

SAP How-to Guide

SAP HANA™

Network Required for SAP HANA System Replication

Applicable Releases:

SAP HANA 1.0

Version 2.0

July 2016

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



The Best-Run Businesses Run SAP™

© Copyright 2016 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 expressly prohibited, as is any decompilation of these components.

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

Document History

Document Version	Description
1.0	First official release
1.2	Fixed error in SQL statement
2.0	SAP HANA SPS11 & SPS12 extensions



Typographic Conventions

Type Style	Description
<i>Example Text</i>	Words or characters quoted from the screen. These include field names, screen titles, pushbuttons labels, menu names, menu paths, and menu options. Cross-references to other documentation
Example text	Emphasized words or phrases in body text, graphic titles, and table titles
<code>Example text</code>	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.
<code>EXAMPLE TEXT</code>	Keys on the keyboard, for example, F2 or ENTER.

Icons





Icon	Description
	Caution
	Note or Important
	Example
	Recommendation or Tip



Table of Contents

- 1. Before you start..... 2
- 2. Network relevant basics 2
 - 2.1 Data and Log Compression.....2
 - 2.2 Replication modes2
 - 2.3 Operation modes3
- 3. Network recommendations..... 5
 - 3.1 Throughput.....5
 - 3.1.1 Operation Mode `delta_datashipping`.....7
 - 3.1.2 Operation Mode `logreplay`.....7
 - 3.2 Latency8

1. Before you start

It is recommended to have read the SAP HANA system replication sections in the following documents:

- *SAP HANA Master Guide*:
http://help.sap.com/hana/SAP_HANA_Master_Guide_en.pdf
- *HowTo guide on System Replication*:
<https://scn.sap.com/docs/DOC-47702>
- *SAP HANA Network Requirements whitepaper*:
<https://scn.sap.com/docs/DOC-63221>

2. Network relevant basics

A network connection is required to transport the data and redo log from the primary site to a secondary site – and in multitier system replication also to a third site. Here are some facts regarding the network required for SAP HANA system replication that are to be taken into consideration.

2.1 Data and Log Compression


Since SAP HANA SPS09 compression can be used 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 redo 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 this content compression is turned off, however a tail compression¹ for the log buffers is done always.

2.2 Replication modes

-  The replication mode is only relevant for the *continuous log shipping* of the transactional redo log buffers – and not for the data shipping (based on HANA in-place snapshots taken on the primary site).

SAP HANA offers different replication modes for replicating the transactional redo log from the primary to the secondary site:

¹ In tail compression the appended *blank* characters are omitted.

- *Synchronous*: Secondary system sends acknowledgement back to primary as soon as data are received and persisted to disk.²
- *Synchronous in-memory*: Secondary system sends acknowledgement back to primary as soon as data is received (this might lead to performance increase depending on disk speed).
- *Asynchronous*: As per design of asynchronous replication, the primary does not wait until the secondary sends an acknowledgement.

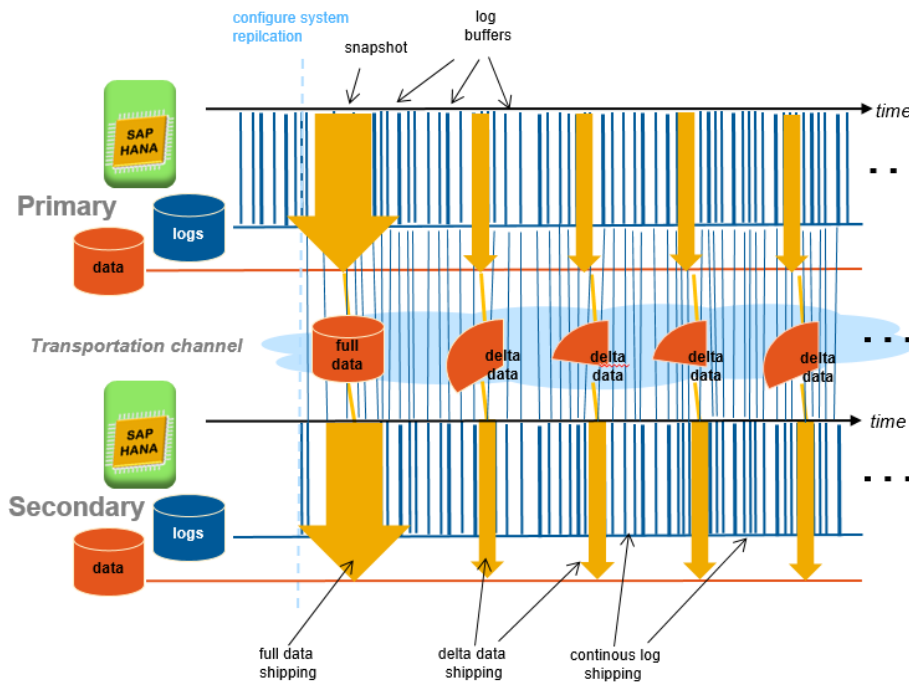
2.3 Operation modes

As of SAP HANA SPS11 there is the additional choice to run SAP HANA system replication in two different operation modes – the “classic” *delta_datashipping* operation mode which implicitly was the only possible mode of operation before SPS11, and the new “HotStandby” operation mode *logreplay*. These operation modes differ among others regarding the network traffic required between the primary and the secondary site.

- *delta_datashipping*: In addition to the continuous redo log shipping the secondary system requests a delta data shipping on regularly (per default every 10 minutes). This is the “classical” operation mode of SAP HANA system replication.
- *logreplay*: In this operation mode pure redo log shipping is done, after the system replication was initially set up with a full data shipping. The amount of data which needs to be transferred to the secondary site is reduced dramatically, because no delta data shipping is required anymore.

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

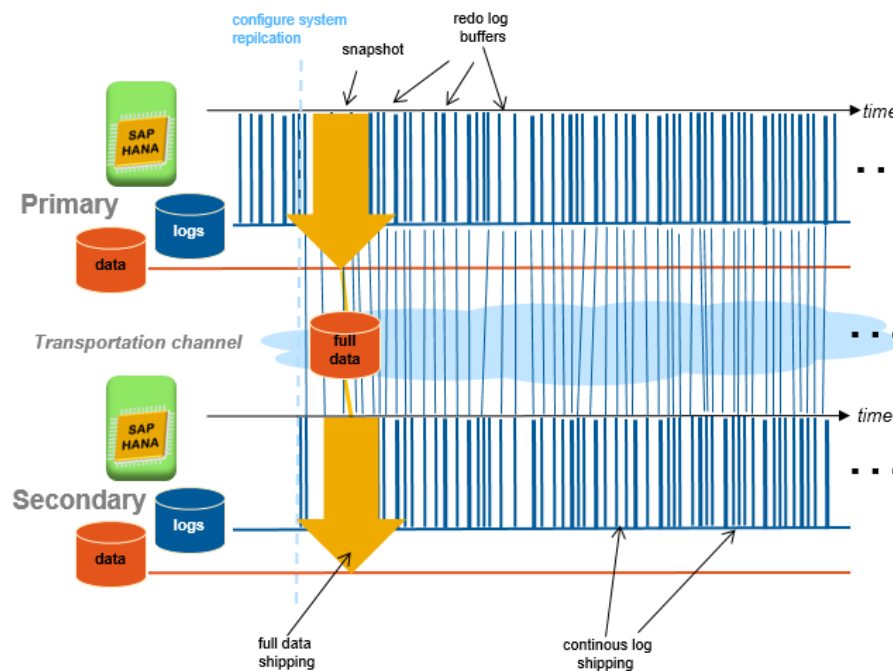
² Additionally (as of SPS08) the *synchronous* replication mode (SYNC) can run with “full sync” enabled. In *full sync* operation, transaction processing on the primary site blocks, when the secondary is currently not connected and newly created redo log buffers cannot be shipped to the secondary site. This behavior ensures that no transaction can be committed locally without shipping the redo log buffers to the secondary site.



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

When system replication is set up initially, a full set of data is sent over the network to the primary. In parallel the shipping of the transactional redo log is started and will continuously run while the replication is active. Occasionally – by default every 10 minutes or when a certain threshold size is met – a delta data shipping takes place.

With **logreplay** operation mode the delta data shipping is not required – only the initial full data shipping is done and the continuous shipping of redo log.



*Operation mode **logreplay**:
Initial full data and redo log shipped to secondary*

There is much less traffic on the network between the participating sites when running in operation mode *logreplay*.

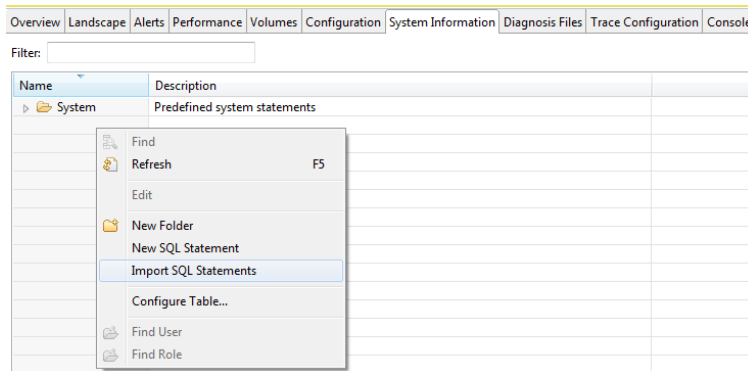
3. Network recommendations

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 acknowledgment from the secondary site arrives.


3.1 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 a 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 and 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 in the "Name" column → *Import SQL Statements*.




Select the “SQL_Statements.zip” file you downloaded from the SAP note:



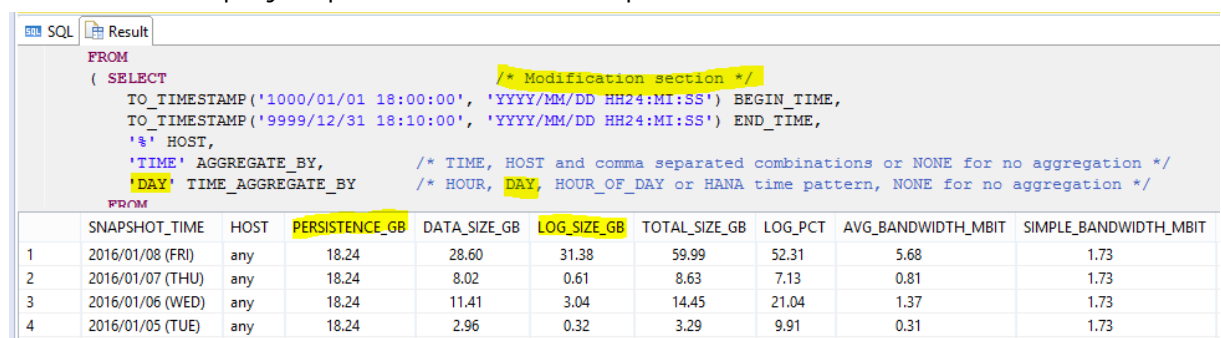
A folder with the corresponding SQL statements will be imported.

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:



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 amount of data written to disk (GB)
LOG_SIZE_GB	Total amount 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_BANDWIDTH_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 amounts of data are shipped over the network. Of special relevance are the above marked values for **PERSISTENCE_GB** and **LOG_SIZE_GB** – as described below.

3.1.1 Operation Mode `delta_datashipping`

For this “classic” operation mode the following applies:

Throughput

It must be possible to transport the size of the persistently stored data within one day from the primary to the secondary.

Here is an example further explaining 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`. (The value represents the currently valid size of the persistently stored data.)

3.1.2 Operation Mode `logreplay`

With this operation mode no delta data shippings are necessary anymore; pure log shipping is done. Thus, 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 above mentioned Bandwidth SQL statement, you receive the above shown result – among others the sizes of log that was created per day (column `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.

3.2 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 replication mode 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). 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 collect the local and the remote *log write wait times* with the SQL statement *HANA_Replication_Overview* (attached to the above mentioned SAP Note 1969700).

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 period between enqueueing and finishing a request.

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

Replication	
SystemReplication	
KeyFigures	
Bandwidth	- Calculation of bandwidth requirements for system replication - Based on write I/O to DATA and size of log backups
Overview	- General system replication information
ParameterDeviations	- Displays statistic server information for SAP HANA parameter deviations between primary and replication side - Sect
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:

AB	REPLICATION_PATH	AB	HOSTS	AB	PORT	AB	KEY	AB	VALUE
	SiteA -> SiteB		Id2131 -> Id2132		30103		Replication mode		SYNC
							Secondary connect time		2016/01/05 14:48:22
							Days since secondary connect time		5.83
							Used persistence size (GB)		23.73
							Log backup size / day (GB)		9.51
							Local log buffer write size (MB)		303464.66
							Shipped log buffer size (MB)		35074.25
							Avg. local log buffer write size (KB)		683.35
							Avg. shipped log buffer size (KB)		80.66
							Avg. local log buffer write time (ms)		3.34
							Avg. log buffer shipping time (ms)		1.88
							Local log buffer write throughput (MB/s)		199.47
							Log buffer shipping throughput (MB/s)		41.74
							Initial data shipping size (MB)		624.00
							Initial data shipping time (s)		3.45
							Last delta data shipping size (MB)		48.00
							Delta data shipping size (MB)		48.00
							Delta data shipping time (s)		0.40
							Delta data shipping throughput (MB/s)		118.41
							Delta data shipping size / day (MB)		n/a

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 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.



www.sap.com/contactsap

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



The Best-Run Businesses Run SAP™