



SAP HANA on NetApp FAS Systems with NFS Configuration Guide

NetApp solutions for SAP

NetApp
August 18, 2025

Table of Contents

SAP HANA on NetApp FAS Systems with NFS Configuration Guide	1
SAP HANA on NetApp FAS systems with NFS Configuration guide	1
SAP HANA Tailored Data Center integration	2
SAP HANA using VMware vSphere	3
Architecture	3
SAP HANA backup	5
SAP HANA disaster recovery	6
Storage sizing	8
Performance considerations	8
Mixed workloads	9
Capacity considerations	10
Configuration of performance test tool	10
Storage sizing process overview	13
Infrastructure setup and configuration	13
Network setup	13
Time synchronization	16
Storage controller setup	16
Host setup	28
SAP HANA installation preparations for NFSv4	34
I/O stack configuration for SAP HANA	35
SAP HANA data volume size	36
SAP HANA software installation	36
Adding additional data volume partitions	39
Where to find additional information	43
Update history	44

SAP HANA on NetApp FAS Systems with NFS Configuration Guide

SAP HANA on NetApp FAS systems with NFS Configuration guide

The NetApp FAS product family has been certified for use with SAP HANA in tailored data center integration (TDI) projects. This guide provides best practices for SAP HANA on this platform with NFS.

Marco Schoen, NetApp

This certification is currently only valid for the following models:

- FAS2750, FAS2820, FAS8300, FAS50, FAS8700, FAS70, FAS9500, FAS90
- A complete list of NetApp certified storage solutions for SAP HANA can be found at the [Certified and Supported SAP HANA Hardware Directory](#).

This document describes the ONTAP configuration requirements for the NFS version 3 (NFSv3) protocol or the NFS version 4 (NFSv4.1) protocol.



Only NFS versions 3 or 4.1 are supported. NFS versions 1, 2, 4.0, and 4.2 aren't supported.



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 cause performance degradation or unexpected behavior and should only be performed if advised by NetApp support.

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

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

The following table shows the supported combinations for NFS versions, NFS locking, and the required isolation implementations, depending on the SAP HANA database configuration.

For SAP HANA single-host systems or multiple hosts without Host Auto-Failover, NFSv3 and NFSv4 are supported.

For SAP HANA multiple host systems with Host Auto-Failover, NetApp only supports NFSv4, while using NFSv4 locking as an alternative to a server-specific STONITH (SAP HANA HA/DR provider) implementation.

SAP HANA	NFS Version	NFS Locking	SAP HANA HA/DR Provider
SAP HANA single host, multiple hosts without Host Auto-Failover	NFSv3	Off	n/a
	NFSv4	On	n/a
SAP HANA multiple hosts with Host Auto-Failover	NFSv3	Off	Server-specific STONITH implementation mandatory
	NFSv4	On	Not required



A server-specific STONITH implementation is not part of this guide. Contact your server vendor for such an implementation.

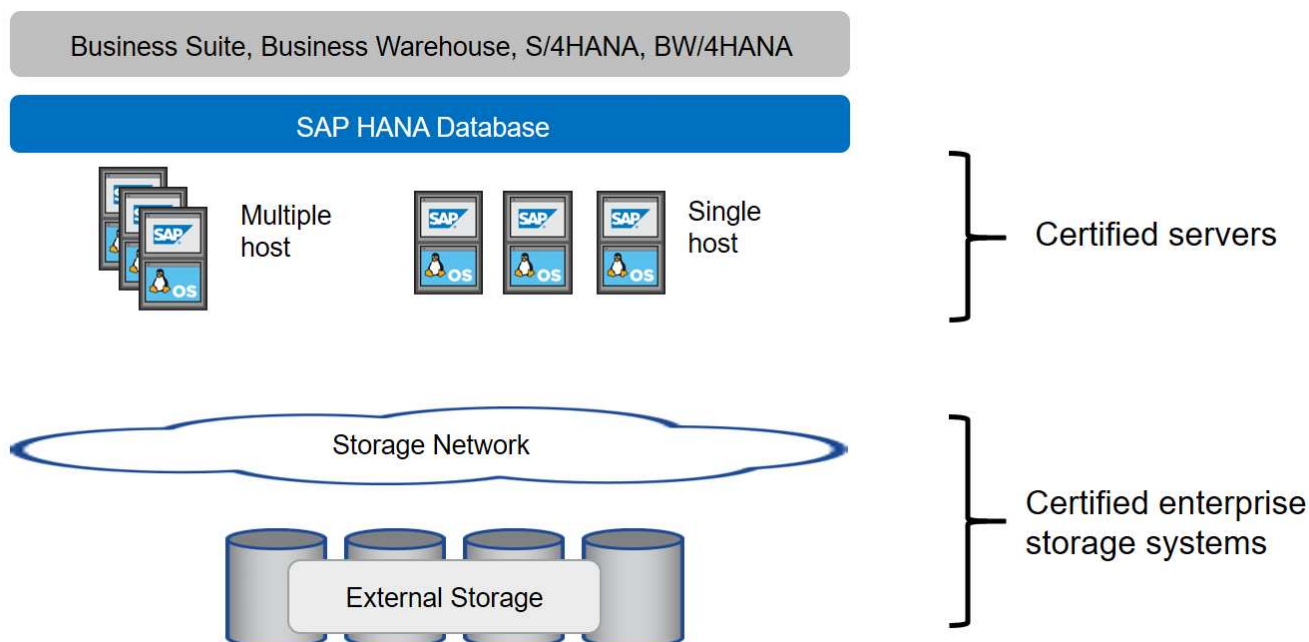
This document covers configuration recommendations for SAP HANA running on physical servers and on virtual servers that use VMware vSphere.



Always 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 TDI program using both NFS (NAS) and FC (SAN) protocols. They can be deployed in any of the current SAP HANA scenarios such as SAP Business Suite on HANA, S/4HANA, BW/4HANA, or SAP Business Warehouse on HANA in either single- host or multiple-host configurations. Any server that is certified for use with SAP HANA can be combined with NetApp certified storage solutions. See the following figure for an architecture overview.



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

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

SAP HANA using VMware vSphere

There are several options to connect the storage to virtual machines (VMs). The preferred one is to connect the storage volumes with NFS directly out of the guest operating system. Using this option, the configuration of hosts and storages do not differ between physical hosts and VMs.

NFS datastores or VVOL datastores with NFS are supported as well. For both options, only one SAP HANA data or log volume must be stored within the datastore for production use cases.

This document describes the recommended setup with direct NFS mounts from the guest OS.

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 storage controllers by using a redundant 10GbE or faster network infrastructure. Data communication between SAP HANA hosts and storage controllers is based on the NFS protocol.

A redundant switching infrastructure is recommended to provide fault-tolerant SAP HANA host- to- storage connectivity in case of switch or network interface card (NIC) failure. The switches might aggregate individual port performance with port channels in order to appear as a single logical entity at the host level.


Different models of the FAS system product family can be mixed and matched at the storage layer to allow for growth and differing performance and capacity needs. The maximum number of SAP HANA hosts that can be attached to the storage system is defined by the SAP HANA performance requirements and the model of NetApp controller used. The number of required disk shelves is only 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 high availability (HA) pair.




The architecture can be scaled in two dimensions:

- By attaching additional SAP HANA hosts and/or storage capacity to the existing storage, if the storage controllers provide enough performance to meet the current SAP key performance indicators (KPIs)
- By adding more storage systems with additional storage capacity for the additional SAP HANA hosts

The following figure shows an example configuration in which more SAP HANA hosts are attached to the storage controllers. In this example, more disk shelves are necessary to fulfill both the capacity and performance requirements of 16 SAP HANA hosts. Depending on the total throughput requirements, additional 10GbE (or faster) connections to the storage controllers must be added.



Independent of the deployed FAS system, the SAP HANA landscape can also be scaled by adding any of the certified storage controllers to meet the desired node density (the following figure).



SAP HANA backup


The ONTAP software present on all NetApp storage controllers provides a built-in mechanism to back up SAP HANA databases while in operation with no effect on performance. Storage-based NetApp Snapshot backups are a fully supported and integrated backup solution available for SAP HANA single containers and for SAP HANA Multitenant Database Container (MDC) systems with a single tenant or multiple tenants.

Storage-based Snapshot backups are implemented by using the NetApp SnapCenter plug-in for SAP HANA.

This allows users to create consistent storage-based Snapshot backups by using the interfaces provided natively by SAP HANA databases. SnapCenter registers each of the Snapshot backups into the SAP HANA backup catalog. Therefore, the backups taken by SnapCenter are visible within SAP HANA Studio and Cockpit where they can be selected directly for restore and recovery operations.

NetApp SnapMirror technology allows Snapshot copies that were created on one storage system to be replicated to a secondary backup storage system that is controlled by SnapCenter. Different backup retention policies can then be defined for each of the backup sets on the primary storage and for the backup sets on the secondary storage systems. The SnapCenter Plug-in for SAP HANA automatically 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 also allows the execution of a block integrity check of the SAP HANA database by executing 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 when compared to conventional file-based backups. These advantages include, but are not limited to, the following:

- Faster backup (a few minutes)
- Reduced recovery time objective (RTO) due to a much faster restore time on the storage layer (a few minutes) as well as more frequent backups
- No performance degradation of the SAP HANA database host, network, or storage during backup and recovery operations
- 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 either on the database layer by using SAP HANA 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 solutions, 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 that uses synchronous SnapMirror replication to the local disaster recovery data center and asynchronous SnapMirror to replicate data to the remote disaster recovery data center.


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

Protection against failures of both the primary and the local disaster recovery site is performed by replicating the data to a third remote disaster recovery data center 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 disaster recovery 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 disaster recovery 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 disaster recovery production servers.


Both replication methods allow you to execute disaster recovery workflow testing without influencing the RPO and RTO. FlexClone volumes are created on the storage and are attached to the disaster recovery 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 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 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 the production workload. The data of each site is synchronously replicated to the other location and is available if there is disaster failover.



Storage sizing

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



Contact NetApp or your NetApp partner sales representative to assist you in creating a properly sized storage environment.

Performance considerations

SAP has defined a static set of storage KPIs that 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 only depends on the number of active SAP HANA hosts that are attached to the storage system.



Storage performance KPIs are only mandated for production SAP HANA systems, but you can implement them in all HANA systems.

SAP delivers a performance test tool used to validate the performance of the storage system 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. However, they must be 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 HDDs and solid-state drives (SSDs).

HDDs

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 as well as the capacity requirements of the database. Adding more disk shelves does not increase the maximum amount of SAP HANA hosts a storage controller can support.

Solid-state drives

With 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 single disk shelf and the minimum number of SSDs required per SAP HANA host were determined by running the SAP performance test tool. This test does not consider the actual storage capacity requirements of the hosts. In addition, you must also calculate the capacity requirements to determine the actual storage configuration needed.

- 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 do not increase the maximum amount of SAP HANA hosts a storage controller can support.

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 adequate performance is available for SAP HANA within the mixed workload environment. NetApp also recommends that you use quality of service (QoS) parameters to regulate the effect these other applications could have and to guarantee throughput for SAP HANA applications.

The SAP performance test tool must be used to check if additional SAP HANA hosts can be run on an existing storage controller that is already in use for other workloads. SAP application servers can be safely placed on the same storage controller and/or aggregate as the SAP HANA databases.

Capacity considerations

A detailed description of the capacity requirements for SAP HANA is in the [SAP Note 1900823](#) attached 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 when storage performance is being tested with the SAP performance test tool.

NetApp conducted performance tests to define the optimal values. The following table lists the parameters that must be set within the configuration file of the SAP performance 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 the 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}"
      }
    }, ...
  ]
}

```

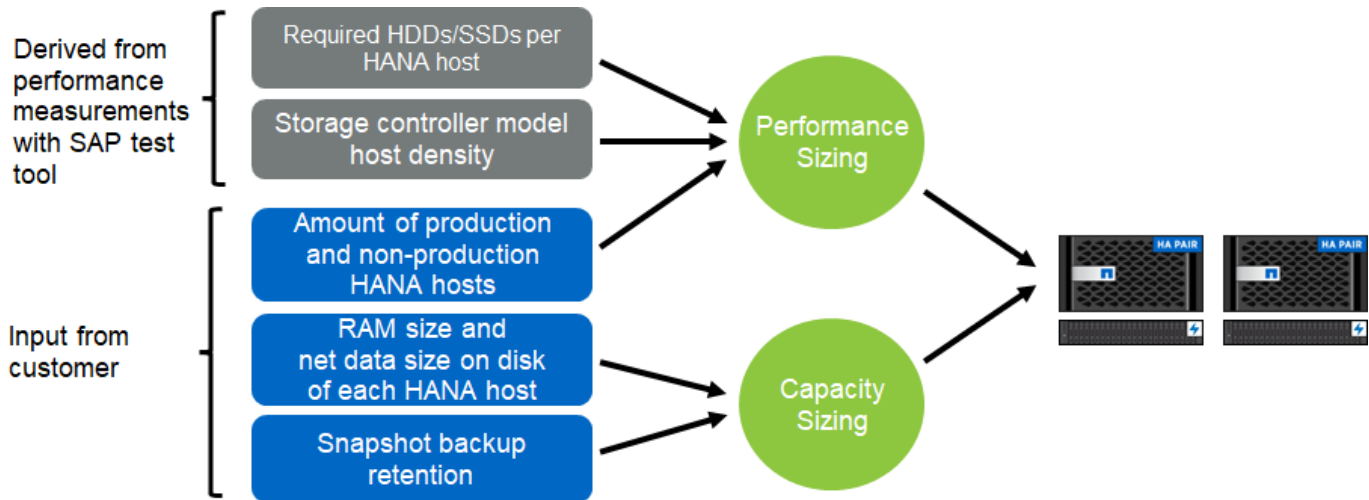
Storage sizing process overview

The number of disks per HANA host and the SAP HANA host density for each storage model were determined with the SAP performance 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 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, 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

Network setup

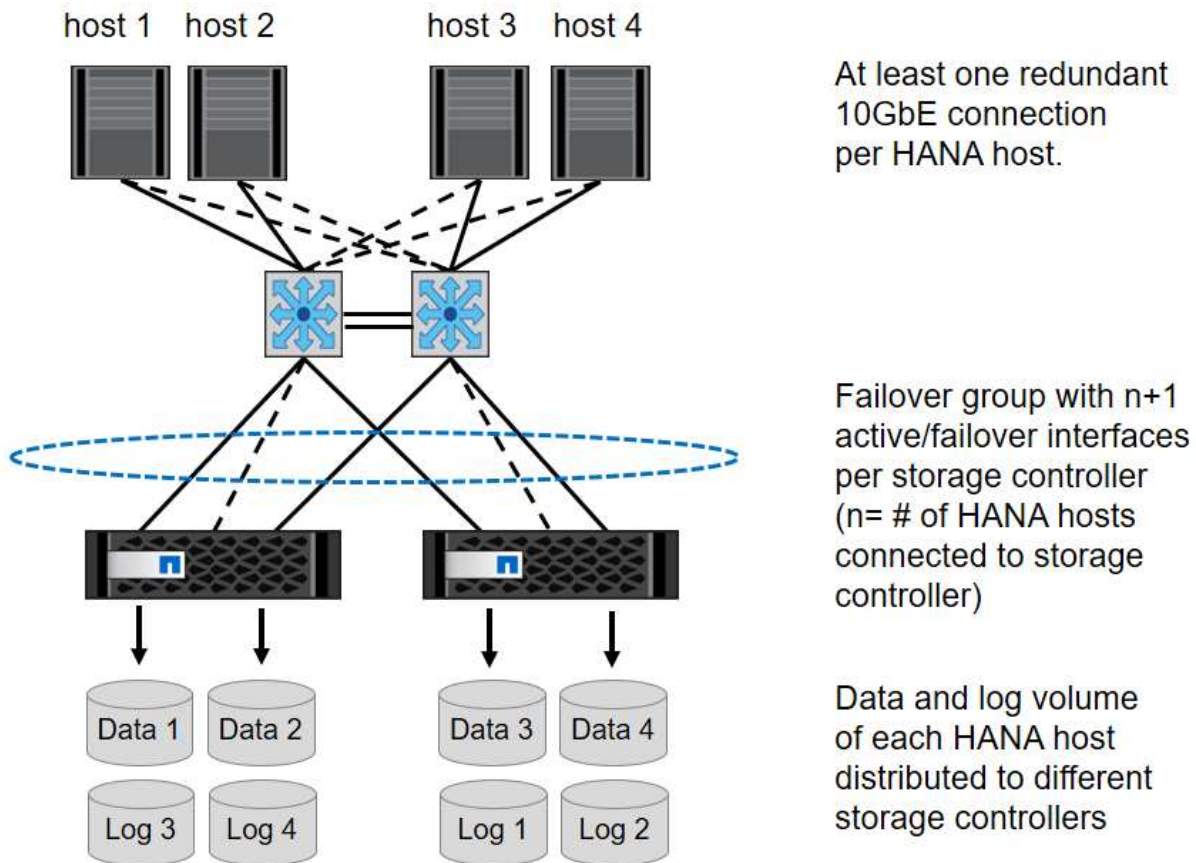
Use the following guidelines when configuring the network:

- A dedicated storage network must be used to connect the SAP HANA hosts to the storage controllers with a 10GbE or faster network.
- Use the same connection speed for storage controllers and SAP HANA hosts. If this is not possible, ensure that the network components between the storage controllers and the SAP HANA hosts are able to handle different speeds. For example, you must provide enough buffer space to allow speed negotiation at the NFS level between storage and hosts. Network components are usually switches, but other components within blade chassis, such as the back plane, must be considered as well.
- Disable flow control on all physical ports used for storage traffic on the storage network switch and host layer.
- Each SAP HANA host must have a redundant network connection with a minimum of 10Gb of bandwidth.
- Jumbo frames with a maximum transmission unit (MTU) size of 9,000 must be enabled on all network components between the SAP HANA hosts and the storage controllers.
- In a VMware setup, dedicated VMXNET3 network adapters must be assigned to each running virtual machine. Check the relevant papers mentioned in the [Introduction](#) for further requirements.
- To avoid interference between each other, use separate network/IO paths for the log and data area.

The following figure shows an example with four SAP HANA hosts attached to a storage controller HA pair using a 10GbE network. Each SAP HANA host has an active-passive connection to the redundant fabric.

At the storage layer, four active connections are configured to provide 10Gb throughput for each SAP HANA host. In addition, one spare interface is configured on each storage controller.

At the storage layer, a broadcast domain with an MTU size of 9000 is configured, and all required physical interfaces are added to this broadcast domain. This approach automatically assigns these physical interfaces to the same failover group. All logical interfaces (LIFs) that are assigned to these physical interfaces are added to this failover group.



In general, it is also possible to use HA interface groups on the servers (bonds) and the storage systems (for example, Link Aggregation Control Protocol [LACP] and ifgroups). With HA interface groups, verify that the load is equally distributed between all interfaces within the group. The load distribution depends on the functionality of the network switch infrastructure.




Depending on the number of SAP HANA hosts and the connection speed used, different numbers of active physical ports are needed.

VMware-specific network setup

Because all data for SAP HANA instances, including performance-critical data and log volumes for the database, is provided through NFS in this solution, proper network design and configuration are crucial. A dedicated storage network is used to separate the NFS traffic from communication and user access traffic between SAP HANA nodes. Each SAP HANA node requires a redundant dedicated network connection with a minimum of 10Gb of bandwidth. Higher bandwidth is also supported. This network must extend end to end

from the storage layer through network switching and computing up to the guest operating system hosted on VMware vSphere. In addition to the physical switching infrastructure, a VMware distributed switch (vDS) is used to provide adequate performance and manageability of network traffic at the hypervisor layer.

The following figure provide a network overview.



Each SAP HANA node uses a dedicated port group on the VMware distributed switch. This port group allows for enhanced quality of service (QoS) and dedicated assignment of physical network interface cards (NICs) on the ESX hosts. To use dedicated physical NICs while preserving HA capabilities if there was a NIC failure, the dedicated physical NIC is configured as an active uplink. Additional NICs are configured as standby uplinks in the teaming and failover settings of the SAP HANA port group. In addition, jumbo frames (MTU 9,000) must be enabled end to end on physical and virtual switches. In addition, turn off flow control on all ethernet ports used for storage traffic on servers, switches, and storage systems. The following figure shows an example of such a configuration.



LRO (large receive offload) must be turned off for interfaces used for NFS traffic. For all other network configuration guidelines, see the respective VMware best practices guides for SAP HANA.

General

Advanced

Security

Traffic shaping

VLAN

Teaming and failover

Monitoring

Traffic filtering and marking

Miscellaneous

Load balancing:

Route based on originating virtual port

Network failure detection:

Link status only

Notify switches:

Yes

Failback:

Yes

Failover order

↑

↓

Active uplinks

dvUplink2

Standby uplinks

dvUplink1

Unused uplinks

Time synchronization

You must synchronize the time between the storage controllers and the SAP HANA database hosts. To do so, set the same time server 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 storage efficiency features in an HDD-based 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 results 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 other applications on a shared-use controller. 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.

Starting with ONTAP 9, QoS is configured on the storage volume level and uses maximum values for throughput (MBps) and the amount of I/O (IOPS).

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.

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 a 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.

Storage configuration


The following overview summarizes the required storage configuration steps. Each step is covered in detail in the subsequent sections. In this section, we assume that the storage hardware is set up and that the ONTAP software is already installed. Also, the connections between the storage ports (10GbE or faster) and the network must already be in place.

1. Check the correct SAS stack 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 LIFs as described in [Logical interface configuration](#).
5. Create volumes within the aggregates as described in [Volume configuration for SAP HANA single-host systems](#) and [Volume configuration for SAP HANA multiple-host systems](#).
6. Set the required volume options as described in [Volume options](#).
7. Set the required options for NFSv3 as described in [NFS configuration for NFSv3](#) or for NFSv4 as described in [NFS configuration for NFSv4](#).


8. Mount the volumes to namespace and set export policies as described in [Mount volumes to namespace and set export policies](#).

Disk shelf connection

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 (100GbE) 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 the disk shelf or drive technology (SSD or HDD) that is used. For FAS2000 series systems, one data aggregate is enough.

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, you must configure two aggregates per controller, independent of which disk shelf or disk technology (SSDs or HDDs) is used. For FAS2000 series systems, one data aggregate is enough.

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 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.

If a QoS profile was automatically created and assigned during new SVM creation, remove the automatically created profile from the SVM to provide the required performance for SAP HANA:

```
vserver modify -vserver <svm-name> -qos-policy-group none
```

Logical interface configuration

For SAP HANA production systems, you must use different LIFs for mounting the data volume and the log volume from the SAP HANA host. Therefore at least two LIFs are required.

The data and log volume mounts of different SAP HANA hosts can share a physical storage network port by using either the same LIFs or by using individual LIFs for each mount.

The maximum number of data and log volume mounts per physical interface are shown in the following table.

Ethernet port speed	10GbE	25GbE	40GbE	100GbE
Maximum number of log or data volume mounts per physical port	3	8	12	30



Sharing one LIF between different SAP HANA hosts might require a remount of data or log volumes to a different LIF. This change avoids performance penalties if a volume is moved to a different storage controller.

Development and test systems can use more data and volume mounts or LIFs on a physical network interface.

For production, development, and test systems, the `/hana/shared` file system can use the same LIF as the data or log volume.

Volume 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.



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



If the data and log volumes are stored on the same controller, access from the server to the storage must be performed with two different LIFs: one LIF to access the data volume and one to access the log volume.



For each SAP HANA DB host, a data volume, a log volume, and a volume for `/hana/shared` are configured. The following table shows an example configuration for single-host SAP HANA 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 following table shows an example of the mount point configuration for a single-host system. To place the home directory of the `sidadm` user on the central storage, the `/usr/sap/SID` file system should be mounted

from the `SID_shared` volume.

Junction Path	Directory	Mount point at HANA host
<code>SID_data_mnt00001</code>	–	<code>/hana/data/SID/mnt00001</code>
<code>SID_log_mnt00001</code>	–	<code>/hana/log/SID/mnt00001</code>
<code>SID_shared</code>	<code>usr-sap</code> <code>shared</code>	<code>/usr/sap/SID</code> <code>/hana/shared</code>

Volume configuration for SAP HANA multiple-host systems

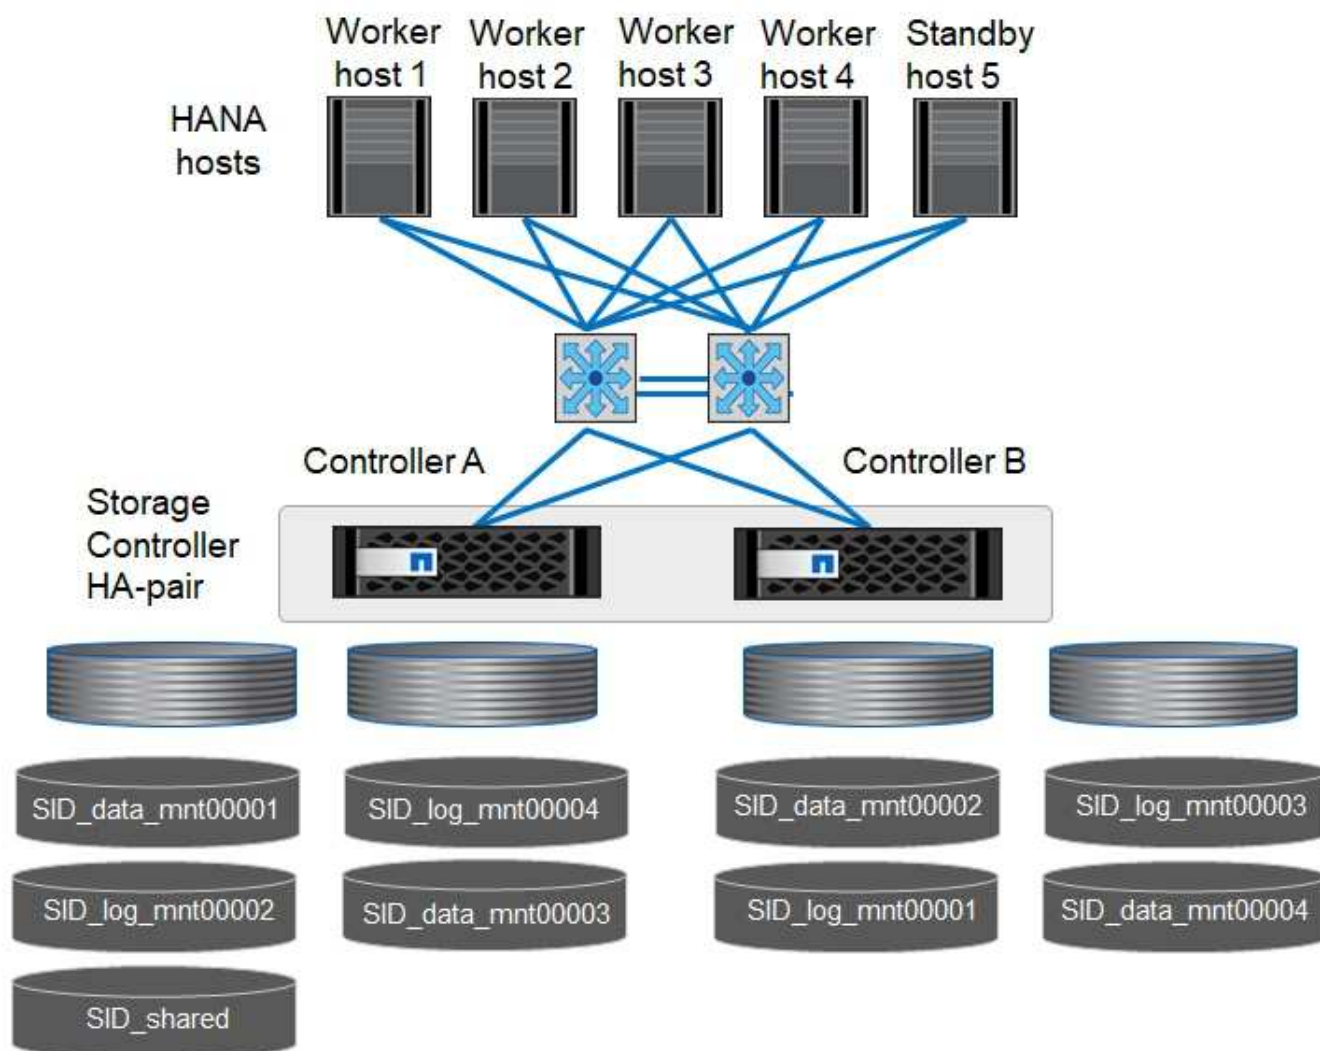
The following figure shows the volume configuration of a 4+1 SAP HANA system. The data and log volumes of each SAP HANA host are distributed to different storage controllers. For example, volume `SID1_data1_mnt00001` is configured on controller A, and volume `SID1_log1_mnt00001` is configured on controller B.



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



If the data and log volumes are stored on the same controller, access from the server to the storage must be performed with two different LIFs: one to access the data volume and one to access the log volume.



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 table shows an example configuration for a multiple-host SAP HANA system with four active hosts.

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 following table shows the configuration and the mount points of a multiple-host system with four active SAP HANA hosts. To place the home directories of the `sidadm` user of each host on the central storage, the

/usr/sap/SID file systems are mounted from the SID_shared volume.

Junction path	Directory	Mount point at SAP HANA host	Note
SID_data_mnt00001	–	/hana/data/SID/mnt00001	Mounted at all hosts
SID_log_mnt00001	–	/hana/log/SID/mnt00001	Mounted at all hosts
SID_data_mnt00002	–	/hana/data/SID/mnt00002	Mounted at all hosts
SID_log_mnt00002	–	/hana/log/SID/mnt00002	Mounted at all hosts
SID_data_mnt00003	–	/hana/data/SID/mnt00003	Mounted at all hosts
SID_log_mnt00003	–	/hana/log/SID/mnt00003	Mounted at all hosts
SID_data_mnt00004	–	/hana/data/SID/mnt00004	Mounted at all hosts
SID_log_mnt00004	–	/hana/log/SID/mnt00004	Mounted at all hosts
SID_shared	shared	/hana/shared/	Mounted at all hosts
SID_shared	usr-sap-host1	/usr/sap/SID	Mounted at host 1
SID_shared	usr-sap-host2	/usr/sap/SID	Mounted at host 2
SID_shared	usr-sap-host3	/usr/sap/SID	Mounted at host 3
SID_shared	usr-sap-host4	/usr/sap/SID	Mounted at host 4
SID_shared	usr-sap-host5	/usr/sap/SID	Mounted at host 5

Volume options

You must verify and set the volume options listed in the following table on all SVMs. For some of the commands, you must switch to the advanced privilege mode within ONTAP.

Action	Command
Disable visibility of Snapshot directory	vol modify -vserver <vserver-name> -volume <volname> -snapdir-access false
Disable automatic Snapshot copies	vol modify -vserver <vserver-name> -volume <volname> -snapshot-policy none
Disable access time update except of the SID_shared volume	set advanced vol modify -vserver <vserver-name> -volume <volname> -atime-update false set admin

NFS configuration for NFSv3

The NFS options listed in the following table must be verified and set on all storage controllers.

For some of the commands shown, you must switch to the advanced privilege mode within ONTAP.

Action	Command
Enable NFSv3	nfs modify -vserver <vserver-name> v3.0 enabled

Action	Command
Set NFS TCP maximum transfer size to 1MB	set advanced nfs modify -vserver <vserver_name> -tcp-max-xfer -size 1048576 set admin



In shared environments with different workloads set the max NFS TCP transfer size to 262144

NFS configuration for NFSv4

The NFS options listed in the following table must be verified and set on all SVMs.

For some of the commands, you must switch to the advanced privilege mode within ONTAP.

Action	Command
Enable NFSv4	nfs modify -vserver <vserver-name> -v4.1 enabled
Set NFS TCP maximum transfer size to 1MB	set advanced nfs modify -vserver <vserver_name> -tcp-max-xfer -size 1048576 set admin
Disable NFSv4 access control lists (ACLs)	nfs modify -vserver <vserver_name> -v4.1-acl disabled
Set NFSv4 domain ID	nfs modify -vserver <vserver_name> -v4-id-domain <domain-name>
Disable NFSv4 read delegation	nfs modify -vserver <vserver_name> -v4.1-read -delegation disabled
Disable NFSv4 write delegation	nfs modify -vserver <vserver_name> -v4.1-write -delegation disabled
Disable NFSv4 numeric ids	nfs modify -vserver <vserver_name> -v4-numeric-ids disabled
Change amount of NFSv4.x session slots optional	set advanced nfs modify -vserver hana -v4.x-session-num-slots <value> set admin



In shared environments with different workloads set the max NFS TCP transfer size to 262144



Please note that disabling numbering ids requires user management as described in [SAP HANA installation preparations for NFSv4](#).



The NFSv4 domain ID must be set to the same value on all Linux servers (/etc/idmapd.conf) and SVMs, as described in [SAP HANA installation preparations for NFSv4](#).



pNFS can be enabled and used.

If SAP HANA multiple-host systems with host auto-failover are being used, the failover parameters need to be adjusted within `nameserver.ini` as shown in the following table. Keep the default retry interval of 10 seconds within these sections.

Section within <code>nameserver.ini</code>	Parameter	Value
failover	normal_retries	9
distributed_watchdog	deactivation_retries	11
distributed_watchdog	takeover_retries	9

Mount volumes to namespace and set export policies

When a volume is created, the volume must be mounted to the namespace. In this document, we assume that the junction path name is the same as the volume name. By default, the volume is exported with the default policy. The export policy can be adapted if required.

Host setup

All the steps described in this section are valid for both SAP HANA environments on physical servers and for SAP HANA running on VMware vSphere.

Configuration parameter for SUSE Linux Enterprise Server

Additional kernel and configuration parameters at each SAP HANA host must be adjusted for the workload generated by SAP HANA.

SUSE Linux Enterprise Server 12 and 15

Starting with SUSE Linux Enterprise Server (SLES) 12 SP1, the kernel parameter must be set in a configuration file in the `/etc/sysctl.d` directory. For example, a configuration file with the name `91-NetApp-HANA.conf` must be created.

```
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.ipv4.tcp_rmem = 4096 131072 16777216
net.ipv4.tcp_wmem = 4096 16384 16777216
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_slow_start_after_idle = 0
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
sunrpc.tcp_max_slot_table_entries = 128
```



Saptune, which is included in SLES for SAP OS versions, can be used to set these values. See [SAP Note 3024346](#) (requires SAP login).

Configuration parameter for Red Hat Enterprise Linux 7.2 or later

You must adjust additional kernel and configuration parameters at each SAP HANA host for the workload generated by SAP HANA.

Starting with Red Hat Enterprise Linux 7.2, you must set the kernel parameters in a configuration file in the `/etc/sysctl.d` directory. For example, a configuration file with the name `91-NetApp-HANA.conf` must be created.

```
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
net.ipv4.tcp_rmem = 4096 131072 16777216
net.ipv4.tcp_wmem = 4096 16384 16777216
net.core.netdev_max_backlog = 300000
net.ipv4.tcp_slow_start_after_idle = 0
net.ipv4.tcp_no_metrics_save = 1
net.ipv4.tcp_moderate_rcvbuf = 1
net.ipv4.tcp_window_scaling = 1
net.ipv4.tcp_timestamps = 1
net.ipv4.tcp_sack = 1
sunrpc.tcp_max_slot_table_entries = 128
```



Since Red Hat Enterprise Linux version 8.6, these settings can also be applied by using RHEL System Roles for SAP (Ansible). See [SAP Note 3024346](#) (requires SAP login).

Create subdirectories in `/hana/shared` volume



The examples show an SAP HANA database with `SID=NF2`.

To create the required subdirectories, take one of the following actions:

- For a single-host system, mount the `/hana/shared` volume and create the `shared` and `usr-sap` subdirectories.

```
sapcc-hana-tst-06:/mnt # mount <storage-hostname>:/NF2_shared /mnt/tmp
sapcc-hana-tst-06:/mnt # cd /mnt/tmp
sapcc-hana-tst-06:/mnt/tmp # mkdir shared
sapcc-hana-tst-06:/mnt/tmp # mkdir usr-sap
sapcc-hana-tst-06:/mnt/tmp # cd ..
sapcc-hana-tst-06:/mnt # umount /mnt/tmp
```

- For a multiple-host system, mount the `/hana/shared` volume and create the `shared` and the `usr-sap` subdirectories for each host.

The example commands show a 2+1 multiple-host HANA system.

```

sapcc-hana-tst-06:/mnt # mount <storage-hostname>:/NF2_shared /mnt/tmp
sapcc-hana-tst-06:/mnt # cd /mnt/tmp
sapcc-hana-tst-06:/mnt/tmp # mkdir shared
sapcc-hana-tst-06:/mnt/tmp # mkdir usr-sap-host1
sapcc-hana-tst-06:/mnt/tmp # mkdir usr-sap-host2
sapcc-hana-tst-06:/mnt/tmp # mkdir usr-sap-host3
sapcc-hana-tst-06:/mnt/tmp # cd ..
sapcc-hana-tst-06:/mnt # umount /mnt/tmp

```

Create mount points



The examples show an SAP HANA database with SID=NF2.

To create the required mount point directories, take one of the following actions:

- For a single-host system, create mount points and set the permissions on the database host.

```

sapcc-hana-tst-06:/ # mkdir -p /hana/data/NF2/mnt00001
sapcc-hana-tst-06:/ # mkdir -p /hana/log/NF2/mnt00001
sapcc-hana-tst-06:/ # mkdir -p /hana/shared
sapcc-hana-tst-06:/ # mkdir -p /usr/sap/NF2
sapcc-hana-tst-06:/ # chmod -R 777 /hana/log/NF2
sapcc-hana-tst-06:/ # chmod -R 777 /hana/data/NF2
sapcc-hana-tst-06:/ # chmod -R 777 /hana/shared
sapcc-hana-tst-06:/ # chmod -R 777 /usr/sap/NF2

```

- For a multiple-host system, create mount points and set the permissions on all worker and standby hosts.

The following example commands are for a 2+1 multiple-host HANA system.

- First worker host:

```

sapcc-hana-tst-06:~ # mkdir -p /hana/data/NF2/mnt00001
sapcc-hana-tst-06:~ # mkdir -p /hana/data/NF2/mnt00002
sapcc-hana-tst-06:~ # mkdir -p /hana/log/NF2/mnt00001
sapcc-hana-tst-06:~ # mkdir -p /hana/log/NF2/mnt00002
sapcc-hana-tst-06:~ # mkdir -p /hana/shared
sapcc-hana-tst-06:~ # mkdir -p /usr/sap/NF2
sapcc-hana-tst-06:~ # chmod -R 777 /hana/log/NF2
sapcc-hana-tst-06:~ # chmod -R 777 /hana/data/NF2
sapcc-hana-tst-06:~ # chmod -R 777 /hana/shared
sapcc-hana-tst-06:~ # chmod -R 777 /usr/sap/NF2

```

- Second worker host:

```

sapcc-hana-tst-07:~ # mkdir -p /hana/data/NF2/mnt00001
sapcc-hana-tst-07:~ # mkdir -p /hana/data/NF2/mnt00002
sapcc-hana-tst-07:~ # mkdir -p /hana/log/NF2/mnt00001
sapcc-hana-tst-07:~ # mkdir -p /hana/log/NF2/mnt00002
sapcc-hana-tst-07:~ # mkdir -p /hana/shared
sapcc-hana-tst-07:~ # mkdir -p /usr/sap/NF2
sapcc-hana-tst-07:~ # chmod -R 777 /hana/log/NF2
sapcc-hana-tst-07:~ # chmod -R 777 /hana/data/NF2
sapcc-hana-tst-07:~ # chmod -R 777 /hana/shared
sapcc-hana-tst-07:~ # chmod -R 777 /usr/sap/NF2

```

- Standby host:

```

sapcc-hana-tst-08:~ # mkdir -p /hana/data/NF2/mnt00001
sapcc-hana-tst-08:~ # mkdir -p /hana/data/NF2/mnt00002
sapcc-hana-tst-08:~ # mkdir -p /hana/log/NF2/mnt00001
sapcc-hana-tst-08:~ # mkdir -p /hana/log/NF2/mnt00002
sapcc-hana-tst-08:~ # mkdir -p /hana/shared
sapcc-hana-tst-08:~ # mkdir -p /usr/sap/NF2
sapcc-hana-tst-08:~ # chmod -R 777 /hana/log/NF2
sapcc-hana-tst-08:~ # chmod -R 777 /hana/data/NF2
sapcc-hana-tst-08:~ # chmod -R 777 /hana/shared
sapcc-hana-tst-08:~ # chmod -R 777 /usr/sap/NF2

```

Mount file systems

Different mount options are used depending on the NFS version and ONTAP release. The following file systems must be mounted to the hosts:

- /hana/data/SID/mnt0000*
- /hana/log/SID/mnt0000*
- /hana/shared
- /usr/sap/SID

The following table shows the NFS versions that must be used for the different file systems for single-host and multiple-host SAP HANA databases.

File systems	SAP HANA single host	SAP HANA multiple hosts
/hana/data/SID/mnt0000*	NFSv3 or NFSv4	NFSv4
/hana/log/SID/mnt0000*	NFSv3 or NFSv4	NFSv4
/hana/shared	NFSv3 or NFSv4	NFSv3 or NFSv4
/usr/sap/SID	NFSv3 or NFSv4	NFSv3 or NFSv4

The following table shows the mount options for the various NFS versions and ONTAP releases. The common parameters are independent of the NFS and ONTAP versions.



SAP LaMa requires the `/usr/sap/SID` directory to be local. Therefore, do not mount an NFS volume for `/usr/sap/SID` if you are using SAP LaMa.

For NFSv3, you must switch off NFS locking to avoid NFS lock cleanup operations if there is a software or server failure.

With ONTAP 9, the NFS transfer size can be configured up to 1MB. Specifically, with 40GbE or faster connections to the storage system, you must set the transfer size to 1MB to achieve the expected throughput values.

Common parameter	NFSv3	NFSv4	NFS transfer size with ONTAP 9	NFS transfer size with ONTAP 8
rw, bg, hard, timeo=600, noatime,	nfsvers=3,nolock,	nfsvers=4.1,lock	rsiz=1048576,wsiz=262144,	rsiz=65536,wsiz=65536,



To improve read performance with NFSv3, NetApp recommends that you use the `nconnect=n` mount option, which is available with SUSE Linux Enterprise Server 12 SP4 or later and RedHat Enterprise Linux (RHEL) 8.3 or later.



Performance tests show that `nconnect=4` provides good read results especially for the data volumes. Log writes might benefit from a lower number of sessions, such as `nconnect=2`. Shared volumes might benefit as well from using the 'nconnect' option. Be aware that the first mount from an NFS server (IP address) defines the amount of sessions being used. Further mounts to the same IP address do not change this even if a different value is used for `nconnect`.



Starting with ONTAP 9.8 and SUSE SLES15SP2 or RedHat RHEL 8.4 or higher, NetApp supports the `nconnect` option also for NFSv4.1.



If `nconnect` is being used with NFSv4.x the amount of NFSv4.x session slots should be adjusted according to the following rule:

Amount of session slots equals `<nconnect value> x 64`.

At the host this will be adjusted by

```
echo options nfs max_session_slots=<calculated value> >
/etc/modprobe.d/nfsclient.conf
```

followed by a reboot. The server side value must be adjusted as well, set the number of session slots as described in [NFS configuration for NFSv4](#).

To mount the file systems during system boot with the `/etc/fstab` configuration file, complete the following steps:

The following example shows a single host SAP HANA database with `SID=NF2` using NFSv3 and an NFS transfer size of 1MB for reads and 256k for writes.

1. Add the required file systems to the `/etc/fstab` configuration file.

```

sapcc-hana-tst-06:/ # cat /etc/fstab
<storage-vif-data01>:/NF2_data_mnt00001 /hana/data/NF2/mnt00001 nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsiz=1048576,wsiz=262144,bg,noa
time,nolock 0 0
<storage-vif-log01>:/NF2_log_mnt00001 /hana/log/NF2/mnt00001 nfs
rw,nfsvers=3,hard,timeo=600,nconnect=2,rsiz=1048576,wsiz=262144,bg,noa
time,nolock 0 0
<storage-vif-data01>:/NF2_shared/usr-sap /usr/sap/NF2 nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsiz=1048576,wsiz=262144,bg,noa
time,nolock 0 0
<storage-vif-data01>:/NF2_shared/shared /hana/shared nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsiz=1048576,wsiz=262144,bg,noa
time,nolock 0 0

```

2. Run `mount -a` to mount the file systems on all hosts.

The next example shows a multiple-host SAP HANA database with `SID=NF2` using NFSv4.1 for data and log file systems and NFSv3 for the `/hana/shared` and `/usr/sap/NF2` file systems. An NFS transfer size of 1MB for reads and 256k for writes is used.

1. Add the required file systems to the `/etc/fstab` configuration file on all hosts.



The `/usr/sap/NF2` file system is different for each database host. The following example shows `/NF2_shared/usr-sap-host1`.

```

sapcc-hana-tst-06:/ # cat /etc/fstab
<storage-vif-data01>:/NF2_data_mnt00001 /hana/data/NF2/mnt00001 nfs
rw,nfsvers=4.1,hard,timeo=600,nconnect=4,rsiz=1048576,wsiz=262144,bg,n
oatime,lock 0 0
<storage-vif-data02>:/NF2_data_mnt00002 /hana/data/NF2/mnt00002 nfs
rw,nfsvers=4.1,hard,timeo=600,nconnect=4,rsiz=1048576,wsiz=262144,bg,n
oatime,lock 0 0
<storage-vif-log01>:/NF2_log_mnt00001 /hana/log/NF2/mnt00001 nfs
rw,nfsvers=4.1,hard,timeo=600,nconnect=2,rsiz=1048576,wsiz=262144,bg,n
oatime,lock 0 0
<storage-vif-log02>:/NF2_log_mnt00002 /hana/log/NF2/mnt00002 nfs
rw,nfsvers=4.1,hard,timeo=600,nconnect=2,rsiz=1048576,wsiz=262144,bg,n
oatime,lock 0 0
<storage-vif-data02>:/NF2_shared/usr-sap-host1 /usr/sap/NF2 nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsiz=1048576,wsiz=262144,bg,noa
time,nolock 0 0
<storage-vif-data02>:/NF2_shared/shared /hana/shared nfs
rw,nfsvers=3,hard,timeo=600,nconnect=4,rsiz=1048576,wsiz=262144,bg,noa
time,nolock 0 0

```

2. Run `mount -a` to mount the file systems on all hosts.

SAP HANA installation preparations for NFSv4

NFS version 4 and higher requires user authentication. This authentication can be accomplished by using a central user management tool such as a Lightweight Directory Access Protocol (LDAP) server or with local user accounts. The following sections describe how to configure local user accounts.

The administration users `<sid>adm`, `<sid>crypt` and the `sapsys` group must be created manually on the SAP HANA hosts and the storage controllers before the installation of the SAP HANA software begins.

SAP HANA hosts

If it doesn't exist, the `sapsys` group must be created on the SAP HANA host. A unique group ID must be chosen that does not conflict with the existing group IDs on the storage controllers.

The users `<sid>adm` and `<sid>crypt` are created on the SAP HANA host. Unique IDs must be chosen that does not conflict with existing user IDs on the storage controllers.

For a multiple-host SAP HANA system, the user and group IDs must be the same on all SAP HANA hosts. The group and user are created on the other SAP HANA hosts by copying the affected lines in `/etc/group` and `/etc/passwd` from the source system to all other SAP HANA hosts.



The NFSv4 domain must be set to the same value on all Linux servers (`/etc/idmapd.conf`) and SVMs. Set the domain parameter "Domain = `<domain-name>`" in the file `/etc/idmapd.conf` for the Linux hosts.

Enable and start the NFS IDMAPD service.

```
systemctl enable nfs-idmapd.service
systemctl start nfs-idmapd.service
```



The latest Linux kernels do not require this step. Warning messages can be safely ignored.

Storage controllers

The user IDs and group ID must be the same on the SAP HANA hosts and the storage controllers. The group and user are created by entering the following commands on the storage cluster:

```
vserver services unix-group create -vserver <vserver> -name <group name>
-id <group id>
vserver services unix-user create -vserver <vserver> -user <user name> -id
<user-id> -primary-gid <group id>
```

Additionally, set the group ID of the UNIX user root of the SVM to 0.

```
vserver services unix-user modify -vserver <vserver> -user root -primary  
-gid 0
```

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 systems used.

NetApp conducted performance tests to define the ideal values. The following table lists the optimal values 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 versions 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 by using the `hdbparam` framework.

```
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/HDB00> hdbparam --paramset  
fileio.max_parallel_io_requests=128  
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/HDB00> hdbparam --paramset  
fileio.async_write_submit_active=on  
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/HDB00> hdbparam --paramset  
fileio.async_read_submit=on  
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/HDB00> hdbparam --paramset  
fileio.async_write_submit_blocks=all
```

Starting with SAP HANA 2.0, `hdbparam` has been deprecated, and the parameters have been moved to `global.ini`. The parameters can be set using SQL commands or SAP HANA Studio. For more details, see SAP note [2399079: Elimination of hdbparam in HANA 2](#). You can also set the parameters within `global.ini` as shown in the following text:

```
nf2adm@stlrx300s8-6: /usr/sap/NF2/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
...
```

Since SAP HANA 2.0 SPS5, the `setParameter.py` script can be used to set the correct parameters:

```
nf2adm@sapcc-hana-tst-06:/usr/sap/NF2/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 data volume size

As the default, SAP HANA uses only one data volume per SAP HANA service. Due to the maximum file size limitation of the file system, we recommend limiting the maximum data volume size.

To do so automatically, set the following parameter in `global.ini` in the section `[persistence]`:

```
datavolume_stripping = true
datavolume_stripping_size_gb = 8000
```

This creates a new data volume after the 8,000GB limit is reached. [SAP note 240005 question 15](#) provides more information.

SAP HANA software installation

The following are requirements for software installation for SAP HANA.

Install on single-host system

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

Install on multiple-host system

To install SAP HANA on a multiple-host system, complete the following steps:

1. Using the SAP `hdbclm` 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-03`) and the standby host (`sapcc-hana-tst-04`).

```
apcc-hana-tst-02:/mnt/sapcc-share/software/SAP/HANA2SPS7-
73/DATA_UNITS/HDB_LCM_LINUX_X86_64 #
./hdbclm --action=install --addhosts=sapcc-hana-tst-03:role=worker,sapcc
-hana-tst-04:role=standby

SAP 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

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

Adding additional data volume partitions

Starting with SAP HANA 2.0 SPS4, you can configure additional data volume partitions, which allows you to configure two or more volumes for the data volume of an SAP HANA tenant database. You can also scale beyond the size and performance limits of a single volume.



Using two or more individual volumes for the data volume is available for SAP HANA single-host and multiple-host systems. You can add additional data volume partitions at any time, but doing so might require a restart of the SAP HANA database.

Enabling additional data volume partitions

1. To enable additional data volume partitions, add the following entry within `global.ini` using SAP HANA Studio or Cockpit in the SYSTEMDB configuration.

```

[customizable_functionalities]
persistence_datavolume_partition_multipath = true

```



Adding the parameter manually to the `global.ini` file requires the restart of the database.

Volume configuration for a single-host SAP HANA system

The layout of volumes for a single-host SAP HANA system with multiple partitions is like the layout for a system with one data volume partition, but with an additional data volume stored on a different aggregate as the log volume and the other data volume. The following table shows an example configuration of an SAP HANA single-host system 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 following table shows an example of the mount point configuration for a single-host system with two data volume partitions.

Junction path	Directory	Mount point at HANA host
SID_data_mnt00001	–	/hana/data/SID/mnt00001
SID_data2_mnt00001	–	/hana/data2/SID/mnt00001
SID_log_mnt00001	–	/hana/log/SID/mnt00001
SID_shared	usr-sap shared	/usr/sap/SID /hana/shared

Create the new data volume and mount it to the namespace using either ONTAP System Manager or the ONTAP cluster command line interface.

Volume configuration for multiple-host SAP HANA system

The layout of volumes for a multiple-host SAP HANA system with multiple partitions is like the layout for a system with one data volume partition, but with an additional data volume stored on a different aggregate as the log volume and the other data volume. The following table shows an example configuration of an SAP HANA multiple-host system with two data volume partitions.

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	Data2 volume: SID_data2_mnt00001
Data and log volumes for node 2	Log volume: SID_log_mnt00002	Data2 volume: SID_data2_mnt00002	Data volume: SID_data_mnt00002	–
Data and log volumes for node 3	–	Data volume: SID_data_mnt00003	Data2 volume: SID_data2_mnt00003	Log volume: SID_log_mnt00003
Data and log volumes for node 4	Data2 volume: SID_data2_mnt00004	Log volume: SID_log_mnt00004	–	Data volume: SID_data_mnt00004
Shared volume for all hosts	Shared volume: SID_shared	–	–	–

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

Junction path	Directory	Mount point at SAP HANA host	Note
SID_data_mnt00001	–	/hana/data/SID/mnt00001	Mounted at all hosts
SID_data2_mnt00001	–	/hana/data2/SID/mnt00001	Mounted at all hosts
SID_log_mnt00001	–	/hana/log/SID/mnt00001	Mounted at all hosts

Junction path	Directory	Mount point at SAP HANA host	Note
SID_data_mnt00002	–	/hana/data/SID/mnt00002	Mounted at all hosts
SID_data2_mnt00002	–	/hana/data2/SID/mnt00002	Mounted at all hosts
SID_log_mnt00002	–	/hana/log/SID/mnt00002	Mounted at all hosts
SID_data_mnt00003	–	/hana/data/SID/mnt00003	Mounted at all hosts
SID_data2_mnt00003	–	/hana/data2/SID/mnt00003	Mounted at all hosts
SID_log_mnt00003	–	/hana/log/SID/mnt00003	Mounted at all hosts
SID_data_mnt00004	–	/hana/data/SID/mnt00004	Mounted at all hosts
SID_data2_mnt00004	–	/hana/data2/SID/mnt00004	Mounted at all hosts
SID_log_mnt00004	–	/hana/log/SID/mnt00004	Mounted at all hosts
SID_shared	shared	/hana/shared/SID	Mounted at all hosts
SID_shared	usr-sap-host1	/usr/sap/SID	Mounted at host 1
SID_shared	usr-sap-host2	/usr/sap/SID	Mounted at host 2
SID_shared	usr-sap-host3	/usr/sap/SID	Mounted at host 3
SID_shared	usr-sap-host4	/usr/sap/SID	Mounted at host 4
SID_shared	usr-sap-host5	/usr/sap/SID	Mounted at host 5

Create the new data volume and mount it to the namespace using either ONTAP System Manager or the ONTAP cluster command line interface.

Host configuration

In addition to the tasks described in the section “[Host setup](#),” you must create the additional mount points and fstab entries for the new additional data volume(s), and you must mount the new volumes.

1. Create additional mount points:

- For a single-host system, create mount points and set the permissions on the database host.

```
sapcc-hana-tst-06:/ # mkdir -p /hana/data2/SID/mnt00001
sapcc-hana-tst-06:/ # chmod -R 777 /hana/data2/SID
```

- For a multiple-host system, create mount points and set the permissions on all worker and standby hosts. The following example commands are for a 2+1 multiple-host HANA system.
 - First worker host:

```
sapcc-hana-tst-06:~ # mkdir -p /hana/data2/SID/mnt00001
sapcc-hana-tst-06:~ # mkdir -p /hana/data2/SID/mnt00002
sapcc-hana-tst-06:~ # chmod -R 777 /hana/data2/SID
```

▪ Second worker host:

```
sapcc-hana-tst-07:~ # mkdir -p /hana/data2/SID/mnt00001
sapcc-hana-tst-07:~ # mkdir -p /hana/data2/SID/mnt00002
sapcc-hana-tst-07:~ # chmod -R 777 /hana/data2/SID
```

▪ Standby host:

```
sapcc-hana-tst-07:~ # mkdir -p /hana/data2/SID/mnt00001
sapcc-hana-tst-07:~ # mkdir -p /hana/data2/SID/mnt00002
sapcc-hana-tst-07:~ # chmod -R 777 /hana/data2/SID
```

2. Add the additional file systems to the `/etc/fstab` configuration file on all hosts. An example for a single-host system using NFSv4.1 is as follows:

```
<storage-vif-data02>:/SID_data2_mnt00001 /hana/data2/SID/mnt00001 nfs
rw,vers=4,
minorversion=1,hard,timeo=600,rsz=1048576,wsz=262144,bg,noatime,lock
0 0
```



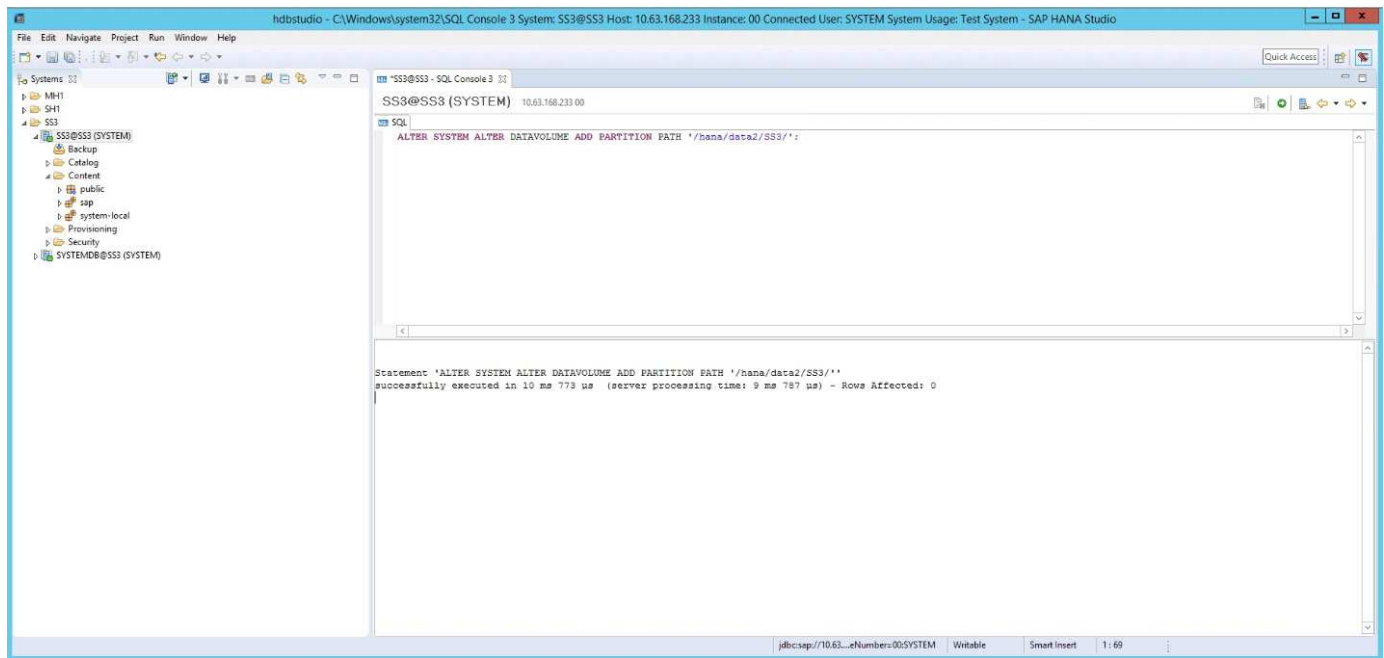
Use a different storage virtual interface for connecting to each data volume to make sure that different TCP sessions are used for each volume. You can also use the `nconnect` mount option if it is available for your OS.

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

Adding an additional data volume partition

Execute the following SQL statement against the tenant database to add an additional data volume partition to your tenant database. Use the path to additional volume(s):

```
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
April 2014	Initial version
August 2014	Updated disk sizing selection and added SSD configuration Added Red Hat Enterprise Linux OS configuration Added SAP HANA storage connector information Added information about VMware configuration
November 2014	Updated storage sizing section
January 2015	Updated storage connector API section Updated aggregate and volume configuration
March 2015	Added new STONITH implementation for SAP HANA SPS9 Added compute node setup and HANA installation section
October 2015	Added NFSv4 support for cDOT Updated sysctl parameter Included I/O parameter for SAP HANA and HWVAL > SPS10
March 2016	Updated capacity sizing Updated mount options for /hana/shared Updated sysctl parameter
February 2017	New NetApp storage systems and disk shelves New features of ONTAP 9 Support for 40GbE New OS releases (SUSE Linux Enterprise Server12 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 (SUSE Linux Enterprise Server 12 SP3 and Red Hat Enterprise Linux 7.4) Additional minor changes SAP HANA 2.0 SPS3
September 2019	New OS releases (SUSE Linux Enterprise Server 12 SP4, SUSE Linux Enterprise Server 15, and Red Hat Enterprise Linux 7.6) Max data volume size Minor changes
December 2019	New NetApp storage systems New OS release SUSE Linux Enterprise Server 15 SP1
March 2020	Support of nconnect for NFSv3 New OS release Red Hat Enterprise Linux 8
May 2020	Introduced multiple data partition features available since SAP HANA 2.0 SPS4

Date	Update summary
June 2020	Additional information about optional functionalities Minor updates
December 2020	Support for nconnect for NFSv4.1 starting with ONTAP 9.8 New operating system releases New SAP HANA version
February 2021	Changes in host network settings and other minor changes
April 2021	VMware vSphere-specific information added
September 2022	New OS-Releases
December 2023	Update of host setup Revised nconnect settings Added information about NFSv4.1 sessions
September 2024	New Storage Systems and Minor Updates
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.