 	A tenant database is copied/moved within the primary system of a system replication. The "cloned" tenant is created in primary system and replication starts after the FINALIZE command was executed.

Further information on SAP HANA tenant database systems – and the described tenant copy / move feature can be found in the *SAP HANA Tenant Databases guide*: Copying and Moving Tenant Databases Between Systems

10.8 Data and Log Volume Encryption

Before HANA 2.0 SPS02 data volume encryption was enabled by setting the `ini` parameter `global.ini/persistence/data_encryption=true`. And depending on the configuration of the `inifile` checker (see section *Monitoring and replicating `ini` parameter changes* in this document) either this change was automatically replicated to the secondary (for the setting `globel.ini/[inifile_checker]=replicate`) or you had to set it manually there as well.

As of HANA 2.0 SPS02 in a system replication configuration, you must change the root keys used for data volume encryption and log volume encryption[42] only in the primary system. The new keys will be propagated to all secondary systems. Also, enable/disable encryption in the primary system only. The setting will be propagated to all secondary systems as well. For details on enabling data and log volume encryption please refer to the SAP HANA Administration Guide – section: *Enable Data and Log Volume Encryption in an Existing SAP HANA Database*.)

10.9 Upgrade and Maintenance

If for some reason you have to stop and restart the primary or the secondary, once the systems are available they will automatically try to get in sync again. There are no manual steps necessary.

Caution

If the system is running with operation mode "`logreplay`" or "`logreplay_readaccess`" please check the section on "Log Retention" in this document to prevent your system from running full or having to do a full data shipping, in case the time the primary could not replicate gets too long.

Note

To avoid a full data shipping after the upgrade to get both systems in sync again, refer to the sections "Data Retention" and "Log Retention". Depending on the settings of the `logshipping_max_retention_size` and `datashipping_snapshot_max_retention_time` parameter settings and the time the upgrade is taking, an optimized resync with delta data or log shipping can be achieved – avoiding a full data shipping.

System Replication with SAP HANA 2.0 requires authentication for data and log shipping channels. The authentication is done using the certificates in the system PKI SSFS store. Thus, there is an additional manual setup step required to exchange certificates in the system PKI SSFS store between primary and secondary site when upgrading to HANA 2.0 (please refer to SAP Note 2369981 as well.)

[42] Please refer to the SAP HANA Administration Guide to find out how to "Change the Root Key for Data Volume Encryption in SAP HANA Cockpit".

10.9.1 "Normal" Upgrade

Since the version of the secondary system must be higher or the same as the one running on the primary system you always must upgrade the secondary first.

The process is straightforward:

1. Upgrade the secondary system and wait until the upgrade is done
   ```
   ./hdblcm --action=update
   ```
2. Upgrade the primary system and wait until the upgrade is done
   ```
   ./hdblcm --action=update
   ```

Caution

If you are upgrading from HANA 1.0 to HANA 2.0. proceed as follows instead:

1. Upgrade the secondary system and wait until the upgrade is done
   ```
   ./hdblcm --action=update --hdbupd_server_nostart
   ```
2. Upgrade the primary system and wait until the upgrade is done
   ```
   ./hdblcm --action=update
   ```
3. Copy the `systemPKI` SSFS key and data file from the primary to the secondary; these files can be found here:
   ```
   /usr/sap/<SID>/SYS/global/security/rsecssfs/data/SSFS_<SID>.DAT
   /usr/sap/<SID>/SYS/global/security/rsecssfs/key/SSFS_<SID>.KEY
   ```
4. Start the secondary system.

After both systems are available again they will get in sync automatically.

10.9.2 Near Zero Downtime Upgrade

System replication can be used for a Near Zero Downtime Upgrade (NZDU), where first the secondary is upgraded, then it takes over operation from the original primary (via takeover), and afterwards the primary is upgraded. In a failback, the original primary is attached as new secondary to re-establish the system replication with changed roles.

Caution

With HANA 2.0 SPS00 the Active/Active (read enabled) operation mode "`logreplay_readaccess`" can be configured. One important restriction is that primary and secondary system must have the exact same HANA version. Thus, during NZDU no read access is possible after the secondary was upgraded and is getting in sync again with the primary before it will take over!

The process, which is described in detail in the SAP HANA Administration Guide looks like this:

1. Set user store entry for automatic repository import at takeover time on primary and secondary by executing the following command, where `<myUser>`[43] requires the necessary privileges to import the repository content of the new version of the software during the takeover:

   ```
   hdbuserstore SET SRTAKEOVER <public hostname>:<sqlport>
   <myUser> <myUsersPasswd>
   ```

   Caution

   With HANA 2.0 SPS01 where SAP HANA Tenant Database Systems are the default database mode, `<sqlport>` must contain the System DB's SQL port **`3xx13`**.

2. Upgrade secondary system

   ```
   ./hdblcm --action=update
   ```

3. Wait until the primary and the secondary systems are in sync as shown in the `M_SERVICE_REPLICATION` view

   Note

   In NZDU updates from SAP HANA single container systems to SAP HANA tenant database systems (=default as of HANA 2.0 SPS01) the multi container migration is done automatically during upgrade. That is, temporarily in this NZDU process, the secondary can get in sync after update with the single container primary, before the takeover is executed. Please refer to SAP Note 2404375 - SAP HANA Platform 2.0 SPS 01 Release Note for details.

4. Stop primary system and perform takeover to the secondary (new primary)

   Caution

   If you are upgrading from HANA 1.0 to HANA 2.0 copy the `systemPKI` SSFS key and data file from this current primary to the new to-be-secondary; these files can be found here:

   ```
   /usr/sap/<SID>/SYS/global/security/rsecssfs/data/SSFS_<SID>.DAT
   /usr/sap/<SID>/SYS/global/security/rsecssfs/key/SSFS_<SID>.KEY
   ```

---

[43] How to create a user `<myUser>` with the privileges required for importing the repository content is shown in the following example:

```
CREATE USER MY_REPO_IMPORT_USER PASSWORD MyRepoUserPW123;
GRANT EXECUTE ON SYS.REPOSITORY_REST TO MY_REPO_IMPORT_USER;
GRANT REPO.READ ON ".REPO_PACKAGE_ROOT" TO MY_REPO_IMPORT_USER;
GRANT REPO. IMPORT TO MY_REPO_IMPORT_USER;
GRANT SELECT ON _SYS_REPO.DELIVERY_UNITS  TO MY_REPO_IMPORT_USER;
GRANT REPO.ACTIVATE_IMPORTED_OBJECTS ON ".REPO_PACKAGE_ROOT" TO MY_REPO_IMPORT_USER;
```

You create this user on the primary from where it is replicated to the secondary automatically.

5. Upgrade the previous primary system without starting the system
   ```
   ./hdblcm --action=update --hdbupd_server_nostart
   ```

   Caution
   Depending on the setting of the INI parameter `datashipping_snapshot_max_retention_time` and the duration between the takeover step (4) and the following registration step (6) a full data shipping or a delta shipping will be necessary to get the systems in sync again. As of HANA 2.0 the default setting is 300 minutes; thus, you have 5 hours' time between the steps (4) and (6).

6. Register the previous primary as secondary
   ```
   ./hdbnsutil -sr_register …
   ```

7. Start the previous primary as secondary

Tip
To achieve a real zero downtime upgrade from the ABAP application server perspective, please have a look at this SAP Note 1913302 (Connectivity suspend of Appserver during takeover).

Tip
For the procedure to upgrade a Multitier System Replication in a NZDU manner, please refer to SAP Note 2386973 (*Near Zero Downtime Upgrades for HANA database 3-tier System Replication*)

10.9.3 Hardware Exchange

As described in SAP Note 1984882 (*Using HANA system replication for Hardware Exchange with Minimum Downtime*) hardware can be exchanged with a minimal downtime using SAP HANA system replication. The procedure is very similar to the described NZDU.

11. Further documentation

11.1 Whitepapers and HowTo guides

SAP Note 2407186 - How-To Guides & Whitepapers For SAP HANA High Availability
SAP HANA Network Requirements whitepaper

SAP HANA Memory Usage explained
SAP HANA Storage Requirements whitepaper
SAP Note 190823 - SAP HANA Storage
SAP HANA Academy

11.2 Official Guides

SAP HANA Administration Guide
SAP HANA Tenant Databases Operations Guide
SAP HANA LCM Tools Reference (hdblcm, ...)
SAP HANA Master Guide
SAP HANA Security Guide
SAP HANA server installation
SAP HANA SQL and System Views Reference
SAP HANA Technical Operations Manual
SAP HANA Troubleshooting and Performance Analysis Guide
SAP HANA guides for previous HANA versions

www.sap.com/contactsap

www.sdn.sap.com/irj/sdn/howtoguides

SAP®
The Best-Run Businesses Run SAP™