



# SAP HANA Business Continuity with HANA System Replication and Red Hat Enterprise Linux Cluster on Cisco Unified Computing System

---

# Contents

<b>Trademarks and Disclaimers .....</b>	<b>2</b>
<b>Document History .....</b>	<b>3</b>
<b>Overview .....</b>	<b>4</b>
Audience .....	4
Architecture .....	4
Network Recommendations .....	5
<b>Setting Up the Cluster .....</b>	<b>15</b>
<b>Updating the Cluster Packages .....</b>	<b>16</b>
<b>Configuring the Server Network.....</b>	<b>18</b>
<b>Initializing the Cluster.....</b>	<b>18</b>
Configure the STONITH Devices .....	20
Test the Cluster Configuration .....	22
Test STONITH .....	23
Create a Cluster IP Address .....	24
Create the SAP HANA Topology Resource .....	24
Create the Clone Resource for the Topology Resource ...	24
Create the SAP HANA Resource .....	25
Create a SAP HANA Master and Slave .....	25
Set Up Constraints .....	26
Results .....	26
Enable STONITH .....	27
<b>Setting Up the Network for System Replication.....</b>	<b>27</b>
<b>Choosing a SAP HANA Synchronization Mode .....</b>	<b>30</b>
<b>Performing Cluster Switchback After a Failure.....</b>	<b>31</b>
Resynchronize the Primary Database.....	31
Activate the Original Master as the Primary System .....	33
Reenable the Secondary Site .....	33
<b>Disabling System Replication.....</b>	<b>34</b>
<b>For More Information.....</b>	<b>35</b>

## Trademarks and Disclaimers

Cisco's trademarks signify Cisco's high quality products and services, and they are valuable assets of the company. Cisco's trademark and brand policies can be found at <http://www.cisco.com/go/logo>. The following is a list of Cisco's trademarks and registered trademarks in the United States and certain other countries. Please note, however, that this listing is not all inclusive, and the absence of any mark from this list does not mean that it is not a Cisco trademark.

<http://www.cisco.com/web/siteassets/legal/trademark.html>

---

## Document History

Version	Date	Author	Document Change
1.0	June 22nd 2015	Ralf Klahr	Initial Draft

---

## Overview

This document is not a step-by-step installation guide but instead provides guidance for setting up SAP HANA system replication. This document refers to specific vendor information, such as the Red Hat SAP HANA replication white paper and other SAP HANA documentation listed at the end of this document.

This document is based on Red Hat Enterprise Linux (RHEL) 6.5.

As of this writing, automatic failover is available only with scale-up (single-system) architecture.

SAP HANA is SAP's implementation of in-memory database technology.

The SAP HANA database takes advantage of the low cost of main memory (RAM), the data processing capabilities of multicore processors, and the fast data access of solid-state disk (SSD) drives relative to traditional hard-disk drives (HDDs) to deliver better performance for analytical and transactional applications. Its multiple-engine query processing environment allows it to support relational data (with both row- and column-oriented physical representations in a hybrid engine) as well as graph and text processing for management of semistructured and unstructured data within the same system. The SAP HANA database is 100 percent compliant with the atomicity, consistency, isolation, and durability (ACID) model.

For more information about SAP HANA, see the SAP help portal, at <http://help.sap.com/hana/>.

## Audience

The intended audience for this document includes sales engineers, field consultants, professional services staff, IT managers, partner engineering staff, and customers deploying SAP HANA.

## Architecture

The solution presented in this document is based on the Cisco Unified Computing System™ (Cisco UCS®) and FlexPod, but it could also be configured with EMC VNX systems. The solution provides a general approach to a single system (scale-up design) with system replication and automatic failover.

The scenarios for this solution can use the following hardware:

- Cisco UCS C460 M4 Rack Server
- Cisco® appliance for SAP HANA Tailored Datacenter Integration (TDI)
- Cisco appliance for NetApp
- Cisco appliance for EMC

Possible failover scenarios are:

- Standalone Cisco UCS C460 M4 > Standalone Cisco UCS C460 M4
- Standalone Cisco UCS C460 M4 > Cisco appliance for SAP HANA TDI
- Cisco appliance for NetApp > Cisco appliance for NetApp
- Cisco appliance for EMC > Cisco appliance for EMC
- Cisco appliance for NetApp or EMC > Cisco appliance for SAP HANA TDI
- Cisco appliance for SAP HANA TDI > Cisco appliance for SAP HANA TDI

## Network Recommendations

In general, SAP does not give a network bandwidth recommendation for the replication network.

The bandwidth required between the two sites depends by the change rate of the database. Therefore, you should start with a 10-Gbps Layer 2 connection for one scale-up system. If you configure a scale-out system, the bandwidth can be increased to the number of nodes multiplied by 10 Gbps.

If the change rate is not high, less bandwidth may be required. In this case, even a switched (Layer 3) connection is possible.

## Single-Node Disaster-Tolerance Configuration

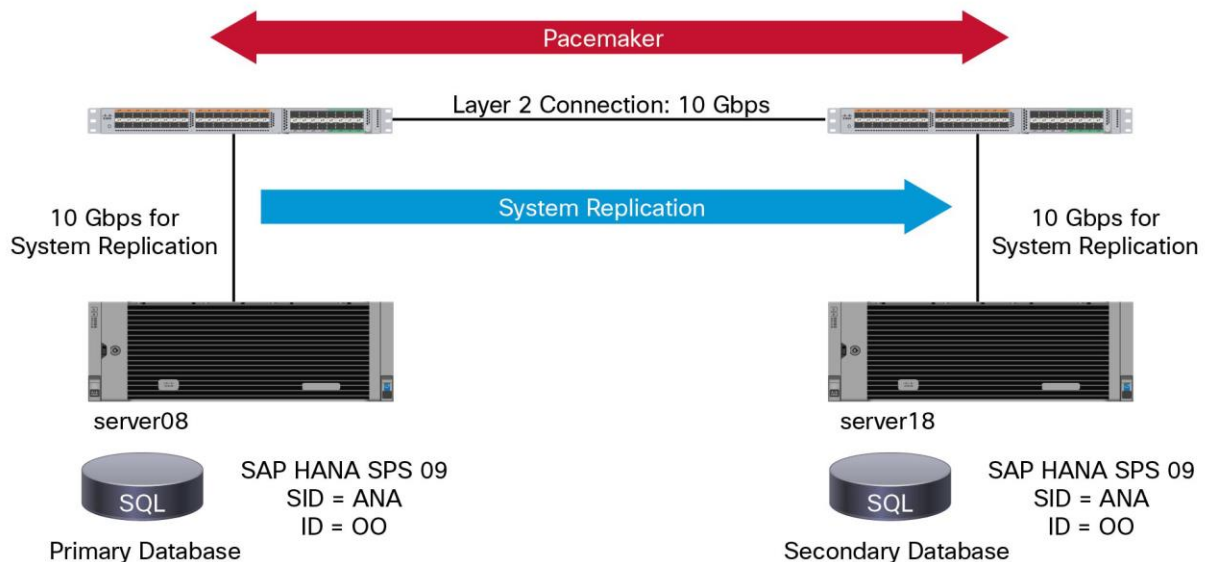
The single-node disaster-tolerance configuration is built on RHEL for SAP HANA with High Availability and Smart Management. It consists of these components:

- RHEL for SAP HANA
- RHEL High-Availability add-on

Red Hat provides a step-by-step description of how to set up a Pacemaker cluster. Please visit <https://access.redhat.com/articles/1466063> to get the latest copy of this document.

Figure 1 shows the design of the current version of the solution.

**Figure 1.** SAP HANA System Replication in the Cluster



---

## Supported Configuration

For the first version of the SAP HANA system replication high-availability solution, the support is limited to the following scenarios and parameters:

- Only 2-node clusters are supported.
- Both SAP HANA SPS 08 and SPS 09 can be used.
- Only SAP HANA scale-up (single device to single device) system replication is supported. SAP HANA scale-out system replication and SAP HANA installations configured with a standby node are not supported.
- Both SAP HANA instances must have the same SAP identifier (SID) and instance number.

No other SAP HANA system (such as SAP Quick Address System [QAS] or TST) can be on the replicating node that needs to be stopped during takeover.

- Only one system replication is supported for the SAP HANA database. Multitier system replication and replication chaining are not supported. (The resource agent can detect whether multitier system replication is configured, but it can't act in this case because doing so could result in an unsupported configuration for SAP HANA system replication [star topology].)
- SAP HANA system replication with multiple components and one database (MCOD) is not supported (<http://scn.sap.com/docs/DOC-59893>). This issue is being tracked at [https://bugzilla.redhat.com/show\\_bug.cgi?id=1200903](https://bugzilla.redhat.com/show_bug.cgi?id=1200903). MCOS is supported only if all databases running on the hosts are replicated, and if the replication is always to the same secondary node. See the Automated SAP Hana System Replication with Pacemaker cookbook.
- All nodes must be in the same network segment (Layer 2).
- Technical users and groups such as <SID>adm must be identically defined on all cluster nodes.
- Name resolution for the cluster nodes and the virtual IP address can be performed locally on all cluster nodes.
- Time on all cluster nodes must be synchronized (using Network Time Protocol [NTP] or some other time synchronization method).
- If the cluster nodes are installed in different data centers or data center areas, the environment must match both the requirements defined by SAP for HANA system replication (see Section 4.2, "Distance Between Data Centers," in the SAP "How to Perform System Replication for SAP HANA" guide) and also the RHEL high-availability add-on stretch cluster requirements, specifically the network latencies between the nodes and the recommended maximum distance.

This proof-of-concept (PoC) design will be the focus of the testing described in this document.

If you need to implement a different scenario, you should define a PoC design with Red Hat.

The resource agent supports SAP HANA for system replication beginning with SAP HANA Release 1.0 SPS 08. SPS 09 with patch level 94 or later is recommended.

## Preparing for Installation

Both nodes for system replication must be installed identically.

First update the nodes with the latest available kernel and OS patches for RHEL 6.5.

```
server08:~ # uname -a
Linux server08 2.6.32-431.el6.x86_64 #1 SMP Sun Nov 10 22:19:54 EST 2013 x86_64
x86_64 x86_64 GNU/Linux

Server18:~ # uname -a
Linux server18 2.6.32-431.el6.x86_64 #1 SMP Sun Nov 10 22:19:56 EST 2013 x86_64
x86_64 x86_64 GNU/Linux
```

## Prerequisites

- You must have access to the SAP Service Marketplace to be able to download SAP installation media, installation guides, etc.
- You must have subscriptions for RHEL for SAP HANA and the RHEL High-Availability add-on for all cluster nodes.
- You must be able to access to the RHEL 6 repository and high-availability and RHEL for SAP HANA add-on repositories with the latest patches. At minimum, the following versions of CMAN, Pacemaker, and Pacemaker Configuration System (PCS) (and their dependencies) are required:
  - CMAN: 3.0.12.1-58.el6.x86\_64
  - Automated SAP Hana System Replication with Pacemaker cookbook
  - libqb-0.16.0-2.el6.x86\_64
  - Pacemaker: 1.1.12-4.el6.x86\_64
  - fence-agents-3.1.5-34.el6.x86\_64
  - PCS: 0.9.90-2.el6.noarch
  - resource-agents-sap-hana-3.9.5-24.el6 (until the package is officially available, it can be obtained from <https://brewweb.devel.redhat.com/buildinfo?buildID=433702>)
- You need a current set of SAP HANA installation media (see the SAP HANA installation guides for the correct list of installation media); the following SAP HANA revisions are recommended: SPS 08 Revision 85\_3 or SPS 09 Revision 94.
- You need the latest version of SAP HOSTAGENT (at least PL206; see <http://service.sap.com/sap/support/>).
- You must have notes/2138524) installed on all nodes (see SAP Note 1031096, Installing Package SAPHOSTAGENT).

## Known Issues

- When the cluster node names are not identical to the host names of the nodes, automatic registration of the former primary replication node as a secondary replication node can fail, and manual intervention using **hdbnsutil** can cause the resource agents to become confused.

This issue was fixed by the following upstream commits:

- <https://github.com/fmherschel/SAPHanaSR/commit/0106a3157fa3d366391878e10134551d3e32aa88>
- <https://github.com/fmherschel/SAPHanaSR/commit/0c51246ab89c9cd87436e43754797afcaebdac93>

---

Those commits are also included in the resource agents shipped with the resource-agents-3.9.5-24.el6 package for RHEL 6.7. For RHEL 6.6, the fixed versions of the resource agents will be released through ZStream errata.

- If SAP HANA system replication is configured with mode **syncmem** after a takeover, the mode will be changed to **sync**. This change is due to limitations in the API provided by SAP HANA. Cisco and the SAP HANA developers are working to get this problem fixed.

As a workaround, manually change the SAP HANA system replication mode back to the old value (see "Changing HANA System Replication Mode on the Fly").

#### Resource Agents

**You must** use at least the versions of the resource agents provided by the resourceagents-3.9.5-12.el6 package. A Red Hat Network (RHN) account is required to download the package.

- [https://rhn.redhat.com/network/software/packages/name\\_overview.pxt?package\\_name=resource-agents-sap-hana&archIdList=520&archLabelList=x86\\_64&search\\_subscribed\\_channels=](https://rhn.redhat.com/network/software/packages/name_overview.pxt?package_name=resource-agents-sap-hana&archIdList=520&archLabelList=x86_64&search_subscribed_channels=)
- [https://access.redhat.com/documentation/en-US/Red\\_Hat\\_Enterprise\\_Linux/6/html/Cluster\\_Administration/s1-install-clust-sw-CA.html](https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Cluster_Administration/s1-install-clust-sw-CA.html)

#### Useful Documentation

See the following for more information:

- SAP Note 2009879: SAP HANA Guidelines for RHEL Operating System
- SAP Note 2013638: SAP HANA Database: Recommended OS Settings for RHEL 6.5
- SAP Note 2136965: SAP HANA Database: Recommended OS Settings for RHEL 6.6
- SAP Note 2001528—Linux: SAP HANA Database SPS 08 Revision 80 (or later) on RHEL 6 or SUSE Enterprise Linux Server (SLES) 11
- SAP Note 2063657: SAP HANA System Replication Takeover Decision Guidelines
- Availability and Scalability: SAP HANA Administration Guide
- How to Perform System Replication for SAP HANA
- Automate SAP HANA System Replication with SLES
- HANA System Replication: Takeover Process
- Configuring the Red Hat High-Availability Add-on with Pacemaker
- Support for Red Hat Enterprise Linux High-Availability Cluster Stretch and Multisite Architectures



## Installing the High-Availability Package

Use these commands to install the high-availability package:

```
[root@server08 sapcd]# cd /sapcd
[root@server18 sapcd]# ls -l rhel*
-rw-r--r-- 1 nobody nobody 3853516800 Oct 29 22:02 rhel-server-6.5-x86_64-
dvd.iso
[root@server08 sapcd]# mkdir RHEL_REPO
[root@server08 sapcd]# mount -o loop ./rhel-server-6.5-x86_64-dvd.iso ./RHEL_REPO
[root@server08 ~]# yum-config-manager --add-repo file:///sapcd/RHEL_REPO
[root@server08 ~]# yum-config-manager --add-repo
file:///sapcd/RHEL_REPO/HighAvailability
[root@server08 ~]# yum-config-manager --enable file:///sapcd/RHEL_REPO
[root@server08 ~]# yum-config-manager --enable
file:///sapcd/RHEL_REPO/HighAvailability
[root@server08 ~]# yum-config-manager --disable rhel-source
[root@server08 ~]# yum clean all
[root@server08 ~]# yum repolist

Loaded plugins: product-id, refresh-packagekit, security, subscription-manager,
versionlock

This system is not registered to Red Hat Subscription Management. You can use
subscription-manager to register.

repo id                repo name
sapcd_RHEL_REPO        added from: file:///sapcd/RHEL_REPO
sapcd_RHEL_REPO_HighAvailability added from:
file:///sapcd/RHEL_REPO/HighAvailability

[root@server08 ~]# yum install rgmanager pcs
```

## Installing the Resource Agents

Download and install the resource agents from RHN (an account is needed).

- resource-agents: <https://rhn.redhat.com/rhn/software/packages/details/Overview.do?pid=1062230>
- resource-agents-sap-hana: <https://rhn.redhat.com/rhn/software/packages/details/Overview.do?pid=1062211>

Install the agents (on all nodes) as follows:

```
[root@server18 Pacemaker]# ls -l resource*
-rw-r--r-- 1365636 Jun 21 resource-agents-3.9.5-12.el6_6.5.x86_64.rpm
-rw-r--r-- 1 61848 Jun 21 resource-agents-sap-hana-3.9.5-12.el6_6.5.x86_64.rpm

[root@server08 Pacemaker]# rpm -Uvh resource*.rpm
Preparing...                               ##### [100%]
 1:resource-agents                          ##### [ 50%]
 2:resource-agents-sap-han##### [100%]
```

Disable the automatic start of Pacemaker:

```
[root@server08 ~]# chkconfig pacemaker off
[root@server18 ~]# chkconfig pacemaker off
```

Installing SAP HANA

Install SAP HANA on all nodes as described in the SAP HANA installation guide.

SAP HANA must be installed with the same SID and instance number on all nodes.

#### server08 (Primary Database)

```
sapcontrol -nr 00 -function GetProcessList
21.06.2015 23:42:57
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GREEN, Running, 2015 06 21 23:37:44, 0:05:13, 44435
hdbnameserver, HDB Nameserver, GREEN, Running, 2015 06 21 23:37:51, 0:05:06,
44507
hdbpreprocessor, HDB Preprocessor, GREEN, Running, 2015 06, 0:04:52, 44690
hdbindexserver, HDB Indexserver, GREEN, Running, 2015 06 21, 0:04:43, 44874
hdbstatisticsserver, HDB Statisticsserver, GREEN, Running, 2015 06 21, , 44877
hdbxsengine, HDB XSEngine, GREEN, Running, 2015 06 21 23:38:14, 0:04:43, 44880
hdbcompileserver, HDB Compileserver, GREEN, Running, 2015 06 21, 0:04:52, 44695
hdbwebdispatcher, HDB WebDispatcher, GREEN, Running, 2015 06 21, 0:04:02, 46062
rhladm@server08:/usr/sap/RH1/HDB00>
```

#### server18 (Secondary Database)

```
sapcontrol -nr 00 -function GetProcessList
21.06.2015 23:45:09
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GREEN, Running, 2015 06 21 23:39:58, 0:05:11, 50291
hdbnameserver, HDB Nameserver, GREEN, Running, 2015 06 21 23:40:03, 0:05:06,
50373
hdbpreprocessor, HDB Preprocessor, GREEN, Running, 2015 06 21, 0:04:56, 50449
hdbindexserver, HDB Indexserver, GREEN, Running, 2015 06 21, 0:04:49, 50587
hdbstatisticsserver, HDB Statisticsserver, GREEN, Running, 2015, 0:04:49, 50590
hdbxsengine, HDB XSEngine, GREEN, Running, 2015 06 21 23:40:20, 0:04:49, 50593
hdbcompileserver, HDB Compileserver, GREEN, Running, 2015 06 21, 0:04:56, 50452
hdbwebdispatcher, HDB WebDispatcher, GREEN, Running, 2015 06 21, 0:04:35, 50829
rhladm@server18:/usr/sap/RH1/HDB00>
```

---

For system replication to work, the SAP HANA **log\_mode** variable must be set to **normal**. You can verify this setting after installation with the following command:

```
hdbsql -u system -i 00 "select value from "SYS"."M_INIFILE_CONTENTS" where
key='log_mode'"
Password:
VALUE
"normal"
1 row selected (overall time 833.880 msec; server time 826.608 msec)
```

```
hdbsql -u system -i 00 "select value from "SYS"."M_INIFILE_CONTENTS" where
key='log_mode'"
Password:
VALUE
"normal"
1 row selected (overall time 694.305 msec; server time 683.936 msec)
```

Make sure that SAP HANA is not configured to automatically start after the system boots, because SAP HANA startup and shutdown will be controlled by the cluster. Verify this setting on all nodes:

```
grep Autostart /usr/sap/RH1/SYS/profile/*
/usr/sap/RH1/SYS/profile/RH1_HDB00_server18:Autostart = 0
rhladm@server18:/usr/sap/RH1/HDB00>
```

```
grep Autostart /usr/sap/RH1/SYS/profile/*
/usr/sap/RH1/SYS/profile/RH1_HDB00_server18:Autostart = 0
rhladm@server18:/usr/sap/RH1/HDB00>
```

In test environments, you may need to reduce the number of log backups that are generated to prevent the log backups from filling the disks.

On both systems, enter this command:

```
vi /hana/shared/RH1/global/hdb/custom/config/global.ini

[persistence]
basepath_datavolumes=/hana/data/RH1
basepath_logvolumes=/hana/log/RH1
log_backup_timeout_s = 0
```

Then restart the database:

```
sapcontrol -nr 00 -function StopSystem
sapcontrol -nr 00 -function StartSystem
```

If SAP HANA system replication uses a separate network for its traffic, you need to configure the system replication hostname resolution as described in “Host Name Resolution for System Replication” in the “SAP HANA Master Guide.”

#### Setting Up the Primary Node

Back up the database as <sidadm> (SAP HANA system replication will not work until backup has been performed):

```
hdbssql -i 00 -u system -p manager "BACKUP DATA USING FILE ('/tmp/backup_1')"  
0 rows affected (overall time 5258.071 msec; server time 5256.815 msec)  
  
rhladm@server08:/hana/shared/RH1/global/hdb/custom/config> ls -l /tmp/backup_1*  
-rw-r----- 1 rhladm sapsys      147456 Mar 15 20:14 /tmp/backup_1_databackup_0_1  
-rw-r----- 1 rhladm sapsys    68927488 Mar 15 20:14 /tmp/backup_1_databackup_1_1  
-rw-r----- 1 rhladm sapsys   143609856 Mar 15 20:14 /tmp/backup_1_databackup_2_1  
-rw-r----- 1 rhladm sapsys  1937190912 Mar 15 20:14 /tmp/backup_1_databackup_3_1  
-rw-r----- 1 rhladm sapsys    77426688 Mar 15 20:14 /tmp/backup_1_databackup_4_1
```

On server08 (primary site), set up system replication on the source site:

```
hdbnsutil -sr_enable --name=DC1  
checking for active nameserver ...  
nameserver is active, proceeding ...  
successfully enabled system as system replication source site  
done.
```

Check the status of the replication:

```
rhladm@server08:/hana > hdbnsutil -sr_state  
checking for active or inactive nameserver ...  
  
System Replication State  
~~~~~  
  
mode: primary  
site id: 1  
site name: DC1  
  
Host Mappings:  
~~~~~  
  
done.  
rhladm@server08:/hana >
```

On server18 (secondary site), before you can register the secondary system, it must be shut down.

The actual status is:

```
python exe/python_support/systemReplicationStatus.py
there are no secondary sites attached
```

**HDB stop**

```
hdbnsutil -sr_register --remoteHost=server08 --remoteInstance=00 --mode=syncmem -
-name=DC2
adding site ...
checking for inactive nameserver ...
nameserver server18:30001 not responding.
collecting information ...
updating local ini files ...
done.
rhladm@server18:/hana>
```

```
python exe/python_support/systemReplicationStatus.py
```

```
| Host      | Port  | Service Name | Volume ID | Site ID | Site Name | Secondary
| Secondary | Secondary | Secondary | Secondary |          |           | Replication
| Replication |          |           |           |           |           | Replication
|          |          |           |           |           |           |           | Host
| Port      | Site ID | Site Name | Active Status | Mode | Status
| Status Details |
| ----- | ----- | ----- | ----- | ----- | ----- | -----
|          |          |          |          |          |          |          |
| server08 | 30005 | statisticserver | 2 | 1 | DC1 | server18
| 30005 | 2 | DC2 | CONNECTION TIMEOUT | UNKNOWN | UNKNOWN
|
| server08 | 30007 | xsengine | 4 | 1 | DC1 | server18
| 30007 | 2 | DC2 | CONNECTION TIMEOUT | UNKNOWN | UNKNOWN
|
| server08 | 30001 | nameserver | 1 | 1 | DC1 | server18
| 30001 | 2 | DC2 | CONNECTION TIMEOUT | UNKNOWN | UNKNOWN
|
| server08 | 30003 | indexserver | 3 | 1 | DC1 | server18
| 30003 | 2 | DC2 | CONNECTION TIMEOUT | UNKNOWN | UNKNOWN
|
status system replication site "DC2": ERROR
```

**HDB start**

**python exe/python\_support/systemReplicationStatus.py**

Secondary Replication	Secondary	Secondary	Secondary	Volume ID	Site ID	Site Name	Secondary Replication	Host	Port
Site ID	Site Name	Active	Status	Mode	Status	Status Details			
server08 30005	30005 2   DC2	statisticsserver   YES	2	1	DC1	server18			
server08 30007	30007 2   DC2	xsengine   YES	4	1	DC1	server18			
server08 30001	30001 2   DC2	nameserver   YES	1	1	DC1	server18			
server08 30003	30003 2   DC2	indexserver   YES	3	1	DC1	server18			

status system replication site "DC2": ACTIVE

**hdbnsutil -sr\_state**

checking for active or inactive nameserver ...

System Replication State

~~~~~

mode: syncmem  
site id: 2  
site name: DC2  
active primary site: 1

re-setup replication: hdbnsutil -sr\_register --name=DC2 --mode=syncmem --remoteHost=server08 --remoteInstance=00

Host Mappings:

~~~~~

server18 -> [DC1] server08  
server18 -> [DC2] server18

done.

rhladm@server18:/usr/sap/RH1/HDB00>

---

## Setting Up the Cluster

Create a monitoring account in SAP HANA on the primary node. This user will be replicated through the SAP HANA system replication process.

```
hdbsql -i 00 -u system -p manager "create user rhelhasync password UCS4hana"

hdbsql -i 00 -u system -p manager "grant DATA ADMIN to rhelhasync"

hdbsql -i 00 -u system -p manager "ALTER USER rhelhasync DISABLE PASSWORD
LIFETIME"
```

Then add a user key **SAPHANA<SID>SR** for this user to the local SAP HANA user store for the root user. Run the following commands on **all** nodes:

User sidadm

```
hdbuserstore SET SAPHANARH1SR localhost:30015 rhelhasync UCS4hana
hdbuserstore SET SAPHANARH1SR localhost:30015 rhelhasync UCS4hana
```

User root

```
/usr/sap/RH1/HDB00/exe/hdbuserstore SET SAPHANARH1SR localhost:30015 rhelhasync
UCS4hana
/usr/sap/RH1/HDB00/exe/hdbuserstore SET SAPHANARH1SR localhost:30015 rhelhasync
UCS4hana
```

Test

```
/usr/sap/RH1/HDB00/exe/hdbsql -U SAPHANARH1SR -i 00 "select distinct REPLICATION_STATUS from
SYS.M_SERVICE_REPLICATION"
```

```
REPLICATION_STATUS
```

```
"ACTIVE"
```

```
1 row selected (overall time 26.169 msec; server time 17.996 msec)
```

Test the password for database access:

```
hdbsql -U rhelhasync "select * from dummy"
DUMMY
"X"
lines 1-2/2 (END)
```

## Updating the Cluster Packages

Download the latest patches from <http://rhn.redhat.com> and store them under /sapcd/Pacemaker. Install the latest updates on all cluster nodes.

```
cd /sapcd/Pacemaker
```

```
ls -l
```

```
-rw-r--r-- 1 37797 Jun 10 02:09 SAPHanaSR-master.zip
-rw-r--r-- 1 107428 Jun 2 09:10 clusterlib-3.0.12.1-68.el6.x86_64.rpm
-rw-r--r-- 1 459468 Jun 2 08:55 cman-3.0.12.1-68.el6.x86_64.rpm
-rw-r--r-- 1 132460 Jun 2 09:27 compat-readline5-5.2-17.1.el6.x86_64.rpm
-rw-r--r-- 1 170136 Jun 2 08:57 fence-agents-3.1.5-48.el6_6.3.x86_64.rpm
-rw-r--r-- 1 1753048 Jun 2 09:14 glib2-2.28.8-4.el6.x86_64.rpm
-rw-r--r-- 1 72120 Jun 2 08:56 libqb-0.16.0-2.el6.x86_64.rpm
-rw-r--r-- 1 459260 Jun 2 08:56 pacemaker-1.1.12-4.el6.x86_64.rpm
-rw-r--r-- 1 195960 Jun 2 09:10 pacemaker-cli-1.1.12-4.el6.x86_64.rpm
-rw-r--r-- 1 81904 Jun 2 09:10 pacemaker-cluster-libs-1.1.12-4.el6.x86_64.rpm
-rw-r--r-- 1 470512 Jun 2 09:10 pacemaker-libs-1.1.12-4.el6.x86_64.rpm
-rw-r--r-- 1 5939160 Jun 2 08:58 pcs-0.9.123-9.el6_6.2.x86_64.rpm
-rw-r--r-- 1 182676 Jun 2 09:24 readline-6.0-4.el6.x86_64.rpm
-rw-r--r-- 1 551272 Jun 2 09:10 ruby-1.8.7.374-4.el6_6.x86_64.rpm
-rw-r--r-- 1 325032 Jun 2 09:22 ruby-irb-1.8.7.374-4.el6_6.x86_64.rpm
-rw-r--r-- 1 1732688 Jun 2 09:14 ruby-libs-1.8.7.374-4.el6_6.x86_64.rpm
-rw-r--r-- 1 389896 Jun 2 09:17 ruby-rdoc-1.8.7.374-4.el6_6.x86_64.rpm
-rw-r--r-- 1 211652 Jun 2 09:10 rubygems-1.3.7-5.el6.noarch.rpm
```

```
rpm -Uvh compat-readline5-5.2-17.1.el6.x86_64.rpm glib2-2.28.8-4.el6.x86_64.rpm
readline-6.0-4.el6.x86_64.rpm
```

```
Preparing... ##### [100%]
 1:glib2 ##### [ 33%]
 2:readline ##### [ 66%]
 3:compat-readline5 ##### [100%]
```

```
rpm -Uvh ruby*
```

```
Preparing... ##### [100%]
 1:ruby-libs ##### [ 20%]
 2:ruby ##### [ 40%]
 3:ruby-irb ##### [ 60%]
 4:ruby-rdoc ##### [ 80%]
 5:rubygems ##### [100%]
```



```
rpm -Uvh pacemaker-* pcs-0.9.123-9.el6_6.2.x86_64.rpm clusterlib-3.0.12.1-68.el6.x86_64.rpm cman-3.0.12.1-68.el6.x86_64.rpm fence-agents-3.1.5-48.el6_6.3.x86_64.rpm
```

```
Preparing... ##### [100%]
 1:clusterlib ##### [ 13%]
 2:pacemaker-libs ##### [ 25%]
 3:pacemaker-cli ##### [ 38%]
 4:pacemaker-cluster-libs ##### [ 50%]
 5:fence-agents ##### [ 63%]
 6:cman ##### [ 75%]
 7:pacemaker ##### [ 88%]
 8:pcs ##### [100%]
```

```
Pacemaker Cluster Manager is already stopped[ OK ]
```

```
Stopping cluster:
```

```
Leaving fence domain... [ OK ]
```

```
Stopping gfs_controld...
```

```
[ OK ]
```

```
Stopping dlm_controld... [ OK ]
```

```
Stopping fenced... [ OK ]
```

```
Stopping cman... [ OK ]
```

```
Waiting for corosync to shutdown:[ OK ]
```

```
Unloading kernel modules... [ OK ]
```

```
Unmounting configfs... [ OK ]
```

```
[root@server18 Pacemaker]#
```

```
Pacemaker Cluster Manager is already stopped[ OK ]
```

```
Stopping cluster:
```

```
Leaving fence domain... [ OK ]
```

```
Stopping gfs_controld... [ OK ]
```

```
Stopping dlm_controld... [ OK ]
```

```
Stopping fenced... [ OK ]
```

```
Stopping cman... [ OK ]
```

```
Unloading kernel modules... [ OK ]
```

```
Unmounting configfs... [ OK ]
```

```
[root@server08 Pacemaker]#
```

## Configuring the Server Network

Use Table 1 to configure the server network.

**Table 1.** Server Network Configuration Settings

Function	server08 Setting	server18 Setting
Management LAN	192.168.76.108 server08m	192.168.76.118 server08m
Cisco Integrated Management Controller (IMC; STONITH) Note: STONITH stands for "shoot the other node in the head."	192.168.76.162 server08-ipmi	192.168.76.165 server08-ipmi
Network File System (NFS) data network	192.168.201.108 server08d	192.168.201.118 server08d
NFS log network	192.168.228.108 server08l	192.168.228.118 server08l
Boot (Preboot Execution Environment [PXE])	192.168.127.108 server08pxe	192.168.127.118 server08pxe
SAP HANA system replication	192.168.222.108 server08sys	192.168.222.118 server08sys
Server-to-server and PCS heartbeat	192.168.220.108 server08s	192.168.220.118 server08s
Access LAN	10.1.1.108 server08	10.1.1.118 server18

## Initializing the Cluster

Make sure that the files `/etc/hosts`, `/etc/passwd`, `/etc/shadow`, and `/etc/group` are copied to the other node.

Run the following command as root on each cluster node to configure the cluster infrastructure and start the cluster:

```
pcs cluster setup --name sap_hana --local server08s server18s
```

The fabric interconnect is not configured to support multicast between the nodes. CMAN communication must be set to unicast.

Remember to increase the `config_version` number by 1 when you change the cluster configuration.

```
[root@server18 cluster]# vi /etc/cluster/cluster.conf
```

```
<cluster config_version="9" name="sap_hana">
  <fence_daemon/>
  <clusternodes>
    <clusternode name="server08s" nodeid="1">
      <fence>
        <method name="pcmk-method">
          <device name="pcmk-redirect" port="server08s"/>
        </method>
      </fence>
    </clusternode>
    <clusternode name="server18s" nodeid="2">
      <fence>
        <method name="pcmk-method">
          <device name="pcmk-redirect" port="server18s"/>
        </method>
      </fence>
    </clusternode>
  </clusternodes>
</cluster>
```

```

    </clusternode>
</clusternodes>
<cmn expected_votes="1" two_node="1" transport="udpu"/>
<fencedevices>
  <fencedevice agent="fence_pcmk" name="pcmk-redirect"/>
</fencedevices>
<rm>
  <failoverdomains/>
  <resources/>
</rm>
</cluster>

```

```

[root@server18 cluster]#
[root@server08 ~]# service pacemaker start
Starting cluster:
  Checking if cluster has been disabled at boot...      [ OK ]
  Checking Network Manager...                          [ OK ]
  Global setup...                                       [ OK ]
  Loading kernel modules...                            [ OK ]
  Mounting configfs...                                  [ OK ]
  Starting cman...                                      [ OK ]
  Waiting for quorum...                                 [ OK ]
  Starting fenced...                                    [ OK ]
  Starting dlm_controld...                             [ OK ]
  Tuning DLM kernel config...                          [ OK ]
  Starting gfs_controld...                             [ OK ]
  Unfencing self...                                    [ OK ]
  Joining fence domain...                              [ OK ]
Starting Pacemaker Cluster Manager                    [ OK ]
[root@server08 ~]#

```

```

[root@server18 ~]# pcs cluster setup --name hanasr --local server08 server18
[root@server18 ~]# service pacemaker start
Starting cluster:
  Checking if cluster has been disabled at boot...      [ OK ]
  Checking Network Manager...                          [ OK ]
  Global setup...                                       [ OK ]
  Loading kernel modules...                            [ OK ]
  Mounting configfs...                                  [ OK ]
  Starting cman...                                      [ OK ]
  Waiting for quorum...                                 [ OK ]
  Starting fenced...                                    [ OK ]
  Starting dlm_controld...                             [ OK ]
  Tuning DLM kernel config...                          [ OK ]

```

```
Starting gfs_controld... [ OK ]
Unfencing self... [ OK ]
Joining fence domain... [ OK ]
Starting Pacemaker Cluster Manager [ OK ]
[root@server18 ~]#
```

```
[root@server18 cluster]# pcs status
Cluster name: sap_hana
WARNING: no stonith devices and stonith-enabled is not false
Last updated: Tue Jun 9 00:37:50 2015
Last change: Tue Jun 9 00:34:45 2015
Stack: cman
Current DC: server08 - partition with quorum
Version: 1.1.11-97629de
2 Nodes configured
0 Resources configured

Online: [ server08 server18 ]

Full list of resources:

[root@server18 cluster]#
```

#### Configure the basic cluster parameters:

```
pcs property set stonith-enabled=false
pcs property set no-quorum-policy="ignore"
pcs resource defaults default-resource-stickiness=1000
pcs resource defaults default-migration-threshold=5000
pcs resource op defaults timeout=600s
```

#### Configure the STONITH Devices

To configure the STONITH devices, collect the Intelligent Platform Management Interface (IPMI) address either from the IMC console or from the Cisco UCS service profile. Figure 2 shows the information for server08, and Figure 3 shows the information for server18.

Figure 2. IPMI Information for server08

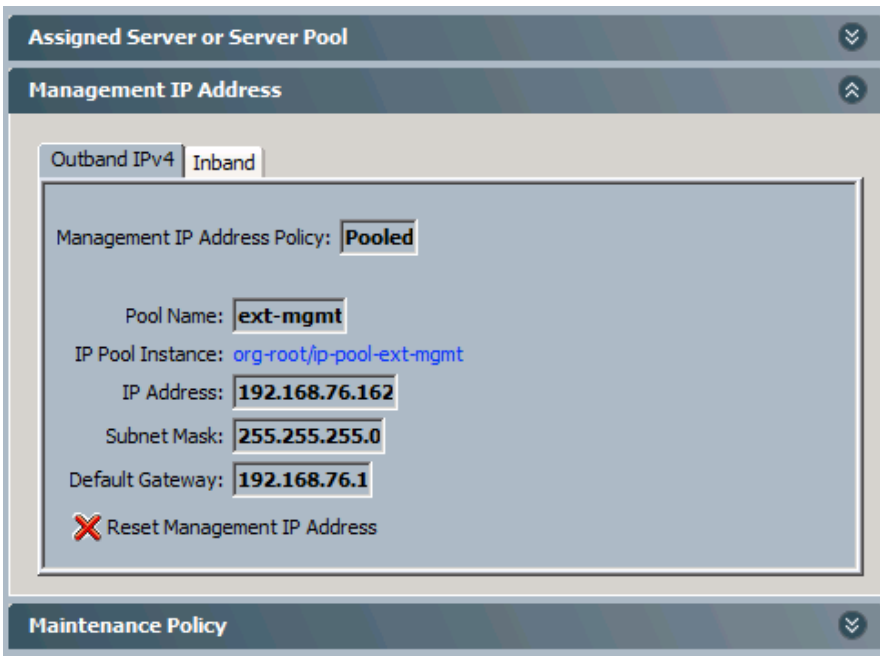
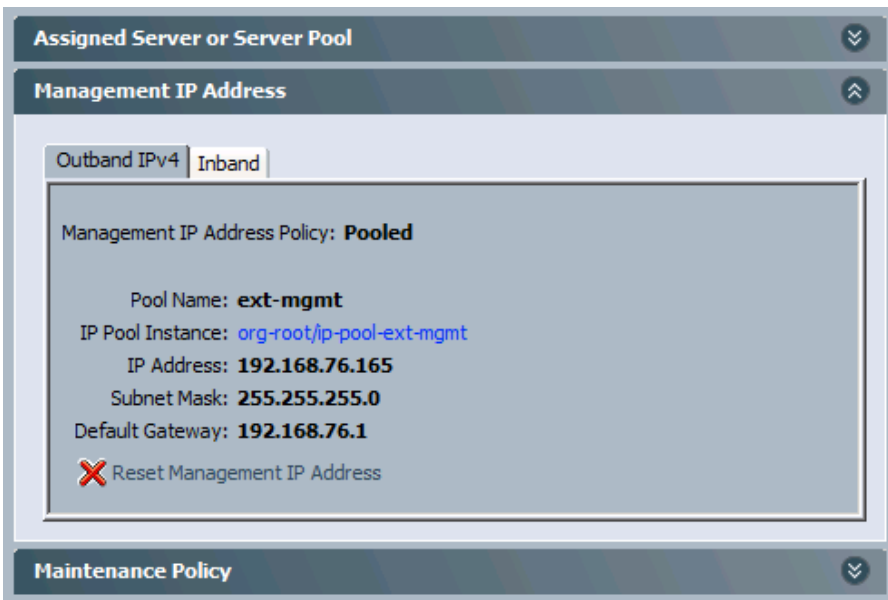


Figure 3. IPMI Information for server18



Add those entries to the /etc/hosts file and synchronize the host file with the other node:

```
vi /etc/hosts
...
..
192.168.76.162  server08-ipmi
192.168.76.165  server18-ipmi
...
...

scp /etc/hosts server18:/etc/
```

Test the IPMI connectivity on **both** nodes:

```
ipmitool -H server18-ipmi -U sapadm -P cisco -I lanplus power status
Chassis Power is on
```

```
ipmitool -H server08-ipmi -U sapadm -P cisco -I lanplus power status
Chassis Power is on
```

Configure the IPMI in the cluster:

```
pcs stonith create st_ipmi_node1 fence_ipmilan ipaddr="server08-ipmi"
login="sapadm" passwd="cisco" pcmk_host_list="server08"
```

```
pcs stonith create st_ipmi_node2 fence_ipmilan ipaddr="server18-ipmi"
login="sapadm" passwd="cisco" pcmk_host_list="server18"
```

```
pcs constraint location st_ipmi_node1 avoids server08
```

```
pcs constraint location st_ipmi_node2 avoids server18
```

Test the Cluster Configuration

Test the cluster configuration:

```
pcs stonith
  st_ipmi_node1 (stonith:fence_ipmilan): Started
  st_ipmi_node2 (stonith:fence_ipmilan): Started
```

```
pcs status
Cluster name: sap_hana
Last updated: Wed Jun 10 00:27:35 2015
Last change: Wed Jun 10 00:23:33 2015
Stack: cman
Current DC: server18 - partition with quorum
Version: 1.1.11-97629de
2 Nodes configured
2 Resources configured
```

---

```
Online: [ server08 server18 ]
```

```
Full list of resources:
```

```
st_ipmi_node1 (stonith:fence_ipmilan): Started server18
st_ipmi_node2 (stonith:fence_ipmilan): Started server08
```

## Test STONITH

Call **ipmi-kill** to make sure that the STONITH IPMI is working correctly.

On server08, enter:

```
pcs stonith fence server18
```

The server18 node must immediately reboot.

Restart Pacemaker when the node is back online.

On server18, enter:

```
pcs stonith fence server08
```

The server08 node must immediately reboot.

Restart Pacemaker when the node is back online.

If both nodes reboot, the STONITH interfaces are working correctly.

```
[root@server08 ~]# pcs status
Cluster name: sap_hana
Last updated: Tue Jun 16 23:02:20 2015
Last change: Wed Jun 10 00:23:33 2015
Stack: cman
Current DC: server18 - partition with quorum
Version: 1.1.11-97629de
2 Nodes configured
2 Resources configured
```

```
Online: [ server08 server18 ]
```

```
Full list of resources:
```

```
st_ipmi_node1 (stonith:fence_ipmilan): Started server18
st_ipmi_node2 (stonith:fence_ipmilan): Started server08
```

```
[root@server08 ~]#
```

## Create a Cluster IP Address

Create a cluster IP address:

```
pcs resource create rsc_ip_SAPHana_RH1_HDB00 IPAddr2 ip="192.168.220.200"
```

```
pcs status
```

```
Cluster name: sap_hana
Last updated: Tue Jun 16 23:08:35 2015
Last change: Tue Jun 16 23:08:21 2015
Stack: cman
Current DC: server18 - partition with quorum
Version: 1.1.11-97629de
2 Nodes configured
3 Resources configured
```

```
Online: [ server08 server18 ]
```

```
Full list of resources:
```

```
st_ipmi_node1 (stonith:fence_ipmilan): Started server18
st_ipmi_node2 (stonith:fence_ipmilan): Started server08
rsc_ip_SAPHana_RH1_HDB00 (ocf::heartbeat:IPAddr2): Started server08
```

```
[root@server08 ~]#
```

## Create the SAP HANA Topology Resource

Create the SAP HANA topology resource:

```
pcs resource create rsc_SAPHanaTopology_RH1_HDB00 SAPHanaTopology \  
SID=RH1 InstanceNumber=00 op start timeout=600 \  
op stop timeout=300 op monitor interval=10 timeout=600
```

## Create the Clone Resource for the Topology Resource

Create the clone resource for the topology resource:

```
pcs resource clone rsc_SAPHanaTopology_RH1_HDB00 meta is-managed=true \  
clone-node-max=1 target-role="Started"
```

```
pcs status
```

```
Cluster name: sap_hana
Last updated: Tue Jun 16 23:16:43 2015
Last change: Tue Jun 16 23:14:33 2015
```



```
Stack: cman
Current DC: server18 - partition with quorum
Version: 1.1.11-97629de
2 Nodes configured
5 Resources configured

Online: [ server08 server18 ]

Full list of resources:

st_ipmi_node1 (stonith:fence_ipmilan):      Started server18
st_ipmi_node2 (stonith:fence_ipmilan):      Started server08
rsc_ip_SAPHana_RH1_HDB00 (ocf::heartbeat:IPaddr2):      Started server08
Clone Set: rsc_SAPHanaTopology_RH1_HDB00-clone [rsc_SAPHanaTopology_RH1_HDB00]
Started: [ server08 server18 ]
```

## Create the SAP HANA Resource

Create the SAP HANA resource.

**Note:** The timeout values shown here for the start, stop, promote, and demote operations are only examples and may need to be adjusted depending on the local SAP HANA setup. For example, large SAP HANA databases can take longer to start, and therefore the start timeout value may need to be increased.

```
pcs resource create rsc_SAPHana_RH1_HDB00 SAPHana \
  SID=RH1 \
  InstanceNumber=00 \
  PREFER_SITE_TAKEOVER=true \
  DUPLICATE_PRIMARY_TIMEOUT=7200 \
  AUTOMATED_REGISTER=false \
  op start timeout=3600 \
  op stop timeout=3600 \
  op promote timeout=3600 \
  op demote timeout=3600 \
  op monitor interval=59 role="Master" timeout=700 \
  op monitor interval=61 role="Slave" timeout=700
```

## Create a SAP HANA Master and Slave

Create a SAP HANA master and slave:

```
pcs resource master msl_rsc_SAPHana_RH1_HDB00 rsc_SAPHana_RH1_HDB00 \ meta is-
managed=true notify=true clone-max=2 clone-node-max=1 \
target-targetrole="Started"
```

## Set Up Constraints

Set up constraints for the system:

```
pcs constraint colocation \  
add rsc_ip_SAPHana_RH1_HDB00 with master msl_rsc_SAPHana_RH1_HDB00 2000  
  
pcs constraint order \  
rsc_SAPHanaTopology_RH1_HDB00-clone then msl_rsc_SAPHana_RH1_HDB00  
symmetrical=false
```

## Results

When you start the cluster (Pacemaker), the SAP HANA database automatically starts too.

```
service pacemaker restart  
service pacemaker restart
```

### **pcs status**

```
Cluster name: sap_hana  
Last updated: Wed Jun 17 00:05:34 2015  
Last change: Wed Jun 17 00:05:32 2015  
Stack: cman  
Current DC: server08 - partition with quorum  
Version: 1.1.11-97629de  
2 Nodes configured  
7 Resources configured
```

```
Online: [ server08 server18 ]
```

Full list of resources:

```
st_ipmi_node1 (stonith:fence_ipmilan): Started server18  
st_ipmi_node2 (stonith:fence_ipmilan): Started server08  
rsc_ip_SAPHana_RH1_HDB00 (ocf::heartbeat:IPaddr2): Started server08  
Clone Set: rsc_SAPHanaTopology_RH1_HDB00-clone [rsc_SAPHanaTopology_RH1_HDB00]  
Started: [ server08 server18 ]  
Master/Slave Set: msl_rsc_SAPHana_RH1_HDB00 [rsc_SAPHana_RH1_HDB00]  
Masters: [ server08 ]  
Slaves: [ server18 ]
```

## Enable STONITH

You **must** enable STONITH in production environments:

```
pcs property set stonith-enabled=true
```

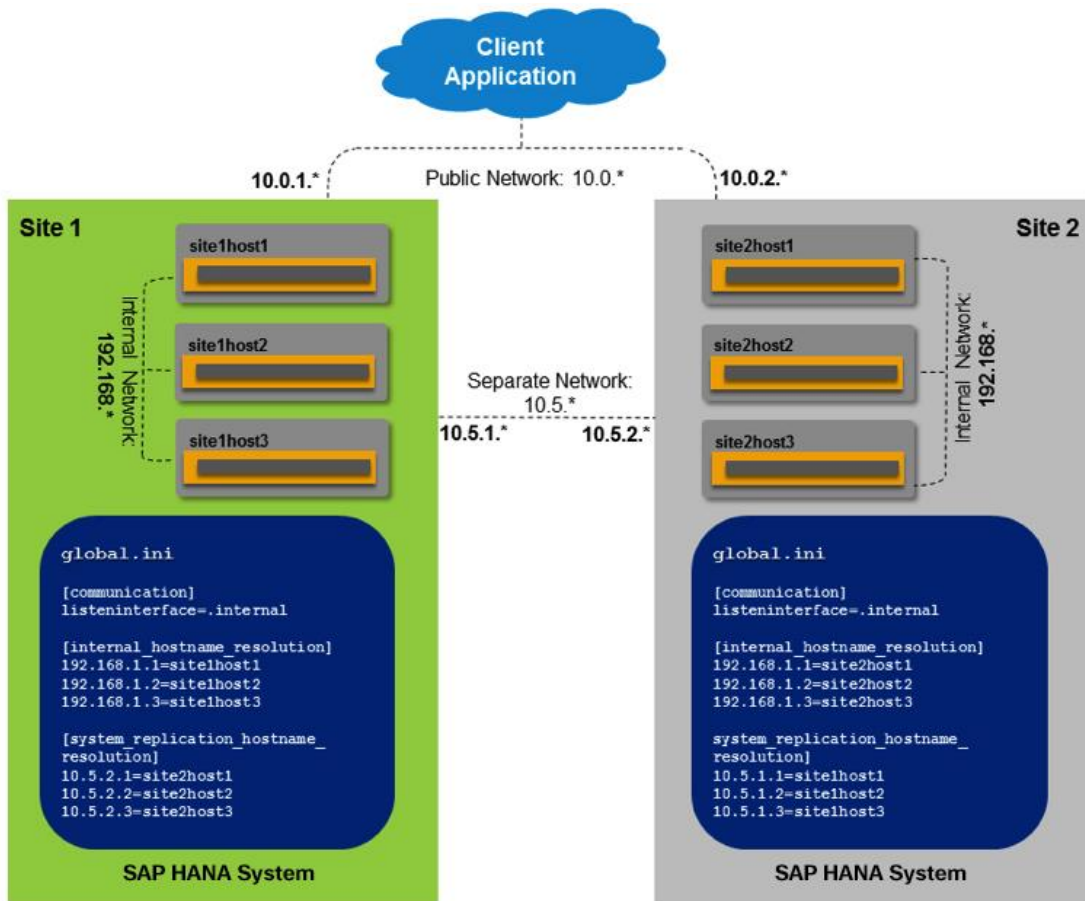
## Setting Up the Network for System Replication

If nothing is configured, SAP HANA uses the access network to synchronize the systems. The solution in this document uses a separate network for system replication, so you should configure the network (Figure 4).

For more information, see:

- [http://help.sap.com/saphelp\\_hanaplatform/helpdata/en/32/c22f81e8c14268a4e5de01cd033e8f/content.htm](http://help.sap.com/saphelp_hanaplatform/helpdata/en/32/c22f81e8c14268a4e5de01cd033e8f/content.htm)
- [http://help.sap.com/saphelp\\_hanaplatform/helpdata/en/0c/6738ab85c64da1aed0fa91c25ed47c/content.htm](http://help.sap.com/saphelp_hanaplatform/helpdata/en/0c/6738ab85c64da1aed0fa91c25ed47c/content.htm)

**Figure 4.** Network Configuration for System Replication



Source: SAP AG

Change the configuration as shown here:

```
server08:/usr/sap/HDB00> netstat -ia
Kernel Interface table
Iface  MTU Met    RX-OK RX-ERR RX-DRP RX-OVR    TX-OK TX-ERR TX-DRP TX-OVR Flg
access  1500 0    17116    0 17085    0      7      0      0      0 BMRU
appl    9000  0      0          0  0        0      0      0      0      0 BM
backup  9000  0    17099      0 17085    0      7      0      0      0 BMRU
datasrc 9000  0    17116      0 17085    0      8      0      0      0 BMRU
eth0    1500  0 96907431    0 17085    0 115121556    0      0      0 BMRU
lo      16436 0 48241697    0  0        0 48241697    0      0      0 LRU
mgmt    1500  0    41529      0 17085    0    2043      0      0      0 BMRU
nfsd    9000  0    9547248    0 17085    0 10164551    0      0      0 BMRU
nfsl    9000  0    1088756    0 17087    0  1515500    0      0      0 BMRU
server  9000  0    1296168    0 17085    0 10105637    0      0      0 BMRU
sysrep  9000 0    17118    0 17085    0      8      0      0      0 BMRU
server08:/usr/sap/RH1/HDB00>
```

#### server08

```
server08:/usr/sap/RH1/HDB00> cdglo
server08:/usr/sap/RH1/SYS/global> cd hdb/custom/config
server08:/usr/sap/RH1/SYS/global/hdb/custom/config> cat global.ini
[persistence]
basepath_datavolumes = /hana/data/RH1
basepath_logvolumes = /hana/log/RH1

[system_replication]
mode = primary
actual_mode = primary
site_id = 1
site_name = SJC-Rack-1
```

#### server18

```
server18:/usr/sap/RH1/HDB00> cdglo
server18:/usr/sap/RH1/SYS/global> cd hdb/custom/config/
server18:/usr/sap/RH1/SYS/global/hdb/custom/config> cat global.ini
[persistence]
basepath_datavolumes = /hana/data/RH1
basepath_logvolumes = /hana/log/RH1

[system_replication]
site_id = 2
```

```
mode = sync
actual_mode = sync
site_name = SJC-Rack-2

[system_replication_site_masters]
1 = server08:30001
```

Reroute the network traffic from the access network to the system replication network.

#### **server08**

```
server08:/usr/sap/RH1/SYS/global/hdb/custom/config> cat global.ini
```

```
[persistence]
basepath_datavolumes = /hana/data/RH1
basepath_logvolumes = /hana/log/RH1

[system_replication]
mode = primary
actual_mode = primary
site_id = 1
site_name = SJC-Rack-1
```

#### **[communication]**

```
listeninterface=.internal
```

#### **[internal\_hostname\_resolution]**

```
192.168.220.101 = server08
192.168.220.111 = server18
```

#### **[system\_replication\_hostname\_resolution]**

```
192.168.222.101 = server08
192.168.222.111 = server18
```

#### **server18**

```
server18:/usr/sap/RH1/SYS/global/hdb/custom/config> cat global.ini
```

```
[persistence]
basepath_datavolumes = /hana/data/RH1
basepath_logvolumes = /hana/log/RH1

[system_replication]
site_id = 2
mode = sync
actual_mode = sync
site_name = SJC-Rack-2
```

```
[system_replication_site_masters]
```

```
1 = server08:30001

[communication]
listeninterface=.internal

[internal_hostname_resolution]
192.168.220.101 = server08
192.168.220.111 = server18

[system_replication_hostname_resolution]
192.168.222.101 = server08
192.168.222.111 = server18
```

Restart SAP HANA.

### Choosing a SAP HANA Synchronization Mode

Several log replication modes are available to send log information to the secondary system. You need to decide which mode to use.

- Synchronous (**mode=sync**): In this mode, the log write operation is considered successful when the log entry has been written to the log volume of the primary and secondary systems. If the connection to the secondary system is lost, the primary system continues transaction processing and writes the changes only to the local disk. No data loss occurs in this scenario as long as the secondary system is connected. Data loss can occur if takeover is performed while the secondary system is disconnected.
- Synchronous in memory (**mode=syncmem**): In this mode, the log write operation is considered successful when the log entry has been written to the log volume of the primary system and transmission of the log has been acknowledged by the secondary system after the log has been copied to memory. If the connection to the secondary system is lost, the primary system continues transaction processing and writes only the changes to the local disk. Data loss can occur if the primary and secondary systems fail at the same time when the secondary system is connected or takeover is performed when the secondary system is disconnected. This option provides better performance, because it is not necessary to wait for disk I/O on the secondary system, but it is more vulnerable to data loss.
- Asynchronous (**mode=async**): In this mode, the primary system sends a redo log buffer to the secondary system asynchronously. The primary system commits a transaction when it has been written to the log file of the primary system and sent to the secondary system through the network. It does not wait for confirmation from the secondary system. This option provides better performance because it is not necessary to wait for log I/O on the secondary system. Database consistency across all services on the secondary system is guaranteed. However, this option is more vulnerable to data loss. Data changes may be lost on takeover.

To set up system replication, you need to perform the configuration steps on the secondary system. You can complete this configuration using the `hdbnsutil` tool, which initializes the topology of the database during installation, or exports, imports, and converts the topology of an existing database. You also can use SAP HANA Studio.

You can use the full synchronization (**fullsync**) option for system replication. When activated, the **fullsync** option helps ensure that a log buffer is shipped to the secondary system before a commit operation occurs on the local primary system.

As of SPS 08, the **fullsync** option can be enabled for **sync** replication (that is, not for **syncmem**). With the **fullsync** option activated, transaction processing occurs on the primary blocks when the secondary system is not currently connected, and newly created log buffers cannot be shipped to the secondary site. This behavior helps ensure that no transaction can be locally committed without shipping the log buffers to the secondary site.

The **fullsync** option can be switched on and off using the following command:

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

The **fullsync** option is not supported by the cluster because configuring this option in a cluster environment makes no sense.

## Performing Cluster Switchback After a Failure

When the cluster detects an error on the primary site, the cluster activates the secondary site automatically.

The cluster cannot switch back the SAP HANA database after the primary system is repaired. This step is manual and is described on the following pages.

This is the result after the cluster switchover:

```
server08 > SAP HANA is down.
```

```
server18 > SAP HANA is active.
```

## Resynchronize the Primary Database

The primary SAP HANA database is down:

```
server08:~ # su - rhladm
server08:/usr/sap/RH1/HDB00> HDB info
USER      PID  PPID %CPU   VSZ   RSS COMMAND
rhladm    24542 24541 12.3 7857084 144916 hdbnsutil -sr_state --sapcontrol=1
rhladm    24646 24542  3.8   3396   192 \_ /bin/sh -c
PATH="/bin:/sbin:/usr/bin:/usr/sbin:/etc" sysctl -n "kernel.shmall"
rhladm    24392 24391  0.1  13884  2740 -sh
rhladm    24592 24392  0.3  12900  1720 \_ /bin/sh /usr/sap/RH1/HDB00/HDB info
rhladm    24662 24592  5.7   4944   876 \_ ps fx -U rhladm -o
user,pid,ppid,pcpu,vsz,rss,args
rhladm    9840 1 0.0 218884 75584 /usr/sap/RH1/HDB00/exe/sapstartsrv
pf=/usr/sap/RH1/SYS/profile/RH1_HDB00_server08 -D -u rhladm
server08:/usr/sap/RH1/HDB00>
```

Enable resynchronization (server08 now becomes the secondary system):

```
server08:~> hdbnsutil -sr_register --remoteHost=server18 --remoteInstance=00 --
mode=sync --name=SJC-Rack-1

adding site ...
```

```
checking for inactive nameserver ...
nameserver server08:30001 not responding.
collecting information ...
registered at 192.168.222.111 (server18)
updating local ini files ...
done.
```

Start the database now:

```
❑ server08:~> sapcontrol -nr 00 -function StartSystem
25.03.2015 15:42:40
StartSystem
OK
```

#### HDB info

```
USER      PID  PPID %CPU  VSZ  RSS COMMAND
rhladm   39342 39341  5.1  50740  4804 python
exe/python_support/landscapeHostConfiguration.py
rhladm   39211 39210 14.4 1079000 104080 python
/usr/sap/RH1/HDB00/exe/python_support/landscapeHostConfiguration.py
rhladm   39402 39211  3.1   300    4  \_ sysctl -n kernel.shmall
rhladm   24392 24391  0.0  13884  2780 -sh
rhladm   39261 24392  0.3  12900  1720 \_ /bin/sh /usr/sap/RH1/HDB00/HDB info
rhladm   39404 39261  5.5  4944   876 \_ ps fx -U rhladm -o
user,pid,ppid,pcpu,vsz,rss,args
rhladm   33980 1  0.0  22024  1512 sapstart
pf=/usr/sap/RH1/SYS/profile/RH1_HDB00_server08
rhladm   34060 33980  0.0 825832 196608 \_
/usr/sap/RH1/HDB00/server08/trace/hdb.sapANA_HDB00 -d -nw -f
/usr/sap/RH1/HDB00/server08/daemon.ini
pf=/usr/sap/RH1/SYS/profile/RH1_HDB00_server08
rhladm   34145 34060  1.4 13646420 1079384 \_ hdbnameserver
rhladm   34290 34060  0.6 9653164 278408 \_ hdbpreprocessor
rhladm   34295 34060  0.4 9460648 267816 \_ hdbcompileserver
rhladm   34441 34060  2.3 10835544 1776224 \_ hdbindexserver
rhladm   34444 34060  2.1 9828516 1481288 \_ hdbstatisticsserver
rhladm   34447 34060  1.3 12396176 715808 \_ hdbxsengine
rhladm   34652 34060  0.1 205464 33000 \_ sapwebdisp_hdb
pf=/usr/sap/RH1/HDB00/server08/wdisp/sapwebdisp.pfl -f
/usr/sap/RH1/HDB00/server08/trace/dev_webdisp
rhladm   9840 1  0.0 219400 75704 /usr/sap/RH1/HDB00/exe/sapstartsrv
pf=/usr/sap/RH1/SYS/profile/RH1_HDB00_server08 -D -u rhladm
server08:~>
```



## Activate the Original Master as the Primary System

Switchover the IP address to server08:

```
server08:~ # crm resource migrate rsc_ip_RH1_HDB00 server08
```

Switchover the database:

```
server08:~ # su - rhladm
server08:/usr/sap/RH1/HDB00> hdbnsutil -sr_takeover
checking local nameserver ...
done.
server08:/usr/sap/RH1/HDB00>
```

The source database (server18) will be shut down now after the takeover process.

## Reenable the Secondary Site

Reenable the secondary site as follows:

```
server18:/var/log # su - rhladm
server18:/usr/sap/RH1/HDB00> hdbnsutil -sr_register --remoteHost=server08 --
remoteInstance=00 --mode=sync --name=SJC-Rack-2
adding site ...
checking for inactive nameserver ...
nameserver server18:30001 not responding.
collecting information ...
registered at 192.168.222.101 (server08)
updating local ini files ...
done.
server18:/usr/sap/RH1/HDB00> sapcontrol -nr 00 -function StartSystem
25.03.2015 16:41:23
StartSystem
OK
server18:/usr/sap/RH1/HDB00> exit
```

## Disabling System Replication

Before disabling system replication, verify that it is enabled.

```
hdbsql -U slehaloc 'select distinct REPLICATION_STATUS from
SYS.M_SERVICE_REPLICATION'
REPLICATION_STATUS
"ACTIVE"
```

Now stop the secondary site.

```
server18
sapcontrol -nr 00 -function StopSystem
25.01.2015 22:26:43
StopSystem
OK
```

```
server08
hdbsql -U slehaloc 'select distinct REPLICATION_STATUS from
SYS.M_SERVICE_REPLICATION'
REPLICATION_STATUS
"ERROR"
```

```
server18
server18:/ > hdbnsutil -sr_unregister
unregistering site ...
nameserver server18:30001 not responding.
nameserver server18:30001 not responding.
checking for inactive nameserver ...
nameserver server18:30001 not responding.
nameserver is shut down, proceeding ...
opening persistence ...
run as transaction master
updating topology for system replication takeover ...
mapped host server08 to server18
sending unregister request to primary site (2) ...
```

```
#####
# CAUTION: You must start the database in order to complete the unregistration!
#####
done.
```

```
server08
server08:/ > hdbnsutil -sr_disable
checking local nameserver:
```

```
checking for inactive nameserver ...
nameserver is running, proceeding ...
done.
server08:/ > hdbnsutil -sr_state
checking for active or inactive nameserver ...

System Replication State
~~~~~

mode: none

done.
server08:/usr/sap/RH1/SYS/global/hdb/custom/config>
```

## For More Information

Consult the following resources for more information:

- SAP HANA takeover process: <http://scn.sap.com/docs/DOC-52345>
- Introduction to SAP HANA high availability: <http://scn.sap.com/docs/DOC-60334>
- SAP guide to performing system replication: <https://scn.sap.com/docs/DOC-47702>
- SAP HANA administration guide: [http://help.sap.com/hana/SAP\\_HANA\\_Administration\\_Guide\\_en.pdf](http://help.sap.com/hana/SAP_HANA_Administration_Guide_en.pdf)
- System replication blog: <http://scn.sap.com/community/hana-in-memory/blog/2013/12/16/sap-hana-system-replication--using-hdbnsutil-sr>
- Red Hat system replication with Pacemaker: <https://access.redhat.com/articles/1466063>
- Cluster SBD fencing: [http://www.linux-ha.org/wiki/SBD\\_Fencing](http://www.linux-ha.org/wiki/SBD_Fencing)
- SAP HANA system replication:  
[http://help.sap.com/saphelp\\_hanaplatform/helpdata/en/09/9caa1959ce4b3fa1144562fa09e163/content.htm](http://help.sap.com/saphelp_hanaplatform/helpdata/en/09/9caa1959ce4b3fa1144562fa09e163/content.htm)
- SAP HANA multinode network configuration:  
[http://help.sap.com/saphelp\\_hanaplatform/helpdata/en/0c/6738ab85c64da1aed0fa91c25ed47c/content.htm](http://help.sap.com/saphelp_hanaplatform/helpdata/en/0c/6738ab85c64da1aed0fa91c25ed47c/content.htm)
- Takeover process: <http://scn.sap.com/docs/DOC-52345>
- Network setup for system replication:
  - [http://help.sap.com/saphelp\\_hanaplatform/helpdata/en/32/c22f81e8c14268a4e5de01cd033e8f/content.htm](http://help.sap.com/saphelp_hanaplatform/helpdata/en/32/c22f81e8c14268a4e5de01cd033e8f/content.htm)
  - [http://help.sap.com/saphelp\\_hanaplatform/helpdata/en/0c/6738ab85c64da1aed0fa91c25ed47c/content.htm](http://help.sap.com/saphelp_hanaplatform/helpdata/en/0c/6738ab85c64da1aed0fa91c25ed47c/content.htm)
- Switching back and forth using SAP HANA system replication: <http://scn.sap.com/community/hana-in-memory/blog/2013/12/16/sap-hana-system-replication--using-hdbnsutil-sr>



---

**Americas Headquarters**  
Cisco Systems, Inc.  
San Jose, CA

**Asia Pacific Headquarters**  
Cisco Systems (USA) Pte. Ltd.  
Singapore

**Europe Headquarters**  
Cisco Systems International BV Amsterdam,  
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com/go/trademarks](http://www.cisco.com/go/trademarks). Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Printed in USA

C11-735382-00 07/15