



SAP HANA on FAS Systems with FCP Configuration Guide

NetApp solutions for SAP

NetApp
August 18, 2025

Table of Contents

- SAP HANA on FAS Systems with FCP Configuration Guide 1
 - SAP HANA on NetApp FAS Systems with Fibre Channel Protocol Configuration Guide 1
 - SAP HANA tailored data center integration 1
 - SAP HANA using VMware vSphere 2
- Architecture 2
 - SAP HANA backup 4
 - SAP HANA disaster recovery 5
- Storage sizing 7
 - Performance considerations 7
 - Mixed workloads 8
 - Capacity considerations 9
 - Configuration of performance test tool 9
 - Storage sizing process overview 11
- Infrastructure setup and configuration 12
 - SAN fabric setup 12
 - Time synchronization 13
 - Storage controller setup 13
 - SAP HANA storage connector API 31
 - Host setup 31
 - I/O stack configuration for SAP HANA 43
 - SAP HANA software installation 44
 - Adding additional data volume partitions for SAP HANA single-host systems 48
- Where to find additional information 51
- Update history 52

SAP HANA on FAS Systems with FCP Configuration Guide

SAP HANA on NetApp FAS Systems with Fibre Channel Protocol Configuration Guide

The NetApp FAS product family has been certified for use with SAP HANA in TDI projects. This guide provides best practices for SAP HANA on this platform for FCP.

Marco Schoen, NetApp

The certification is valid for the following models:

- FAS2750, FAS2820, FAS8300, FAS50, FAS8700, FAS70, FAS9500, FAS90

For a complete list of NetApp's certified storage solutions for SAP HANA, see the [certified and supported SAP HANA hardware directory](#).

This document describes FAS configurations that use the Fibre Channel Protocol (FCP).



The configuration described in this paper is necessary to achieve the required SAP HANA KPIs and the best performance for SAP HANA. Changing any settings or using features not listed herein might result in performance degradation or unexpected behavior and should only be done if advised by NetApp support.


The configuration guides for FAS systems using NFS and NetApp AFF systems can be found using the following links:

- [SAP HANA on NetApp AFF Systems with FCP](#)
- [SAP HANA on NetApp ASA Systems with FCP](#)
- [SAP HANA on NetApp FAS Systems with NFS](#)
- [SAP HANA on NetApp AFF Systems with NFS](#)

In an SAP HANA multiple-host environment, the standard SAP HANA storage connector is used to provide fencing in the event of an SAP HANA host failover. Refer to the relevant SAP notes for operating system configuration guidelines and HANA-specific Linux kernel dependencies. For more information, see [SAP Note 2235581 – SAP HANA Supported Operating Systems](#).

SAP HANA tailored data center integration

NetApp FAS storage controllers are certified in the SAP HANA Tailored Data Center Integration (TDI) program using NFS (NAS) and Fibre Channel (SAN) protocols. They can be deployed in any SAP HANA scenario, such as, SAP Business Suite on HANA, S/4HANA, BW/4HANA or SAP Business Warehouse on HANA in single-host or multiple-host configurations. Any server that is certified for use with SAP HANA can be combined with the certified storage solution. See the following figure for an architecture overview.



For more information regarding the prerequisites and recommendations for productive SAP HANA systems, see the following resource:

- [SAP HANA Tailored Data Center Integration Frequently Asked Questions](#)

SAP HANA using VMware vSphere

There are several options for connecting storage to virtual machines (VMs). The preferred one is to connect the storage volumes with NFS directly out of the guest operating system. This option is described in [SAP HANA on NetApp AFF Systems with NFS](#).

Raw device mappings (RDM), FCP datastores, or VVOL datastores with FCP are supported as well. For both datastore options, only one SAP HANA data or log volume must be stored within the datastore for productive use cases.

For more information about using vSphere with SAP HANA, see the following links:

- [SAP HANA on VMware vSphere - Virtualization - Community Wiki](#)
- [SAP HANA on VMware vSphere Best Practices Guide](#)
- [2161991 - VMware vSphere configuration guidelines - SAP ONE Support Launchpad \(Login required\)](#)


Architecture

SAP HANA hosts are connected to the storage controllers using a redundant FCP infrastructure and multipath software. A redundant FCP switch infrastructure is required to provide fault-tolerant SAP HANA host-to-storage connectivity in case of switch or host bus adapter (HBA) failure. Appropriate zoning must be configured at the switch to allow all HANA hosts to reach the required LUNs on the storage controllers.

Different models of the FAS product family can be used at the storage layer. The maximum number of SAP HANA hosts attached to the storage is defined by the SAP HANA performance requirements. The number of

disk shelves required is determined by the capacity and performance requirements of the SAP HANA systems.


The following figure shows an example configuration with eight SAP HANA hosts attached to a storage HA pair.




This architecture can be scaled in two dimensions:

- By attaching additional SAP HANA hosts and disk capacity to the storage, assuming that the storage controllers can provide enough performance under the new load to meet key performance indicators (KPIs)
- By adding more storage systems and disk capacity for the additional SAP HANA hosts

The following figure shows a configuration example in which more SAP HANA hosts are attached to the storage controllers. In this example, more disk shelves are necessary to meet the capacity and performance requirements of the 16 SAP HANA hosts. Depending on the total throughput requirements, you must add additional FC connections to the storage controllers.



Independent of the deployed FAS system storage model, the SAP HANA landscape can also be scaled by adding more storage controllers, as shown in the following figure.



SAP HANA backup

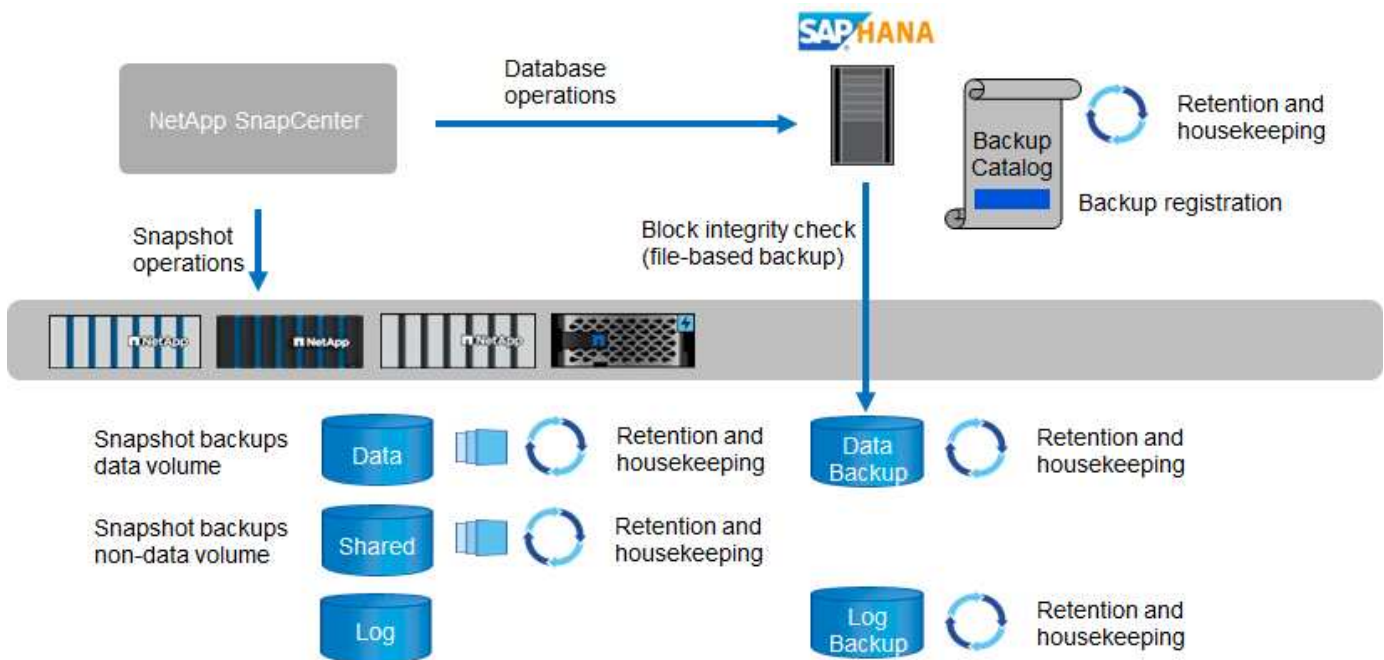
NetApp ONTAP software provides a built-in mechanism to back up SAP HANA databases. Storage-based Snapshot backup is a fully supported and integrated backup solution available for SAP HANA single-container systems and for SAP HANA MDC single-tenant systems.

Storage-based Snapshot backups are implemented by using the NetApp SnapCenter plug-in for SAP HANA, which enables consistent storage-based Snapshot backups by using the interfaces provided by the SAP HANA

database. SnapCenter registers the Snapshot backups in the SAP HANA backup catalog so that the backups are visible within the SAP HANA studio and can be selected for restore and recovery operations.

By using NetApp SnapVault software, the Snapshot copies that were created on the primary storage can be replicated to the secondary backup storage controlled by SnapCenter. Different backup retention policies can be defined for backups on the primary storage and for backups on the secondary storage. The SnapCenter Plug-in for SAP HANA Database manages the retention of Snapshot copy-based data backups and log backups including the housekeeping of the backup catalog. The SnapCenter Plug-in for SAP HANA Database also enables the execution of a block-integrity check of the SAP HANA database by performing a file-based backup.

The database logs can be backed up directly to the secondary storage by using an NFS mount, as shown in the following figure.



Storage-based Snapshot backups provide significant advantages compared to file-based backups. Those advantages include the following:

- Faster backup (few minutes)
- Faster restore on the storage layer (a few minutes)
- No effect on the performance of the SAP HANA database host, network, or storage during backup
- Space-efficient and bandwidth-efficient replication to secondary storage based on block changes

For detailed information about the SAP HANA backup and recovery solution using SnapCenter, see [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#).

SAP HANA disaster recovery

SAP HANA disaster recovery can be performed on the database layer by using SAP system replication or on the storage layer by using storage-replication technologies. The following section provides an overview of disaster recovery solutions based on storage replication.

For detailed information about the SAP HANA disaster recovery solution using SnapCenter, see [TR-4646: SAP HANA Disaster Recovery with Storage Replication](#).

Storage replication based on SnapMirror

The following figure shows a three-site disaster recovery solution, using synchronous SnapMirror replication to the local DR datacenter and asynchronous SnapMirror to replicate data to the remote DR datacenter.

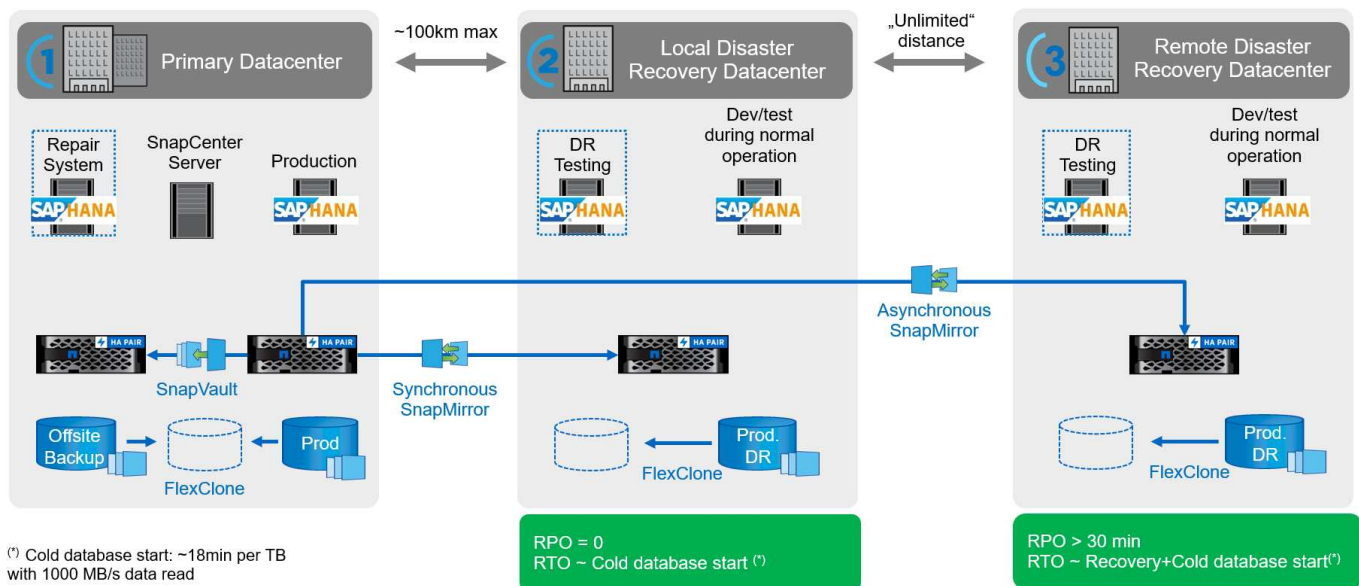
Data replication using synchronous SnapMirror provides an RPO of zero. The distance between the primary and the local DR datacenter is limited to around 100km.

Protection against failures of both the primary and the local DR site is performed by replicating the data to a third remote DR datacenter using asynchronous SnapMirror. The RPO depends on the frequency of replication updates and how fast they can be transferred. In theory, the distance is unlimited, but the limit depends on the amount of data that must be transferred and the connection that is available between the data centers. Typical RPO values are in the range of 30 minutes to multiple hours.

The RTO for both replication methods primarily depends on the time needed to start the HANA database at the DR site and load the data into memory. With the assumption that the data is read with a throughput of 1000MBps, loading 1TB of data would take approximately 18 minutes.

The servers at the DR sites can be used as dev/test systems during normal operation. In the case of a disaster, the dev/test systems would need to be shut down and started as DR production servers.


Both replication methods allow you to execute DR workflow testing without influencing the RPO and RTO. FlexClone volumes are created on the storage and are attached to the DR testing servers.



Synchronous replication offers StrictSync mode. If the write to secondary storage is not completed for any reason, the application I/O fails, thereby ensuring that the primary and secondary storage systems are identical. Application I/O to the primary resumes only after the SnapMirror relationship returns to the InSync status. If the primary storage fails, application I/O can be resumed on the secondary storage after failover with no loss of data. In StrictSync mode, the RPO is always zero.

Storage replication based on NetApp MetroCluster

The following figure shows a high-level overview of the solution. The storage cluster at each site provides local high availability and is used for production workloads. The data at each site is synchronously replicated to the other location and is available in case of disaster failover.



Storage sizing

The following section provides an overview of performance and capacity considerations for sizing a storage system for SAP HANA.



Contact your NetApp or NetApp partner sales representative to support the storage sizing process and to create a properly sized storage environment.

Performance considerations

SAP has defined a static set of storage KPIs. These KPIs are valid for all production SAP HANA environments independent of the memory size of the database hosts and the applications that use the SAP HANA database. These KPIs are valid for single-host, multiple-host, Business Suite on HANA, Business Warehouse on HANA, S/4HANA, and BW/4HANA environments. Therefore, the current performance sizing approach depends on only the number of active SAP HANA hosts that are attached to the storage system.



Storage performance KPIs are required only for production SAP HANA systems.

SAP delivers a performance test tool, which must be used to validate the storage performance for active SAP HANA hosts attached to the storage.

NetApp tested and predefined the maximum number of SAP HANA hosts that can be attached to a specific storage model, while still fulfilling the required storage KPIs from SAP for production-based SAP HANA systems.



The storage controllers of the certified FAS product family can also be used for SAP HANA with other disk types or disk back-end solutions, as long as they are supported by NetApp and fulfill SAP HANA TDI performance KPIs. Examples include NetApp Storage Encryption (NSE) and NetApp FlexArray technology.

This document describes disk sizing for SAS hard disk drives and solid-state drives.

Hard disk drives

A minimum of 10 data disks (10k RPM SAS) per SAP HANA node is required to fulfill the storage performance KPIs from SAP.



This calculation is independent of the storage controller and disk shelf used.

Solid-state drives

With solid-state drives (SSDs), the number of data disks is determined by the SAS connection throughput from the storage controllers to the SSD shelf.

The maximum number of SAP HANA hosts that can be run on a disk shelf and the minimum number of SSDs required per SAP HANA host were determined by running the SAP performance test tool.

- The 12Gb SAS disk shelf (DS224C) with 24 SSDs supports up to 14 SAP HANA hosts, when the disk shelf is connected with 12Gb.
- The 6Gb SAS disk shelf (DS2246) with 24 SSDs supports up to 4 SAP HANA hosts.

The SSDs and the SAP HANA hosts must be equally distributed between both storage controllers.

The following table summarizes the supported number of SAP HANA hosts per disk shelf.

	6Gb SAS shelves (DS2246) fully loaded with 24 SSDs	12Gb SAS shelves (DS224C) fully loaded with 24 SSDs
Maximum number of SAP HANA hosts per disk shelf	4	14



This calculation is independent of the storage controller used. Adding more disk shelves does not increase the maximum number of SAP HANA hosts that a storage controller can support.

NS224 NVMe shelf

One NVMe SSDs (data) supports up to 2 SAP HANA hosts.

The SSDs and the SAP HANA hosts must be equally distributed between both storage controllers.

Mixed workloads

SAP HANA and other application workloads running on the same storage controller or in the same storage aggregate are supported. However, it is a NetApp best practice to separate SAP HANA workloads from all other application workloads.

You might decide to deploy SAP HANA workloads and other application workloads on either the same storage controller or the same aggregate. If so, you must make sure that enough performance is always available for SAP HANA within the mixed workload environment. NetApp also recommends that you use quality of service (QoS) parameters to regulate the impact these other applications could have on SAP HANA applications.

The SAP HCMT test tool must be used to check if additional SAP HANA hosts can be run on a storage controller that is already used for other workloads. However, SAP application servers can be safely placed on the same storage controller and aggregate as the SAP HANA databases.

Capacity considerations

A detailed description of the capacity requirements for SAP HANA is in the [SAP Note 1900823](#) white paper.



The capacity sizing of the overall SAP landscape with multiple SAP HANA systems must be determined by using SAP HANA storage sizing tools from NetApp. Contact NetApp or your NetApp partner sales representative to validate the storage sizing process for a properly sized storage environment.

Configuration of performance test tool

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. These parameters must also be set for the performance test tool from SAP (fsperf) when the storage performance is tested by using the SAP test tool.

Performance tests were conducted by NetApp to define the optimal values. The following table lists the parameters that must be set within the configuration file of the SAP test tool.

Parameter	Value
max_parallel_io_requests	128
async_read_submit	on
async_write_submit_active	on
async_write_submit_blocks	all

For more information about the configuration of SAP test tool, see [SAP note 1943937](#) for HWCCT (SAP HANA 1.0) and [SAP note 2493172](#) for HCMT/HCOT (SAP HANA 2.0).

The following example shows how variables can be set for the HCMT/HCOT execution plan.

```
...{
    "Comment": "Log Volume: Controls whether read requests are
submitted asynchronously, default is 'on'",
    "Name": "LogAsyncReadSubmit",
    "Value": "on",
    "Request": "false"
},
{
    "Comment": "Data Volume: Controls whether read requests are
submitted asynchronously, default is 'on'",
    "Name": "DataAsyncReadSubmit",
    "Value": "on",
    "Request": "false"
},
{
    "Comment": "Log Volume: Controls whether write requests can be
submitted asynchronously",
    "Name": "LogAsyncWriteSubmitActive",
```

```

        "Value": "on",
        "Request": "false"
    },
    {
        "Comment": "Data Volume: Controls whether write requests can be
submitted asynchronously",
        "Name": "DataAsyncWriteSubmitActive",
        "Value": "on",
        "Request": "false"
    },
    {
        "Comment": "Log Volume: Controls which blocks are written
asynchronously. Only relevant if AsyncWriteSubmitActive is 'on' or 'auto'
and file system is flagged as requiring asynchronous write submits",
        "Name": "LogAsyncWriteSubmitBlocks",
        "Value": "all",
        "Request": "false"
    },
    {
        "Comment": "Data Volume: Controls which blocks are written
asynchronously. Only relevant if AsyncWriteSubmitActive is 'on' or 'auto'
and file system is flagged as requiring asynchronous write submits",
        "Name": "DataAsyncWriteSubmitBlocks",
        "Value": "all",
        "Request": "false"
    },
    {
        "Comment": "Log Volume: Maximum number of parallel I/O requests
per completion queue",
        "Name": "LogExtMaxParallelIoRequests",
        "Value": "128",
        "Request": "false"
    },
    {
        "Comment": "Data Volume: Maximum number of parallel I/O requests
per completion queue",
        "Name": "DataExtMaxParallelIoRequests",
        "Value": "128",
        "Request": "false"
    }, ...

```

These variables must be used for the test configuration. This is usually the case with the predefined execution plans SAP delivers with the HCMT/HCOT tool. The following example for a 4k log write test is from an execution plan.

```

...
{
  "ID": "D664D001-933D-41DE-A904F304AEB67906",
  "Note": "File System Write Test",
  "ExecutionVariants": [
    {
      "ScaleOut": {
        "Port": "${RemotePort}",
        "Hosts": "${Hosts}",
        "ConcurrentExecution": "${FSConcurrentExecution}"
      },
      "RepeatCount": "${TestRepeatCount}",
      "Description": "4K Block, Log Volume 5GB, Overwrite",
      "Hint": "Log",
      "InputVector": {
        "BlockSize": 4096,
        "DirectoryName": "${LogVolume}",
        "FileOverwrite": true,
        "FileSize": 5368709120,
        "RandomAccess": false,
        "RandomData": true,
        "AsyncReadSubmit": "${LogAsyncReadSubmit}",
        "AsyncWriteSubmitActive":
"${LogAsyncWriteSubmitActive}",
        "AsyncWriteSubmitBlocks":
"${LogAsyncWriteSubmitBlocks}",
        "ExtMaxParallelIoRequests":
"${LogExtMaxParallelIoRequests}",
        "ExtMaxSubmitBatchSize": "${LogExtMaxSubmitBatchSize}",
        "ExtMinSubmitBatchSize": "${LogExtMinSubmitBatchSize}",
        "ExtNumCompletionQueues":
"${LogExtNumCompletionQueues}",
        "ExtNumSubmitQueues": "${LogExtNumSubmitQueues}",
        "ExtSizeKernelIoQueue": "${ExtSizeKernelIoQueue}"
      }
    }, ...
  ]
}

```

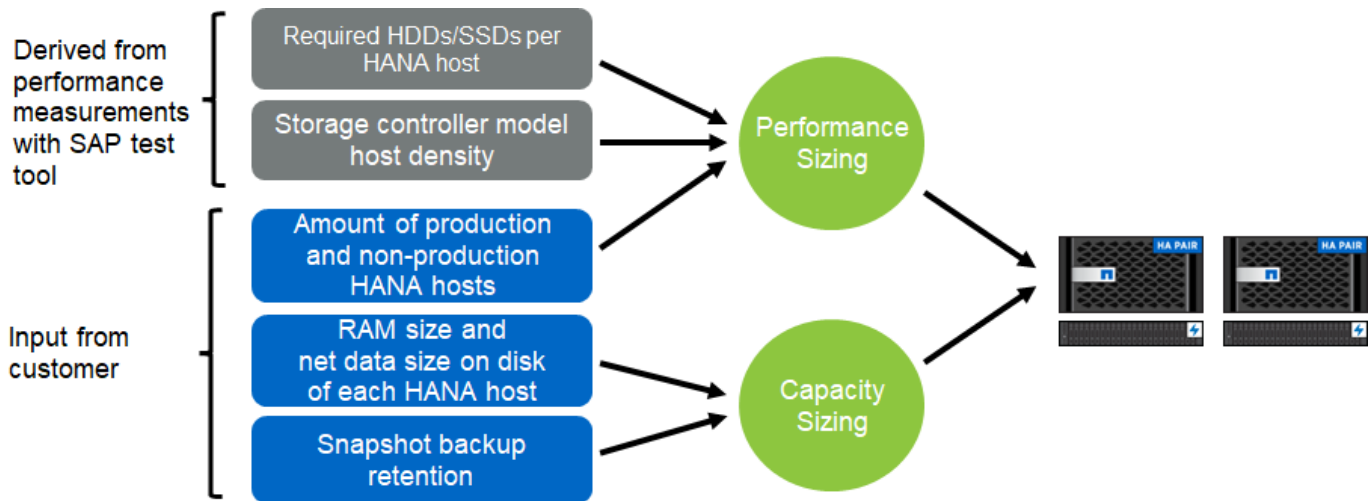
Storage sizing process overview

The number of disks per HANA host and the SAP HANA host density for each storage model were determined with the SAP HANA test tool.

The sizing process requires details such as the number of production and nonproduction SAP HANA hosts, the RAM size of each host, and the backup retention period of the storage-based Snapshot copies. The number of SAP HANA hosts determines the storage controller and the number of disks required.

The size of the RAM, the net data size on the disk of each SAP HANA host, and the Snapshot copy backup retention period are used as inputs during capacity sizing.

The following figure summarizes the sizing process.



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 policy 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 [\[Disk shelf connection\]](#).
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

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.

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

hana

Access protocol

SMB/CIFS, NFS

ISCSI

FC

NVMe

Enable FC

Configure FC ports

Nodes	1a	1b	1c	1d
a400-sapcc-01				
a400-sapcc-02				

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

vsadmin

Password

Confirm password

Add a network interface for storage VM management.

Node

a400-sapcc-01

IP address

10.10.10.10

Subnet mask

255.255.255.0

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:

21

```
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 igroup 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	<code>/hana/data/SID1/mnt00001</code>	Mounted using <code>/etc/fstab</code> 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_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_data_mnt00002/FC5_data_mnt00002 -size 1t
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_data2_mnt00002/FC5_data2_mnt00002 -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
lun create -path /vol/FC5_log_mnt00002/FC5_log_mnt00002 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
lun create -path /vol/FC5_log2_mnt00002/FC5_log2_mnt00002 -size 260g
-ostype linux -space-reserve disabled -space-allocation disabled -class
regular
```

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

```
lun igroup 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_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_data_mnt00002/FC5_data_mnt00002 -igroup HANA-FC5
lun map -path /vol/FC5_data2_mnt00002/FC5_data2_mnt00002 -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
lun map -path /vol/FC5_log_mnt00002/FC5_log_mnt00002 -igroup HANA-FC5
lun map -path /vol/FC5_log2_mnt00002/FC5_log2_mnt00002 -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          FC5_log2_mnt00002    /dev/sdbb
host21        FCP          500g    cDOT
svm1          FC5_log_mnt00002    /dev/sdba
host21        FCP          500g    cDOT
svm1          FC5_log2_mnt00001    /dev/sdaz
host21        FCP          500g    cDOT
svm1          FC5_log_mnt00001    /dev/sday
host21        FCP          500g    cDOT
svm1          FC5_data2_mnt00002    /dev/sdax
host21        FCP          1t      cDOT
svm1          FC5_data_mnt00002    /dev/sdaw
host21        FCP          1t      cDOT
svm1          FC5_data2_mnt00001    /dev/sdav
host21        FCP          1t      cDOT
svm1          FC5_data_mnt00001    /dev/sdau
host21        FCP          1t      cDOT
svm1          FC5_log2_mnt00002    /dev/sdat
host21        FCP          500g    cDOT
svm1          FC5_log_mnt00002    /dev/sdas
host21        FCP          500g    cDOT
svm1          FC5_log2_mnt00001    /dev/sdar
host21        FCP          500g    cDOT
svm1          FC5_log_mnt00001    /dev/sdaq
host21        FCP          500g    cDOT
svm1          FC5_data2_mnt00002    /dev/sdap
host21        FCP          1t      cDOT
svm1          FC5_data_mnt00002    /dev/sdao
host21        FCP          1t      cDOT
svm1          FC5_data2_mnt00001    /dev/sdan
host21        FCP          1t      cDOT
svm1          FC5_data_mnt00001    /dev/sdam
host21        FCP          1t      cDOT
svm1          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.

```
pvccreate /dev/mapper/hana-FC5_data_mnt00001
pvccreate /dev/mapper/hana-FC5_data2_mnt00001
pvccreate /dev/mapper/hana-FC5_log_mnt00001
pvccreate /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 FC5_data_mnt00001-vol
mkfs.xfs 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.

```
pvccreate /dev/mapper/hana-FC5_data_mnt00001
pvccreate /dev/mapper/hana-FC5_data2_mnt00001
pvccreate /dev/mapper/hana-FC5_data_mnt00002
pvccreate /dev/mapper/hana-FC5_data2_mnt00002
pvccreate /dev/mapper/hana-FC5_log_mnt00001
pvccreate /dev/mapper/hana-FC5_log2_mnt00001
pvccreate /dev/mapper/hana-FC5_log_mnt00002
pvccreate /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_mnt00001
lvcreate --extents 100%FREE -i 2 -I 256k --name vol FC5_data_mnt00002
lvcreate --extents 100%FREE -i 2 -I 64k --name vol FC5_log_mnt00002
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.

```
mkfs.xfs FC5_data_mnt00001-vol
mkfs.xfs FC5_data_mnt00002-vol
mkfs.xfs FC5_log_mnt00001-vol
mkfs.xfs FC5_log_mnt00002-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/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt00001
sapcc-hana-tst:/ # mkdir -p /hana/data/FC5/mnt00002
sapcc-hana-tst:/ # mkdir -p /hana/log/FC5/mnt00002
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@stlrx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.max_parallel_io_requests=128
SS3adm@stlrx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.async_write_submit_active=on
SS3adm@stlrx300s8-6:/usr/sap/SS3/HDB00> hdbparam --paramset
fileio.async_read_submit=on
SS3adm@stlrx300s8-6:/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@stlrx300s8-6:/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 hdb1cm 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.

```

/hdb1cm --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/XSACSAPWEBIDE07_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
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/';
```



Where to find additional information

To learn more about the information described in this document, refer to the following documents and/or websites:

- [SAP HANA Software Solutions](#)
- [TR-4646: SAP HANA Disaster Recovery with Storage Replication](#)
- [TR-4614: SAP HANA Backup and Recovery with SnapCenter](#)
- [TR-4667: Automating SAP System Copies Using the SnapCenter SAP HANA Plug-In](#)
- [NetApp Documentation Centers](#)

<https://www.netapp.com/support-and-training/documentation/>

- [SAP Certified Enterprise Storage Hardware for SAP HANA](#)

<https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/>

- [SAP HANA Storage Requirements](#)

<https://www.sap.com/documents/2024/03/146274d3-ae7e-0010-bca6-c68f7e60039b.html>

- [SAP HANA Tailored Data Center Integration Frequently Asked Questions](#)

<https://www.sap.com/documents/2016/05/e8705aae-717c-0010-82c7-eda71af511fa.html>

- [SAP HANA on VMware vSphere Wiki](#)

https://help.sap.com/docs/SUPPORT_CONTENT/virtualization/3362185751.html

- [SAP HANA on VMware vSphere Best Practices Guide](#)

https://www.vmware.com/docs/sap_hana_on_vmware_vsphere_best_practices_guide-white-paper

Update history

The following technical changes have been made to this solution since its original publication.

Date	Update summary
February 2015	Initial version
October 2015	Included I/O parameters for SAP HANA and HWVAL SPS 10 and later
February 2016	Updated capacity sizing
February 2017	New NetApp storage systems and disk shelves New features of ONTAP 9 New OS releases (SLES12 SP1 and Red Hat Enterprise Linux 7.2) New SAP HANA release
July 2017	Minor updates
September 2018	New NetApp storage systems New OS releases (SLES12 SP3 and Red Hat Enterprise Linux 7.4) Additional minor updates SAP HANA 2.0 SPS3
September 2019	New OS releases Minor updates
April 2020	Introduced multiple data partition features available since SAP HANA 2.0 SPS4
June 2020	Additional information about optional functionalities Minor updates
February 2021	Linux LVM support New NetApp storage systems New OS releases (SLES15SP2, RHEL 8)
April 2021	VMware vSphere-specific information added
September 2022	New OS-Releases
September 2024	New Storage Systems
February 2025	New Storage System
July 2025	Minor updates

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—without 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.