



Infrastructure setup and configuration

NetApp solutions for SAP

NetApp
December 09, 2025

Table of Contents

| | |
|---|----|
| Infrastructure setup and configuration | 1 |
| SAN fabric setup | 1 |
| Time synchronization | 2 |
| Storage controller setup | 2 |
| Storage efficiency | 2 |
| NetApp FlexGroup Volumes | 2 |
| NetApp Volume and Aggregate Encryption | 2 |
| Quality of Service | 2 |
| NetApp FabricPool | 3 |
| Configure storage | 4 |
| Disk shelf connections | 4 |
| Aggregate configuration | 5 |
| Storage virtual machine configuration | 7 |
| Logical interface configuration | 7 |
| Initiator groups | 9 |
| SAP HANA storage connector API | 19 |
| Host setup | 19 |
| Configure multipathing | 19 |
| I/O stack configuration for SAP HANA | 31 |
| SAP HANA software installation | 32 |
| Install on single-host system | 32 |
| Install on multiple-host system | 32 |
| Adding additional data volume partitions for SAP HANA single-host systems | 36 |
| Enabling additional data volume partitions | 37 |
| Volume and LUN configuration | 37 |
| Host configuration | 37 |
| Adding an additional datavolume partition | 38 |

Infrastructure setup and configuration


The following sections provide SAP HANA infrastructure setup and configuration guidelines and describes all the steps needed to set up an SAP HANA system. Within these sections, the following example configurations are used:

- HANA system with SID=FC5
 - SAP HANA single and multiple host using Linux logical volume manager (LVM)
 - SAP HANA single host using SAP HANA multiple partitions

SAN fabric setup

Each SAP HANA server must have a redundant FCP SAN connection with a minimum of 8Gbps bandwidth. For each SAP HANA host attached to a storage controller, at least 8Gbps of bandwidth must be configured at the storage controller.

The following figure shows an example with four SAP HANA hosts attached to two storage controllers. Each SAP HANA host has two FCP ports connected to the redundant fabric. At the storage layer, four FCP ports are configured to provide the required throughput for each SAP HANA host.



In addition to the zoning on the switch layer, you must map each LUN on the storage system to the hosts that connect to this LUN. Keep the zoning on the switch simple; that is, define one zone set in which all host HBAs can see all controller HBAs.

Time synchronization

You must synchronize the time between the storage controllers and the SAP HANA database hosts. The same time server must be set for all storage controllers and all SAP HANA hosts.

Storage controller setup

This section describes the configuration of the NetApp storage system. You must complete the primary installation and setup according to the corresponding ONTAP setup and configuration guides.

Storage efficiency

Inline deduplication, cross- volume inline deduplication, inline compression, and inline compaction are supported with SAP HANA in an SSD configuration.

Enabling the storage efficiency features in an HDD configuration is not supported.

NetApp FlexGroup Volumes

The usage of NetApp FlexGroup Volumes is not supported for SAP HANA. Due to the architecture of SAP HANA the usage of FlexGroup Volumes does not provide any benefit and may result in performance issues.

NetApp Volume and Aggregate Encryption

The use of NetApp Volume Encryption (NVE) and NetApp Aggregate Encryption (NAE) are supported with SAP HANA.

Quality of Service

QoS can be used to limit the storage throughput for specific SAP HANA systems or non-SAP applications on a shared controller.

Production and Dev/Test

One use case would be to limit the throughput of development and test systems so that they cannot influence production systems in a mixed setup.

During the sizing process, you should determine the performance requirements of a nonproduction system. Development and test systems can be sized with lower performance values, typically in the range of 20% to 50% of a production-system KPI as defined by SAP.

Large write I/O has the biggest performance effect on the storage system. Therefore, the QoS throughput limit should be set to a percentage of the corresponding write SAP HANA storage performance KPI values in the data and log volumes.

Shared Environments

Another use case is to limit the throughput of heavy write workloads, especially to avoid that these workloads have an impact on other latency sensitive write workloads.

In such environments it is best practice to apply a non-shared throughput ceiling QoS group-policy to each LUN within each Storage Virtual Machine (SVM) to restrict the max throughput of each individual storage object to the given value. This reduces the possibility that a single workload can negatively influence other workloads.

To do so, a group-policy needs to be created using the CLI of the ONTAP cluster for each SVM:

```
qos policy-group create -policy-group <policy-name> -vserver <vserver name> -max-throughput 1000MB/s -is-shared false
```

and applied to each LUN within the SVM. Below is an example to apply the policy group to all existing LUNs within an SVM:

```
lun modify -vserver <vserver name> -path * -qos-policy-group <policy-name>
```

This needs to be done for every SVM. The name of the QoS police group for each SVM needs to be different. For new LUNs, the policy can be applied directly:

```
lun create -vserver <vserver_name> -path /vol/<volume_name>/<lun_name> -size <size> -ostype <e.g. linux> -qos-policy-group <policy-name>
```

It is recommended to use 1000MB/s as maximum throughput for a given LUN. If an application requires more throughput, multiple LUNs with LUN striping shall be used to provide the needed bandwidth. This guide provides an example for SAP HANA based on Linux LVM in section [Host Setup](#).



The limit applies also to reads. Therefore use enough LUNs to fulfil the required SLAs for SAP HANA database startup time and for backups.

NetApp FabricPool

NetApp FabricPool technology must not be used for active primary file systems in SAP HANA systems. This includes the file systems for the data and log area as well as the /hana/shared file system. Doing so results in unpredictable performance, especially during the startup of an SAP HANA system.

Using the “snapshot-only” tiering policy is possible as well as using FabricPool in general at a backup target such as SnapVault or SnapMirror destination.



Using FabricPool for tiering Snapshot copies at primary storage or using FabricPool at a backup target changes the required time for the restore and recovery of a database or other tasks such as creating system clones or repair systems. Take this into consideration for planning your overall lifecycle- management strategy, and check to make sure that your SLAs are still being met while using this function.

FabricPool is a good option for moving log backups to another storage tier. Moving backups affects the time needed to recover an SAP HANA database. Therefore, the option “tiering-minimum-cooling-days” should be set to a value that places log backups, which are routinely needed for recovery, on the local fast storage tier.


Configure storage

The following overview summarizes the required storage configuration steps. Each step is covered in more detail in the subsequent sections. Before initiating these steps, complete the storage hardware setup, the ONTAP software installation, and the connection of the storage FCP ports to the SAN fabric.


1. Check the correct disk shelf configuration, as described in [Disk shelf connections](#).
2. Create and configure the required aggregates, as described in [Aggregate configuration](#).
3. Create a storage virtual machine (SVM), as described in [Storage virtual machine configuration](#).
4. Create logical interfaces (LIFs), as described in [Logical interface configuration](#).
5. Create initiator groups (igroups) with worldwide names (WWNs) of HANA servers as described in the section [xref:/bp/hana-fas-fc-storage-controller-setup.html#initiator-groups](#) [Initiator groups](#).
6. Create and configure volumes and LUNs within the aggregates as described in the section [Single Host Setup](#) for single hosts
or in section [Multiple Host Setup](#) for multiple hosts

Disk shelf connections

With HDDs, a maximum of two DS2246 disk shelves or four DS224C disk shelves can be connected to one SAS stack to provide the required performance for the SAP HANA hosts, as shown in the following figure. The disks within each shelf must be distributed equally to both controllers of the HA pair.




With SSDs, a maximum of one disk shelf can be connected to one SAS stack to provide the required performance for the SAP HANA hosts, as shown in the following figure. The disks within each shelf must be distributed equally to both controllers of the HA pair. With the DS224C disk shelf, quad-path SAS cabling can also be used but is not required.



NVMe disk shelves

Each NS224 NVMe disk shelf is connected with two 100GbE ports per controller, as shown in the following figure. The disks within each shelf must be distributed equally to both controllers of the HA pair.




Aggregate configuration

In general, you must configure two aggregates per controller, independent of which disk shelf or disk technology (SSD or HDD) is used. This step is necessary so that you can use all available controller resources. For FAS 2000 series systems, one data aggregate is sufficient.

Aggregate configuration with HDDs


The following figure shows a configuration for eight SAP HANA hosts. Four SAP HANA hosts are attached to each storage controller. Two separate aggregates, one at each storage controller, are configured. Each aggregate is configured with $4 \times 10 = 40$ data disks (HDDs).



Aggregate configuration with SDD-only systems

In general, two aggregates per controller must be configured, independently of which disk shelf or disk technology (SSDs or HDDs) is used.

The following figure shows a configuration of 12 SAP HANA hosts running on a 12Gb SAS shelf configured with ADPv2. Six SAP HANA hosts are attached to each storage controller. Four separate aggregates, two at each storage controller, are configured. Each aggregate is configured with 11 disks with nine data and two parity disk partitions. For each controller, two spare partitions are available.



Storage virtual machine configuration

Multiple-host SAP landscapes with SAP HANA databases can use a single SVM. An SVM can also be assigned to each SAP landscape if necessary in case they are managed by different teams within a company. The screenshots and command outputs in this document use an SVM named `hana`.

Logical interface configuration

Within the storage cluster configuration, one network interface (LIF) must be created and assigned to a dedicated FCP port. If, for example, four FCP ports are required for performance reasons, four LIFs must be created. The following figure shows a screenshot of the eight LIFs that were configured on the SVM.

The screenshot shows the ONTAP System Manager interface. The left sidebar is collapsed, and the main area is divided into three sections:

- IPspaces:** Shows a table with a single row for 'Cluster' under 'Broadcast domains'.
- Broadcast domains:** Shows a table with four rows: 'Cluster' (9000 MTU, IPspace: Cluster), 'Default' (1500 MTU, IPspace: Default), 'NFS' (9000 MTU, IPspace: Default), and 'NFS2' (9000 MTU, IPspace: Default).
- Network interfaces:** Shows a table with eight rows of network interface data, including columns for Name, Status, Storage VM, IPspace, Address, Current node, Current port, Portset, Protocols, and Throughput.

During SVM creation with ONTAP 9 System Manager, all the required physical FCP ports can be selected, and one LIF per physical port is created automatically.

The following figure depicts the creation of SVM and LIFs with ONTAP System Manager.

☰ **NetApp** ONTAP System Manager | a400-sapcc

Search actions, objects, and pages ? < > 👤

Dashboard

Insights

Storage ^

Overview

Volumes

LUNs

NVMe namespaces

Consistency groups

Shares

Qtrees

Quotas

Storage VMs

Tiers

Network

Events & jobs

Protection

Hosts

Cluster

Add storage VM

Storage VM name:

Access protocol: SMB/CIFS, NFS iSCSI FC NVMe

Enable FC

Configure FC ports ?

| Nodes | 1a | 1b | 1c | 1d |
|---------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| a400-sapcc-01 | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| a400-sapcc-02 | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |

Storage VM administration

Enable maximum capacity limit
The maximum capacity that all volumes in this storage VM can allocate. [Learn More](#)

Manage administrator account

User name:

Password:

Confirm password:

Add a network interface for storage VM management.

Node:

IP address: Subnet mask:

Save Cancel

Initiator groups

An igroup can be configured for each server or for a group of servers that require access to a LUN. The igroup configuration requires the worldwide port names (WWPNs) of the servers.

Using the `sanlun` tool, run the following command to obtain the WWPNs of each SAP HANA host:

```
stlrx300s8-6:~ # sanlun fcp show adapter
/sbin/udevadm
/sbin/udevadm

host0 ..... WWPN:2100000e1e163700
host1 ..... WWPN:2100000e1e163701
```



The `sanlun` tool is part of the NetApp Host Utilities and must be installed on each SAP HANA host. More details can be found in section [Host setup](#).

The initiator groups can be created using the CLI of the ONTAP Cluster.

```
lun igrup create -igroup <igroup name> -protocol fcp -ostype linux
-initiator <list of initiators> -vserver <SVM name>
```

Single host


This section describes the configuration of the NetApp storage system specific to SAP HANA single-host systems

Volume and LUN configuration for SAP HANA single-host systems

The following figure shows the volume configuration of four single-host SAP HANA systems. The data and log volumes of each SAP HANA system are distributed to different storage controllers. For example, volume `SID1_data_mnt00001` is configured on controller A and volume `SID1_log_mnt00001` is configured on controller B. Within each volume, a single LUN is configured.



If only one storage controller of a high-availability (HA) pair is used for the SAP HANA systems, data volumes and log volumes can also be stored on the same storage controller.



For each SAP HANA host, a data volume, a log volume, and a volume for /hana/shared are configured. The following table shows an example configuration with four SAP HANA single-host systems.

| Purpose | Aggregate 1 at Controller A | Aggregate 2 at Controller A | Aggregate 1 at Controller B | Aggregate 2 at Controller B |
|---|---------------------------------|---------------------------------|---------------------------------|---------------------------------|
| Data, log, and shared volumes for system SID1 | Data volume: SID1_data_mnt00001 | Shared volume: SID1_shared | – | Log volume: SID1_log_mnt00001 |
| Data, log, and shared volumes for system SID2 | – | Log volume: SID2_log_mnt00001 | Data volume: SID2_data_mnt00001 | Shared volume: SID2_shared |
| Data, log, and shared volumes for system SID3 | Shared volume: SID3_shared | Data volume: SID3_data_mnt00001 | Log volume: SID3_log_mnt00001 | – |
| Data, log, and shared volumes for system SID4 | Log volume: SID4_log_mnt00001 | – | Shared volume: SID4_shared | Data volume: SID4_data_mnt00001 |

The next table shows an example of the mount point configuration for a single-host system.

| LUN | Mount point at HANA host | Note |
|--------------------|--------------------------|--------------------------------|
| SID1_data_mnt00001 | /hana/data/SID1/mnt00001 | Mounted using /etc/fstab entry |

| LUN | Mount point at HANA host | Note |
|-------------------|--------------------------|--------------------------------|
| SID1_log_mnt00001 | /hana/log/SID1/mnt00001 | Mounted using /etc/fstab entry |
| SID1_shared | /hana/shared/SID1 | Mounted using /etc/fstab entry |



With the described configuration, the `/usr/sap/SID1` directory in which the default home directory of user SID1adm is stored, is on the local disk. In a disaster recovery setup with disk-based replication, NetApp recommends creating an additional LUN within the `SID1_shared` volume for the `/usr/sap/SID1` directory so that all file systems are on the central storage.

Volume and LUN configuration for SAP HANA single-host systems using Linux LVM

The Linux LVM can be used to increase performance and to address LUN size limitations. The different LUNs of an LVM volume group should be stored within a different aggregate and at a different controller. The following table shows an example for two LUNs per volume group.



It is not necessary to use LVM with multiple LUNs to fulfil the SAP HANA KPIs, but it is recommended

| Purpose | Aggregate 1 at Controller A | Aggregate 2 at Controller A | Aggregate 1 at Controller B | Aggregate 2 at Controller B |
|--|------------------------------------|---|--------------------------------------|----------------------------------|
| Data, log, and shared volumes for LVM based system | Data volume: SID1_data_mnt00001 | Shared volume: SID1_shared Log2 volume: SID1_log2_mnt00001 | Data2 volume: SID1_data2_mnt00001 | Log volume: SID1_log_mnt00001 |



With the described configuration, the `/usr/sap/SID1` directory in which the default home directory of user SID1adm is stored, is on the local disk. In a disaster recovery setup with disk-based replication, NetApp recommends creating an additional LUN within the `SID1_shared` volume for the `/usr/sap/SID1` directory so that all file systems are on the central storage.

Volume options

The volume options listed in the following table must be verified and set on all volumes used for SAP HANA.

| Action | ONTAP 9 |
|--|---|
| Disable automatic Snapshot copies | <code>vol modify -vserver <vserver-name> -volume <volname> -snapshot-policy none</code> |
| Disable visibility of Snapshot directory | <code>vol modify -vserver <vserver-name> -volume <volname> -snapdir-access false</code> |

Creating LUNs, volumes, and mapping LUNs to initiator groups

You can use NetApp ONTAP System Manager to create storage volumes and LUNs and the map them to the igroups of the servers and the ONTAP CLI. This guide describes the usage of the CLI.

Creating LUNs, volumes, and mapping LUNs to igroups using the CLI

This section shows an example configuration using the command line with ONTAP 9 for a SAP HANA single host system with SID FC5 using LVM and two LUNs per LVM volume group:

1. Create all necessary volumes.

```
vol create -volume FC5_data_mnt00001 -aggregate aggr1_1 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log_mnt00001 -aggregate aggr1_2 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_data2_mnt00001 -aggregate aggr1_2 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log2_mnt00001 -aggregate aggr1_1 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_shared -aggregate aggr1_1 -size 512g -state
online -policy default -snapshot-policy none -junction-path /FC5_shared
-encrypt false -space-guarantee none
```

2. Create all LUNs.

```
lun create -path /vol/FC5_data_mnt00001/FC5_data_mnt00001 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_data2_mnt00001/FC5_data2_mnt00001 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log_mnt00001/FC5_log_mnt00001 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log2_mnt00001/FC5_log2_mnt00001 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
```

3. Create the initiator group for all ports belonging to sythe hosts of FC5.

```
lun igroup create -igroup HANA-FC5 -protocol fcp -ostype linux
-initiator 10000090fadcc5fa,10000090fadcc5fb -vserver hana
```

4. Map all LUNs to created initiator group.

```
lun map -path /vol/FC5_data_mnt00001/FC5_data_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_data2_mnt00001/FC5_data2_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_log_mnt00001/FC5_log_mnt00001 -igroup HANA-FC5
lun map -path /vol/FC5_log2_mnt00001/FC5_log2_mnt00001 -igroup HANA-FC5
```

Multiple hosts

This section describes the configuration of the NetApp storage system specific to SAP HANA multiple-hosts systems


Volume and LUN configuration for SAP HANA multiple-host systems

The following figure shows the volume configuration of a 4+1 multiple-host SAP HANA system. The data volumes and log volumes of each SAP HANA host are distributed to different storage controllers. For example, the volume `SID_data_mnt00001` is configured on controller A and the volume `SID_log_mnt00001` is configured on controller B. One LUN is configured within each volume.

The `/hana/shared` volume must be accessible by all HANA hosts and is therefore exported by using NFS. Even though there are no specific performance KPIs for the `/hana/shared` file system, NetApp recommends using a 10Gb Ethernet connection.



If only one storage controller of an HA pair is used for the SAP HANA system, data and log volumes can also be stored on the same storage controller.



For each SAP HANA host, a data volume and a log volume are created. The /hana/shared volume is used by all hosts of the SAP HANA system. The following figure shows an example configuration for a 4+1 multiple-host SAP HANA system.

| Purpose | Aggregate 1 at Controller A | Aggregate 2 at Controller A | Aggregate 1 at Controller B | Aggregate 2 at Controller B |
|---------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| Data and log volumes for node 1 | Data volume: SID_data_mnt00001 | — | Log volume: SID_log_mnt00001 | — |
| Data and log volumes for node 2 | Log volume: SID_log_mnt00002 | — | Data volume: SID_data_mnt00002 | — |
| Data and log volumes for node 3 | — | Data volume: SID_data_mnt00003 | — | Log volume: SID_log_mnt00003 |
| Data and log volumes for node 4 | — | Log volume: SID_log_mnt00004 | — | Data volume: SID_data_mnt00004 |
| Shared volume for all hosts | Shared volume: SID_shared | — | — | — |

The next table shows the configuration and the mount points of a multiple-host system with four active SAP HANA hosts.

| LUN or Volume | Mount point at SAP HANA host | Note |
|------------------------|------------------------------|---|
| LUN: SID_data_mnt00001 | /hana/data/SID/mnt00001 | Mounted using storage connector |
| LUN: SID_log_mnt00001 | /hana/log/SID/mnt00001 | Mounted using storage connector |
| LUN: SID_data_mnt00002 | /hana/data/SID/mnt00002 | Mounted using storage connector |
| LUN: SID_log_mnt00002 | /hana/log/SID/mnt00002 | Mounted using storage connector |
| LUN: SID_data_mnt00003 | /hana/data/SID/mnt00003 | Mounted using storage connector |
| LUN: SID_log_mnt00003 | /hana/log/SID/mnt00003 | Mounted using storage connector |
| LUN: SID_data_mnt00004 | /hana/data/SID/mnt00004 | Mounted using storage connector |
| LUN: SID_log_mnt00004 | /hana/log/SID/mnt00004 | Mounted using storage connector |
| Volume: SID_shared | /hana/shared/SID | Mounted at all hosts using NFS and /etc/fstab entry |



With the described configuration, the `/usr/sap/SID` directory in which the default home directory of user SIDadm is stored is on the local disk for each HANA host. In a disaster recovery setup with disk-based replication, NetApp recommends creating four additional subdirectories in the `SID_shared` volume for the `/usr/sap/SID` file system so that each database host has all its file systems on the central storage.

Volume and LUN configuration for SAP HANA multiple-host systems using Linux LVM

The Linux LVM can be used to increase performance and to address LUN size limitations. The different LUNs of an LVM volume group should be stored within a different aggregate and at a different controller. The following table shows an example for two LUNs per volume group for a 2+1 SAP HANA multiple host system.



It is not necessary to use LVM to combine several LUN to fulfil the SAP HANA KPIs, but it is recommended.

| Purpose | Aggregate 1 at Controller A | Aggregate 2 at Controller A | Aggregate 1 at Controller B | Aggregate 2 at Controller B |
|---------------------------------|--------------------------------|--------------------------------|----------------------------------|----------------------------------|
| Data and log volumes for node 1 | Data volume: SID_data_mnt00001 | Log2 volume: SID_log2_mnt00001 | Log volume: SID_log_mnt00001 | Data2 volume: SID_data2_mnt00001 |
| Data and log volumes for node 2 | Log2 volume: SID_log2_mnt00002 | Data volume: SID_data_mnt00002 | Data2 volume: SID_data2_mnt00002 | Log volume: SID_log_mnt00002 |
| Shared volume for all hosts | Shared volume: SID_shared | – | – | – |

Volume options

The volume options listed in the following table must be verified and set on all volumes used for SAP HANA.

| Action | ONTAP 9 |
|--|--|
| Disable automatic Snapshot copies | vol modify -vserver <vserver-name> -volume <volname> -snapshot-policy none |
| Disable visibility of Snapshot directory | vol modify -vserver <vserver-name> -volume <volname> -snapdir-access false |

Creating LUNs, volumes, and mapping LUNs to initiator groups

You can use NetApp ONTAP System Manager to create storage volumes and LUNs and the map them to the igroups of the servers and the ONTAP CLI. This guide describes the usage of the CLI.

Creating LUNs, volumes, and mapping LUNs to igroups using the CLI

This section shows an example configuration using the command line with ONTAP 9 for a 2+1 SAP HANA multiple host system with SID FC5 using LVM and two LUNs per LVM volume group.

1. Create all necessary volumes.

```

vol create -volume FC5_data_mnt00001 -aggregate aggr1_1 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log_mnt00002 -aggregate aggr2_1 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log_mnt00001 -aggregate aggr1_2 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_data_mnt00002 -aggregate aggr2_2 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_data2_mnt00001 -aggregate aggr1_2 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log2_mnt00002 -aggregate aggr2_2 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_log2_mnt00001 -aggregate aggr1_1 -size 280g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_data2_mnt00002 -aggregate aggr2_1 -size 1200g
-snapshot-policy none -foreground true -encrypt false -space-guarantee
none
vol create -volume FC5_shared -aggregate aggr1_1 -size 512g -state
online -policy default -snapshot-policy none -junction-path /FC5_shared
-encrypt false -space-guarantee none

```

2. Create all LUNs.

```
lun create -path /vol/FC5_data_mnt0001/FC5_data_mnt0001 -size 1t  
-ostype linux -space-reserve disabled -space-allocation disabled -class  
regular  
lun create -path /vol/FC5_data2_mnt0001/FC5_data2_mnt0001 -size 1t  
-ostype linux -space-reserve disabled -space-allocation disabled -class  
regular  
lun create -path /vol/FC5_data_mnt0002/FC5_data_mnt0002 -size 1t  
-ostype linux -space-reserve disabled -space-allocation disabled -class  
regular  
lun create -path /vol/FC5_data2_mnt0002/FC5_data2_mnt0002 -size 1t  
-ostype linux -space-reserve disabled -space-allocation disabled -class  
regular  
lun create -path /vol/FC5_log_mnt0001/FC5_log_mnt0001 -size 260g  
-ostype linux -space-reserve disabled -space-allocation disabled -class  
regular  
lun create -path /vol/FC5_log2_mnt0001/FC5_log2_mnt0001 -size 260g  
-ostype linux -space-reserve disabled -space-allocation disabled -class  
regular  
lun create -path /vol/FC5_log_mnt0002/FC5_log_mnt0002 -size 260g  
-ostype linux -space-reserve disabled -space-allocation disabled -class  
regular  
lun create -path /vol/FC5_log2_mnt0002/FC5_log2_mnt0002 -size 260g  
-ostype linux -space-reserve disabled -space-allocation disabled -class  
regular
```

3. Create the igroup for all servers belonging to system FC5.

```
lun igrup create -igroup HANA-FC5 -protocol fcp -ostype linux  
-initiator 10000090fadcc5fa,10000090fadcc5fb,  
10000090fadcc5c1,10000090fadcc5c2, 10000090fadcc5c3,10000090fadcc5c4  
-vserver hana
```

4. Map all LUNs to the created igroup.

```
lun map -path /vol/FC5_data_mnt0001/FC5_data_mnt0001 -igroup HANA-FC5
lun map -path /vol/FC5_data2_mnt0001/FC5_data2_mnt0001 -igroup HANA-FC5
lun map -path /vol/FC5_data_mnt0002/FC5_data_mnt0002 -igroup HANA-FC5
lun map -path /vol/FC5_data2_mnt0002/FC5_data2_mnt0002 -igroup HANA-FC5
lun map -path /vol/FC5_log_mnt0001/FC5_log_mnt0001 -igroup HANA-FC5
lun map -path /vol/FC5_log2_mnt0001/FC5_log2_mnt0001 -igroup HANA-FC5
lun map -path /vol/FC5_log_mnt0002/FC5_log_mnt0002 -igroup HANA-FC5
lun map -path /vol/FC5_log2_mnt0002/FC5_log2_mnt0002 -igroup HANA-FC5
```

SAP HANA storage connector API

A storage connector is required only in multiple-host environments that have failover capabilities. In multiple-host setups, SAP HANA provides high-availability functionality so that an SAP HANA database host can fail over to a standby host. In this case, the LUNs of the failed host are accessed and used by the standby host. The storage connector is used to make sure that a storage partition can be actively accessed by only one database host at a time.

In SAP HANA multiple-host configurations with NetApp storage, the standard storage connector delivered by SAP is used. The “SAP HANA FC Storage Connector Admin Guide” can be found as an attachment to [SAP note 1900823](#).

Host setup

Before setting up the host, NetApp SAN Host Utilities must be downloaded from the [NetApp Support](#) site and installed on the HANA servers. The Host Utility documentation includes information about additional software that must be installed depending on the FCP HBA used.

The documentation also contains information about multipath configurations that are specific to the Linux version used. This document covers the required configuration steps for SLES 15 and Red Hat Enterprise Linux 7.6 or higher, as described in the [Linux Host Utilities 7.1 Installation and Setup Guide](#).

Configure multipathing



Steps 1 to 6 must be performed on all worker and standby hosts in the SAP HANA multiple-host configuration.

To configure multipathing, complete the following steps:

1. Run the Linux `rescan-scsi-bus.sh -a` command on each server to discover new LUNs.
2. Run the `sanlun lun show` command and verify that all required LUNs are visible. The following example

shows the `sanlun lun show` command output for a 2+1 multiple-host HANA system with two data LUNs and two log LUNs. The output shows the LUNs and the corresponding device files, such as LUN `SS3_data_mnt00001` and the device file `/dev/sdag`. Each LUN has eight FC paths from the host to the storage controllers.

```
sapcc-hana-tst:~ # sanlun lun show
controller(7mode/E-Series)/
host          lun                                device
vserver(cDOT/FlashRay)      lun-pathname          filename
adapter      protocol    size      product
-----
-----
svm1
host21      FCP        500g     FC5_log2_mnt00002      /dev/sdbb
svm1
host21      FCP        500g     FC5_log_mnt00002      /dev/sdba
svm1
host21      FCP        500g     FC5_log2_mnt00001      /dev/sdaz
host21      FCP        500g     cDOT
svm1
host21      FCP        500g     FC5_log_mnt00001      /dev/sday
host21      FCP        500g     cDOT
svm1
host21      FCP        1t      cDOT
svm1
host21      FCP        1t      FC5_data_mnt00002      /dev/sdax
host21      FCP        1t      cDOT
svm1
host21      FCP        1t      cDOT
svm1
host21      FCP        1t      FC5_data2_mnt00001      /dev/sdaw
host21      FCP        1t      cDOT
svm1
host21      FCP        1t      cDOT
host21      FCP        1t      FC5_data2_mnt00001      /dev/sdav
host21      FCP        1t      cDOT
svm1
host21      FCP        1t      cDOT
host21      FCP        500g     FC5_log2_mnt00002      /dev/sdat
host21      FCP        500g     cDOT
svm1
host21      FCP        500g     FC5_log_mnt00002      /dev/sdas
host21      FCP        500g     cDOT
svm1
host21      FCP        500g     FC5_log2_mnt00001      /dev/sdar
host21      FCP        500g     cDOT
svm1
host21      FCP        500g     FC5_log_mnt00001      /dev/sdaq
host21      FCP        500g     cDOT
svm1
host21      FCP        1t      FC5_data2_mnt00002      /dev/sdap
host21      FCP        1t      cDOT
svm1
host21      FCP        1t      cDOT
svm1
host21      FCP        1t      FC5_data_mnt00002      /dev/sdao
host21      FCP        1t      cDOT
svm1
host21      FCP        1t      FC5_data2_mnt00001      /dev/sdan
host21      FCP        1t      cDOT
svm1
host21      FCP        1t      FC5_data_mnt00001      /dev/sdam
host21      FCP        500g     FC5_log2_mnt00002      /dev/sdal
host20      FCP        500g     cDOT
```

| | | | | |
|--------|-----|------|--------------------|-----------|
| svm1 | | | FC5_log_mnt00002 | /dev/sdak |
| host20 | FCP | 500g | cDOT | |
| svm1 | | | FC5_log2_mnt00001 | /dev/sdaj |
| host20 | FCP | 500g | cDOT | |
| svm1 | | | FC5_log_mnt00001 | /dev/sdai |
| host20 | FCP | 500g | cDOT | |
| svm1 | | | FC5_data2_mnt00002 | /dev/sdah |
| host20 | FCP | 1t | cDOT | |
| svm1 | | | FC5_data_mnt00002 | /dev/sdag |
| host20 | FCP | 1t | cDOT | |
| svm1 | | | FC5_data2_mnt00001 | /dev/sdaf |
| host20 | FCP | 1t | cDOT | |
| svm1 | | | FC5_data_mnt00001 | /dev/sdae |
| host20 | FCP | 1t | cDOT | |
| svm1 | | | FC5_log2_mnt00002 | /dev/sdad |
| host20 | FCP | 500g | cDOT | |
| svm1 | | | FC5_log_mnt00002 | /dev/sdac |
| host20 | FCP | 500g | cDOT | |
| svm1 | | | FC5_log2_mnt00001 | /dev/sdab |
| host20 | FCP | 500g | cDOT | |
| svm1 | | | FC5_log_mnt00001 | /dev/sdaa |
| host20 | FCP | 500g | cDOT | |
| svm1 | | | FC5_data2_mnt00002 | /dev/sdz |
| host20 | FCP | 1t | cDOT | |
| svm1 | | | FC5_data_mnt00002 | /dev/sdy |
| host20 | FCP | 1t | cDOT | |
| svm1 | | | FC5_data2_mnt00001 | /dev/sdx |
| host20 | FCP | 1t | cDOT | |
| svm1 | | | FC5_data_mnt00001 | /dev/sdw |
| host20 | FCP | 1t | cDOT | |

3. Run the `multipath -r` and `multipath -ll` command to get the worldwide identifiers (WWIDs) for the device file names.



In this example, there are eight LUNs.

```
sapcc-hana-tst:~ # multipath -r
sapcc-hana-tst:~ # multipath -ll
3600a098038314e63492b59326b4b786d dm-7 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`-- policy='service-time 0' prio=50 status=active
 |- 20:0:4:2 sdaf 65:240 active ready running
 |- 20:0:5:2 sdx 65:112 active ready running
 |- 21:0:4:2 sdav 66:240 active ready running
```

```

`- 21:0:6:2 sdan 66:112 active ready running
3600a098038314e63492b59326b4b786e dm-9 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:4 sdah 66:16 active ready running
|- 20:0:5:4 sdz 65:144 active ready running
|- 21:0:4:4 sdax 67:16 active ready running
`- 21:0:6:4 sdap 66:144 active ready running
3600a098038314e63492b59326b4b786f dm-11 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:6 sdaj 66:48 active ready running
|- 20:0:5:6 sdab 65:176 active ready running
|- 21:0:4:6 sdaz 67:48 active ready running
`- 21:0:6:6 sdar 66:176 active ready running
3600a098038314e63492b59326b4b7870 dm-13 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:8 sdal 66:80 active ready running
|- 20:0:5:8 sdad 65:208 active ready running
|- 21:0:4:8 sdbb 67:80 active ready running
`- 21:0:6:8 sdat 66:208 active ready running
3600a098038314e63532459326d495a64 dm-6 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:1 sdae 65:224 active ready running
|- 20:0:5:1 sdw 65:96 active ready running
|- 21:0:4:1 sdau 66:224 active ready running
`- 21:0:6:1 sdam 66:96 active ready running
3600a098038314e63532459326d495a65 dm-8 NETAPP,LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:3 sdag 66:0 active ready running
|- 20:0:5:3 sdy 65:128 active ready running
|- 21:0:4:3 sdaw 67:0 active ready running
`- 21:0:6:3 sdao 66:128 active ready running
3600a098038314e63532459326d495a66 dm-10 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:5 sdai 66:32 active ready running

```

```
| - 20:0:5:5 sdaa 65:160 active ready running
| - 21:0:4:5 sday 67:32 active ready running
` - 21:0:6:5 sdaq 66:160 active ready running
3600a098038314e63532459326d495a67 dm-12 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
| - 20:0:4:7 sdak 66:64 active ready running
| - 20:0:5:7 sdac 65:192 active ready running
| - 21:0:4:7 sdba 67:64 active ready running
` - 21:0:6:7 sdas 66:192 active ready running
```

4. Edit the `/etc/multipath.conf` file and add the WWIDs and alias names.



The example output shows the content of the `/etc/multipath.conf` file, which includes alias names for the four LUNs of a 2+1 multiple-host system. If there is no `multipath.conf` file available, you can create one by running the following command: `multipath -T > /etc/multipath.conf`.

```

sapcc-hana-tst:/ # cat /etc/multipath.conf
multipaths {
    multipath {
        wwid      3600a098038314e63492b59326b4b786d
        alias    svm1-FC5_data2_mnt00001
    }
    multipath {
        wwid      3600a098038314e63492b59326b4b786e
        alias    svm1-FC5_data2_mnt00002
    }
    multipath {
        wwid      3600a098038314e63532459326d495a64
        alias    svm1-FC5_data_mnt00001
    }
    multipath {
        wwid      3600a098038314e63532459326d495a65
        alias    svm1-FC5_data_mnt00002
    }
    multipath {
        wwid      3600a098038314e63492b59326b4b786f
        alias    svm1-FC5_log2_mnt00001
    }
    multipath {
        wwid      3600a098038314e63492b59326b4b7870
        alias    svm1-FC5_log2_mnt00002
    }
    multipath {
        wwid      3600a098038314e63532459326d495a66
        alias    svm1-FC5_log_mnt00001
    }
    multipath {
        wwid      3600a098038314e63532459326d495a67
        alias    svm1-FC5_log_mnt00002
    }
}

```

5. Run the `multipath -r` command to reload the device map.
6. Verify the configuration by running the `multipath -ll` command to list all the LUNs, alias names, and active and standby paths.



The following example output shows the output of a 2+1 multiple-host HANA system with two data and two log LUNs.

```

sapcc-hana-tst:~ # multipath -ll
hsvm1-FC5_data2_mnt00001 (3600a098038314e63492b59326b4b786d) dm-7
NETAPP, LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:2 sdaf 65:240 active ready running
|- 20:0:5:2 sdx 65:112 active ready running
|- 21:0:4:2 sdav 66:240 active ready running
`- 21:0:6:2 sdan 66:112 active ready running
svm1-FC5_data2_mnt00002 (3600a098038314e63492b59326b4b786e) dm-9
NETAPP, LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:4 sdah 66:16 active ready running
|- 20:0:5:4 sdz 65:144 active ready running
|- 21:0:4:4 sdax 67:16 active ready running
`- 21:0:6:4 sdap 66:144 active ready running
svm1-FC5_data_mnt00001 (3600a098038314e63532459326d495a64) dm-6
NETAPP, LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:1 sdae 65:224 active ready running
|- 20:0:5:1 sdw 65:96 active ready running
|- 21:0:4:1 sdau 66:224 active ready running
`- 21:0:6:1 sdam 66:96 active ready running
svm1-FC5_data_mnt00002 (3600a098038314e63532459326d495a65) dm-8
NETAPP, LUN C-Mode
size=1.0T features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:3 sdag 66:0 active ready running
|- 20:0:5:3 sdy 65:128 active ready running
|- 21:0:4:3 sdaw 67:0 active ready running
`- 21:0:6:3 sdao 66:128 active ready running
svm1-FC5_log2_mnt00001 (3600a098038314e63492b59326b4b786f) dm-11
NETAPP, LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:6 sdaj 66:48 active ready running
|- 20:0:5:6 sdab 65:176 active ready running
|- 21:0:4:6 sdaz 67:48 active ready running
`- 21:0:6:6 sdar 66:176 active ready running

```

```

svm1-FC5_log2_mnt00002 (3600a098038314e63492b59326b4b7870) dm-13
NETAPP, LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:8 sdal 66:80 active ready running
|- 20:0:5:8 sdad 65:208 active ready running
|- 21:0:4:8 sdbb 67:80 active ready running
`- 21:0:6:8 sdat 66:208 active ready running
svm1-FC5_log_mnt00001 (3600a098038314e63532459326d495a66) dm-10
NETAPP, LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:5 sdai 66:32 active ready running
|- 20:0:5:5 sdaa 65:160 active ready running
|- 21:0:4:5 sday 67:32 active ready running
`- 21:0:6:5 sdaq 66:160 active ready running
svm1-FC5_log_mnt00002 (3600a098038314e63532459326d495a67) dm-12
NETAPP, LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
`--+ policy='service-time 0' prio=50 status=active
|- 20:0:4:7 sdak 66:64 active ready running
|- 20:0:5:7 sdac 65:192 active ready running
|- 21:0:4:7 sdba 67:64 active ready running
`- 21:0:6:7 sdas 66:192 active ready running

```

Single host setup

This chapter describes the setup of an SAP HANA single host using Linux LVM.

LUN configuration for SAP HANA single-host systems

At the SAP HANA host, volume groups and logical volumes need to be created and mounted, as indicated in the following table.

| Logical volume/LUN | Mount point at SAP HANA host | Note |
|--------------------------|------------------------------|--------------------------------|
| LV: FC5_data_mnt0000-vol | /hana/data/FC51/mnt00001 | Mounted using /etc/fstab entry |
| LV: FC5_log_mnt00001-vol | /hana/log/FC5/mnt00001 | Mounted using /etc/fstab entry |
| LUN: FC5_shared | /hana/shared/FC5 | Mounted using /etc/fstab entry |

 With the described configuration, the `/usr/sap/FC5` directory in which the default home directory of user `FC5adm` is stored, is on the local disk. In a disaster recovery setup with disk-based replication, NetApp recommends creating an additional LUN within the `FC5_shared` volume for the `/usr/sap/FC5` directory so that all file systems are on the central storage.

Create LVM volume groups and logical volumes

1. Initialize all LUNs as a physical volume.

```
pvcreate /dev/mapper/hana-FC5_data_mnt00001
pvcreate /dev/mapper/hana-FC5_data2_mnt00001
pvcreate /dev/mapper/hana-FC5_log_mnt00001
pvcreate /dev/mapper/hana-FC5_log2_mnt00001
```

2. Create the volume groups for each data and log partition.

```
vgcreate FC5_data_mnt00001 /dev/mapper/hana-FC5_data_mnt00001
/dev/mapper/hana-FC5_data2_mnt00001
vgcreate FC5_log_mnt00001 /dev/mapper/hana-FC5_log_mnt00001
/dev/mapper/hana-FC5_log2_mnt00001
```

3. Create a logical volume for each data and log partition. Use a stripe size that is equal to the number of LUNs used per volume group (in this example, it is two) and a stripe size of 256k for data and 64k for log. SAP only supports one logical volume per volume group.

```
lvcreate --extents 100%FREE -i 2 -I 256k --name vol FC5_data_mnt00001
lvcreate --extents 100%FREE -i 2 -I 64k --name vol FC5_log_mnt00001
```

4. Scan the physical volumes, volume groups, and vol groups at all other hosts.

```
modprobe dm_mod
pvscan
vgscan
lvscan
```



If these commands do not find the volumes, a restart is required.

To mount the logical volumes, the logical volumes must be activated. To activate the volumes, run the following command:

```
vgchange -a y
```

Create file systems

Create the XFS file system on all data and log logical volumes and the hana shared LUN.

```
mkfs.xfs /dev/mapper/FC5_data_mnt00001-vol
mkfs.xfs /dev/mapper/FC5_log_mnt00001-vol
mkfs.xfs /dev/mapper/svm1-FC5_shared
```

Create mount points

Create the required mount point directories, and set the permissions on the database host:

```
sapcc-hana-tst:/ # mkdir -p /hana/data/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/shared
sapcc-hana-tst:/ # chmod -R 777 /hana/log/FC5
sapcc-hana-tst:/ # chmod -R 777 /hana/data/FC5
sapcc-hana-tst:/ # chmod 777 /hana/shared
```

Mount file systems

To mount file systems during system boot using the `/etc/fstab` configuration file, add the required file systems to the `/etc/fstab` configuration file:

```
# cat /etc/fstab
/dev/mapper/hana-FC5_shared /hana/shared xfs defaults 0 0
/dev/mapper/FC5_log_mnt00001-vol /hana/log/FC5/mnt00001 xfs
    relatime,inode64 0 0
/dev/mapper/FC5_data_mnt00001-vol /hana/data/FC5/mnt00001 xfs
    relatime,inode64 0 0
```



The XFS file systems for the data and log LUNs must be mounted with the `relatime` and `inode64` mount options.

To mount the file systems, run the `mount -a` command at the host.

Multiple hosts setup

This chapter describes the setup of a 2+1 SAP HANA multiple host system as example.

LUN configuration for SAP HANA multiple-hosts systems

At the SAP HANA host, volume groups and logical volumes need to be created and mounted, as indicated in the following table.

| Logical volume (LV) or volume | Mount point at SAP HANA host | Note |
|-------------------------------|------------------------------|---|
| LV: FC5_data_mnt00001-vol | /hana/data/FC5/mnt00001 | Mounted using storage connector |
| LV: FC5_log_mnt00001-vol | /hana/log/FC5/mnt00001 | Mounted using storage connector |
| LV: FC5_data_mnt00002-vol | /hana/data/FC5/mnt00002 | Mounted using storage connector |
| LV: FC5_log_mnt00002-vol | /hana/log/FC5/mnt00002 | Mounted using storage connector |
| Volume: FC5_shared | /hana/shared | Mounted at all hosts using NFS and /etc/fstab entry |

With the described configuration, the `/usr/sap/FC5` directory in which the default home directory of user FC5adm is stored, is on the local disk for each HANA host. In a disaster recovery setup with disk-based replication, NetApp recommends creating four additional subdirectories in the `FC5_shared` volume for the `/usr/sap/FC5` file system so that each database host has all its file systems on the central storage.

Create LVM volume groups and logical volumes

1. Initialize all LUNs as a physical volume.

```

pvcreate /dev/mapper/hana-FC5_data_mnt00001
pvcreate /dev/mapper/hana-FC5_data2_mnt00001
pvcreate /dev/mapper/hana-FC5_data_mnt00002
pvcreate /dev/mapper/hana-FC5_data2_mnt00002
pvcreate /dev/mapper/hana-FC5_log_mnt00001
pvcreate /dev/mapper/hana-FC5_log2_mnt00001
pvcreate /dev/mapper/hana-FC5_log_mnt00002
pvcreate /dev/mapper/hana-FC5_log2_mnt00002

```

2. Create the volume groups for each data and log partition.

```

vgcreate FC5_data_mnt00001 /dev/mapper/hana-FC5_data_mnt00001
/dev/mapper/hana-FC5_data2_mnt00001
vgcreate FC5_data_mnt00002 /dev/mapper/hana-FC5_data_mnt00002
/dev/mapper/hana-FC5_data2_mnt00002
vgcreate FC5_log_mnt00001 /dev/mapper/hana-FC5_log_mnt00001
/dev/mapper/hana-FC5_log2_mnt00001
vgcreate FC5_log_mnt00002 /dev/mapper/hana-FC5_log_mnt00002
/dev/mapper/hana-FC5_log2_mnt00002

```

3. Create a logical volume for each data and log partition. Use a stripe size that is equal to the number of LUNs used per volume group (in this example, it is two) and a stripe size of 256k for data and 64k for log. SAP only supports one logical volume per volume group.

```
lvcreate --extents 100%FREE -i 2 -I 256k --name vol FC5_data_mnt0001
lvcreate --extents 100%FREE -i 2 -I 256k --name vol FC5_data_mnt0002
lvcreate --extents 100%FREE -i 2 -I 64k --name vol FC5_log_mnt0002
lvcreate --extents 100%FREE -i 2 -I 64k --name vol FC5_log_mnt0001
```

4. Scan the physical volumes, volume groups, and vol groups at all other hosts.

```
modprobe dm_mod
pvscan
vgscan
lvscan
```



If these commands do not find the volumes, a restart is required.

To mount the logical volumes, the logical volumes must be activated. To activate the volumes, run the following command:

```
vgchange -a y
```

Create file systems

Create the XFS file system on all data and log logical volumes.

```
mkfs.xfs /dev/mapper/FC5_data_mnt0001-vol
mkfs.xfs /dev/mapper/FC5_data_mnt0002-vol
mkfs.xfs /dev/mapper/FC5_log_mnt0001-vol
mkfs.xfs /dev/mapper/FC5_log_mnt0002-vol
```

Create mount points

Create the required mount point directories, and set the permissions on all worker and standby hosts:

```
sapcc-hana-tst:/ # mkdir -p /hana/data/FC5/mnt0001
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt0001
sapcc-hana-tst:/ # mkdir -p /hana/data/FC5/mnt0002
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt0002
sapcc-hana-tst:/ # mkdir -p /hana/shared
sapcc-hana-tst:/ # chmod -R 777 /hana/log/FC5
sapcc-hana-tst:/ # chmod -R 777 /hana/data/FC5
sapcc-hana-tst:/ # chmod 777 /hana/shared
```

Mount file systems

To mount the /hana/shared file systems during system boot using the /etc/fstab configuration file, add the /hana/shared file system to the /etc/fstab configuration file of each host.

```
sapcc-hana-tst:/ # cat /etc/fstab
<storage-ip>:/hana_shared /hana/shared nfs rw,vers=3,hard,timeo=600,
intr,noatime,nolock 0 0
```



All the data and log file systems are mounted through the SAP HANA storage connector.

To mount the file systems, run the `mount -a` command at each host.

I/O stack configuration for SAP HANA

Starting with SAP HANA 1.0 SPS10, SAP introduced parameters to adjust the I/O behavior and optimize the database for the file and storage system used.

NetApp conducted performance tests to define the ideal values. The following table lists the optimal values as inferred from the performance tests.

| Parameter | Value |
|---------------------------|-------|
| max_parallel_io_requests | 128 |
| async_read_submit | on |
| async_write_submit_active | on |
| async_write_submit_blocks | all |

For SAP HANA 1.0 up to SPS12, these parameters can be set during the installation of the SAP HANA database as described in SAP Note [2267798 – Configuration of the SAP HANA Database during Installation Using hdbparam](#).

Alternatively, the parameters can be set after the SAP HANA database installation using the `hdbparam` framework.

```
SS3adm@stlx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.max_parallel_io_requests=128
SS3adm@stlx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.async_write_submit_active=on
SS3adm@stlx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.async_read_submit=on
SS3adm@stlx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.async_write_submit_blocks=all
```

Starting with SAP HANA 2.0, `hdbparam` is deprecated and the parameters have been moved to the `global.ini` file. The parameters can be set by using SQL commands or SAP HANA Studio. For more

information, see SAP Note [2399079 - Elimination of hdbparam in HANA 2](#). The parameters can be also set within the `global.ini` file.

```
SS3adm@stlx300s8-6:/usr/sap/SS3/SYS/global/hdb/custom/config> cat
global.ini
...
[fileio]
async_read_submit = on
async_write_submit_active = on
max_parallel_io_requests = 128
async_write_submit_blocks = all
...
```

With SAP HANA 2.0 SPS5 and later, you can use the `setParameter.py` script to set the parameters mentioned above.

```
fc5adm@sapcc-hana-tst-03:/usr/sap/FC5/HDB00/exe/python_support>
python setParameter.py
-set=SYSTEM/global.ini/fileio/max_parallel_io_requests=128
python setParameter.py -set=SYSTEM/global.ini/fileio/async_read_submit=on
python setParameter.py
-set=SYSTEM/global.ini/fileio/async_write_submit_active=on
python setParameter.py
-set=SYSTEM/global.ini/fileio/async_write_submit_blocks=all
```

SAP HANA software installation

Below are the requirements for SAP HANA software installation.

Install on single-host system

SAP HANA software installation does not require any additional preparation for a single-host system.

Install on multiple-host system



The following installation procedure is based on SAP HANA 1.0 SPS12 or later.

Before beginning the installation, create a `global.ini` file to enable use of the SAP storage connector during the installation process. The SAP storage connector mounts the required file systems at the worker hosts during the installation process. The `global.ini` file must be available in a file system that is accessible from all hosts, such as the `/hana/shared/SID` file system.

Before installing SAP HANA software on a multiple-host system, the following steps must be completed:

1. Add the following mount options for the data LUNs and the log LUNs to the `global.ini` file:

- `relatime` and `inode64` for the data and log file system

2. Add the WWIDs of the data and log partitions. The WWIDs must match the alias names configured in the `/etc/multipath.conf` file.

The following output shows an example of a 2+1 multiple-host setup in which the system identifier (SID) is SS3.

```
stlrx300s8-6:~ # cat /hana/shared/global.ini
[communication]
listeninterface = .global
[persistence]
basepath_datavolumes = /hana/data/SS3
basepath_logvolumes = /hana/log/SS3
[storage]
ha_provider = hdb_ha.fcClient
partition_*_*_prtype = 5
partition_*_data__mountoptions = -o relatime,inode64
partition_*_log__mountoptions = -o relatime,inode64,nobarrier
partition_1_data__wwid = hana-SS3_data_mnt00001
partition_1_log__wwid = hana-SS3_log_mnt00001
partition_2_data__wwid = hana-SS3_data_mnt00002
partition_2_log__wwid = hana-SS3_log_mnt00002
[system_information]
usage = custom
[trace]
ha_fcclient = info
stlrx300s8-6:~ #
```

If LVM is used, the needed configuration is different. The example below shows a 2+1 multiple-host setup with SID=FC5.

```

sapcc-hana-tst-03:/hana/shared # cat global.ini
[communication]
listeninterface = .global
[persistence]
basepath_datavolumes = /hana/data/FC5
basepath_logvolumes = /hana/log/FC5
[storage]
ha_provider = hdb_ha.fcClientLVM
partition_*_*_prtype = 5
partition_*_data_mountOptions = -o relatime,inode64
partition_*_log_mountOptions = -o relatime,inode64
partition_1_data_lvmname = FC5_data_mnt00001-vol
partition_1_log_lvmname = FC5_log_mnt00001-vol
partition_2_data_lvmname = FC5_data_mnt00002-vol
partition_2_log_lvmname = FC5_log_mnt00002-vol
sapcc-hana-tst-03:/hana/shared #

```

Using the SAP `hdblcm` installation tool, start the installation by running the following command at one of the worker hosts. Use the `addhosts` option to add the second worker (sapcc-hana-tst-06) and the standby host (sapcc-hana-tst-07).

The directory where the prepared the `global.ini` file has been stored is included with the `storage_cfg` CLI option (`--storage_cfg=/hana/shared`).

Depending on the OS version being used, it might be necessary to install python 2.7 before installing the SAP HANA database.

```

/hdblcm --action=install --addhosts=sapcc-hana-tst
-06:role=worker:storage_partition=2,sapcc-hana-tst-07:role=standby
--storage_cfg=/hana/shared/

```

AP HANA Lifecycle Management - SAP HANA Database 2.00.073.00.1695288802

Scanning software locations...

Detected components:

SAP HANA AFL (incl.PAL,BFL,OFL) (2.00.073.0000.1695321500) in
/mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_AFL_LINUX_X86_64/packages
SAP HANA Database (2.00.073.00.1695288802) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-73/DATA_UNITS/HDB_SERVER_LINUX_X86_64/server
SAP HANA Database Client (2.18.24.1695756995) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_CLIENT_LINUX_X86_64/SAP_HANA_CLIENT/client
SAP HANA Studio (2.3.75.000000) in /mnt/sapcc-
share/software/SAP/HANA2SPS7-73/DATA_UNITS/HDB_STUDIO_LINUX_X86_64/studio

SAP HANA Local Secure Store (2.11.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/HANA_LSS_24_LINUX_X86_64/packages
 SAP HANA XS Advanced Runtime (1.1.3.230717145654) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/XSA_RT_10_LINUX_X86_64/packages
 SAP HANA EML AFL (2.00.073.0000.1695321500) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/HDB_EML_AFL_10_LINUX_X86_64/packages
 SAP HANA EPM-MDS (2.00.073.0000.1695321500) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/SAP_HANA_EPM-MDS_10/packages
 Automated Predictive Library (4.203.2321.0.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/PAAPL4_H20_LINUX_X86_64/apl-4.203.2321.0-hana2sp03-linux_x64/installer/packages
 GUI for HALM for XSA (including product installer) Version 1 (1.015.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/XSA_CONTENT_10/XSACALMPIUI15_0.zip
 XSAC FILEPROCESSOR 1.0 (1.000.102) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/XSA_CONTENT_10/XSACFILEPROC00_102.zip
 SAP HANA tools for accessing catalog content, data preview, SQL console, etc. (2.015.230503) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/XSAC_HRTT_20/XSACHRTT15_230503.zip
 Develop and run portal services for customer applications on XSA (2.007.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/XSA_CONTENT_10/XSACPORTALSERV07_0.zip
 The SAP Web IDE for HANA 2.0 (4.007.0) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/XSAC_SAP_WEB_IDE_20/XSAC SAPWEBIDE07_0.zip
 XS JOB SCHEDULER 1.0 (1.007.22) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/XSA_CONTENT_10/XSACSERVICES07_22.zip
 SAPUI5 FESV6 XSA 1 - SAPUI5 1.71 (1.071.52) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/XSA_CONTENT_10/XSACUI5FESV671_52.zip
 SAPUI5 FESV9 XSA 1 - SAPUI5 1.108 (1.108.5) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/XSA_CONTENT_10/XSACUI5FESV9108_5.zip
 SAPUI5 SERVICE BROKER XSA 1 - SAPUI5 Service Broker 1.0 (1.000.4) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/XSA_CONTENT_10/XSACUI5SB00_4.zip
 XSA Cockpit 1 (1.001.37) in /mnt/sapcc-share/software/SAP/HANA2SPS7-73/DATA_UNITS/XSA_CONTENT_10/XSACXSACOCKPIT01_37.zip

SAP HANA Database version '2.00.073.00.1695288802' will be installed.

Select additional components for installation:

| Index | Components | Description |
|--------------------------|-----------------|--|
| <hr/> | | |
| <hr/> | | |
| 1 | all | All components |
| 2 | server | No additional components |
| 3 | client | Install SAP HANA Database Client version |
| 2.18.24.1695756995 | | |
| 4 | lss | Install SAP HANA Local Secure Store version |
| 2.11.0 | | |
| 5 | studio | Install SAP HANA Studio version 2.3.75.000000 |
| 6 | xs | Install SAP HANA XS Advanced Runtime version |
| 1.1.3.230717145654 | | |
| 7 | afl | Install SAP HANA AFL (incl.PAL,BFL,OFL) version 2.00.073.0000.1695321500 |
| 8 | eml | Install SAP HANA EML AFL version |
| 2.00.073.0000.1695321500 | | |
| 9 | epmmds | Install SAP HANA EPM-MDS version |
| 2.00.073.0000.1695321500 | | |
| 10 | sap_afl_sdk_apl | Install Automated Predictive Library version |
| 4.203.2321.0.0 | | |

Enter comma-separated list of the selected indices [3,4]: 2,3

Verify that the installation tool installed all selected components at all worker and standby hosts.

Adding additional data volume partitions for SAP HANA single-host systems

Starting with SAP HANA 2.0 SPS4, additional data volume partitions can be configured. This feature allows you to configure two or more LUNs for the data volume of an SAP HANA tenant database and to scale beyond the size and performance limits of a single LUN.

 It is not necessary to use multiple partitions to fulfil the SAP HANA KPIs. A single LUN with a single partition fulfils the required KPIs.

 Using two or more individual LUNs for the data volume is only available for SAP HANA single-host systems. The SAP storage connector required for SAP HANA multiple-host systems does only support one device for the data volume.

You can add more data volume partitions at any time but it might require a restart of the SAP HANA database.

Enabling additional data volume partitions

To enable additional data volume partitions, complete the following steps:

1. Add the following entry within the `global.ini` file:

```
[customizable_functionalities]
persistence_datavolume_partition_multipath = true
```

2. Restart the database to enable the feature. Adding the parameter through the SAP HANA Studio to the `global.ini` file by using the Systemdb configuration prevents the restart of the database.

Volume and LUN configuration

The layout of volumes and LUNs is similar to the layout of a single host with one data volume partition, but with an additional data volume and LUN stored on a different aggregate as log volume and the other data volume. The following table shows an example configuration of an SAP HANA single-host systems with two data volume partitions.

| Aggregate 1 at Controller A | Aggregate 2 at Controller A | Aggregate 1 at Controller B | Aggregate 2 at Controller B |
|-----------------------------------|------------------------------|------------------------------------|---------------------------------|
| Data volume: SID_data_mnt00001 | Shared volume: SID_shared | Data volume: SID_data2_mnt00001 | Log volume: SID_log_mnt00001 |

The next table shows an example of the mount point configuration for a single-host system with two data volume partitions.

| LUN | Mount point at HANA host | Note |
|--------------------|--------------------------|--------------------------------|
| SID_data_mnt00001 | /hana/data/SID/mnt00001 | Mounted using /etc/fstab entry |
| SID_data2_mnt00001 | /hana/data2/SID/mnt00001 | Mounted using /etc/fstab entry |
| SID_log_mnt00001 | /hana/log/SID/mnt00001 | Mounted using /etc/fstab entry |
| SID_shared | /hana/shared/SID | Mounted using /etc/fstab entry |

Create the new data LUNs by using either ONTAP System Manager or the ONTAP CLI.

Host configuration

To configure a host, complete the following steps:

1. Configure multipathing for the additional LUNs, as described in section 0.
2. Create the XFS file system on each additional LUN belonging to the HANA system.

```
stlrx300s8-6:/ # mkfs.xfs /dev/mapper/hana-FC5_data2_mnt00001
```

3. Add the additional file system/s to the /etc/fstab configuration file.



The XFS file systems for the data LUN must be mounted with the `relatime` and `inode64` mount options. The XFS file systems for the log LUN must be mounted with the `relatime`, `inode64`, and `nobarrier` mount options.

```
stlrx300s8-6:/ # cat /etc/fstab
/dev/mapper/hana-FC5_shared /hana/shared xfs defaults 0 0
/dev/mapper/hana-FC5_log_mnt00001 /hana/log/FC5/mnt00001 xfs
relatime,inode64 0 0
/dev/mapper/hana-FC5_data_mnt00001 /hana/data/FC5/mnt00001 xfs
relatime,inode64 0 0
/dev/mapper/hana-FC5_data2_mnt00001 /hana/data2/FC5/mnt00001 xfs
relatime,inode64 0 0
```

4. Create the mount points and set the permissions on the database host.

```
stlrx300s8-6:/ # mkdir -p /hana/data2/FC5/mnt00001
stlrx300s8-6:/ # chmod -R 777 /hana/data2/FC5
```


5. To mount the file systems, run the `mount -a` command.

Adding an additional datavolume partition

To add an additional datavolume partition to your tenant database, complete the following step:

1. Execute the following SQL statement against the tenant database. Each additional LUN can have a different path.

```
ALTER SYSTEM ALTER DATAVOLUME ADD PARTITION PATH '/hana/data2/SID/';
```



Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—with prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.