

High Availability Guide
SAP Business One, version for SAP HANA
Document Version: 3.0.5 – 2023-11-06

PUBLIC

Setting Up SAP HANA Database High Availability for SAP Business One 10.0 for Automatic Failover

Applicable Release: SAP Business One 10.0, version for SAP HANA



Typographic Conventions

Type Style	Description
<i>Example</i>	Words or characters quoted from the screen. These include field names, screen titles, pushbuttons labels, menu names, menu paths, and menu options. Textual cross-references to other documents.
Example	Emphasized words or expressions.
EXAMPLE	Technical names of system objects. These include report names, program names, transaction codes, table names, and key concepts of a programming language when they are surrounded by body text, for example, SELECT and INCLUDE.
<code>Example</code>	Output on the screen. This includes file and directory names and their paths, messages, names of variables and parameters, source text, and names of installation, upgrade and database tools.
Example	Exact user entry. These are words or characters that you enter in the system exactly as they appear in the documentation.
<Example>	Variable user entry. Angle brackets indicate that you replace these words and characters with appropriate entries to make entries in the system.
EXAMPLE	Keys on the keyboard, for example, F2 or ENTER .

Document History

Version	Date	Change
1.0	2015-04-16	First version.
1.0.1	2015-12-23	<ul style="list-style-type: none">• Section 1.2: Prerequisite for the virtual IP address.• Section 2.3: Installation instructions for SAP Business One, version for SAP HANA.
2.0.1	2016-08-03	<ul style="list-style-type: none">• Updated the document for SAP Business One 9.2, version for SAP HANA.• Support SLES 11 SP4.
2.0.2	2016-08-25	<ul style="list-style-type: none">• Updated the installation part.• Added the upgrading part.
2.1.0	2017-06-01	<ul style="list-style-type: none">• Section 1.1: Updated the high-availability landscape.• Section 1.2: Added configuration instructions for SBD devices, and removed crm_b1.txt in the Others part.
2.1.1	2017-09-29	Section 2.4: Updated the Note in Step 2.
2.1.2	2018-01-24	Added high-availability configuration instructions for SUSE 12 SP1.
2.1.3	2018-02-09	Added SAP Note 2601083 .
2.1.4	2018-4-30	<ul style="list-style-type: none">• Supports SAP Business One 9.3, version for SAP HANA.• Updated introduction.• Updated software and hardware requirements.• Minor revisions.
2.1.5	2018-06-07	Support SLES 12 SP3.
3.0.0	2020-07-16	<ul style="list-style-type: none">• Support SAP Business One 10.0, version for SAP HANA.• Support SLES 15 SP1 for SAP applications.
3.0.1	2021-04-22	Support SLES 15 SP2 for SAP applications.
3.0.2	2021-09-20	<ul style="list-style-type: none">• Removed restriction on SID.• Minor revisions.• Added instruction to section 2.2.1 for upgrade with script delivered with version 3.0.0 and 3.0.1 of this guide.• Added section 4.1: Manual Tasks Without Failover.
3.0.3	2022-09-22	Added instruction on fixing compatibility issue of Python support scripts with Python 3 in SAP HANA 2 Revision 059.04.
3.0.4	2023-01-17	Added two new sections: <ul style="list-style-type: none">• 4.6 Maintenance of Operating SBD Devices in a Live Environment

Version	Date	Change
		<ul style="list-style-type: none">• 4.7 Requesting Technical Support
3.0.5	2023-11-06	<ul style="list-style-type: none">• Support SLES 15 SP3 for SAP applications.• Support SLES 15 SP4 for SAP applications.

Contents

1	Introduction.....	6
1.1	Prerequisites	7
1.2	Related Documentation	18
2	Setup.....	20
2.1	Installation	20
2.1.1	Installing Software in the SAP HANA Box	20
2.1.2	Executing the High-Availability Setup Script	23
2.1.3	Installing SAP Business One, version for SAP HANA.....	24
2.2	Upgrade	26
2.2.1	Upgrading SAP HANA on Host 1 and Host 2.....	26
2.2.2	Upgrading SAP Business One Server Components, version for SAP HANA, on Central Server	27
3	Actions After Failover	28
3.1	Normal Failover	28
3.2	Failover Due to Disk Space Availability	30
4	Daily Maintenance.....	32
4.1	Manual Tasks Without Failover.....	32
4.2	Tenant Database Backup.....	32
4.3	Status Check on Cluster.....	33
4.4	Status Check on Resources.....	33
4.5	Status Check on SAP HANA Replication Function	33
4.6	Maintenance of Operating SBD Devices in a Live Environment	34
4.7	Requesting Technical Support	34

1 Introduction

To ensure business continuity and near-zero data loss in case of planned or unplanned downtime, we provide a high availability solution for the SAP HANA® database server for the SAP Business One® 10.0, version for SAP HANA® application. The high availability solution is made possible by the well-designed system replication mechanism of SAP HANA.

In a highly available environment, you have a primary (active) system and a secondary (standby) system. The secondary system is set up as an exact copy of the primary system and there is a continuous replication of all persisted data to the secondary system.

You can configure high availability for automatic failover. In an automatic failover, the former primary system switches to the secondary system automatically in the event of failure. An automatic failover includes both the analysis and transaction functions of SAP Business One 10.0, version for SAP HANA.

An automatic failover is implemented through clustering of SAP HANA on SUSE® Linux Enterprise Server (SLES), for SAP applications. With SUSE Linux Enterprise Server for SAP Applications, you can even build an OS-level cluster and identify the SAP HANA database as a cluster resource.

This guide provides instructions on how to set up, upgrade and maintain a high availability environment for the SAP HANA database for SAP Business One 10.0, version for SAP HANA, for automatic failover.

Note

Due to unpredictable factors, SAP has reservations about using this solution as part of your disaster recovery plan. For general guidance on disaster recovery, please contact SAP Support.

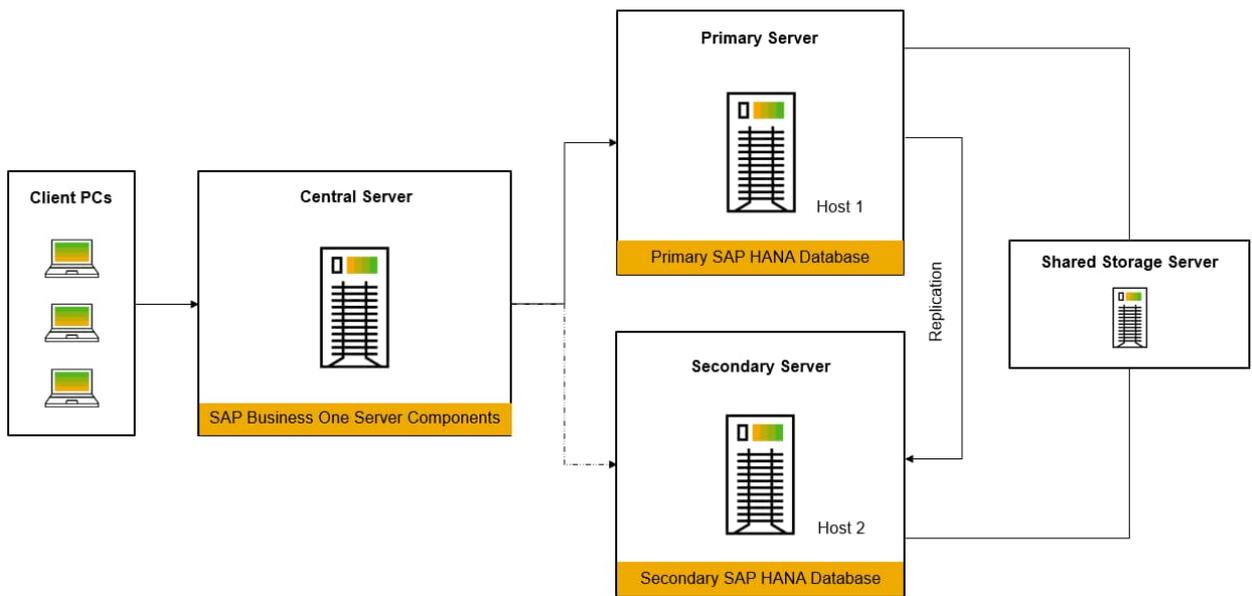
The images and data in this guide are from SAP internal systems, sample data, or demo systems. Any resemblance to real data is purely coincidental.

Landscape

The following figure illustrates the landscape of the high availability environment.

All Linux-based SAP Business One server components are installed on the central server.

We assume the SAP HANA instance on Host 1 as the primary system, and the instance on Host 2 as the secondary system.



1.1 Prerequisites

Before setting up the high availability environment, you must ensure that you have met the following prerequisites.

Database Servers and Network Requirements

- You have installed SUSE Linux Enterprise Server 15 SP2 for SAP Applications and applied the latest updates, including the latest patch of the SAPHanaSR resource agent (0.154.0-4.11.2 or later).
- You have the password of the `root` user.

Caution

You must define the same password for `root` on both SAP HANA servers.

All commands are executed by the `root` user, unless otherwise stated, and are displayed in this style:

`<Command>`.

- Both SAP HANA database servers have been certified for SAP HANA.
- You have configured three IP addresses: two for the SAP HANA database servers and one as a virtual IP address. In addition, the three IP addresses have the same default gateway and are on the same subnet.

The virtual IP address will be managed by the cluster resource manager after the high-availability setup script `ha15.sh` is executed. Do not bind this IP address with any particular network card.

- Both SAP HANA database servers connect to the same switch.
- For hostnames:
 - The hostname of each database server must be in lower case. If you have already defined the hostname in upper case, follow the instructions in SAP Note [1780950](#) to change the hostname to lower case.
 - Each database server can resolve the hostname of the other server. To ensure this, on each database server, add the mapping of the hostname and IP address of the other server to the `/etc/hosts/` file.
- Both database servers connect to a shared storage server of one block through an internal storage area network. The block device is used as an SBD (storage-based death) device and needs to meet the following requirements:
 - No partition exists on the block.
 - Its size is no less than 10 M.

To check if the SBD device meets the requirements, execute this command: `fdisk -l`

```
Disk /dev/mapper/vgsbd-test: 12 MiB, 12582912 bytes, 24576 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

For more information about SBD and STONITH, see http://clusterlabs.org/doc/crm_fencing.html and http://www.linux-ha.org/wiki/SBD_Fencing.

To configure an SBD device, proceed as follows:

1. Configure the central server as an iSCSI Target Server:

Before configuring the central server, make sure that you have met the following prerequisites:

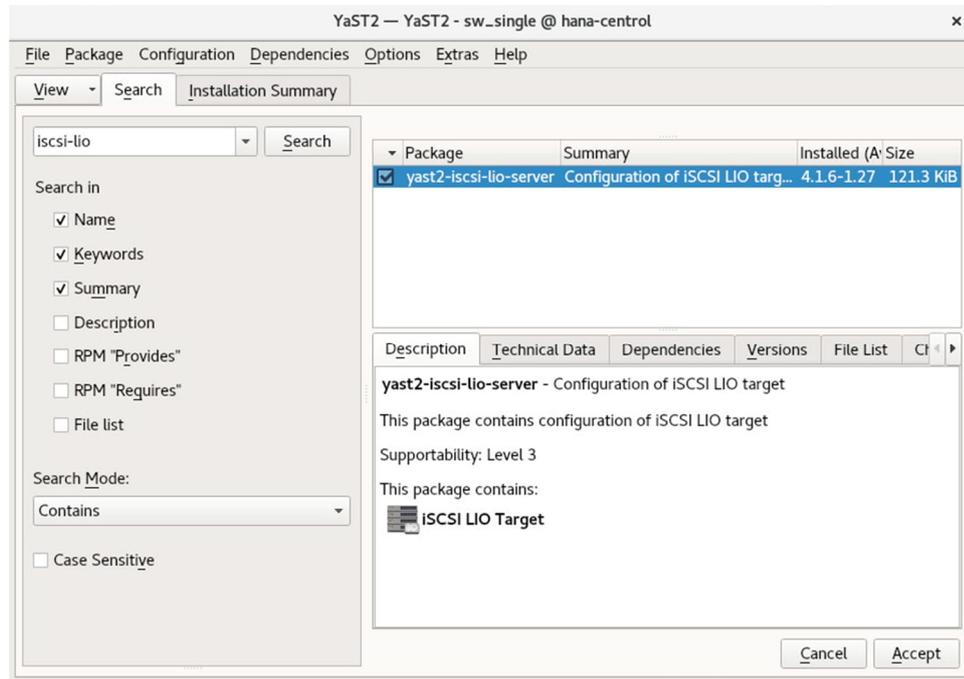
- You have shared storage, or you can use the central server as the shared storage server.
- You must dedicate one small partition as the SBD device.
- The shared storage server should not be virtualized by Hyper-V.
- The SBD device must not make use of host-based RAID.
- The SBD device must not reside on a DRBD instance.
- The SBD device can be shared between different clusters, as long as no more than 255 nodes share the device.

Then follow the below steps to configure the central server as iSCSI Target Server:

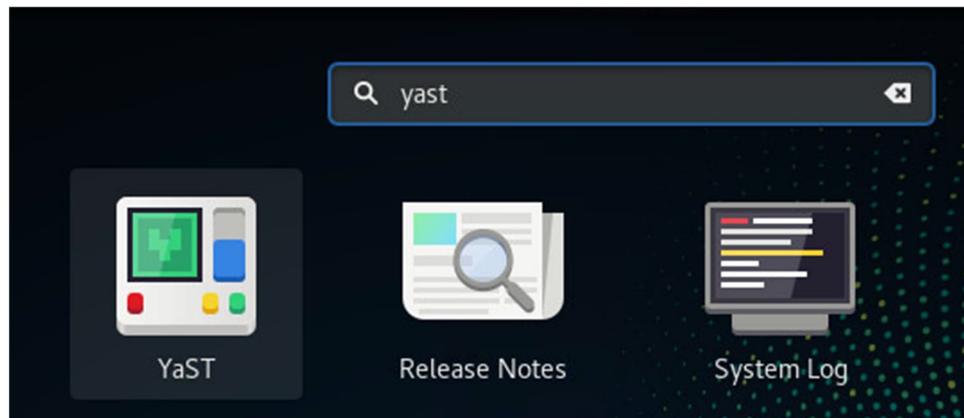
1. Log on to the central server as iSCSI Target Server. Execute command `fdisk -l` to list the SBD device which you prepared before. Here we use:

```
/dev/mapper/hanavg-sbd_test_1
```

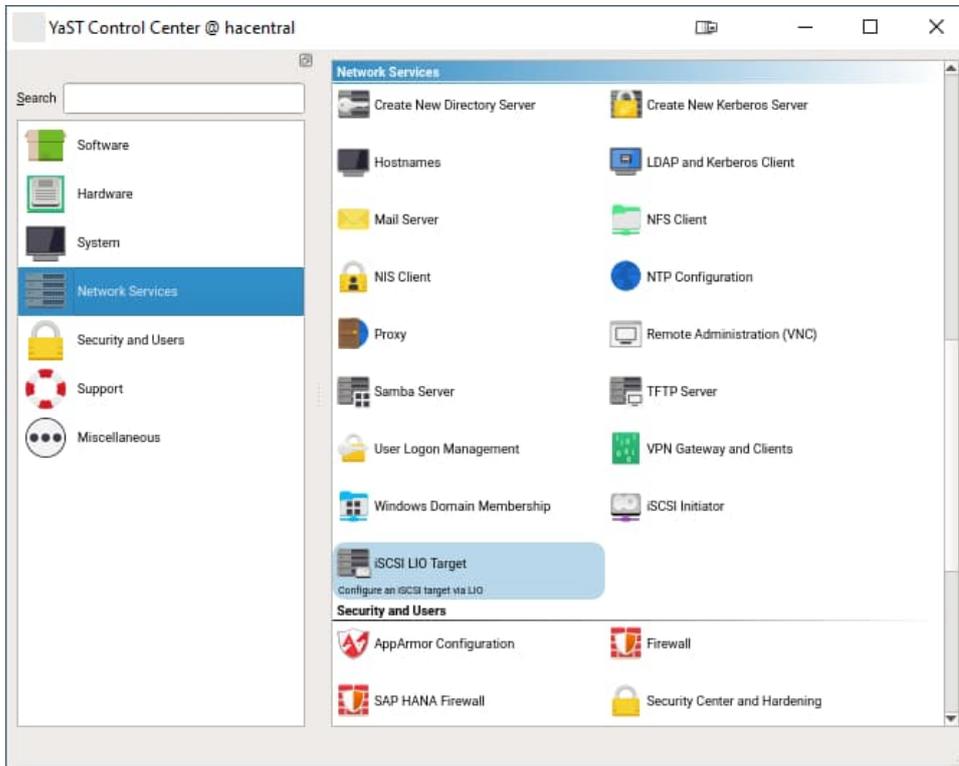
You can copy the device path to be used later.



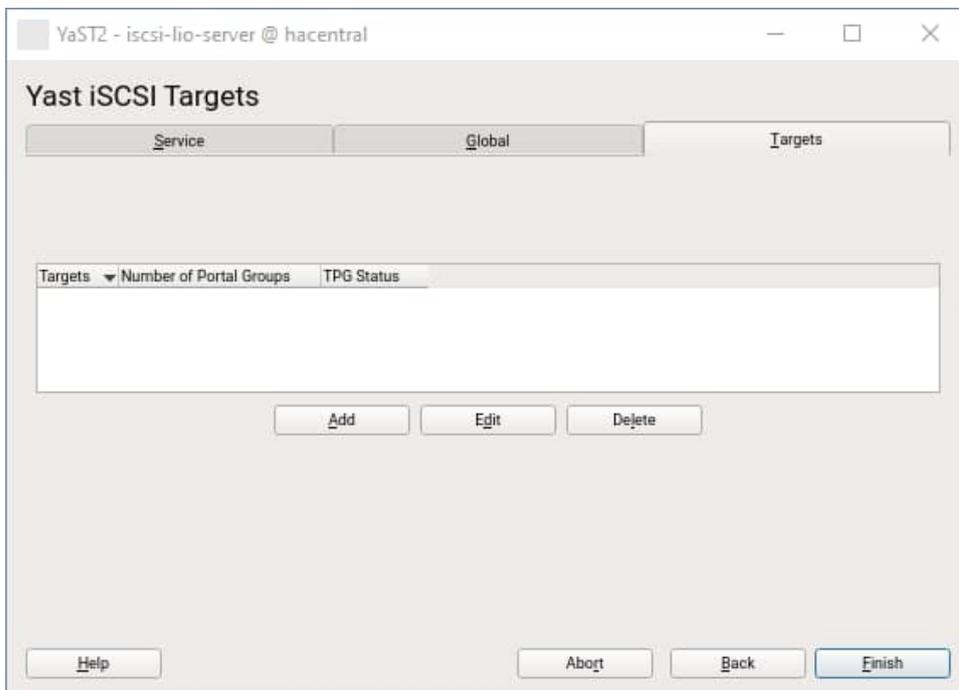
2. Input `yast2`.



- Go to *Network Service/iSCSI Target*.



- On the *Targets* tab, choose *Add* to add a new target.



5. On the *Identifier* tab, enter the iSCSI target name you want and choose *Add*.

YaST2 - iscsi-lio-server @ hacentral

Add iSCSI Target

Target **Identifier** **Portal Grp**

iqn.2020-05.com.example 0d5a 1

IP Address: **Port Number**

3260

Bind all IP addresses

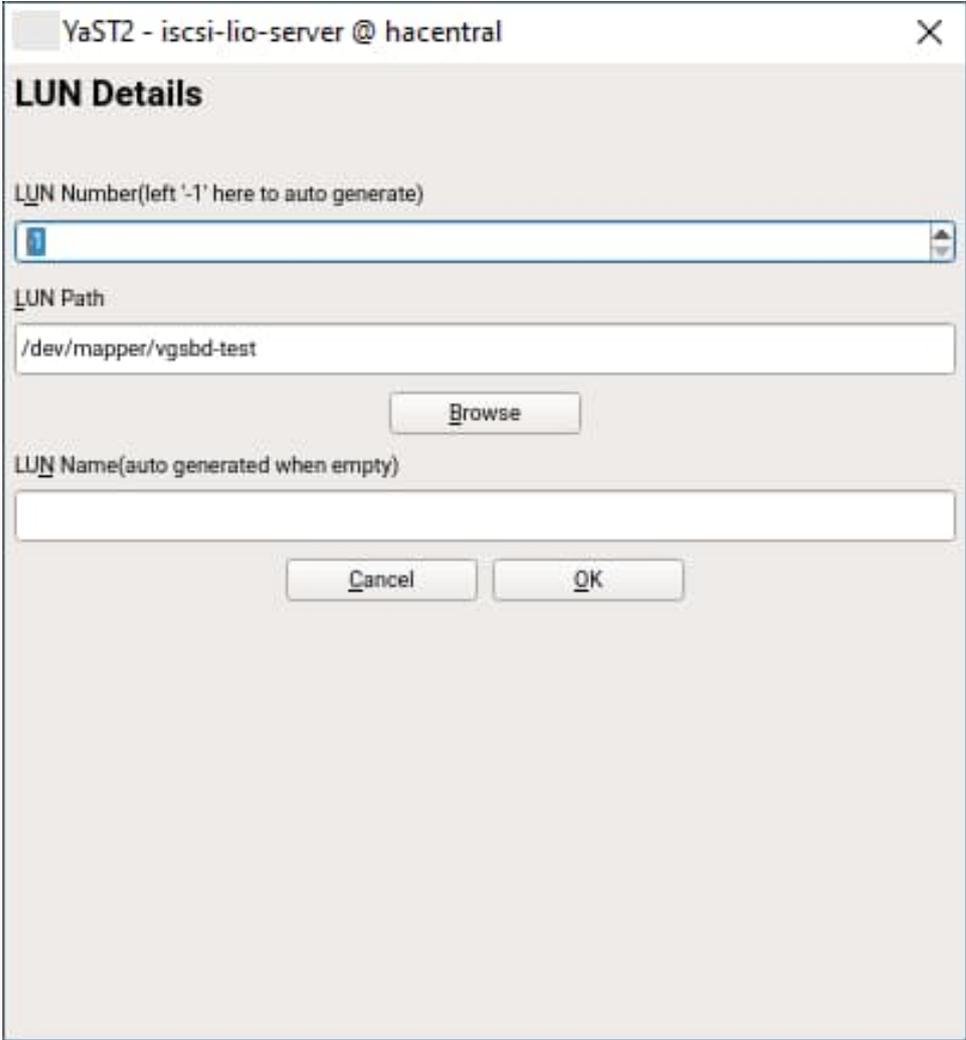
Use Login Authentication

LUN	Name	Path
-----	------	------

Add **Delete**

Help **Abort** **Back** **Next**

- On the *Path* tab, input the devices which you copied before with command `fdisk -l`. Then choose *OK* to add the LUN.

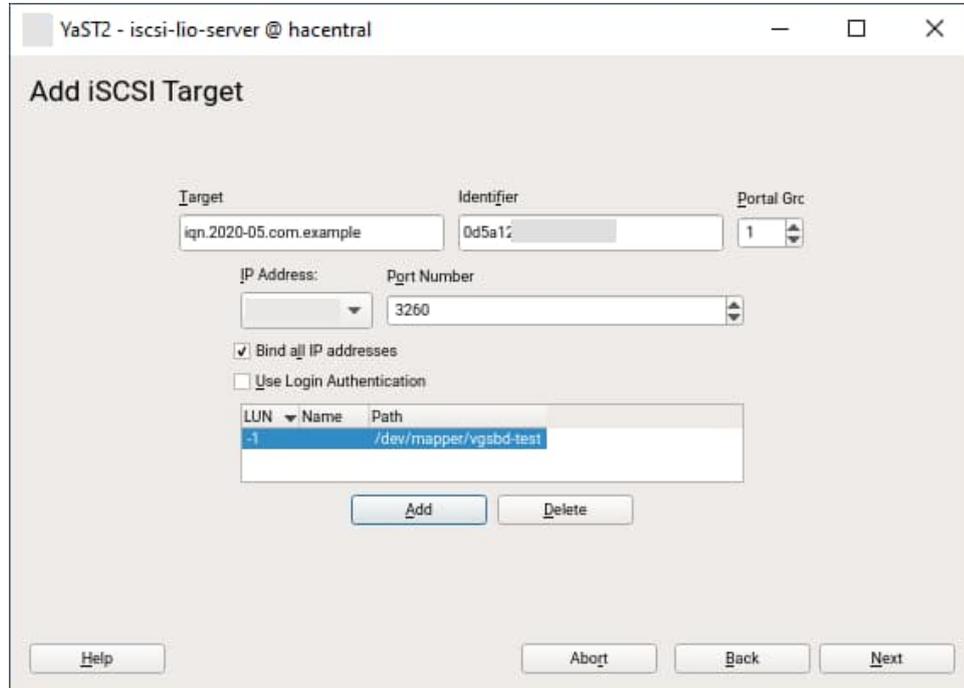


The image shows a YaST2 dialog box titled "YaST2 - iscsi-lio-server @ hacentral" with a close button (X) in the top right corner. The dialog is titled "LUN Details" and contains the following fields and controls:

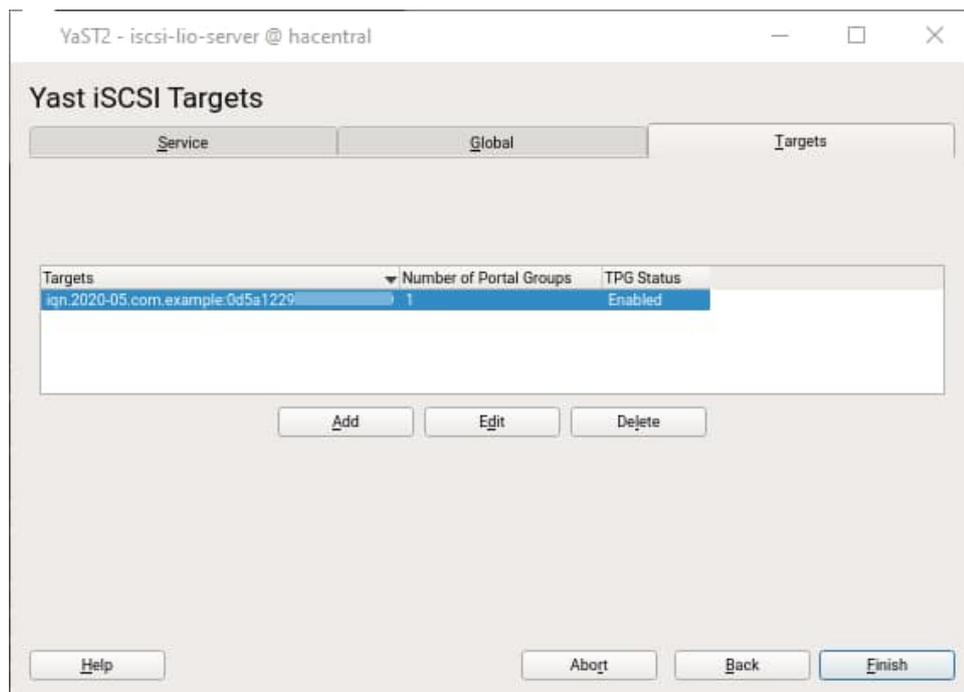
- LUN Number**(left '-1' here to auto generate): A text input field containing the value "1".
- LUN Path**: A text input field containing the path "/dev/mapper/vgsbd-test". Below this field is a "Browse" button.
- LUN Name**(auto generated when empty): An empty text input field.

At the bottom of the dialog, there are two buttons: "Cancel" and "OK".

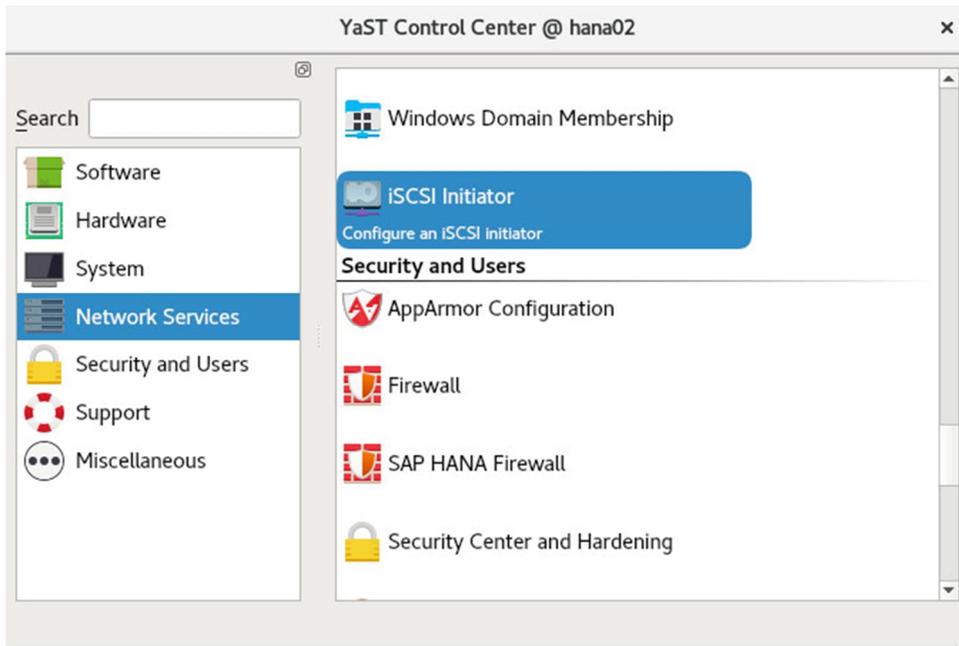
7. Keep the default setting and choose *Next*.



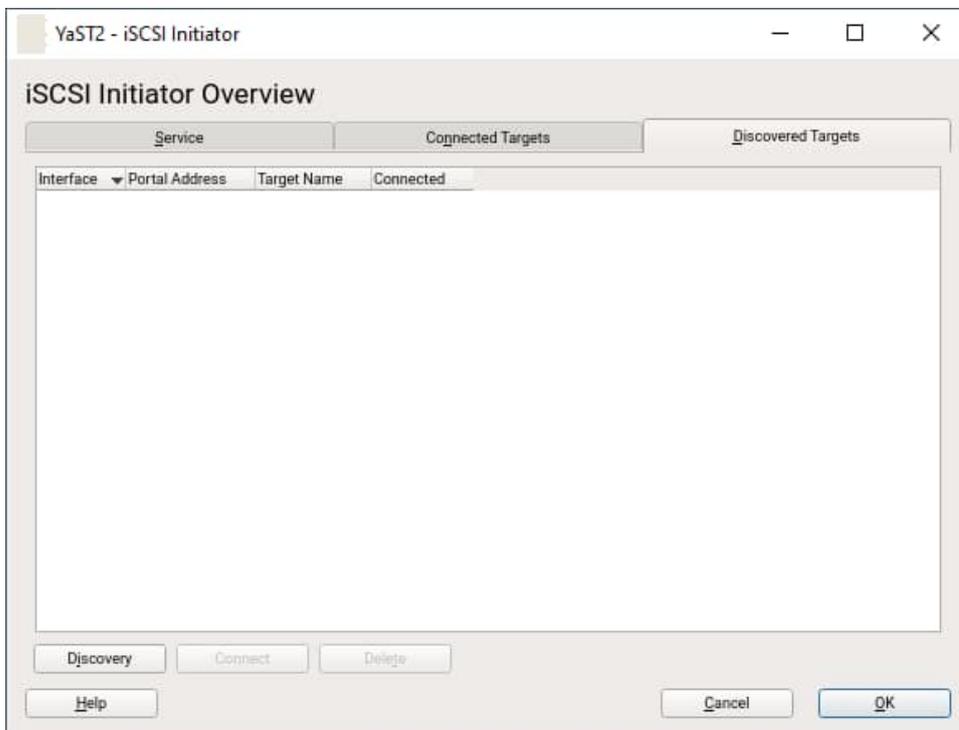
8. Choose *Finish* to complete the configuration.



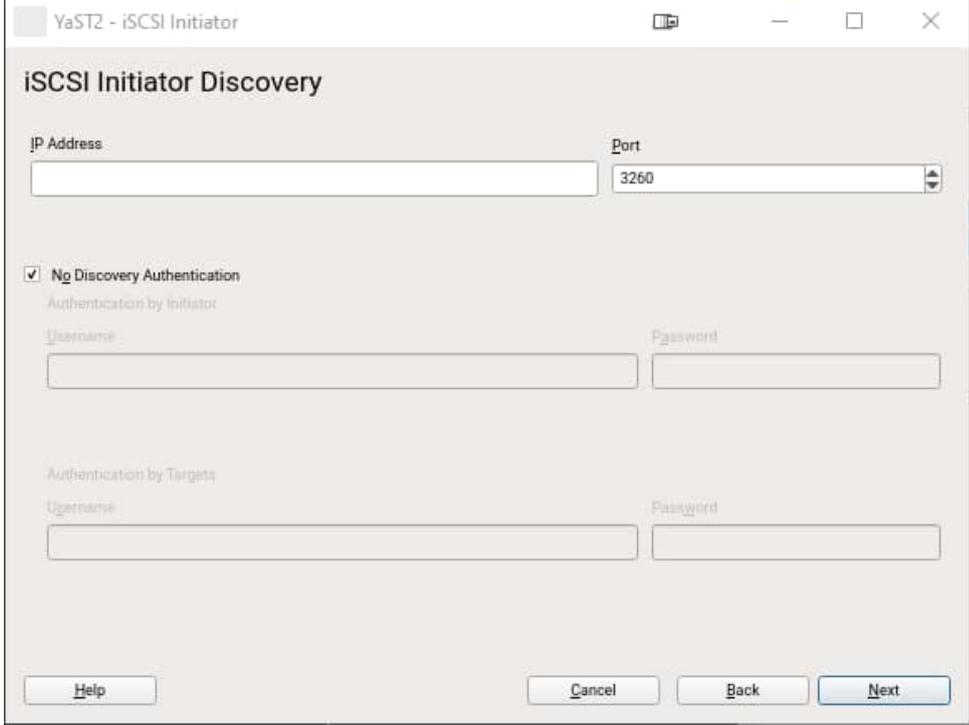
2. Configure the iSCSI initiator on SAP HANA high availability servers.
 1. Input **yast2**.
 2. Go to *Network Service/iSCSI Target*.



3. On the *Discovered Targets* tab, choose *Discovery*.

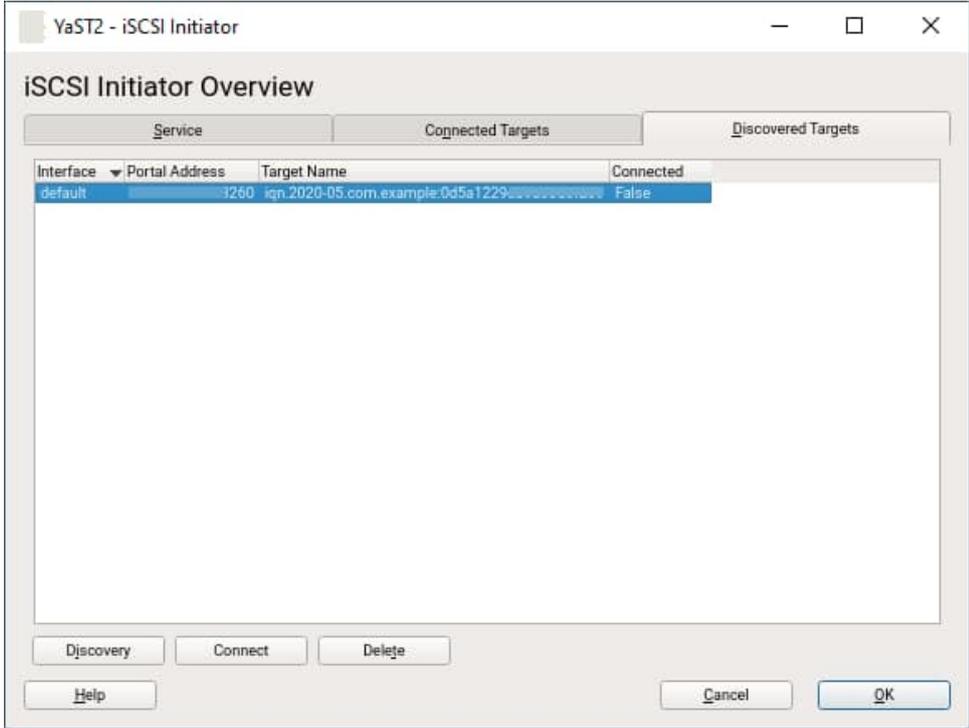


4. Input the iSCSI Target Server IP address and choose *Next*.



The screenshot shows the 'iSCSI Initiator Discovery' window in YaST2. It features a title bar 'YaST2 - iSCSI Initiator' and a main title 'iSCSI Initiator Discovery'. There are two input fields: 'IP Address' (empty) and 'Port' (set to 3260). Below these is a checked checkbox 'No Discovery Authentication'. Underneath, there are two sections for authentication: 'Authentication by Initiator' and 'Authentication by Targets', each with 'Username' and 'Password' fields. At the bottom, there are buttons for 'Help', 'Cancel', 'Back', and 'Next'.

5. Choose *Connect* to connect to the server.

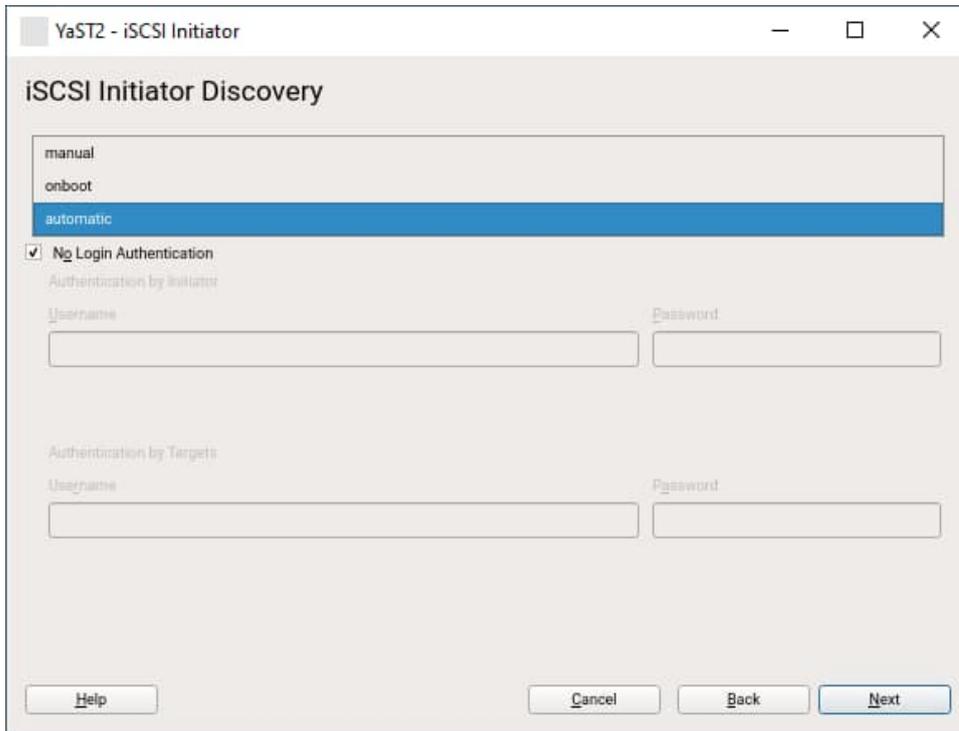


The screenshot shows the 'iSCSI Initiator Overview' window in YaST2. It has a title bar 'YaST2 - iSCSI Initiator' and a main title 'iSCSI Initiator Overview'. The window is divided into three tabs: 'Service', 'Connected Targets', and 'Discovered Targets'. The 'Discovered Targets' tab is active, showing a table with the following data:

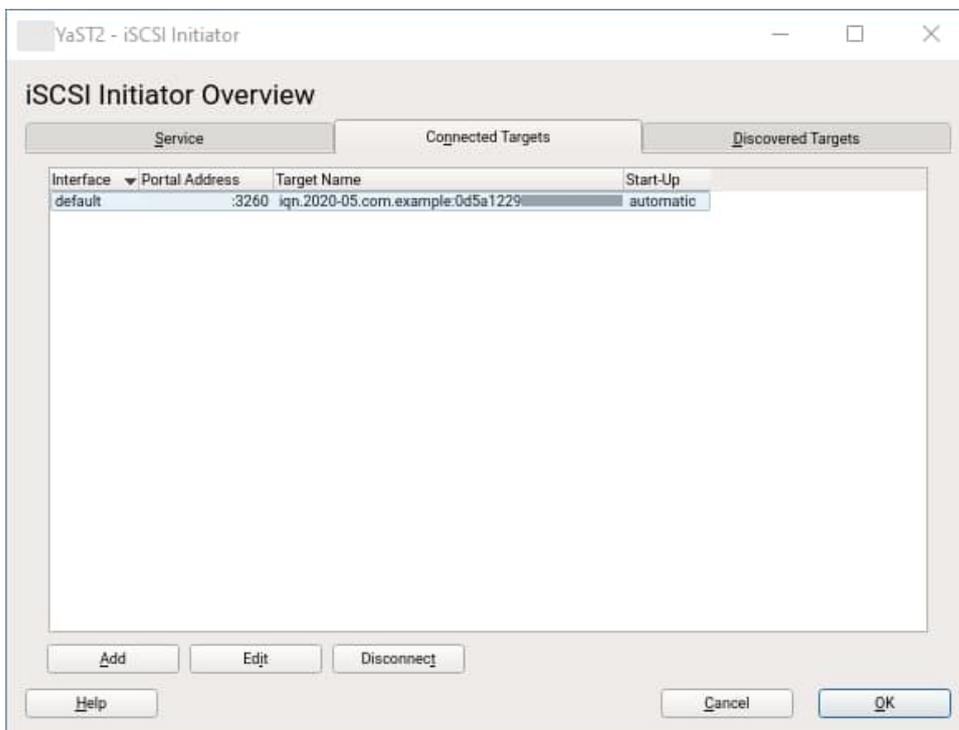
Interface	Portal Address	Target Name	Connected
default	1260	iqn.2020-05.com.example:0d5a12292b0b000000000000	False

At the bottom of the window, there are buttons for 'Discovery', 'Connect', and 'Delete'. Below these are buttons for 'Help', 'Cancel', and 'OK'.

- On the *Startup* tab, choose *automatic* and then *Next*.



- The value for the *Connected* tab changes to *True*. Choose *OK* to complete.



- Execute the command `fdisk -l` to check iSCSI devices status on the SAP HANA servers. You should get 1 additional device as listed below:

dev/sdb

```
Disk /dev/sdb: 12 MiB, 12582912 bytes, 24576 sectors
Disk model: IBLOCK
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 33550336 bytes
```

Time Synchronization

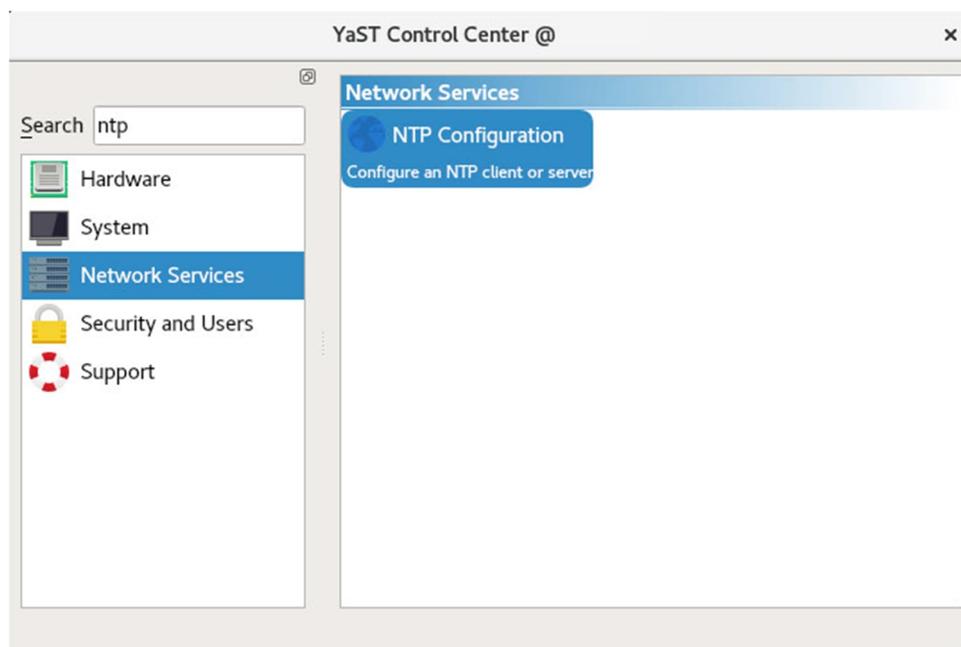
It is important to keep both database servers synchronized in time. We highly recommend that you use an additional NTP server to synchronize the time on the database servers; if you do not have an NTP server, you must be sure to check regularly that the time is synchronized.

Note

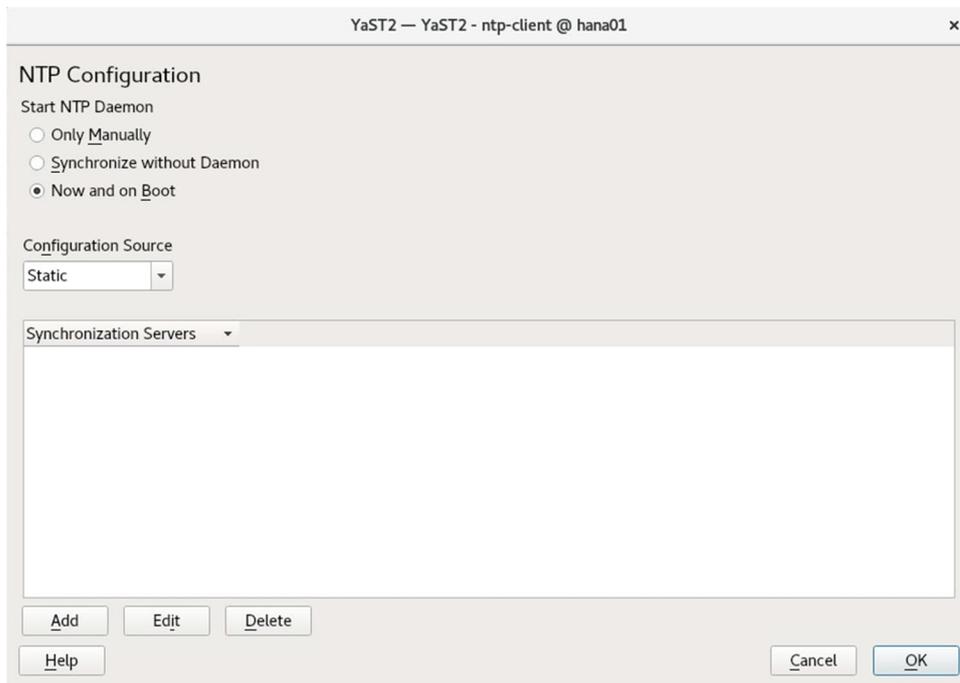
You can use the central server (for SAP Business One server components) as the NTP server.

If you use an NTP server, ensure the following points:

- Both database servers can ping the NTP server.
- On both database servers, do the following:
 - Update the NTP server IP address in Yast:



- Choose to start NTP Daemon *Now and on Boot*, then add one or more NTP servers and test them.



Supported Versions

This high availability solution is supported for SAP Business One 9.1 PLO6 and higher, version for SAP HANA and for new customers only. We require that you be familiar with the installation and upgrade processes of SAP Business One, version for SAP HANA. Additional instructions for this high-availability solution are given below in [Installing SAP Business One, version for SAP HANA](#).

Others

- To ensure the system administrator can receive automatic notification via email if an SAP HANA failover happens, ensure that you have done the following:
 - Configured the mail settings on both servers to ensure the servers can send email.
 - Created a distribution list or a group mail account for relevant system administrators.
- You have uploaded the following files (shipped with this guide) to Host 1:
 - `ha.conf`
 - `ha15.sh`

1.2 Related Documentation

For more information, refer to the following documentation:

Documentation	Remarks
<i>SAP HANA Administrator's Guide</i>	Available at http://help.sap.com/hana_platform
<i>SUSE Linux Enterprise High Availability Extension HA guide</i>	Available at https://www.suse.com/documentation
<i>SAP Business One Administrator's Guide, version for SAP HANA</i>	Available in the SAP Business One product packages: \Documentation\SystemSetup\B1_for_SAP_HANA_Admin_Guide.pdf
Important SAP Notes	
2144869	Central SAP Note for this high availability solution.
2154175	An issue with SBO Mailer on 9.1 PL06; a workaround provided.
2601083	How to modify the hostnames/IP addresses of the SAP HANA servers after configuring the high availability cluster.

2 Setup

2.1 Installation

To set up the high-availability environment, you must install some additional software other than SAP HANA. After that, you can execute a script to set up system replication between both SAP HANA database servers and configure the cluster resources.

2.1.1 Installing Software in the SAP HANA Box

Before executing the high-availability setup script, ensure that you have installed the following software on both database servers:

1. SAP HANA components:
 - o SAP HANA database 2.0 SPS 04 or higher

Caution

Both SAP HANA instances must have the same system ID and instance number. In addition, you must define the same password for `<sid>adm` and `SYSTEM` for both SAP HANA instances.

- o SAP HANA AFL
- o 64-bit SAP HANA database client for Linux

For more information, see the *SAP HANA Server Installation and Update Guide* at http://help.sap.com/hana_platform.

Note

- o You must apply for an SAP HANA license according to the contract for Host 1 (primary system). After system replication is set up between the two database servers, the primary system replicates all relevant license information to the secondary system.
- o If you are using SAP HANA 2 Revision 059.04, find the file `python_support_fixes_hana_2.00.059.04.tar.gz` that is shipped with this guide and extract it to the Python support directory.

Run the command below as a root user:

```
tar -xzf /path/to/the/archive/python_support_fixes_hana_059.04.tar.gz -C
/usr/sap/<SID>/SYS/exe/hdb/python_support/ && chown -R <sid>adm:sapsys
/usr/sap/<SID>/SYS/exe/hdb/python_support/
```

2. SAP host agent

After installing the above SAP HANA components, install the SAP host agent, as follows:

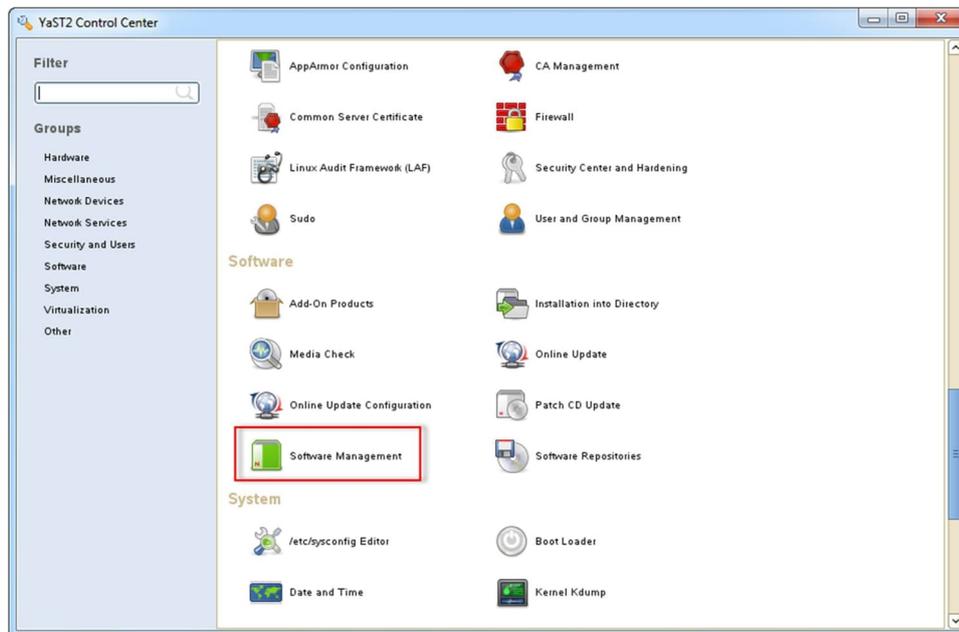
1. Navigate to the following folder: `/usr/sap/<SID>/SYS/global/hdb/saphostagent_setup`.
2. Install the SAP host agent by executing the following command: `./saphostexec -install`

The agent is installed to `/usr/sap/hostctrl`.

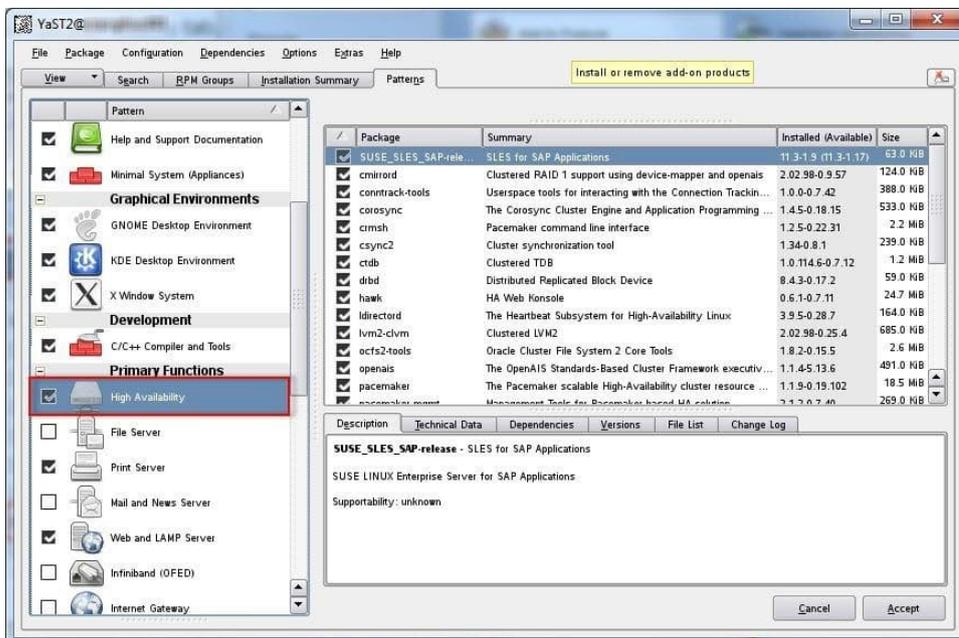
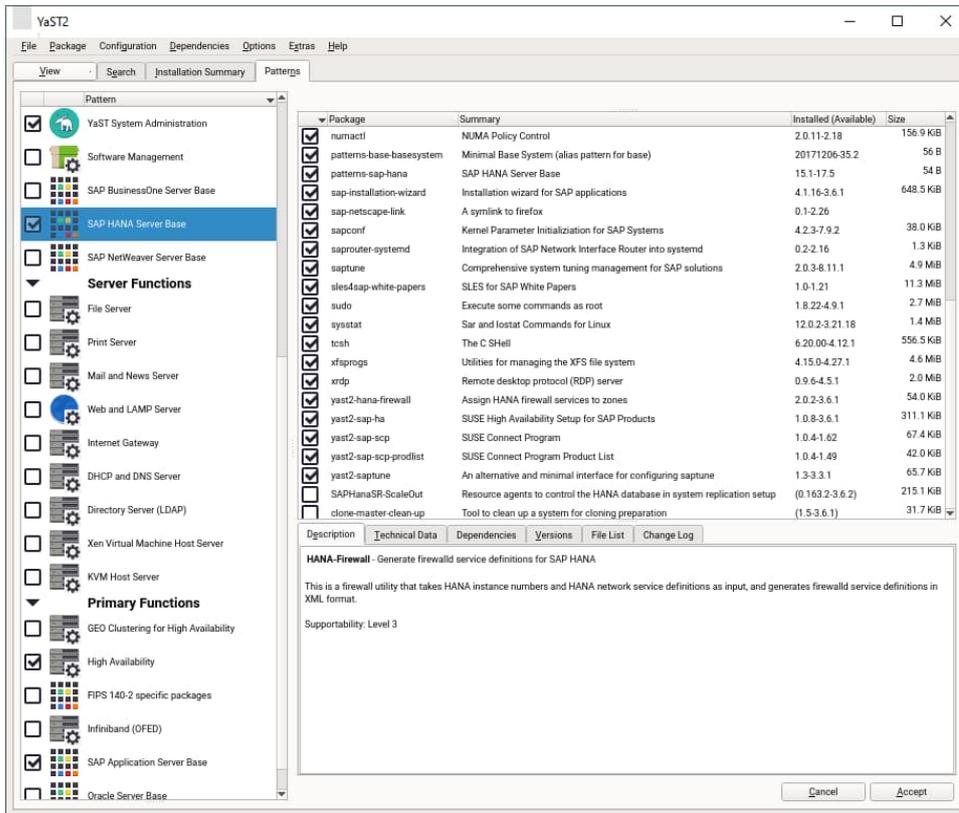
3. SUSE Linux Enterprise High Availability Extension package

To install the package, do one of the following:

- o Execute the following command: `zypper in -t pattern ha_sles sap-hana`
- o Install the package from YaST, as follows:
 1. Open the YaST control center.
 2. Choose *Software Management*.



3. From the *View* dropdown list, choose *Patterns*.
4. On the *Patterns* tab, in the left pane, select *SAP HANA Server Base* and *High Availability*.



5. Choose *Accept*.
4. SAP SAPHanaSR resource agent

To install the agent, execute the following command (make sure you have registered first):

```
zypper install SAPHana
```

To install the agent manually with the download, execute the following command:

```
zypper install SAPHanaSR-0.xxx-xx.noarch.rpm
```

5. saptune

Use saptune to optimize SAP HANA performance:

```
zypper in saptune
```

```
saptune solution apply HANA
```

```
saptune daemon start
```

```
saptune solution enabled
```

2.1.2 Executing the High-Availability Setup Script

To set up the high-availability environment, on Host 1, execute the `ha15.sh` script in one of the following ways:

- [Recommended] Read the required parameters from a configuration file, as below:
 1. Update the information in the configuration file `ha.conf` to the actual values of your system.

i Note

If you create or modify the configuration file on Windows, you must convert all line breaks (end of line, EOL) to the Unix/Linux format. To do so, use either of the following methods:

- o Before transferring the file from Windows to Linux, convert the line breaks to the Unix/Linux format. For example, in Notepad++®, from the menu bar, choose *Edit* → *EOL Conversion* → *UNIX/OSX Format*.
 - o On Linux, in the folder where you store the shell file, run the command `dos2unix <Shell File Name>`.
2. Place `ha.conf` in the same folder as `ha15.sh`.
 3. In the folder where the file `ha15.sh` is located, run `sh ha15.sh -f`.
- Read the required parameters directly with the following command line:

```
sh ha15.sh <Option 1> <Parameter 1> <Option 2> <Parameter 2>... ..
```

To find out which parameters are required and can be used, check the default configuration file `ha.conf` or execute `sh ha15.sh -h`.

The setup script asks you to review and confirm the configuration parameters before making the configuration; a progress bar is displayed during the configuration.

After the configuration, please back up all the tenant database(s) manually to ensure that the high availability cluster can work properly.

```
Checking for configuration file
Configuration file is ha.conf

*****
* SAP HANA SYSTEM ID: NDB *
* SAP HANA SYSTEMDB's SYSTEM user password: ██████████ *
* SAP HANA system admin: ndbadm *
* SAP HANA instance number: 00 *
* Local hostname: hana1 *
* Local IP address: ██████████ *
* Remote server hostname: hana2 *
* Remote server IP address: ██████████ *
* Remote server user root password: ██████████ *
* Virtual IP address: ██████████ *
* SBDs: *
* Notification email address: ██████████ *
*****

Confirm above information? (yes/no) yes

HA configuration progress:[#####]100%
*****
Configuration for
SAP HANA database high availability for SAP Business One 10.0, version for SAP HANA,
has completed.

Check status by " crm status -AD "
*****
```

Results

- A log file SBO_HANA_configure_<datetime>.log is generated in the /tmp folder.
- Apart from other resources, the SAP HANA database is added as a resource to the cluster.
- If the primary system is not available, an automatic failover is performed.

2.1.3 Installing SAP Business One, version for SAP HANA

The installation inside a high availability environment has some slight yet crucial differences from a regular installation, as described in the procedure below.

i Note

This high availability solution is supported for SAP Business One 9.1 PL06 and higher, version for SAP HANA and for new customers only. Existing customers are not supported until further notice.

Prerequisite

- System replication has been set up between both SAP HANA database servers. The currently primary SAP HANA server is `Host 1`.

To check the replication status, follow the instructions in SAP HANA Replication Status Check.

- You have installed the `nfs-kernel-server` on the central server.

To install the `nfs-kernel-server`, execute the following command:

```
zypper install nfs-kernel-server
systemctl enable nfsserver.service
systemctl start nfsserver.service
```

Procedure

Install SAP Business One Server Components, version for SAP HANA, on Central Server

1. On the central server, use the installation package (for example, release 9.2 PLO3) to install all the SAP Business One server components except the Backup Service.

In this step, use the virtual IP address to connect to the SAP HANA instance.

2. Install the Backup Service.

In this step, edit the SAP HANA database servers for backup, use the Host 1 IP address in the *Network Address* field, and DO NOT use the virtual IP address.



Post-Installation Actions

1. In the System Landscape Directory, on the *Services* tab, modify the backup service:
 1. Add a new SAP HANA server instance; the service name should be the virtual IP address.
 2. Delete the old SAP HANA service instance of which the service name is the Host 1 IP address.
2. In the System Landscape Directory, go to the *DB Instances* section on the *DB Instances and Companies* tab, and delete the database server.

Edit the backup path if it displays as *None*.

3. Mount the backup directories on Host 2:

1. On the central server, modify the file `/etc/exports` as follows:

```
/hana/shared/backup_service/backups <host 1 IP
address>(rw,no_root_squash,sync,no_subtree_check) <host 2 IP
address>(rw,no_root_squash,sync,no_subtree_check)
/tmp/backup_service <host 1 IP address>
(rw,no_root_squash,sync,no_subtree_check) <host 2 IP
address>(rw,no_root_squash,sync,no_subtree_check)
```

i Note

You can provide simultaneous access to a range of machines by specifying a network and a netmask. For example, if you want to allow access to all the machines with IP addresses between 192.168.0.0 and 192.168.0.255, you can modify the file `/etc/exports` as follows:

```
/hana/shared/backup_service/backups 192.168.0.0/255.255.255.0
(rw,no_root_squash,sync,no_subtree_check)
/tmp/backup_service 192.168.0.0/255.255.255.0
(rw,no_root_squash,sync,no_subtree_check)
```

For more information about setting up an NFS server, go to the following website:

<http://nfs.sourceforge.net/nfs-howto/ar01s03.html>.

2. Restart `nfsserver`:

```
systemctl start nfs
```

3. Mount the directories on Host 2:

```
mount <central server IP>:/hana/shared/backup_service/backups
/hana/shared/backup_service/backups
mount <central server IP>:/tmp/backup_service /tmp/backup_service
```

2.2 Upgrade

This chapter shows you how to upgrade your high availability environment.

2.2.1 Upgrading SAP HANA on Host 1 and Host 2

Prerequisites

If you have referred to version 3.0.0 or 3.0.1 of this guide and used the `ha15.sh` file delivered with the guide, please do the following beforehand:

1. In the SAP HANA studio, double-click the relevant SAP HANA system to access the SAP HANA administrative settings.
2. Modify the SAP HANA parameter for the SYSTEMDB database:

```
global.ini/[system_replication]/logshipping_max_retention_size = 10240
```

Procedure

1. Use `crm status -AD` to display the primary and secondary hosts.
2. Enable the maintenance mode as described in section 4.1.
3. Update the SAP HANA components, first on the secondary host, and then the primary host, in the following order:

SAP HANA database and SAP HANA AFLs
64-bit SAP HANA database client for Linux



Caution

During the upgrade, the SLD may be out of service for some time.

4. Disable the maintenance mode as described in section 4.1.

2.2.2 Upgrading SAP Business One Server Components, version for SAP HANA, on Central Server

To upgrade the SAP Business One 10.0, version for SAP HANA components on the central server to a higher version, perform the normal upgrade procedures. For more information, see *SAP Business One Administrator's Guide, version for SAP HANA*.



Note

If you want to upgrade SAP Business One, version for SAP HANA from 9.3 to 10.0, you need to first uninstall the SAP Business One components on Host 1 and Host 2, and then upgrade all the SAP Business One server components on the central server.

However, before uninstallation, you must first stop running the resource monitor in the cluster resource manager following below steps (on any one of the HANA hosts with `root` user):

1. Execute below command to stop b1 resource monitor:

```
crm resource stop b1
```

After several minutes, check the status of b1 resource with the command `crm status`.

The status you get should be `stop`.

2. Remove the b1 monitor with the command below:

```
crm configure delete b1
```

3 Actions After Failover

3.1 Normal Failover

If the primary system becomes unavailable and fails over to the secondary system, the system administrator will receive a notification e-mail according to the configuration. After recovering the previously primary server, the administrator must perform the following tasks to restore the cluster.

Note

For more information about troubleshooting, refer to "Part IV. Troubleshooting and Reference", SUSE Linux Enterprise High Availability Extension guide at <https://www.suse.com/documentation>.

1. On either `Host 1` or `Host 2`, execute the command `crm cluster status` to check the status of the cluster:
 - If the status is `Stopped`, check the status of the SBD device by running `sbd -d <Device Path> list` and do the following:
 - o If the status is `exit` or `clear`, execute the command `crm cluster start` to restart the service.
 - o If the status is `reset`, proceed as follows:
 1. Run the following commands to clear the `reset` status.

```
sbd -d <Device Path> message <Server Name> clear
```
 2. Restart the crm cluster:

```
crm cluster start
```

Example

There is one SBD device (`/dev/sdb`). The command is `sbd -d /dev/sdb list` and the result is as follows:

```
0      host1      reset      host2
1      host2      clear
```

To clear the `reset` status of the SBD device on `Host 1`, execute the following command:

```
sbd -d /dev/sdb message host1 clear
```

- If the result is `Running`, execute this command to ensure all resources are indeed available: `crm status -AD`.

The output should be similar to the following:

```

Node List:
* Online: [ hana1 hana2 ]

Active Resources:
* stonith-sbd (stonith:external/sbd): Started hana1
* Clone Set: cln_SAPHanaTopology_SLE_HDB00 [rsc_SAPHanaTopology_SLE_HDB00]:
* Started: [ hana1 hana2 ]
* Clone Set: msl_SAPHana_SLE_HDB00 [rsc_SAPHana_SLE_HDB00] (promotable):
* Masters: [ hana1 ]
* Slaves: [ hana2 ]
* rsc_ip_SLE_HDB00 (ocf::heartbeat:IPaddr2): Started hana1
* rsc_mail (ocf::heartbeat:MailTo): Started hana1
* Clone Set: cln_diskfull_threshold [sysinfo]:
* Started: [ hana1 hana2 ]

Node Attributes:
* Node: hana1:
* arch : x86_64
* cpu_cores : 8
* cpu_info : Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz
* cpu_load : 0.67,
* cpu_speed : 4399.99
* free_swap : 2050
* hana_data_free : 24576
* hana_log_free : 24576
* hana_ndb_clone_state : PROMOTED
* hana_ndb_op_mode : logreplay
* hana_ndb_remoteHost : hana2
* hana_ndb_roles : 4:P:master1:master:worker:master
* hana_ndb_site : hana1
* hana_ndb_smode : syncmem
* hana_ndb_sync_state : PRIM
* hana_ndb_version : 2.00.056.00.1624618329
* hana_ndb_vhost : hana1
* lpa_ndb_lpt : 1632125956
* master-rsc_SAPHana_SLE_HDB00 : 150
* os : Linux-5.3.18-59.16-default
* ram_free : 3400
* ram_total : 128850
* root_free : 55296
* Node: hana2:
* arch : x86_64
* cpu_cores : 8
* cpu_info : Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz
* cpu_load : 0.37
* cpu_speed : 4399.99
* free_swap : 2050
* hana_data_free : 22528
* hana_log_free : 22528
* hana_ndb_clone_state : DEMOTED
* hana_ndb_op_mode : logreplay
* hana_ndb_remoteHost : hana1
* hana_ndb_roles : 4:S:master1:master:worker:master
* hana_ndb_site : hana2
* hana_ndb_smode : syncmem
* hana_ndb_sync_state : SOK
* hana_ndb_version : 2.00.056.00.1624618329
* hana_ndb_vhost : hana2
* lpa_ndb_lpt : 30
* master-rsc_SAPHana_SLE_HDB00 : 100
* os : Linux-5.3.18-59.16-default
* ram_free : 1650
* ram_total : 128850
* root_free : 55296

```

2. Ensure the recovered (previously primary) SAP HANA on `Host 1` is registered as secondary on `Host 2` (currently primary), as follows:
 1. On `Host 1`, switch to the `<sid>adm` user.
 2. Check the SAP HANA status: `hdbnsutil -sr_state` and then do either of the following:
 - o If the output is as below, it means `Host 1` is still (erroneously) registered as the primary system.


```
mode: primary site id: 1
site name: host1
```

 Stop the system and then register it as the secondary system on `Host 2`:


```
HDB stop
hdbnsutil -sr_register --name=<host 1 hostname> --remoteHost=<host 2
hostname> --remoteInstance=<Instance Number> --replicationMode=syncmem
```
 - o If the output is as below, it means `Host 1` has already been registered as the secondary system and you can proceed to the next step.


```
mode: syncmem
site id: 2
site name: host1
active primary site: 1
```
 3. Start the system: `HDB start`



Caution

If you encounter any error during the registration or system startup, contact SAP Support.

3. After several minutes, check the SAP HANA replication status (`Host 1` as the secondary system):


```
crm_attribute -N host1 -G -n hana_ndb_sync_state -l reboot -q,
```

 or check the value of `hana_ndb_sync_state` with the command line `crm status -AD`.

You may get one of the following results:

```
SFAIL: System replication is still in process. You may wait for a few more minutes and check again.
SOK: System replication is complete.
```

3.2 Failover Due to Disk Space Availability

To prevent a primary server from running out of disk space, and thus being unable to provide database service, the primary SAP HANA will automatically fail over to the secondary node when the primary server has less than 8GB of free space.

When the failover happens, you will be notified by email. Please do the following on the previously primary server to recover the cluster:

1. Free up the partition where `/hana/data` and `/hana/log` are located, and the `/root` directory space.
2. Restart the `crm` cluster on the former primary server:

```
crm cluster stop
```

`crm cluster start`

3. Register the former primary SAP HANA database as secondary on the currently primary node.
4. Start the former primary SAP HANA database manually: `HDB start`.

4 Daily Maintenance

System administrators should perform the following tasks on the high availability cluster as part of their daily maintenance work.

4.1 Manual Tasks Without Failover

If you plan to do some manual maintenance tasks, for example, restart the OS, patch the OS, or restart the SAP HANA instance without unexpected behaviors or failover, please enable the maintenance mode:

```
crm maintenance on
```

All the resources will be defined as unmanaged. Check the crm status as below:

```
crm status
```

```
*** Resource management is DISABLED ***
The cluster will not attempt to start, stop or recover services

Node List:
 * Online: [ hana1 hana2 ]

Full List of Resources:
 * stonith-sbd (stonith:external/sbd): Started hana2 (unmanaged)
 * Clone Set: cln_SAPHanaTopology_SLE_HDB00 [rsc_SAPHanaTopology_SLE_HDB00] (unmanaged):
 * rsc_SAPHanaTopology_SLE_HDB00 (ocf::suse:SAPHanaTopology): Started hana2 (unmanaged)
 * rsc_SAPHanaTopology_SLE_HDB00 (ocf::suse:SAPHanaTopology): Started hana1 (unmanaged)
 * Clone Set: msl_SAPHana_SLE_HDB00 [rsc_SAPHana_SLE_HDB00] (promotable) (unmanaged):
 * rsc_SAPHana_SLE_HDB00 (ocf::suse:SAPHana): Master hana2 (unmanaged)
 * rsc_SAPHana_SLE_HDB00 (ocf::suse:SAPHana): Slave hana1 (unmanaged)
 * rsc_ip_SLE_HDB00 (ocf::heartbeat:IPaddr2): Started hana2 (unmanaged)
 * rsc_mail (ocf::heartbeat:MailTo): Started hana2 (unmanaged)
 * Clone Set: cln_diskfull_threshold [sysinfo] (unmanaged):
 * sysinfo (ocf::pacemaker:SysInfo): Started hana2 (unmanaged)
 * sysinfo (ocf::pacemaker:SysInfo): Started hana1 (unmanaged)
```

After you finish your tasks, please disable the maintenance mode:

```
crm maintenance off
```

4.2 Tenant Database Backup

The newly created tenant databases on the primary server must be backed up first, so that they can be replicated to the secondary SAP HANA instance. Otherwise, the replication status will be `SFAIL`, and the failover to the secondary node will be suspended when the primary node is down.

4.3 Status Check on Cluster

1. Login to Host 1.
2. Execute the command `crm cluster status`. The status should be active/Running.
If the status is Stopped, execute `crm cluster start` to bring the node online:
3. Repeat the above steps on Host 2.

4.4 Status Check on Resources

On either Host 1 or Host 2, execute the command `crm_mon -r` to check the status of the cluster resources dynamically. The result should be similar to the following:

```
Stack: corosync
Current DC: hana1 (version 2.0.1+20190417.13d370ca9-3.9.1-2.0.1+20190417.13d370ca9) - partition with quorum
Last updated: Sun May 31 20:14:42 2020
Last change: Sun May 31 20:14:13 2020 by root via crm_attribute on hana1

2 nodes configured
7 resources configured

Online: [ hana1 hana2 ]

Full list of resources:

stonith-sbd (stonith:external/sbd): Started hana1
Clone Set: c1n_SAPHanaTopology_SLE_HDB00 [rsc_SAPHanaTopology_SLE_HDB00]
Started: [ hana1 hana2 ]
Clone Set: ms1_SAPHana_SLE_HDB00 [rsc_SAPHana_SLE_HDB00] (promotable)
Masters: [ hana1 ]
Slaves: [ hana2 ]
rsc_ip_SLE_HDB00 (ocf::heartbeat:IPaddr2): Started hana1
rsc_mail (ocf::heartbeat:MailTo): Started hana1
```

4.5 Status Check on SAP HANA Replication Function

To check the status of replication between the primary and secondary systems, in the SAP HANA studio, in SYSTEMDB, in the administration view of Host 1, on the *Landscapes* tab, on the *System Replication* sub-tab, check the *REPLICATION_STATUS* column. The screenshot below demonstrates the normal status of the system replication:

The screenshot shows the SAP HANA Studio interface. On the left, a tree view displays the 'HANA HA Cluster' structure, including nodes for 'hana1' and 'hana2', and 'VIP'. The main window shows the 'System Replication' tab with a table of replication details.

AB	HOST	AB	SECONDARY_HOST	AB	REPLICATION_MO...	AB	REPLICATION_STATUS
	hana1		hana2		SYNCMEM		ACTIVE
	hana1		hana2		SYNCMEM		ACTIVE
	hana1		hana2		SYNCMEM		ACTIVE

4.6 Maintenance of Operating SBD Devices in a Live Environment

If you want to perform maintenance on operating SBD devices without unexpected behaviors or failover, for example, replace, restart, or patch the iSCSI Target Server, or move it to another endpoint, please follow the steps below:

1. Enable the maintenance mode on either one of the two servers:

```
crm maintenance on
```

2. Disable and stop the cluster on both servers:

```
crm cluster disable
```

```
crm cluster stop
```

3. Change the iSCSI Target Server as described in [1.1 Prerequisites](#).

4. Create new SBD devices on either one of the two servers:

```
sbd -d /dev/sdb create
```

5. Check if the UUIDs on both servers are identical:

```
sdb -d /dev/sdb dump
```

6. start and enable the cluster on both servers:

```
crm cluster start
```

```
crm cluster enable
```

7. Disable the maintenance mode on either one of the two servers:

```
crm maintenance off
```

4.7 Requesting Technical Support

If you need technical support from SAP, please run the command `crm report`, and then find and upload the file `hb_report-<Timestamp>.tar.bz2` to your support ticket.

```
hana1:/home # crm report
INFO: hana1# The report is saved in ./hb_report-Mon-16-Jan-2023.tar.bz2
INFO: hana1# Report timespan: 01/16/23 04:36:00 - 01/16/23 16:36:51
INFO: hana1# Thank you for taking time to create this report.
```


www.sap.com/contactsap

© 2023 SAP SE or an SAP affiliate company. 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 SE or an SAP affiliate company.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies. Please see <http://www.sap.com/corporate-en/legal/copyright/index.epx#trademark> for additional trademark information and notices.

Material Number:

The SAP logo, consisting of the letters "SAP" in white on a blue rectangular background.