



**Hewlett Packard**  
Enterprise

# **HPE Nimble Storage All-Flash Array for Use with SAP HANA Deployment Guide**

# Contents

<b>About the Author.....</b>	<b>4</b>
<b>Introduction.....</b>	<b>5</b>
Audience.....	5
<b>Architecture Overview.....</b>	<b>6</b>
Hardware Components.....	6
Software and Firmware Components.....	7
<b>Deployment.....</b>	<b>8</b>
Servers.....	8
Brocade SAN.....	8
Customer Private Network.....	8
Management Network.....	8
HPE Nimble Storage Array.....	8
Array Overview.....	8
Configure the Network.....	9
Configure Initiator Groups and WWPNS.....	10
Create Volumes.....	11
Sizing Considerations.....	14
Operating System.....	14
Install OS.....	15
Configure OS.....	15
Configure Multipathing.....	16
File Systems for SAP HANA.....	17
Configure the SAP HANA global.ini File.....	18
Test Performance and Verify Parameter Configuration.....	18
<b>Version History.....</b>	<b>20</b>

© Copyright 2018 Hewlett Packard Enterprise Development LP. All rights reserved worldwide.

## Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

## Acknowledgments

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated. Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

## Publication Date

Tuesday June 26, 2018 06:24:42

## Document ID

luk1469542685605

## Support

All documentation and knowledge base articles are available on HPE InfoSight at <https://infosight.hpe.com>. To register for HPE InfoSight, click the *Create Account* link on the main page.

Email: [support@nimblestorage.com](mailto:support@nimblestorage.com)

For all other general support contact information, go to <https://www.nimblestorage.com/customer-support/>.

## About the Author

### Craig Sullivan



*Master Technical Marketing Engineer*

Hewlett Packard Enterprise

Craig Sullivan has over 20 years of experience with SAP Basis and technical architecture. For the first 10 years of his career, Craig was a Basis administrator of global SAP implementations for several companies and consulting firms. In 2010, he brought his skills to the storage industry to develop products and solutions for SAP applications, including SAP HANA. Craig has applied his many years of experience with enterprise applications and storage technologies to helping customers integrate HPE Nimble Storage solutions into their SAP and Splunk environments.

# Introduction

The SAP HANA software is the next generation of the SAP technology platform. Originally designed to support large-scale analytics, it has been adapted to support SAP business applications. The primary innovations of SAP HANA are an in-memory database and an analytics engine that can load the entire database into memory. This capability enables rapid processing of multiple terabytes of data and dramatically reduces the time required to perform queries and complex analytics.

Although SAP HANA is an in-memory database, it still requires a persistent storage area to maintain data between restarts and after unplanned server shutdowns. SAP has implemented a certification program to ensure that hardware vendors meet the performance, scalability, and high-availability requirements of the SAP HANA software.

SAP HANA is delivered to customers in two ways: as an appliance and through the SAP HANA tailored data center integration (TDI) program. The appliance model provides a preinstalled and configured hardware platform for running SAP HANA. Hardware vendors are required to configure these systems and to allow SAP to test them to guarantee that the hardware meets SAP HANA key performance indicators. The TDI model enables customers to combine server, network, and storage resources that have been individually certified by SAP to support SAP HANA.

SAP HANA can be deployed in two configurations, scale-up and scale-out:

- **Scale-up systems** are single-server systems that run the SAP HANA database and application. In the appliance model, these systems typically contain internal storage for data persistence and some flash storage for log retention. In the TDI model, scale-up systems can be deployed with shared storage.
- **Scale-out systems** are multiple servers with a shared storage infrastructure that can support larger volumes of data by partitioning the SAP HANA database between the servers. Currently, only SAP Business Warehouse (BW) is supported on scale-out configurations.

HPE Nimble Storage arrays have achieved the SAP HANA certification for its adaptive flash arrays and for the AF40, AF60, and AF80 all-flash arrays. Customers and partners can verify the HPE Nimble Storage certification by consulting the [Certified and Supported SAP HANA Hardware Directory](#).

## Audience

Hewlett Packard Enterprise (HPE) customers and partners can use the information in this guide to install and configure the infrastructure that supports SAP HANA in a TDI implementation. For the deployment to meet the performance and availability requirements of SAP HANA, it is important to follow the guidelines and recommendations in the guide. Customers running SAP HANA can deploy either HPE Nimble Storage adaptive flash or all-flash arrays that have achieved the SAP certification, but the guide focuses on the deployment of AF-Series (all-flash) arrays.

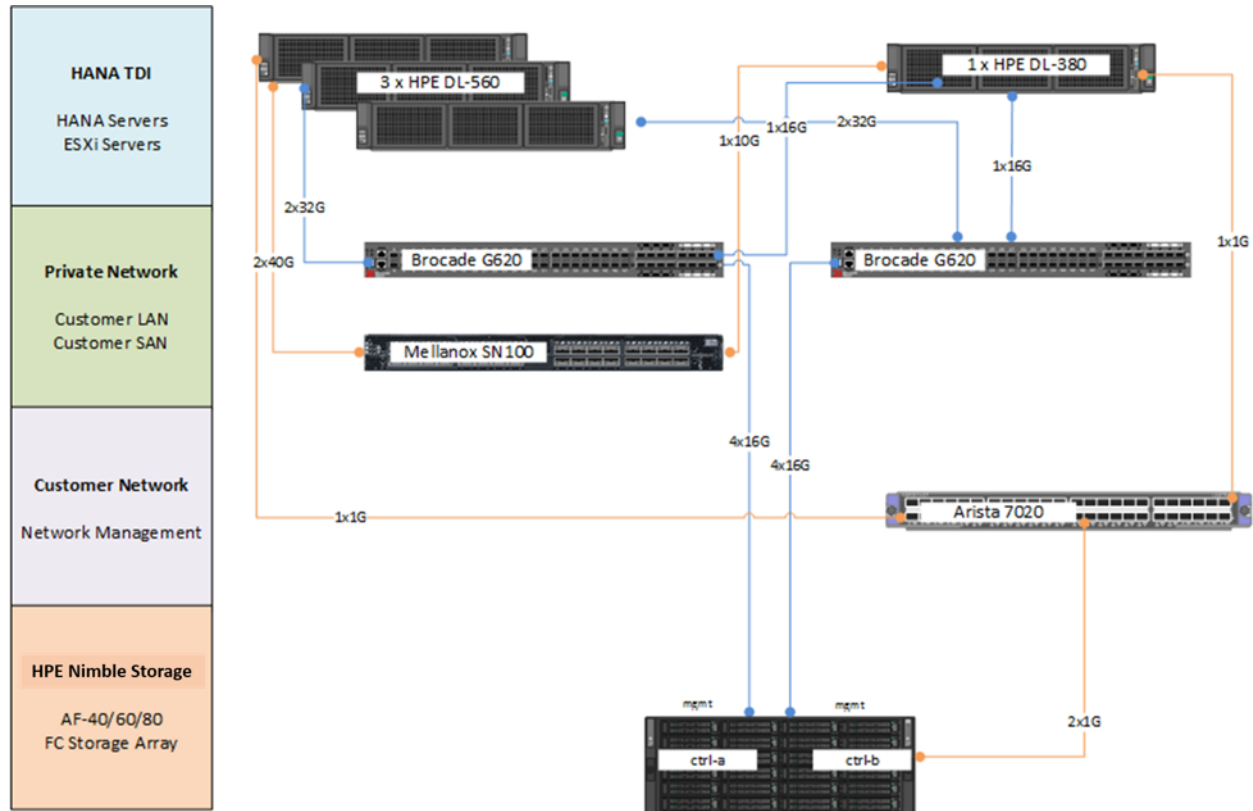
The guide assumes that the audience has a working knowledge of the following products and technologies:

- HPE Nimble Storage all-flash arrays
- Network design for Ethernet connectivity and Fibre Channel (FC) protocol
- UNIX-based operating systems (OSs)
- SAP HANA architecture and implementation
- NFS protocol

# Architecture Overview

The following infrastructure was tested for the solution that uses all-flash arrays with SAP HANA.

**Figure 1: Infrastructure layout**



## Hardware Components

The tested infrastructure includes the following hardware components.

### SAP HANA TDI

- HPE DL-560 servers
- HPE DL-380 server
- ESX 6.5

### Customer Network

- Brocade SAN switches
- Arista Ethernet Switch for Management Access
- Mellanox Ethernet Switch for Private Network

### HPE Nimble Storage All-Flash Array

- AF40 (AF-40-QT4F-11T; includes both controllers)

- 8 FC ports @ 16 Gbps per controller
- 24 x 480 GB SSDs

Depending on the storage requirements of your SAP HANA deployment, you can use other all-flash array models that are certified by SAP. The following table describes the technical specifications of these arrays.

**Table 1: Technical specifications of HPE Nimble Storage all-flash arrays certified for use with SAP HANA**

AF-Series Arrays	AF40	AF60	AF80	Scale-out 4 X AF80
Raw capacity (TB/TiB)	11-184/10-167	11-553/10-502	23-1106/21-1005	4423/4023
Usable capacity (TB/TiB)	8-136/7-124	8-407/7-370	17-815/15-741	3260/2965
Effective capacity (TB/TiB) <sup>2</sup>	40-682/36-620	40-2037/36-1853	82-4075/75-3706	16303/14827
Maximum number of expansion shelves	1	2	2	8
RAID level	Triple+ Parity			
On-board iSCSI/mgmt. 1 Gb/10 Gb ports per array	4	4	4	16
Optional iSCSI 1 Gb ports per array	4, 8, 12, 16, 20, 24	4, 8, 12, 16, 20, 24	4, 8, 12, 16, 20, 24	96
Optional iSCSI 10 Gb ports per array	4, 8, 12, 16, 20, 24	4, 8, 12, 16, 20, 24	4, 8, 12, 16, 20, 24	96
Optional FC 8 Gb/16 Gb ports per array	4, 8, 12, 16, 20, 24	4, 8, 12, 16, 20, 24	4, 8, 12, 16, 20, 24	96
Maximum power requirement (watts/kVA)	800/0.889	850/0.944	1200/1.333	4800/5.332
Thermal (BTU)	2624	2788	3936	15744

## Software and Firmware Components

The solution was tested with the following software and firmware versions.

**Table 2: Software and firmware versions**

Vendor	Model	Version
HPE	DL-560	U34 v1.36
HPE	DL-380	P89 v2.00
Brocade	G620	V8.0.1b
Arista	7020	4.20.5.2F
SUSE	SLES for SAP Applications	12 SP3
HPE	HPE Nimble Storage AF40	5.0.3.0-562630-opt
Mellanox	SN100M	3.6.4120
SAP	HANA	2.0 SP3

# Deployment

You must install and configure all of the hardware and software components of the SAP HANA TDI infrastructure addressed in this guide—switches, servers, hypervisor, storage array, and OS. For the OS configuration, you must modify some settings on the physical servers and virtual machines (VMs), create file systems, and configure the global system parameters for SAP HANA.

## Servers

Server settings for the HPE DL-560 were configured based on HPE guides. The DL-380 did not run HANA workloads but hosted the shared file system served through NFS as well as the VMware vCenter Server appliance to manage the ESXi environment.

## Brocade SAN

For high availability, FC connections were split between the two Brocade switches. Zones were created by using industry best practices.

## Customer Private Network

The Mellanox switch acted as a private-only network switch for inter-node communication and for storage networking through NFS. Two separate subnets were used: 192.168.1.0/20 and 172.16.1.0/20.

## Management Network

An Arista 7020 switch was used for management and client network access to the servers as well as to the storage system.

## HPE Nimble Storage Array

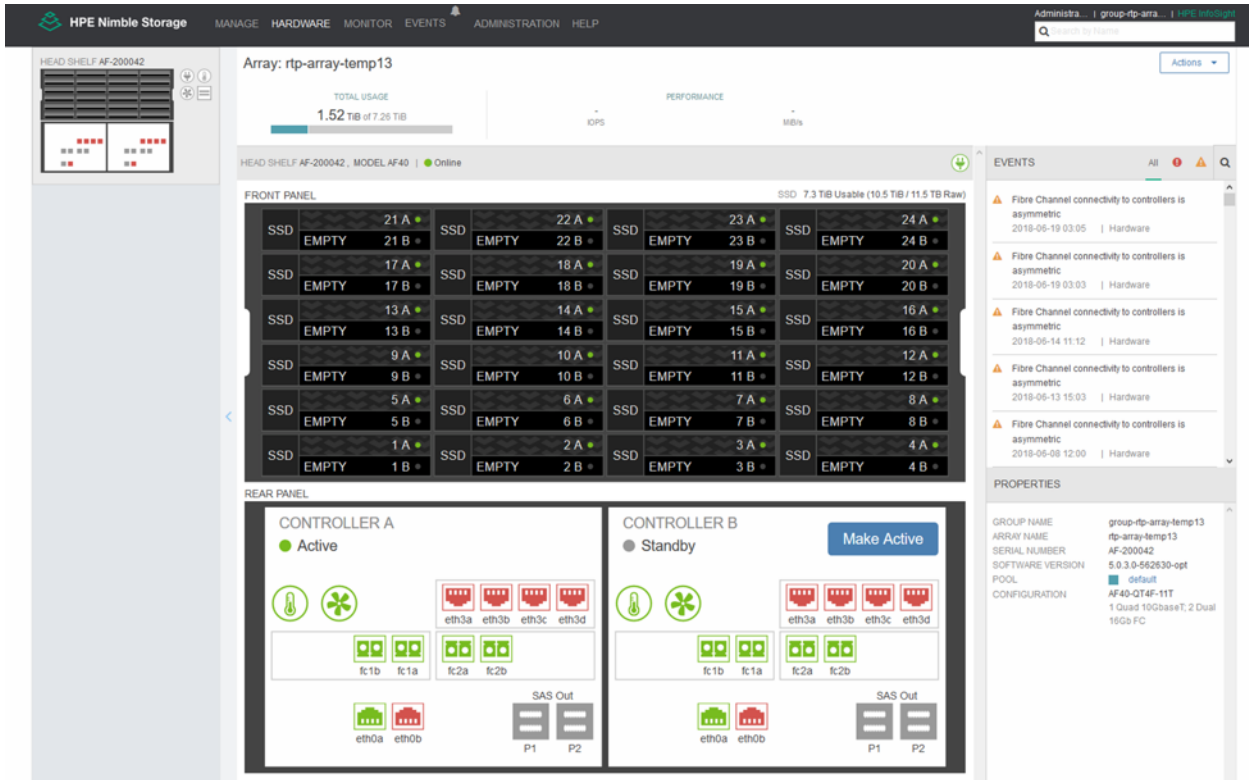
The RAID configuration for the HPE Nimble Storage array is by default a triple-parity configuration. No additional configuration for the storage is necessary.

For more information about the RAID configuration in HPE Nimble Storage arrays, see [HPE Nimble Storage All Flash Arrays](#).

## Array Overview

The HPE Nimble Storage AF-Series array is an all-flash array with dual active-standby storage controllers. It is configured with 24 480 GB SSD drives and has 2 1/10 Gbps onboard Ethernet adapters for management access. It has two quad-port 16 G FC target cards per controller. The storage array has a SAN connection through the Brocade switches. The array uses native FC.



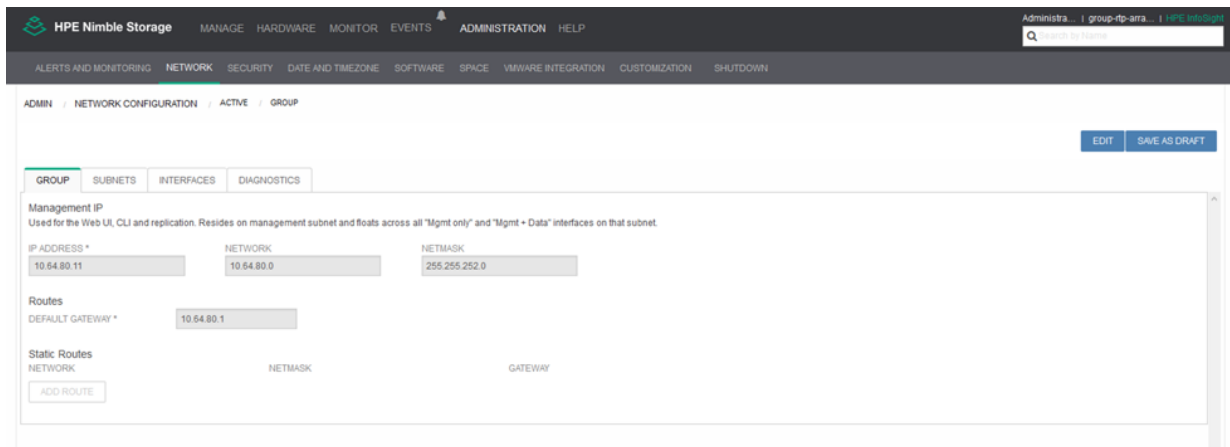
**Figure 2: HPE Nimble Storage array UI view**

## Configure the Network

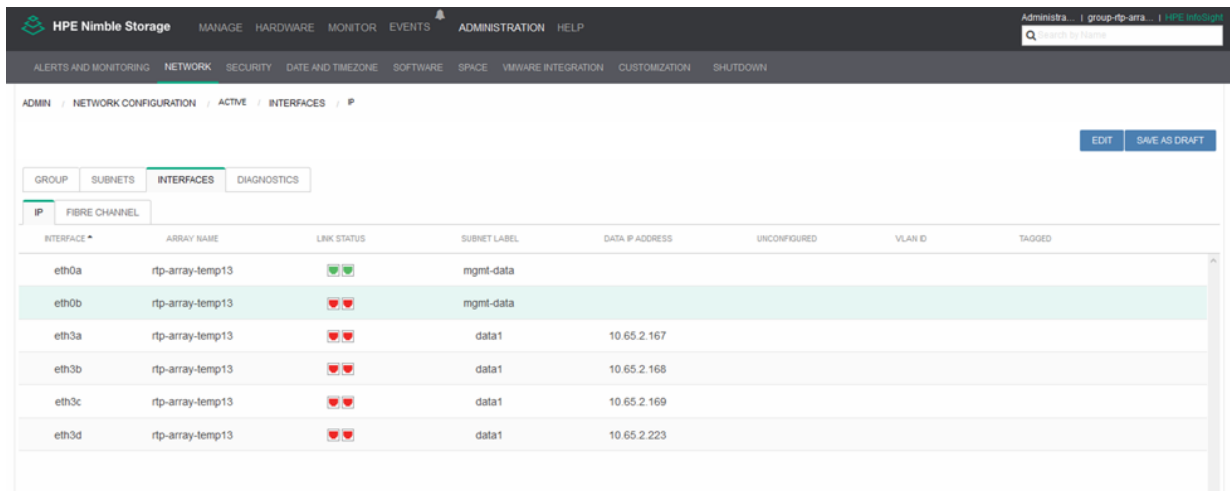
The HPE Nimble Storage array is configured with a single management Ethernet port for access, but you can configure additional ports to support replication. Replication was not configured or tested for the SAP HANA certification process.

### Procedure

- 1 During the array setup process, configure IP addresses for management access.



- 2 Use separate network interfaces for replication and management access.



ADMIN / NETWORK CONFIGURATION / ACTIVE / INTERFACES / P

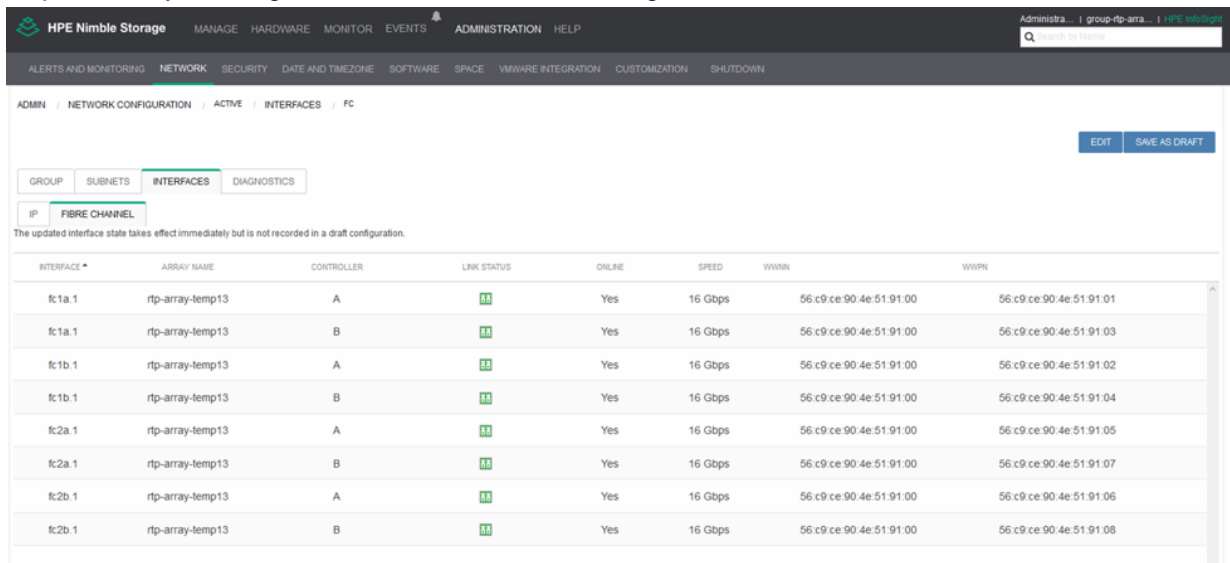
EDIT SAVE AS DRAFT

GROUP SUBNETS INTERFACES DIAGNOSTICS

IP FIBRE CHANNEL

INTERFACE *	ARRAY NAME	LINK STATUS	SUBNET LABEL	DATA IP ADDRESS	UNCONFIGURED	VLAN ID	TAGGED
eth0a	rtp-array-temp13	✓✓	mgmt-data				
eth0b	rtp-array-temp13	✓✓	mgmt-data				
eth3a	rtp-array-temp13	✓✓	data1	10.65.2.167			
eth3b	rtp-array-temp13	✓✓	data1	10.65.2.168			
eth3c	rtp-array-temp13	✓✓	data1	10.65.2.169			
eth3d	rtp-array-temp13	✓✓	data1	10.65.2.223			

Each controller in the array has two quad-port FC target adapters. The FC adapters support up to 16 Gbps, but they autonegotiate based on the fabric configuration.



ADMIN / NETWORK CONFIGURATION / ACTIVE / INTERFACES / FC

EDIT SAVE AS DRAFT

GROUP SUBNETS INTERFACES DIAGNOSTICS

IP FIBRE CHANNEL

The updated interface state takes effect immediately but is not recorded in a draft configuration.

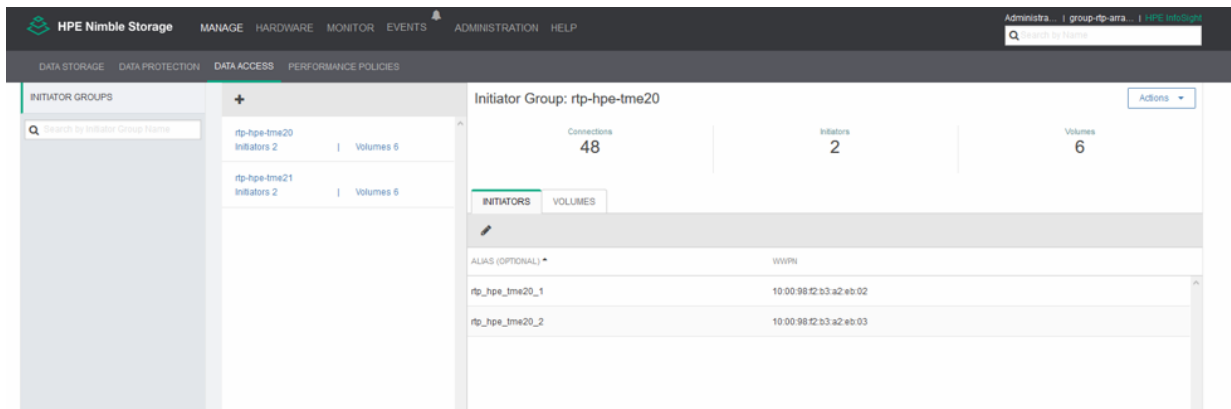
INTERFACE *	ARRAY NAME	CONTROLLER	LINK STATUS	ONLINE	SPEED	WWPN	WWPN
fc1a.1	rtp-array-temp13	A	✓	Yes	16 Gbps	56:c9:ce:90:4e:51:91:00	56:c9:ce:90:4e:51:91:01
fc1a.1	rtp-array-temp13	B	✓	Yes	16 Gbps	56:c9:ce:90:4e:51:91:00	56:c9:ce:90:4e:51:91:03
fc1b.1	rtp-array-temp13	A	✓	Yes	16 Gbps	56:c9:ce:90:4e:51:91:00	56:c9:ce:90:4e:51:91:02
fc1b.1	rtp-array-temp13	B	✓	Yes	16 Gbps	56:c9:ce:90:4e:51:91:00	56:c9:ce:90:4e:51:91:04
fc2a.1	rtp-array-temp13	A	✓	Yes	16 Gbps	56:c9:ce:90:4e:51:91:00	56:c9:ce:90:4e:51:91:05
fc2a.1	rtp-array-temp13	B	✓	Yes	16 Gbps	56:c9:ce:90:4e:51:91:00	56:c9:ce:90:4e:51:91:07
fc2b.1	rtp-array-temp13	A	✓	Yes	16 Gbps	56:c9:ce:90:4e:51:91:00	56:c9:ce:90:4e:51:91:06
fc2b.1	rtp-array-temp13	B	✓	Yes	16 Gbps	56:c9:ce:90:4e:51:91:00	56:c9:ce:90:4e:51:91:08

## Configure Initiator Groups and WWPNs

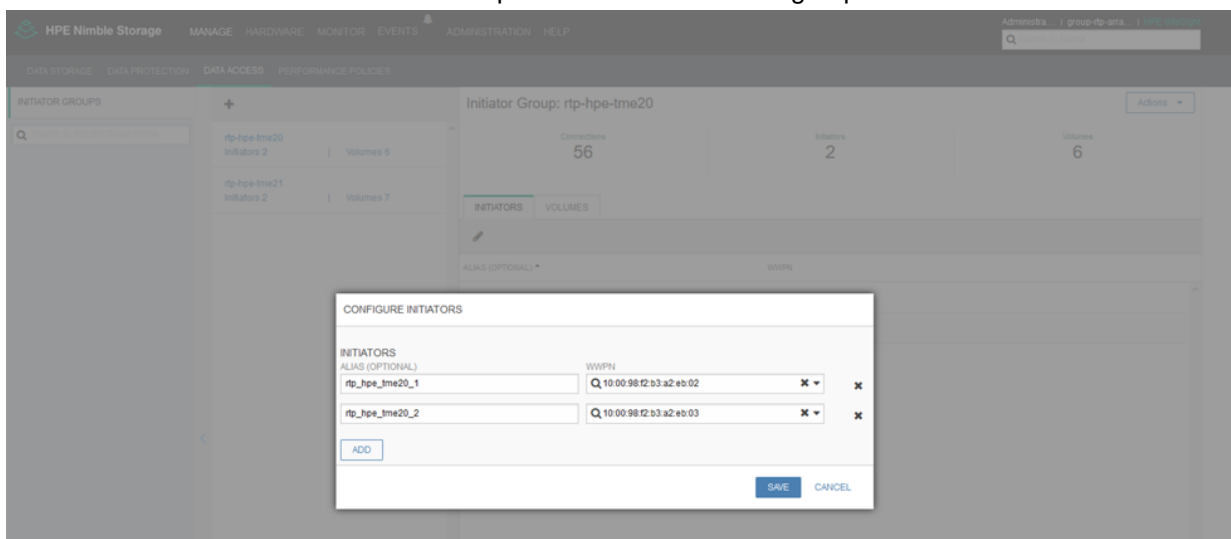
SAP HANA data LUNs and log LUNs are exported to all compute node initiators. Therefore, all nodes can access data LUNs and log LUNs in case of a failover.

### Procedure

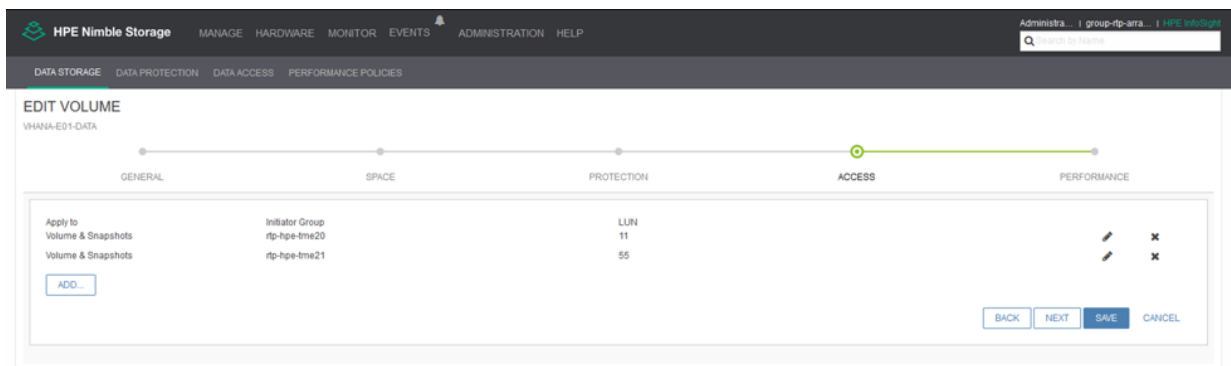
- 1 Create an initiator group for each host in the SAP HANA landscape.



## 2 Add all initiators on each SAP HANA compute node to the initiator group.



## 3 Add all initiators to the volumes.



## Create Volumes

Each HANA node has LUNs for HANA:

- Data (390 GB)

- Log (290 GB)

The size of HANA data and log LUNs depends on the size of server memory. Review SAP guidelines for storage provisioning.

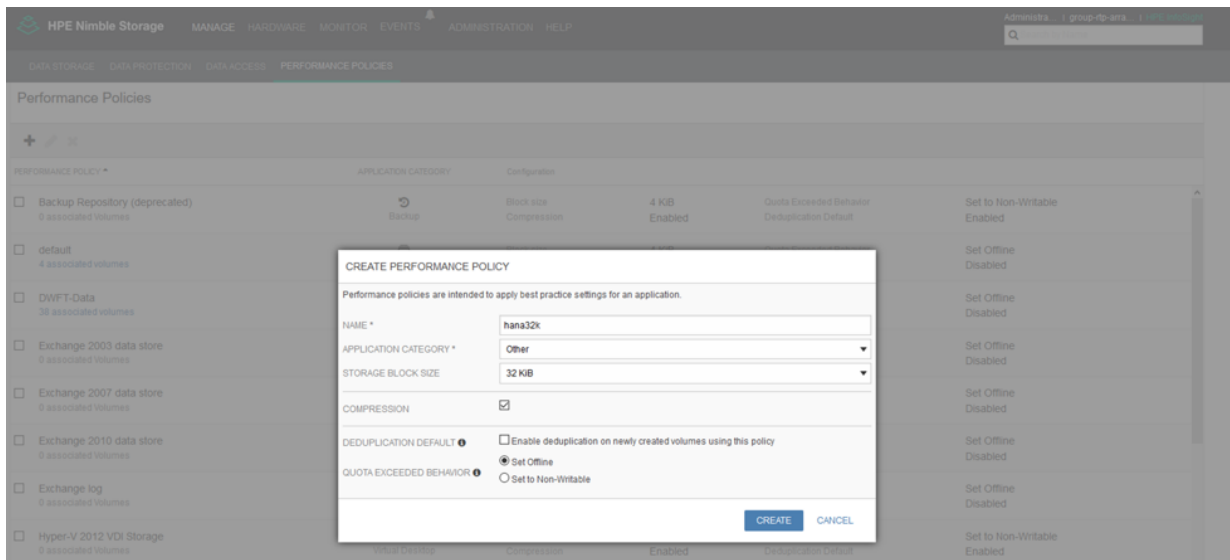
Data volumes are created by using a custom performance policy (`hana4k`) with a 4 KB block size. Log volumes are created by using a custom performance policy (`hana32k`) with a 32 KB block size.

In the ESXi environment, each server has a VMFS datastore for the OS, and all VMs share this datastore for the OS files.

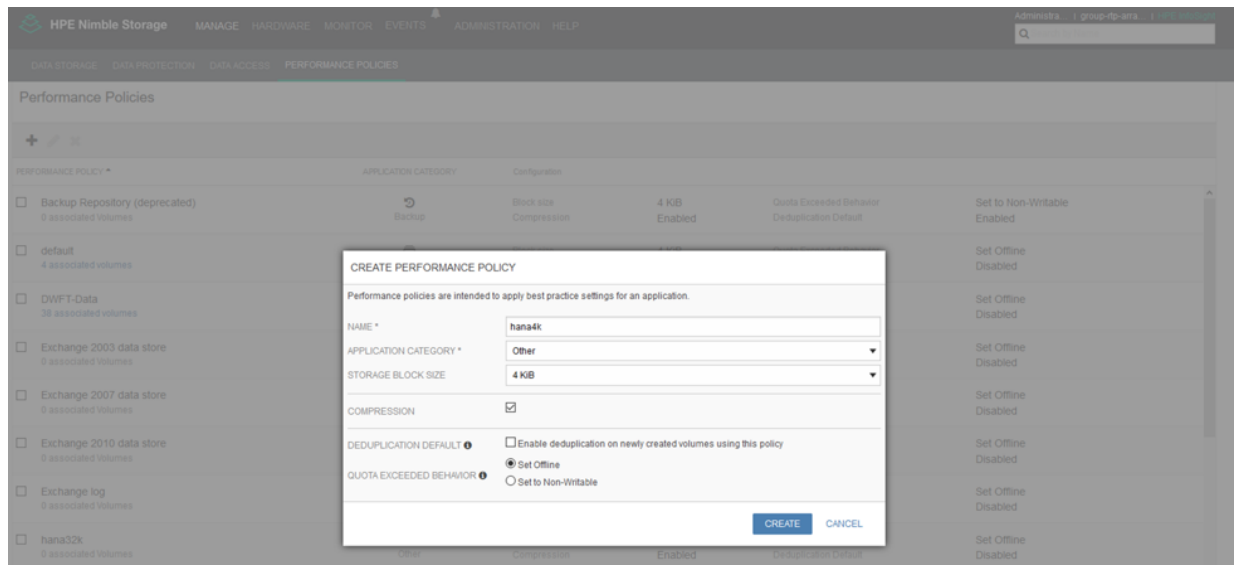
Individual VMFS datastores are created for data volumes and for log volumes. Data volumes are created by using the `hana4k` performance policy, and log volumes are created by using the `hana32k` performance policy. VMware best practices are followed to create individual disks in those datastores, including the use of separate paravirtual SCSI adapters for each disk.

## Procedure

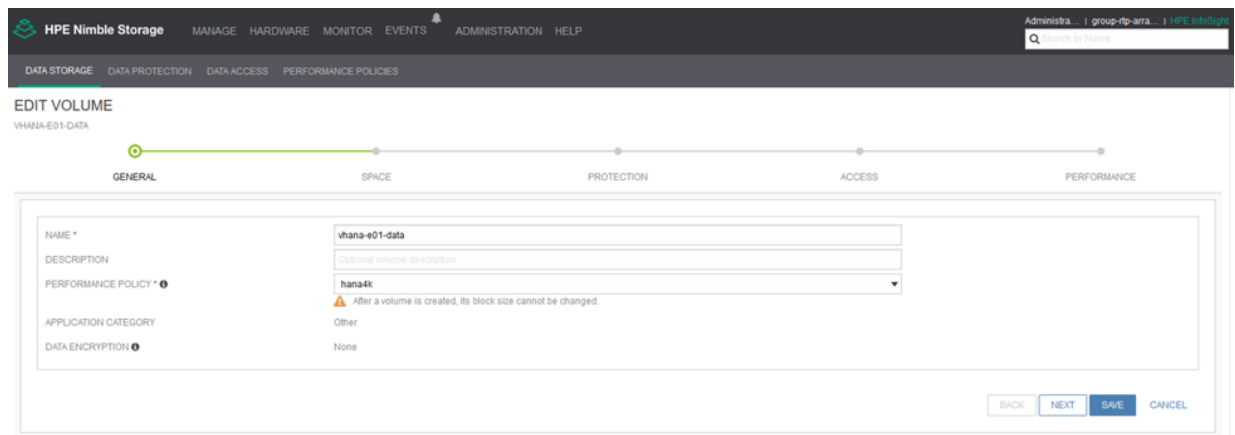
- 1 From the HPE Nimble Storage array UI, navigate to **Manage > Performance Policies** and choose **Add** to create the `hana32k` performance policy.



- 2 From **Manage > Performance Policies**, choose **Add** again to create the `hana4k` performance policy.



- 3 Use the hana4k and the hana32k performance policies to create the data and log volumes, respectively. Consult the SAP documentation to determine the correct size for each volume.



- 4 Add all initiator groups for the SAP HANA compute nodes to the volume access list.

The screenshot displays the HPE Nimble Storage management console. At the top, there's a navigation bar with tabs like MANAGE, HARDWARE, MONITOR, EVENTS, and ADMINISTRATION. Below this, a breadcrumb trail shows 'VOLUMES > default > vhana-e01-data'. The main area is divided into several sections: APPLICATION (Other), VOLUME USAGE (36.1 GiB of 97.7 GiB), PERFORMANCE (0 Total IOPS, 0 Total MB/s, 0 ms R Latency, 0.2 ms W Latency), and DATA PROTECTION. The 'DATA ACCESS' tab is selected, showing a table of connections. The table has columns for INITIATOR, INITIATOR WWPN, and TARGET WWPN. It lists four connections for two initiators. A right-hand panel shows the 'CONFIGURATION' section with details like PERFORMANCE POLICY (hana4k), THIRD PARTY MANAGEMENT (None), SERIAL NUMBER, TARGET WWNN, COMPRESSION (Enabled), DEDUPLICATION (Disabled), VOLUME COLLECTION (Not Protected), and REPLICATION PARTNER (N/A).

## Sizing Considerations

Extensive testing has established the number of SAP HANA nodes that each member of the all-flash array product family can support; the results of these tests are presented in the following table. Capacity for persistent storage is flexible, and it is based on the size of the SSDs that are installed in the array.

**Table 3: Number of SAP HANA nodes supported by each HPE Nimble Storage array model**

All-Flash Array Model	AF40	AF60	AF80
Node count	8	12	16

It is possible to combine competing workloads on an HPE Nimble Storage array that is used for SAP HANA. However, you must be cautious and plan and size your storage capacity carefully to ensure that production SAP HANA systems work as expected. HPE Nimble Storage arrays do provide additional capacity and throughput options to serve other applications, depending on the number of production SAP HANA nodes that are installed.

SAP Business Suite and other applications have workload characteristics that can affect the performance of SAP HANA in production environments. For this reason, you cannot predict which or how many applications can be installed on the same storage array without compromising the performance of SAP HANA. If you plan to deploy other SAP applications with SAP HANA, consult the HPE Nimble Storage team to determine the correct configuration and sizing for your storage.

## Operating System

For the tests, the SUSE Linux OS was installed on the servers of the SAP HANA deployment. If you are implementing SAP HANA in a virtual environment with HPE Nimble Storage arrays, you must install the HPE Nimble Storage Connection Manager (NCM) on the ESXi host to support the virtual environment.

NCM makes changes to the ESXi environment to optimize integration with HPE Nimble Storage arrays and provides the plugin that enables you to manage VMware datastores on the array directly from VMware vCenter. You must connect all datastores to the ESXi host by using the FC protocol. To install NCM, follow the instructions in the HPE Nimble Storage [VMware Integration Guide](#).

## Install OS

For installation of both the physical servers and the virtual servers, follow the [SAP note 2205917](#) from the SAP Service Marketplace.

## Configure OS

On both physical servers and VMs, modifications were made to the OS to support the SAP HANA performance requirements. This was done by using a custom udev rules file called `99-nimble-tune.rules` placed in the `/etc/udev/rules.d` directory. The file was adjusted accordingly based on the physical or virtual configuration.

```
##
ACTION!="add|change", GOTO="nimble_tuning_end"
SUBSYSTEM!="block", GOTO="nimble_tuning_end"
KERNEL!="sd*|dm-*", GOTO="nimble_tuning_end"
KERNEL=="dm-*", ENV{DM_UUID}!="*6c9ce9*", GOTO="nimble_tuning_end"
ENV{DEVTYPE}=="partition", GOTO="nimble_tuning_end"

#Comment the below line to enable these rules for devices attached to VM as
VMDK disks
#KERNEL=="sd*", ENV{ID_SERIAL}!="*6c9ce9*", GOTO="nimble_tuning_end"
#Uncomment the below line to enable these rules for devices attached to VM as
VMDK disks
KERNEL=="sd*", ENV{ID_SERIAL}!="*6c9ce9*", ENV{ID_VENDOR}!="VMware",
GOTO="nimble_tuning_end"

# Below rules are to assist in easily setting block device configurations
# suiting user application needs.
# Please uncomment the lines beginning with ATTR to enable these rules
# and run "udevadm control --reload-rules" and "udevadm trigger" to apply for
all Nimble devices.

# set max sectors_kb to max hw sectors_kb.
ATTR{queue/max_sectors_kb}="4096"
# set nr_requests to 512.
ATTR{queue/nr_requests}="512"
# set scheduler to noop.
ATTR{queue/scheduler}="noop"
# disable add_random.
ATTR{queue/add_random}="0"
# disable rotational.
ATTR{queue/rotational}="0"
# set rq_affinity to 2.
ATTR{queue/rq_affinity}="2"
# set read ahead
ATTR{queue/read_ahead_kb}="4096"

LABEL="nimble_tuning_end"
```

On VMs, queue depth is controlled by the paravirtual SCSI adapter settings on the hypervisor. For detailed instructions on how to increase queue depth on VMs, see the [VMware KB article 2053145](#).

**Procedure**

- 1 To support maximum performance, update the LUN queue depth for devices on the physical hardware by adding the following line to the `/etc/modprobe.d/99-local.conf` file:

```
options lpfc lpfc_lun_queue_depth=256
```

- 2 On the ESXi host, increase the value of the `fnic_max_qdepth` parameter by using the following command:

```
# esxcli system module parameters set -p lpfc_lun_queue_depth=128 -m lpfc
```

- 3 For the VMs, create a file called `pvscsi.conf` with the desired queue depth settings in the `/etc/modprobe.d` directory.

```
# cat /etc/modprobe.d/pvscsi.conf
options vmw_pvscsi cmd_per_lun=128 ring_pages=32
```

- 4 After you create the file, run **vmware-config-tools.pl** for the VM settings to take effect.

**Configure Multipathing**

To configure multipathing support for the HPE Nimble Storage array, create the `/etc/multipath.conf` file in the following way:

```
defaults {
    user_friendly_names yes
}
blacklist {
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
    devnode "^hd[a-z] [[0-9]*]"
    device {
        vendor "*"
        product "*"
    }
}
blacklist_exceptions {
    device {
        vendor "Nimble"
        product "Server"
    }
}
devices {
    device {
        vendor "Nimble"
        product "Server"
        no_path_retry 20
        fast_io_fail_tmo 5
        rr_weight priorities
        path_grouping_policy group_by_prio
        rr_min_io_rq 10
        failback 10
        path_selector "round-robin 0"
        path_checker "tur"
        hardware_handler "1 alua"
        prio "alua"
    }
}
```



## File Systems for SAP HANA

Every SAP HANA node in a scale-out environment must have file systems for the data, the logs, and the shared storage directory for the SAP HANA instance.

### Create the Data and Log File Systems

You must create file systems for data and logs on the SAP HANA nodes. The data and log file systems are local file systems that are connected through the FC protocol. For a virtual environment, the file systems are disks stored in a datastore that is connected to the ESXi host through the FC protocol.

#### Procedure

- 1 Create mount points for SAP HANA `/hana/data` and `/hana/log` on the local host.
- 2 Create the physical volumes by using the LUNs that you mapped to the host.
- 3 Use default options to create volume groups and logical volumes.
- 4 Format the logical volumes with XFS.
- 5 Add an entry in the `/etc/fstab` file on each host so that the file systems are mounted automatically.
- 6 Mount the file systems.

Example of `/etc/fstab` entries for the data and log file systems:

```
/dev/vghana01data/lvhana01data /hana/data/NS1 xfs defaults 0 0
/dev/vghana01log/lvhana01log /hana/log/NS1 xfs defaults 0 0
```

### Create the Shared File System

You must create a file system for the shared storage directory on the SAP HANA nodes. The shared file system should be mounted to each host by using NFS. To connect the shared file system with the array, HPE recommends that you use a Linux-based HA cluster (for example, two physical or virtual compute nodes) to create an NFS server.

#### Procedure

- 1 Create a mount point for `/hana/shared` on the NFS host.
- 2 Create physical volumes by using the LUNs that you mapped to the host.
- 3 Use default options to create the volume groups and logical volumes.
- 4 Format the logical volumes with XFS.
- 5 Add an entry in the `/etc/fstab` file on the NFS host so that the file system is mounted automatically.
- 6 Mount the file system.

Example of `/etc/fstab` entries on the NFS server:

```
/dev/vgshared/lvshared /hana/shared xfs defaults 0 0
```

### Mount the Shared File System on the SAP HANA Nodes

You must export the shared file system from the NFS server and mount it on the SAP HANA nodes.

#### Procedure

- 1 Export the shared file system to the SAP HANA compute nodes from the NFS server.
- 2 Create a mount point on each SAP HANA node for the shared file system.
- 3 Create an entry in the `/etc/fstab` file so that the file system is mounted automatically.
- 4 Mount the shared file system on the SAP HANA compute nodes.

Example of `/etc/fstab` entries for the shared file system:

```
hana-appe01-sto:/hana/shared /hana/shared nfs
rw,bg,vers=3,hard,rsiz=65536,wsiz=65536,nointr,actimeo=0,nolock 0 0
```

## Configure the SAP HANA global.ini File

To support the SAP HANA host autofailover capability, configure the `global.ini` file. You can specify the storage by using a worldwide identifier (WWID) or a logical volume manager (LVM) descriptor. For the SAP HANA certification, the LVM configuration was used.

To install and configure your system, follow the guidelines in the *SAP HANA Fibre Channel Storage Connector Admin Guide*. The latest version of the guide and the SAP HANA storage white paper are attached to SAP Note 1900823 on the [SAP Support Portal](#).

The following `global.ini` file applies to the environment used for the SAP HANA certification.

Example for 2+1 configuration with three nodes:

```
[communication]
internal_network = 192.168.1/20
listeninterface = .internal

[internal_hostname_resolution]
192.168.1.86 = hana01
192.168.1.87 = hana02
192.168.1.88 = hana03

[persistence]
log_mode = overwrite
basepath_logbackup = /backup/log
basepath_databackup = /backup/data
enable_auto_log_backup = no
log_backup_timeout_s = 7200
basepath_datavolumes = /hana/data
basepath_logvolumes = /hana/log

[storage]
ha_provider = hdb_ha.fcClientLVM
partition_*_*_mountoptions = -t xfs
partition_*_*_prtype = 5
partition_1_data_lvmname = vghana01data-lvhana01data
partition_1_log_lvmname = vghana01data-lvhana01data
partition_2_data_lvmname = vghana02data-lvhana02data
partition_2_log_lvmname = vghana02data-lvhana02data

[table_placement]
same_num_partitions = true
method = 2
prefix = /
bw_schema = SAPABAP1,SAPABAP1SHD

[trace]
ha_fcclientlvm = info
```

## Test Performance and Verify Parameter Configuration

For the performance tests to establish the maximum number of SAP HANA nodes that can be supported by a single storage array, SAP provided the test script, `fsperf_mult_106.sh`. The Hardware Certification

Check Tool version used for testing was 2.00.012.05.1524717104. The following table compares the default values of the SAP HANA **hdb** parameters to the tested values.

**Table 4: SAP HANA hdb parameters**

Parameter	hdbparam Default Value	Tested Value
fileio.async_read_submit	on	on
fileio.async_write_submit_blocks	all	all
fileio.async_write_submit_active	auto	auto
fileio.max_submit_batch_size	64	64
fileio.min_submit_batch_size	16	16
fileio.max_parallel_io_requests	64	64 (data) 128 (log)
fileio.size_kernel_io_queue	512	512
fileio.num_completion_queues	1	1
fileio.num_submit_queues	1	1

To confirm that your infrastructure complies with SAP HANA performance requirements, verify the parameter values and, if necessary, adjust them to the tested values.

Example:

```
hdbparam --paramset fileio  
[/hana/log/NS1/mnt00001/hdb00001/].max_parallel_io_requests=64  
hdbparam --paramset fileio  
[/hana/data/NS1/mnt00001/hdb00001/].max_parallel_io_requests=128
```

You can also set the values by using only the usage type and not the full path.

Example:

```
hdbparam --paramset fileio [LOG]. max_parallel_io_requests=64  
hdbparam --paramset fileio [DATA]. max_parallel_io_requests=128
```

For more information about the correct configuration of these parameters, see SAP Note 2267798 on the [SAP Support Portal](#).

## Version History

Version	Release Date	Description
2.0	June 2018	Updates to support new AF-Series models
1.2	June 2017	Addition of AF1000 as an SAP HANA certified array and branding updates
1.1	March 2017	Updates to terminology, the all-flash array tech specs, and the list of arrays certified by SAP
1.0	September 2016	Initial release