



Technical Report

SAP HANA on NetApp E-Series with Fibre Channel Protocol

Configuration Guide

Adolf Hohl, Bernd Herth, NetApp
July 2015 | TR-4296

Abstract

This document describes how to deploy HANA on NetApp® E-Series E5600 storage controllers using the Fibre Channel Protocol. It describes configurations to consolidate multiple independent HANA instances as well as the deployment of HANA systems with multiple nodes.

TABLE OF CONTENTS

1	Introduction	4
1.1	SAP HANA Tailored Datacenter Integration (TDI)	4
2	Architecture	5
2.1	SAP HANA Backup	6
2.2	Storage for Single-Node vs. Multinode HANA	7
3	Storage Sizing	8
3.1	Performance Considerations	8
3.2	Capacity Considerations	8
4	Infrastructure Setup and Configuration	9
4.1	Fibre Channel Setup	9
4.2	Network Setup	10
4.3	Storage System Setup	11
4.4	SAP HANA Storage Connector API	16
4.5	HANA Compute Node Setup	16
4.6	HANA Installation	21
5	Examples for a Typical Setup	22
5.1	Example 2+1 Node HANA Scale-Out	22
	References	26
	Version History	26

LIST OF TABLES

Table 1)	HANA single-node system storage requirements	8
Table 2)	HANA scale-out 3+1 nodes block storage requirements	9
Table 3)	HANA scale-out 3+1 nodes shared storage requirements	9
Table 4)	Volume configuration	24

LIST OF FIGURES

Figure 1)	SAP HANA TDI	4
Figure 2)	Example setup: three single-node HANA systems	5
Figure 3)	Scaling options	6
Figure 4)	Backup architecture from E-Series to NetApp FAS	7
Figure 5)	Fibre Channel connection	10
Figure 6)	Typical networking connectivity	11

Figure 7) HANA storage layout.....	12
Figure 8) Preferred owner.	13
Figure 9) Host mapping.....	13
Figure 10) Host definition wizard.....	14
Figure 11) Host definition: add WWNs.	15
Figure 12) Define host groups.	16
Figure 13) Identify WWN for each LUN.	19
Figure 14) Disk pool layout.....	23
Figure 15) DDP and volume layout.	24

1 Introduction

The NetApp E-Series and EF flash array product family has been certified for SAP HANA in the Tailored Datacenter Integration (TDI) program. This enterprise storage platform is characterized by the NetApp E-Series storage platform running the SANtricity® storage operating system. This certified setup utilizes Fibre Channel Protocol (FCP) to connect the storage to SAP HANA compute nodes. The file system used here is XFS.

The certification is valid for the following models:

- E2700, E5400, E5500, E5600
- EF540, EF550, EF560

NetApp E-Series and EF flash array (in the following named E-Series) are certified for the SAP HANA TDI program, and thus the conditions of this program apply. This program gives customers greater choice when deploying HANA. However, it is the customer's responsibility to provide sufficient resources. To verify this, SAP provides a toolset, the SAP Hardware Configuration Check Tool (HWCCT).

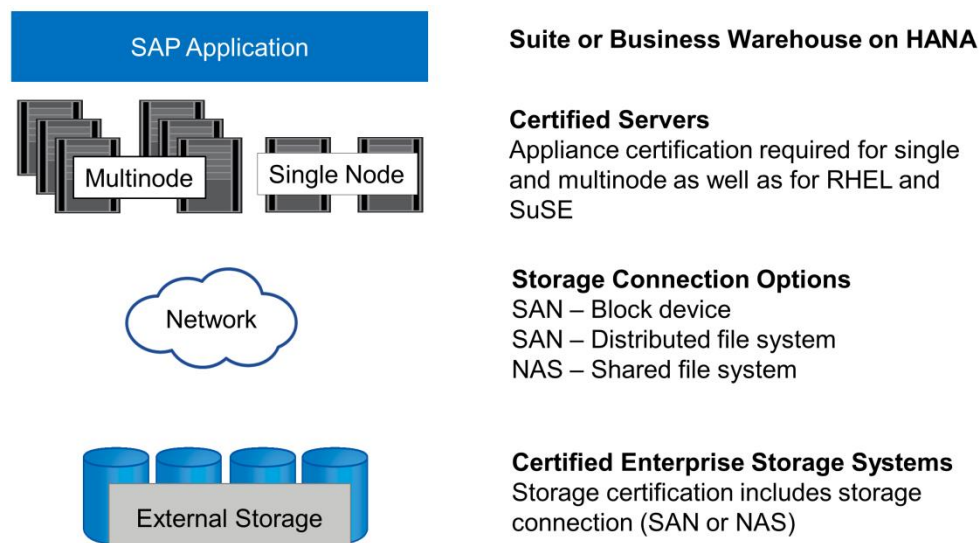
This document gives customers a basis to plan and implement HANA configurations using NetApp E-Series. The configurations shown here were validated using the SAP Hardware Certification toolset for use with different numbers of HANA nodes/systems. These configurations can be implemented as documented or act as a starting point for further customization. This document focuses on the consolidation of HANA nodes.

1.1 SAP HANA Tailored Datacenter Integration (TDI)

As part of the TDI process, the storage hardware vendor is required to test and certify the storage family for use with SAP HANA TDI implementations. As part of this certification process, the vendor must measure key performance indicators (KPIs) specific to SAP HANA by using the test tool provided by SAP. In addition, the vendor must validate an SAP HANA high-availability (HA) scale-out implementation, including failover tests.

The storage certification is combined with a storage connection, which can be either SAN or NAS, as shown in Figure 1. A certified storage solution can be used for SAP Business Suite on HANA or SAP Business Warehouse on HANA in single-node or multinode configurations. Any server that is certified for use with SAP HANA can be combined with the certified storage solution.

Figure 1) SAP HANA TDI.



For more information regarding the prerequisites and recommendations, refer to the following resources:

- [SAP HANA Tailored Datacenter Integration Frequently Asked Questions](#)
- [SAP HANA Tailored Datacenter Integration Overview Presentation](#)
- [SAP HANA Storage Requirements](#)

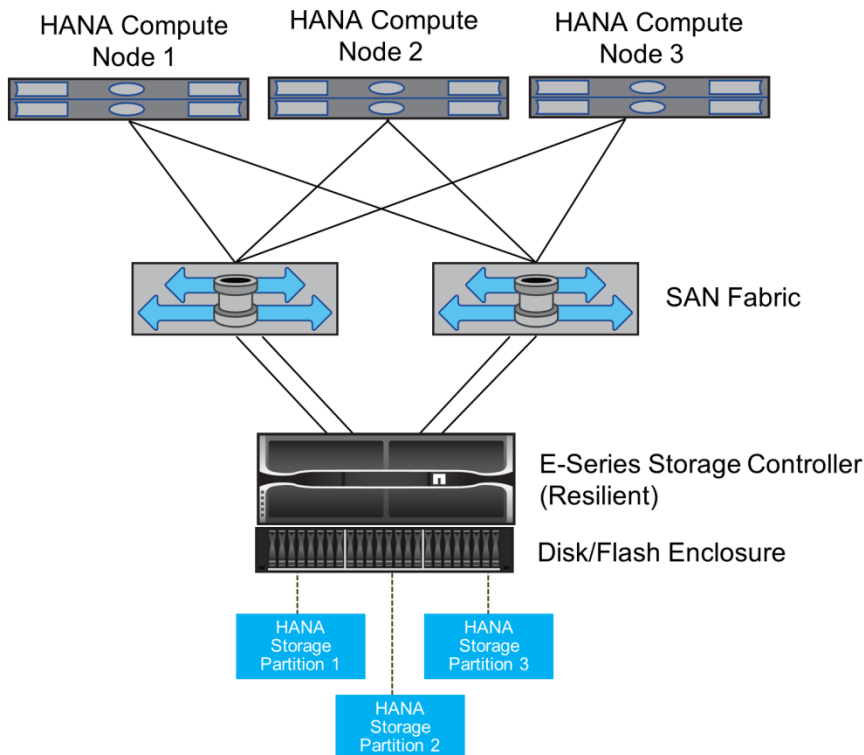
2 Architecture

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

Different models of the E-Series product family can be used at the storage layer. The maximum number of SAP HANA nodes attached to the storage is defined by the SAP HANA performance requirements. The number of disks required is determined by the capacity and performance requirements of the SAP HANA systems. The capacity requirements depend on the amount of SAP HANA nodes and the RAM size of each node. The storage partitions of the SAP HANA nodes are distributed across the storage controllers.

Note: The storage and SAP HANA node-to-node communication is separated, and only the storage communication is illustrated in Figure 2.

Figure 2) Example setup: three single-node HANA systems.



The architecture can be scaled in two dimensions:

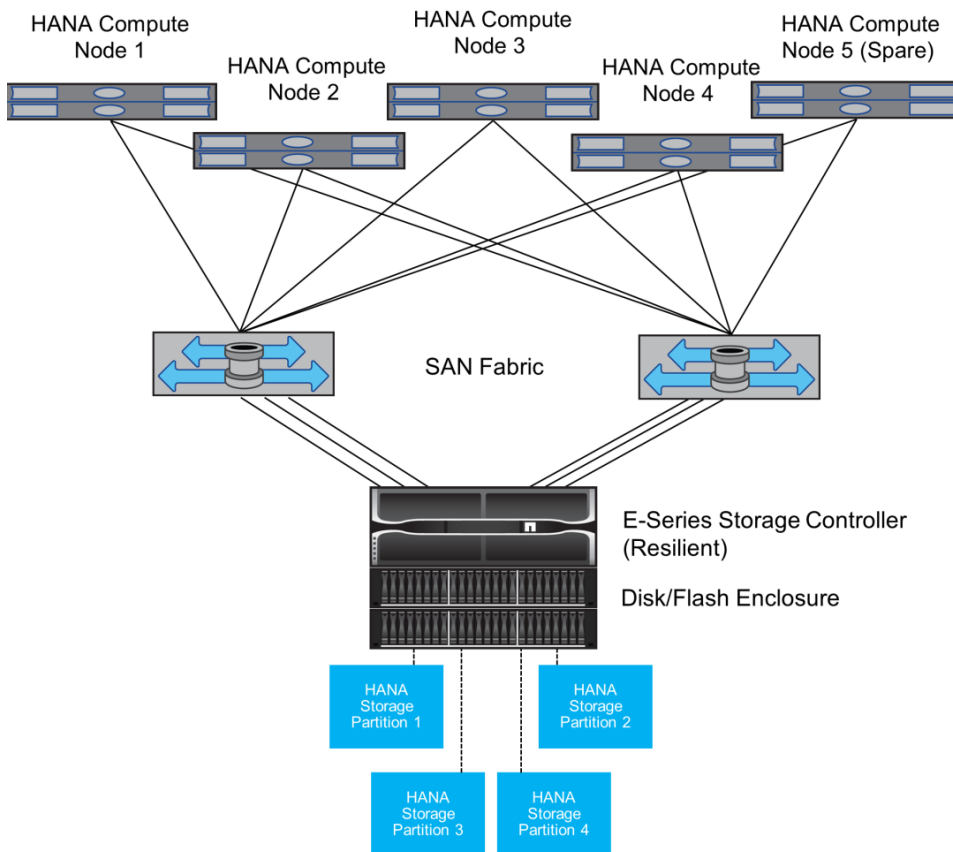
- By attaching additional SAP HANA nodes and disk capacity to the storage, as long as the storage controllers provide enough performance to meet the KPIs

- By adding more storage systems and disk capacity for the additional SAP HANA nodes

Figure 3 shows a scenario where:

- Additional SAP HANA nodes have been added to the fabric switches.
- An additional disk shelf has been added to the controller to provide required additional capacity.
- Additional Fibre Channel connections have been added between controller and switches to support the additional FC bandwidth and throughput requirements.

Figure 3) Scaling options.



2.1 SAP HANA Backup

The scope of this document does not include capacity required for SAP HANA backups. However, we would like to give a brief overview about the backup choices available. The SAP HANA database provides different methods for database backups:

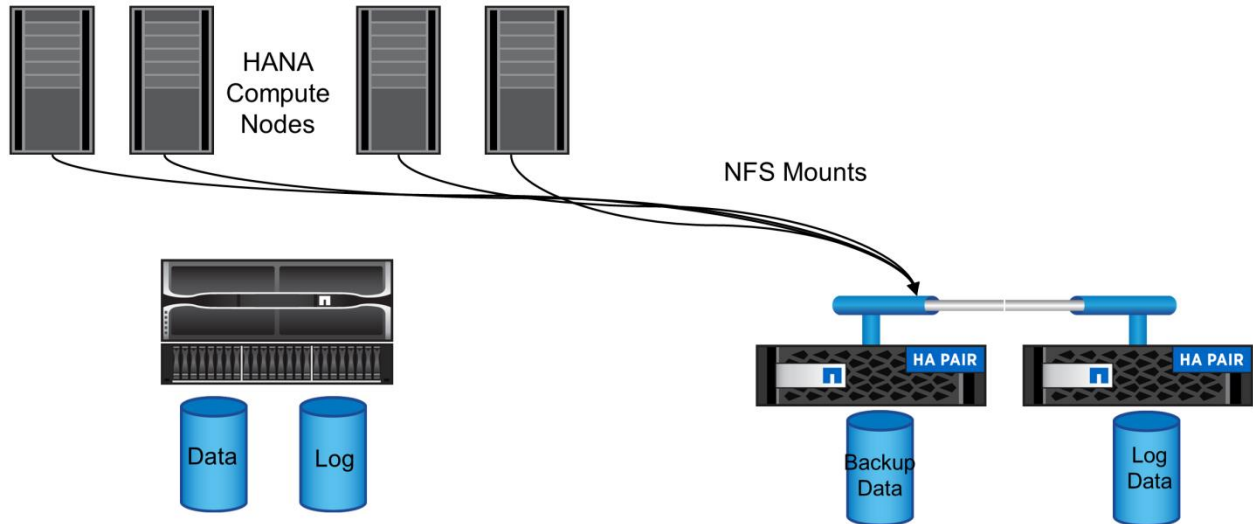
- Backup dump to disk (single-node SAP HANA instances) or shared file system (multinode SAP HANA systems)
- Backup piped to a backup server using backint-interface (single-node/multinode SAP HANA)
- Storage-based Snapshot[®] backups (single-node/multinode SAP HANA)

Note: When this document was published, a convenient end-to-end storage-based Snapshot backup was available for the FAS product line only.

In the scope of this document, we outline the backup dump to a shared file system residing on a NetApp FAS storage system.

Here we recommend mounting an NFS share from the NFS storage system directly to the SAP HANA nodes to store the SAP HANA backups. SAP HANA log backups are stored in the same manner. Refer to Figure 4 for an illustration.

Figure 4) Backup architecture from E-Series to NetApp FAS.



2.2 Storage for Single-Node vs. Multinode HANA

From a storage perspective the HANA system requires four storage containers:

- **SAP HANA storage partition.** Each server can own one storage partition at a time. Each SAP HANA storage partition typically consists of a separate DATA volume and LOG volume (two volumes altogether). At any given point in time only a single SAP HANA node must have exclusive write access to its storage partition.
 - **DATA volume.** Usually provided by means of an exported volume (NAS) or LUN/volume group (SAN).
 - **LOG volume.** Same as DATA volume.
- **SAP HANA shared.** An area that holds a file system for executables, traces, configuration data, and more. Shared access is required in multinode SAP HANA configurations. Typically provided by means of a shared NFS export.
- **SAP HANA backup.** A container for file-based backups and archive redo logs. A multinode SAP HANA configuration requires shared access to the member nodes. Typically provided by a shared NFS export.

For a single-node setup, all four areas may be implemented as LUNs served by the E-Series. The usage of an external storage for backup is required, as recommended previously.

This document focuses on SAP HANA storage partitions (DATA and LOG) provided by the NetApp E-Series. Consider a NetApp FAS storage system to provide sufficient NFS storage for SAP HANA shared and SAP HANA backup containers.

3 Storage Sizing

This section describes the performance and capacity considerations that should be observed to obtain an accurate storage sizing.

3.1 Performance Considerations

SAP has defined a static set of storage performance KPIs. These KPIs are valid for all SAP HANA environments independent of the memory size of the database nodes and the applications that use the SAP HANA database. These KPIs are valid for single-node, multinode, Business Suite on HANA, and Business Warehouse environments. Therefore, the current performance sizing approach depends only on the number of SAP HANA nodes that are attached to the storage system.

SAP delivers a performance test tool that must be used to validate the storage performance for the number of SAP HANA nodes attached to the storage.

The storage vendor defines the maximum number of SAP HANA nodes that can be attached to a specific storage model based on the KPIs from SAP.

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

3.2 Capacity Considerations

From a capacity requirement SAP HANA demands a storage size depending on the RAM of the SAP HANA database. The usable capacity of the dynamic disk pools (DDP) needs to be sized at least 2.5 times the RAM size of the related HANA nodes running on this storage. This includes all of the following:

- 0.5 × RAM for log, max 0.5TB for nodes with more than 1TB RAM
- 1 × RAM for data
- 1 × RAM for the `/hana/shared` file system

Based on the SAP HANA system configuration as single node or scale out, different storage setups are required.

HANA Single-Node System

For a HANA single-node system, Table 1 lists the minimum storage requirements based on the HANA compute nodes RAM size.

Table 1) HANA single-node system storage requirements.

HANA Partition	HANA Compute Nodes with 512GB RAM	HANA Compute Nodes with 1TB RAM	HANA Compute Nodes with 2TB RAM	HANA Compute Nodes with 4TB RAM
Data partition	512GB	1TB	2TB	4TB
Log partition	256GB	512GB	512GB	512GB
HANA shared	512GB	1TB	2TB	4TB
Total	1.25TB	2.5TB	4.5TB	8.5TB

Note: This table does not include storage for data and log backup.

HANA Shared

In the case of a single-node setup, HANA shared can be implemented as an XFS file system using an additional LUN served from the E-Series-based storage.

HANA Backup

The SAP HANA backup storage has to hold two backup areas:

- Space for one or more complete data backups. As a minimum the size of the data area, that is, 1xRAM, is required.
- Additional space for the log backups. The size depends on the data change rate of the SAP HANA database and should be larger 0.5xRAM.

In the case of HANA single-node systems, the HANA backup area can be implemented on an additional LUN served from the E-Series, but it is strongly recommended to store the file-based HANA backups on a separate shared backup server. Table 1 does not include any backup space.

HANA Scale-Out System

As mentioned previously, only the DATA and LOG partitions can be stored on the E-Series for an SAP HANA scale-out system, while SAP HANA shared and SAP HANA backup need to be provided by a shared file system such as NFS.

Table 2) HANA scale-out 3+1 nodes block storage requirements.

HANA Partition 3+1 Node Scale-Out	HANA Compute Nodes with 512GB RAM	HANA Compute Nodes with 1TB RAM	HANA Compute Nodes with 2TB RAM	HANA Compute Nodes with 4TB RAM
Data partition	1.5TB	3TB	6TB	12TB
Log partition	0.75TB	1.5TB	1.5TB	1.5TB
Total	2.25TB	4.5TB	7.5TB	13.5TB

Table 3) HANA scale-out 3+1 nodes shared storage requirements.

HANA NFS Storage	HANA Compute Nodes with 512GB RAM	HANA Compute Nodes with 1TB RAM	HANA Compute Nodes with 2TB RAM	HANA Compute Nodes with 4TB RAM
HANA shared	1.5TB	3TB	6TB	12TB

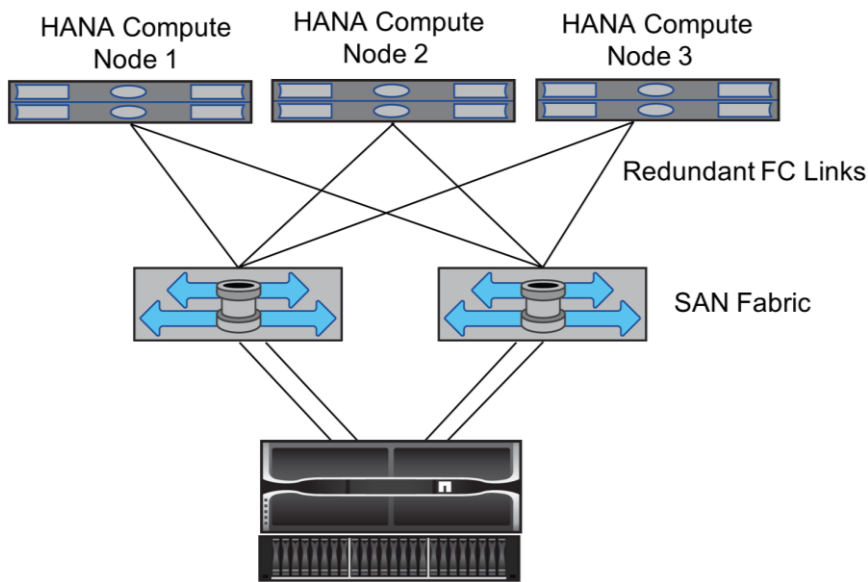
4 Infrastructure Setup and Configuration

4.1 Fibre Channel Setup

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

Figure 5 shows an example with three SAP HANA nodes attached to two storage controllers. Each SAP HANA node has two FCP ports connected to the redundant fabric. At the storage layer, four FCP ports are configured to provide the required throughput for each SAP HANA node.

Figure 5) Fibre Channel connection.



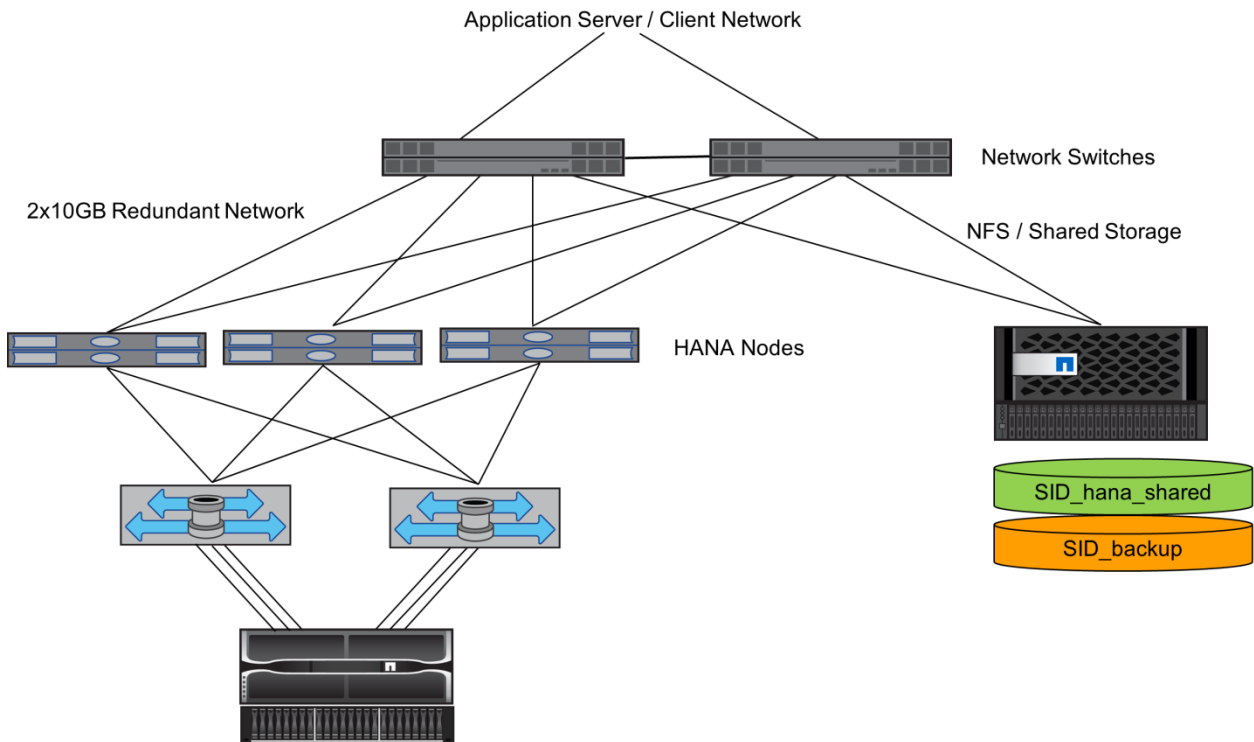
In addition to the zoning on the switch layer, it is required to map each LUN on the storage system to the hosts that should be entitled to connect to this LUN. Therefore, the zoning on the switch should be kept simple, that is, defining one zone set where all node HBAs see all controller HBAs.

4.2 Network Setup

In order to attach the SAP HANA shared file system and SAP HANA backup share using NFS, it is recommended to provide a dedicated storage network to separate the SAP HANA server communication from the storage network.

A typical redundant network setup is using two 10Gb connections from each SAP HANA node to a redundant switch, as illustrated in Figure 6.

Figure 6) Typical networking connectivity.



4.3 Storage System Setup

This section describes the configuration of the NetApp E-Series storage system. As a prerequisite, the primary installation and setup have been done according to the corresponding documentation. All further configurations should be applied with the SANtricity tool.

Dynamic Disk Pools


Multiple measurements with an E55xx/E56xx system with 24, 60, and 120 disks (each 900GB) have shown that best results in regard to the SAP HANA KPI have been reached using DDPs built with 15 disks (12 in the case of a 24-disk shelf).

LUN Creation and Assignment

A single volume/LUN is created out of a DDP. As a best practice, the naming convention for a LUN should include the SAP HANA SID and the storage partition. For example, SID_data, SID_log, SID1_data, and so on for different single-node systems.

Distributing the LUNs of one system to different DDPs is recommended. The following example shows a setup with two single-node HANA systems on a 24-disk E5524 system.

Figure 7) HANA storage layout.



In this example two DDPs have been created, and the data and log devices (LUNs) of the two single-node HANA systems have been created, each data and log LUN on a different DDP.

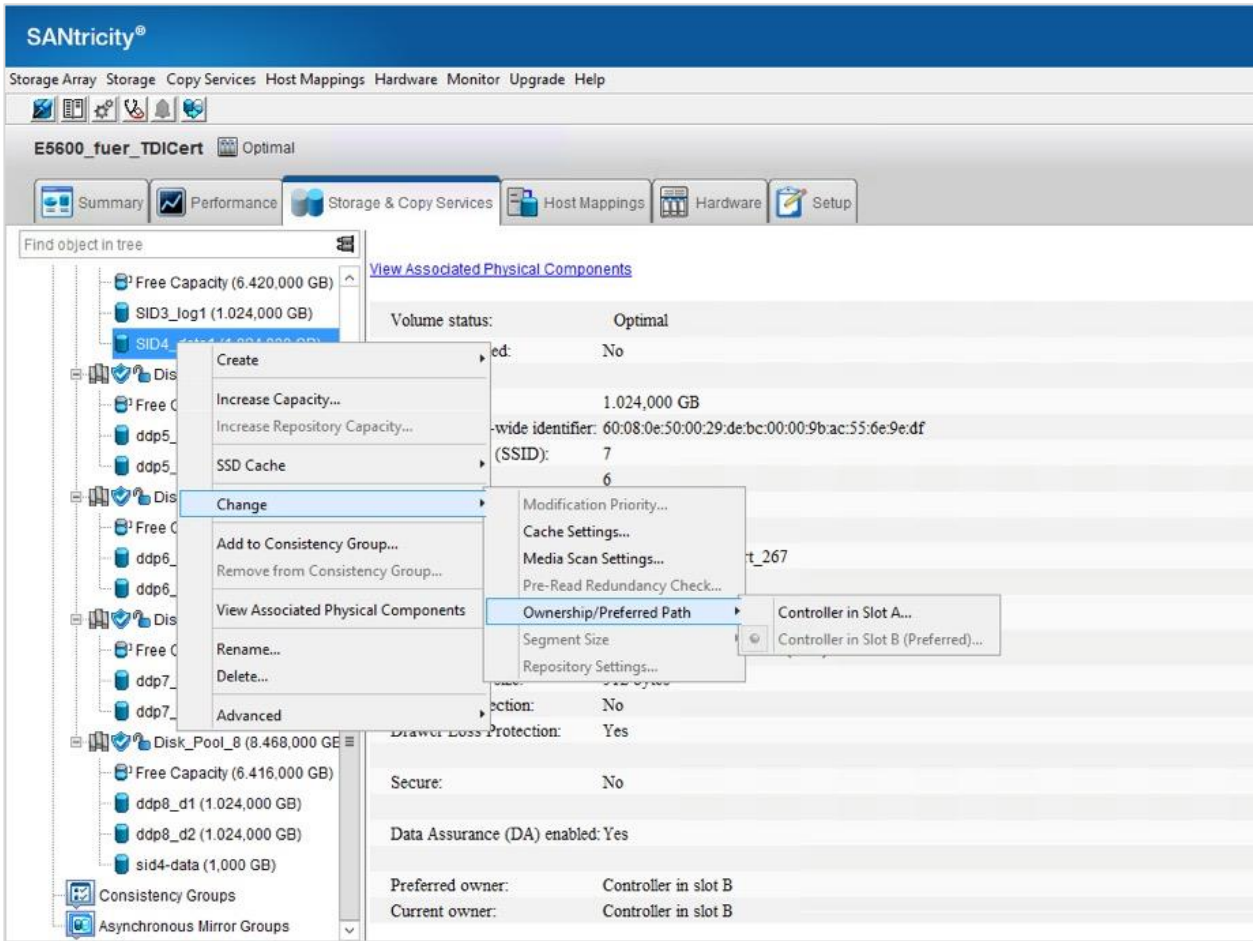
Note: In this example the HANA shared and HANA backup volumes are provisioned outside the E-Series as NFS shares and are not shown in Figure 7.

Controller Assignment

Each LUN within a DDP is accessible by both controllers. The definition of Ownership/Preferred Path is used to assign the preferred owner. This selects the controller that serves the LUN under normal operations. Distributing the ownership of the LUNs evenly between the two controllers is recommended to balance the load between both controllers.

Figure 8 shows how to use the SANtricity Storage Manager to set each LUN's preferred owner.

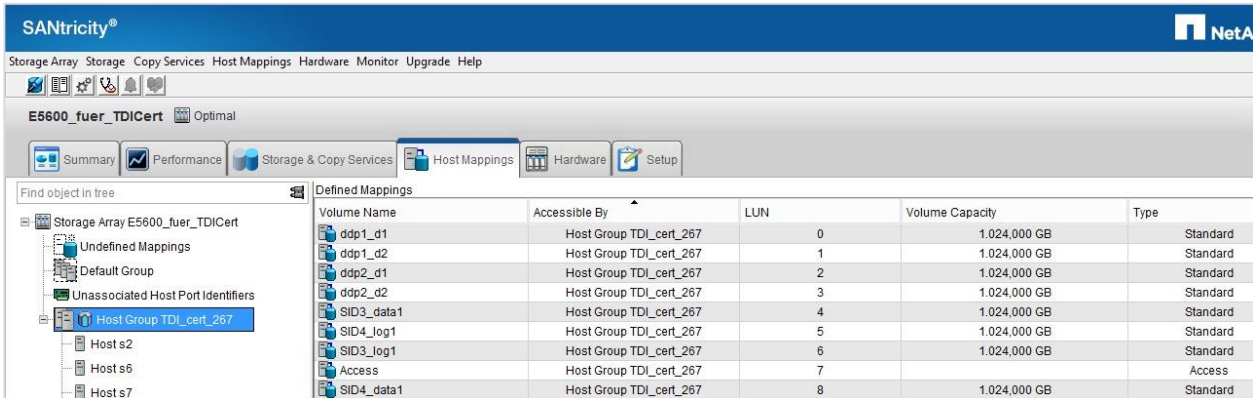
Figure 8) Preferred owner.



Host Mappings

To access a LUN from a host, a “host mapping” needs to be defined. It is a best practice to assign a LUN to all hosts of a HANA multinode configuration (or to the one node of a single-node HANA) with the help of “host groups” using the SANtricity Storage Manager.

Figure 9) Host mapping.



Host Mapping Workflow

The following steps are recommended to create a host mapping.

1. Identify the Linux HBA port names.

Log in to the Linux system and issue the command:

```
stlrx300s8-1:~ # systool -c fc_host -v|grep port_name
port_name           = "0x2100000e1e14ff10"
port_name           = "0x2100000e1e14ff11"
stlrx300s8-1:~ #
```

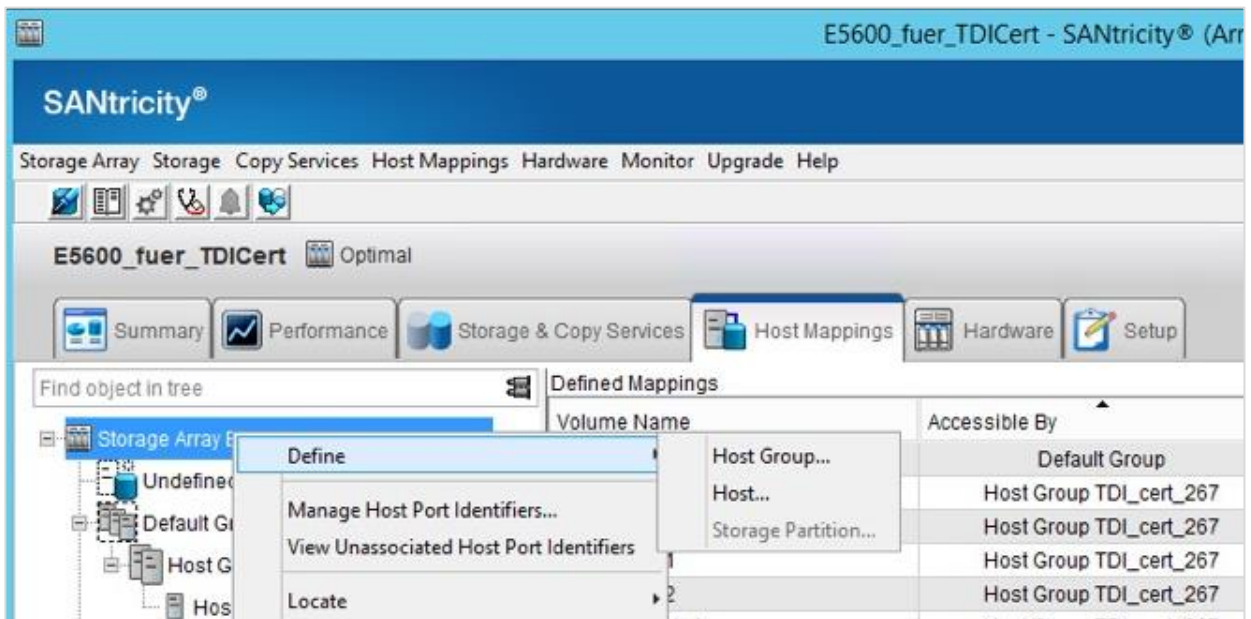
This lists the WWNs.

2. Create a host definition in SANtricity.

In SANtricity, define the host with its WWNs. To better identify this host, you may use descriptive names and aliases for the host HBAs.

Start the Host definition wizard in SANtricity from the Host Mappings tab by right-clicking the Storage Array and selecting Define → Host.

Figure 10) Host definition wizard.



This starts the wizard in which you can define the descriptive name for the host and adds the WWN for this host. In case SAN zoning has already been configured, the WWN might be visible at the E-Series and can be selected from a list of unassociated host port identifiers. Otherwise, you can define them manually, as shown in Figure 11. In Figure 11, Linux (DM-MP) is selected as the host type.

Note: The wizard allows you to add the newly defined host to a host group. The workflow describes this step separately.

Figure 11) Host definition: add WWNs.

The host communicates with the storage array through its host bus adapters (HBAs) or its iSCSI initiators where each physical port has a unique host port identifier. In this step, select or create an identifier, give it an alias or user label, then add it to the list to be associated with host bla.

[How do I match a host port identifier to a host?](#)

Choose a method for adding a host port identifier to a host:

☐ Add by selecting a known unassociated host port identifier

Known unassociated host port identifier:

- There are no known unassociated host port identifiers -

☒ Add by creating a new host port identifier:

New host port identifier (16 characters required):

210000314567e456

Alias (30 characters maximum):

HANA-1-hbab

Add ▼ Remove ▲

Host port identifiers to be associated with the host:

Host Port Identifier	Alias / User Label
21:00:00:31:45:67:e4:55	HANA-1-haba

< Back Next > Cancel Help

3. Define host groups.

Using host groups is optional, but is very helpful in administering LUN assignments for SAP HANA scale-out environments.

Start the Host definition wizard in SANtricity from the Host Mappings tab by right-clicking the Storage Array and selecting Define → Host Groups.

As shown in Figure 12, select a descriptive name for the host group. Add the hosts that have not already been assigned to a host group to the newly created host group.

Figure 12) Define host groups.

E5600_fuer_TDICert - Define Host Group

NetApp

A host group is a collection of hosts that share access to volumes within a storage partition on a storage array. You can create a new host group by providing a host group name and selecting the hosts that will be part of the host group.

Note: Use the Define Host wizard to create a new host that does not currently exist.

[Learn more about host groups](#)

Provide host group name

Enter new host group name (30 characters maximum):

HANA-SID-hosts

Add Hosts to new host group

Select hosts to add:

Name	Operating System

Add >

< Remove

Hosts in group:

Name	Operating System
HANA-node-1	Linux (DM-MP)

OK Cancel Help

4. Define LUN mapping.

Select the host group, right-click, and select Add LUN Mapping. This opens the Add LUN Mapping dialog box, in which all currently unmapped LUNs can be added to the host group.

4.4 SAP HANA Storage Connector API

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

In SAP HANA multinode configurations with NetApp storage, the standard storage connector delivered by SAP is used. The “SAP HANA Fibre Channel Storage Connector Admin Guide” can be found as an attachment to SAP note 1900823.

4.5 HANA Compute Node Setup

The HANA TDI certification requires using certified compute hardware only.

Operating System Version

The operating system used for this test configuration is SuSE SLES 11 SP3. At the time of writing this report, the alternative use of Red Hat RHEL 6.5 is permitted.

Operating System Configuration

The operating system needs to be configured according to the following SAP documents:

- [SAP HANA Master Guide](#)
- [SAP HANA Server Installation and Update Guide](#)

As the time of writing, the following SAP notes focus on operating system–specific configuration requirements:

- 1514967 SAP HANA: Central Note
- 1310037 SUSE LINUX Enterprise Server 11: Installation notes
- 1954788 SAP HANA DB: Recommended OS settings for SLES 11/SLES for SAP Applications 11 SP3
- 2009879 SAP HANA Guidelines for Red Hat Enterprise Linux (RHEL) Operating System
- 2013638 SAP HANA DB: Recommended OS settings for RHEL 6.5

Zoning

All nodes must have access to all LUNs. Zoning on the switch has to reflect this. In addition, the storage controller must be configured such that all volumes are visible to all required SAP HANA hosts (refer to section “Host Mapping Workflow”).

Multipathing

Multipathing has to be enabled to aggregate traffic and make sure of proper failover behavior. To enable multipathing, complete the following steps.

1. Enable the multipath daemon.

```
stlrx300s8-1:~ # chkconfig multipathd on
stlrx300s8-1:~ # chkconfig multipathd
multipathd on
```

2. Configure the multipath configuration file `/etc/multipath.conf`.

Prepare the `multipath.conf` file with the required E-Series specific settings.

These settings are within section `devices{}`.

```
...
devices {
    device {
        vendor "NETAPP"
        product "INF-01-00"
        path_grouping_policy "group_by_prio"
        path_selector "service-time 0"
        path_checker "rdac"
        features "2 pg_init_retries 50"
        hardware_handler "1 rdac"
        prio "rdac"
        failback "immediate"
        no_path_retry 30
        #polling_interval 5
    }
}
...
```

The device attribute `features 2 pg_init_retries 50` and `no_path_retry 30` are used in combination with the `global.ini` settings `partition_*_*__prtype = 5` (see next paragraph).

3. Run the Linux `rescan-scsi-bus.sh` command on each server to discover new LUNs. Alternatively reboot the server.
4. Restart the multipath daemon.

```
stlrx300s8-1:~ # service multipathd restart
Shutting down multipathd
Starting multipathd
```

done
done

5. Reload the LUNs by using the Linux command `multipath -r`.

```
stlrx300s8-1:~ # multipath -r
...
Jun 19 04:22:06 | 360030057013a20d01a68328209c16300: ignoring map
reload: 360080e500029debc00009bac556e9edf undef NETAPP,INF-01-00
size=1.0T features='3 pg_init_retries 50 queue_if_no_path' hwhandler='1 rdac' wp=undef
|+- policy='service-time 0' prio=6 status=undef
| |- 7:0:1:8 sdaa 65:160 active ready running
| |- 7:0:2:8 sdar 66:176 active ready running
| |- 8:0:1:8 sdcq 69:224 active ready running
| ` 8:0:2:8 sddh 70:240 active ready running
`+- policy='service-time 0' prio=1 status=undef
| |- 7:0:0:8 sdj 8:144 active ghost running
| |- 7:0:6:8 sdbi 67:192 active ghost running
| |- 8:0:0:8 sdbz 68:208 active ghost running
| ` 8:0:6:8 sddy 128:0 active ghost running
reload: 360080e500029e2d00000dcf2556e9ec8 undef NETAPP,INF-01-00
size=1.0T features='3 pg_init_retries 50 queue_if_no_path' hwhandler='1 rdac' wp=undef
|+- policy='service-time 0' prio=6 status=undef
| |- 7:0:0:9 sdk 8:160 active ready running
| |- 7:0:6:9 sdbj 67:208 active ready running
| |- 8:0:0:9 sdca 68:224 active ready running
| ` 8:0:6:9 sddz 128:16 active ready running
`+- policy='service-time 0' prio=1 status=undef
| |- 7:0:1:9 sdab 65:176 active ghost running
| |- 7:0:2:9 sdas 66:192 active ghost running
| |- 8:0:1:9 sdcr 69:240 active ghost running
| ` 8:0:2:9 sddi 71:0 active ghost running
```

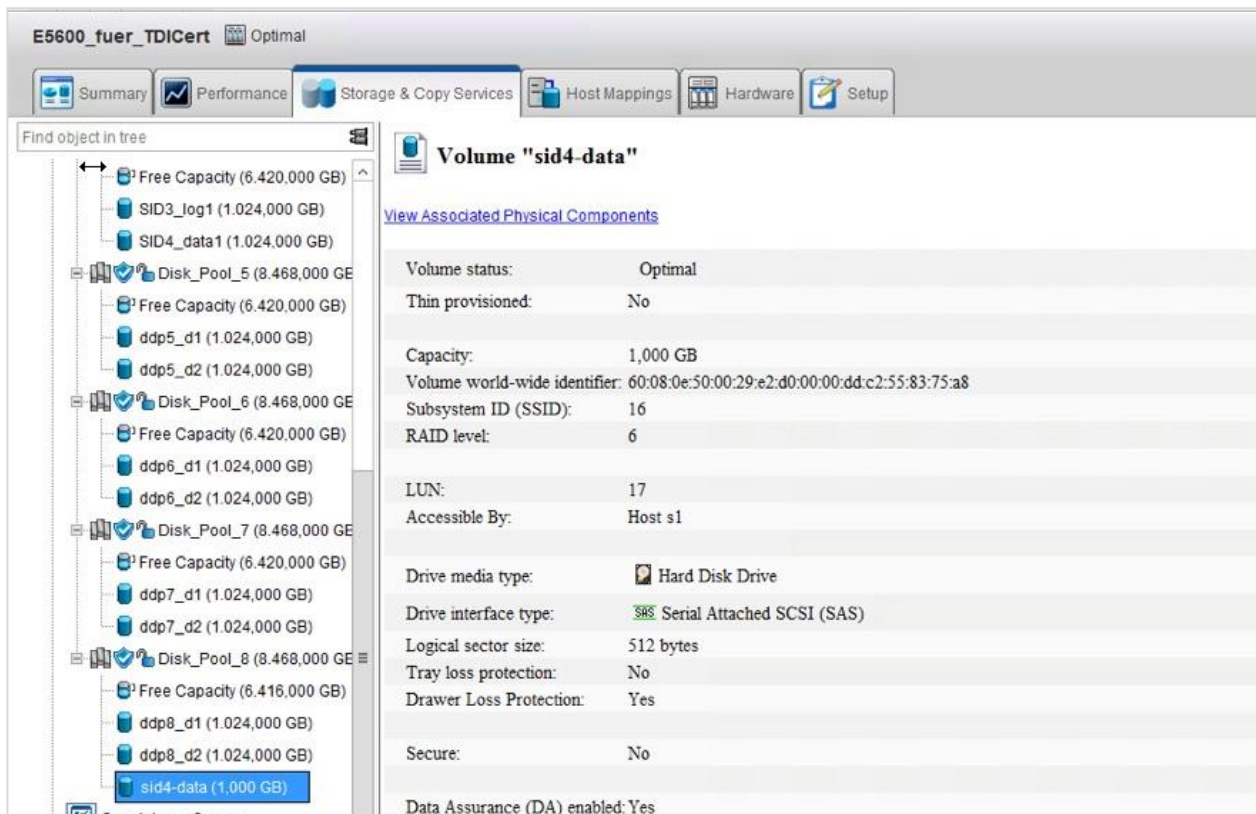
6. Identify all LUNs without the alias definition by running the Linux command `multipath -ll`.

A list of all loaded LUNs appears.

```
Stlrx300s8-1:~ # multipath -ll
...
...
360080e500029e2d00000ddc2558375a8 dm-16 NETAPP,INF-01-00
size=1.0G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='1
rdac' wp=rw
|+- policy='service-time 0' prio=6 status=active
| |- 7:0:0:17 sdei 128:160 active ready running
| |- 7:0:6:17 sdel 128:208 active ready running
| |- 8:0:0:17 sdem 128:224 active ready running
| ` 8:0:6:17 sdep 129:16 active ready running
`+- policy='service-time 0' prio=1 status=enabled
| |- 7:0:1:17 sdej 128:176 active ghost running
| |- 7:0:2:17 sdek 128:192 active ghost running
| |- 8:0:1:17 sden 128:240 active ghost running
| ` 8:0:2:17 sdeo 129:0 active ghost running
...
...
```

The new LUNs are listed with their WWN IDs. You may use SANtricity to match the new WWNs with the exported LUN's WWN. Select a LUN, and SANtricity displays its attribute on the right. The WWN can be found under "Volume world-wide identifier."

Figure 13) Identify WWN for each LUN.



Note: Linux prefixes “3” in front of the WWN.

7. Add an alias for each identified LUN to the /etc/multipath.conf.

```

multipaths {
...
multipath {
    wwid 360080e500029e2d00000dcee556e9de9
    alias FC1_data
}
multipath {
    wwid 360080e500029debc00009ba9556e9e36
    alias FC1_log
}
...
}

```

The section `multipaths {}` is used to define aliases for the WWID that simplify the configuration and handling of the SAP HANA configuration file `global.ini`. All LUNs that need to be used on this node should be defined.

It is recommended to use descriptive names including the SAP HANA SID and the storage partition type and number in case of a multinode SAP HANA system:

- FC1_data for the data partition of a single-node SAP HANA system with SID “FC1”
- FC1_data1, FC1_log1, FC1_data2,... for the partitions for a multimode SAP HANA system

In a HANA scale-out installation, the multipath configuration must include all data and log devices on each of the nodes that are part of the SAP HANA scale-out system.

Create File Systems

To create the XFS file system on each LUN belonging to the HANA system, complete the following step:

1. Take one of the following actions:

- For a single-node system, create the XFS file system on the data LUN, the log LUN, and the /hana/shared LUN.

```
stlrx300s8-4:/ # mkfs.xfs /dev/mapper/FC1_data
stlrx300s8-4:/ # mkfs.xfs /dev/mapper/FC1_log
stlrx300s8-4:/ # mkfs.xfs /dev/mapper/FC1_shared
```

- For a multinode system, create the XFS file system on all data and log LUNs.

```
stlrx300s8-4:/ # mkfs.xfs /dev/mapper/FC1_data1
stlrx300s8-4:/ # mkfs.xfs /dev/mapper/FC1_log1
stlrx300s8-4:/ # mkfs.xfs /dev/mapper/FC1_data2
stlrx300s8-4:/ # mkfs.xfs /dev/mapper/FC1_log2
```

Create Mount Points

To create the required mount point directories, complete the following step:

1. Take one of the following actions:

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

```
stlrx300s8-5:/ # mkdir -p /hana/data/FC1/mnt00001
stlrx300s8-5:/ # mkdir -p /hana/log/FC1/mnt00001
stlrx300s8-5:/ # mkdir -p /hana/shared
```

```
stlrx300s8-5:/ # chmod -R 777 /hana/log/FC1
stlrx300s8-5:/ # chmod -R 777 /hana/data/FC1
stlrx300s8-5:/ # chmod 777 /hana/shared
```

- For a multinode system, set the permissions and create mount points on all worker and standby nodes.

Note: The example commands show a 2+1 multinode HANA system.

```
stlrx300s8-5:/ # mkdir -p /hana/data/FC1/mnt00001
stlrx300s8-5:/ # mkdir -p /hana/log/FC1/mnt00001
stlrx300s8-5:/ # mkdir -p /hana/data/FC1/mnt00002
stlrx300s8-5:/ # mkdir -p /hana/log/FC1/mnt00002
stlrx300s8-5:/ # mkdir -p /hana/shared
```

```
stlrx300s8-5:/ # chmod -R 777 /hana/log/FC1
stlrx300s8-5:/ # chmod -R 777 /hana/data/FC1
stlrx300s8-5:/ # chmod 777 /hana/shared
```

Mount File Systems

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

1. Take one of the following actions:

- For a single-node system, add the required file systems to the /etc/fstab configuration file.

```
stlrx300s8-5:/ # cat /etc/fstab
/dev/mapper/FC1_shared /hana/shared/FC1 xfs default 0 0
/dev/mapper/FC1_log /hana/log/FC1/mnt00001 xfs default 0 0
/dev/mapper/FC1_data /hana/data/FC1/mnt00001 xfs default 0 0
```

- For a multinode system, add the /hana/shared file system to the /etc/fstab configuration file.

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

```
stlrx300s8-5:/ # cat /etc/fstab
192.168.173.101:/hana_shared /hana/shared nfs
rw,vers=3,hard,timeo=600,rsz=65536,wsz=65536,intr,actimeo=0,noatime,nolock 0 0
```

Note: In a multimode configuration `/hana/shared` needs to be provisioned through NFS.

2. Run `mount -a` to mount the file systems.

NTP Configuration

NTP is required and needs to be configured.

4.6 HANA Installation

Install on Single-Node System

For a single-node system, the SAP HANA software installation does not require any additional preparation.

Install on Multinode System

Note: The following installation procedure is based on SAP HANA SPS9.

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

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

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

The following output shows an example of a 2+1 multinode setup where the system identifier (SID) is FC1.

```
stlrx300s8-4:/ # cat /hana/shared/FC1/install-cfg-4FC1/global.ini
[communication]
listeninterface = .global

[persistence]
basepath_datavolumes = /hana/data/FC1
basepath_logvolumes = /hana/log/FC1

[system_information]
usage = test

[storage]
ha_provider = hdb_ha.fcClient
partition_*__prtype = 5
partition_1_data__wwid = FC1_data1
partition_1_log__wwid = FC1_log1
partition_2_data__wwid = FC1_data2
partition_2_log__wwid = FC1_log2
```

2. 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 (`stlrx300s8-5`) and the standby host (`stlrx300s8-3`).

Note: The directory where the prepared `global.ini` file has been stored is included with the `storage_cfg` command line option (`--storage_cfg=/hana/shared/FC1`).

```
stlrx300s8-4:/ # ./hdblcm --action=install --addhosts=stlrx300s8-
5:role=worker:storage_partition=2,stlrx300s8-3:role=standby
--storage_cfg=/hana/shared/FC1
```

```
SAP HANA Lifecycle Management - SAP HANA 1.00.090.00.1416514886
*****

Scanning Software Locations...
Detected components:
  SAP HANA Database (1.00.090.00.1416514886) in
/mnt/hwval/HANAuSUIToH/51049219/DATA_UNITS/HDB_SERVER_LINUX_X86_64/server
  SAP HANA AFL (Misc) (1.00.090.00.1416514886) in
/mnt/hwval/HANAuSUIToH/51049219/DATA_UNITS/HDB_AFL_LINUX_X86_64/packages
  SAP HANA LCAPPS (1.00.090.000.449616) in
/mnt/hwval/HANAuSUIToH/51049219/DATA_UNITS/HANA_LCAPPS_10_LINUX_X86_64/packages
  SAP TRD AFL FOR HANA (1.00.090.00.1416514886) in
/mnt/hwval/HANAuSUIToH/51049219/DATA_UNITS/HDB_TRD_AFL_LINUX_X86_64/packages
  SAP HANA Database Client (1.00.090.00.1416514886) in
/mnt/hwval/HANAuSUIToH/51049219/DATA_UNITS/HDB_CLIENT_LINUX_X86_64/client
  SAP HANA Smart Data Access (1.00.3.005.0) in
/mnt/hwval/HANAuSUIToH/51049219/DATA_UNITS/SAP_HANA_SDA_10/packages
Truncated ....
```

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

5 Examples for a Typical Setup

5.1 Example 2+1 Node HANA Scale-Out

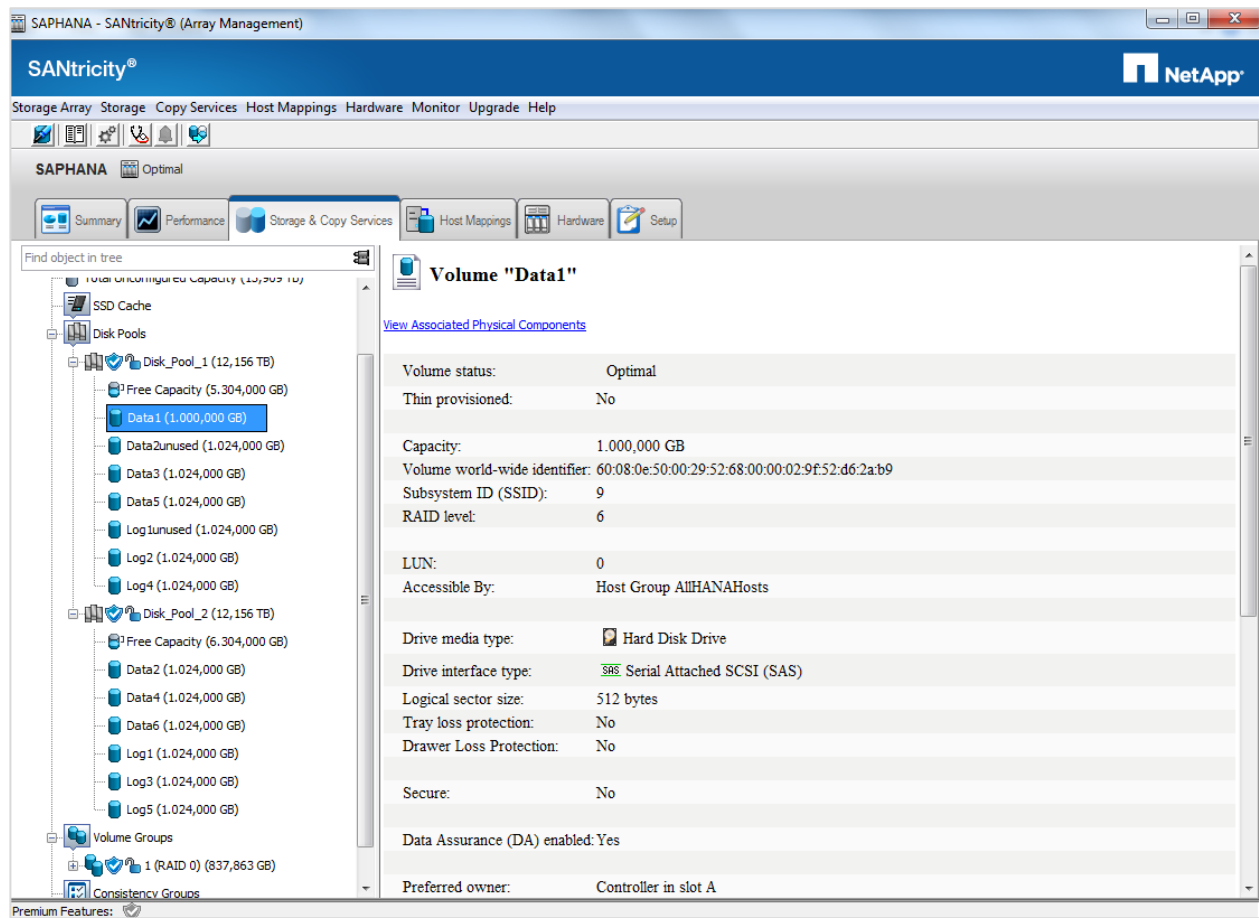
This section lists the details to set HANA up with multiple worker nodes on NetApp E-Series storage systems. It covers the setup of the storage API for HANA to provide high availability and failover capabilities.

Volume Layout

In this setup each HANA node has a dedicated log and data volume. The mount point `/hana/shared` is served by an independent NetApp FAS storage controller and is mounted on all HANA nodes. Because the default location for backup (data and log backup) lies within the `/hana/shared` volume, the NFS server also has to deal with the DATA and LOG backup.

This setup uses NetApp SANtricity DDP. The combination of DDP and NetApp E-Series storage is ideal for high-performance computing (HPC) and other big data storage environments. DDP features a next-generation architecture that minimizes the impact of drive failures and returns the system to optimal conditions up to eight times more quickly than traditional RAID.

Figure 14) Disk pool layout.



The number of disk pools depends on the capacity requirements of the HANA nodes and the number of disks that are attached to the storage system. In this example two DDPs, each with 15 disks, have been created. Each disk pool contains two volumes, which are attached to active HANA compute nodes. Data and log partitions are distributed in such a way that no HANA node has its data and log files in the same disk pool together.

Disk Pool Settings

Each disk pool (Disk_Pool_1, Disk_Pool_2) contains 15 disks in this setup, and each volume provided by this disk pool uses RAID 6 data redundancy schemes.

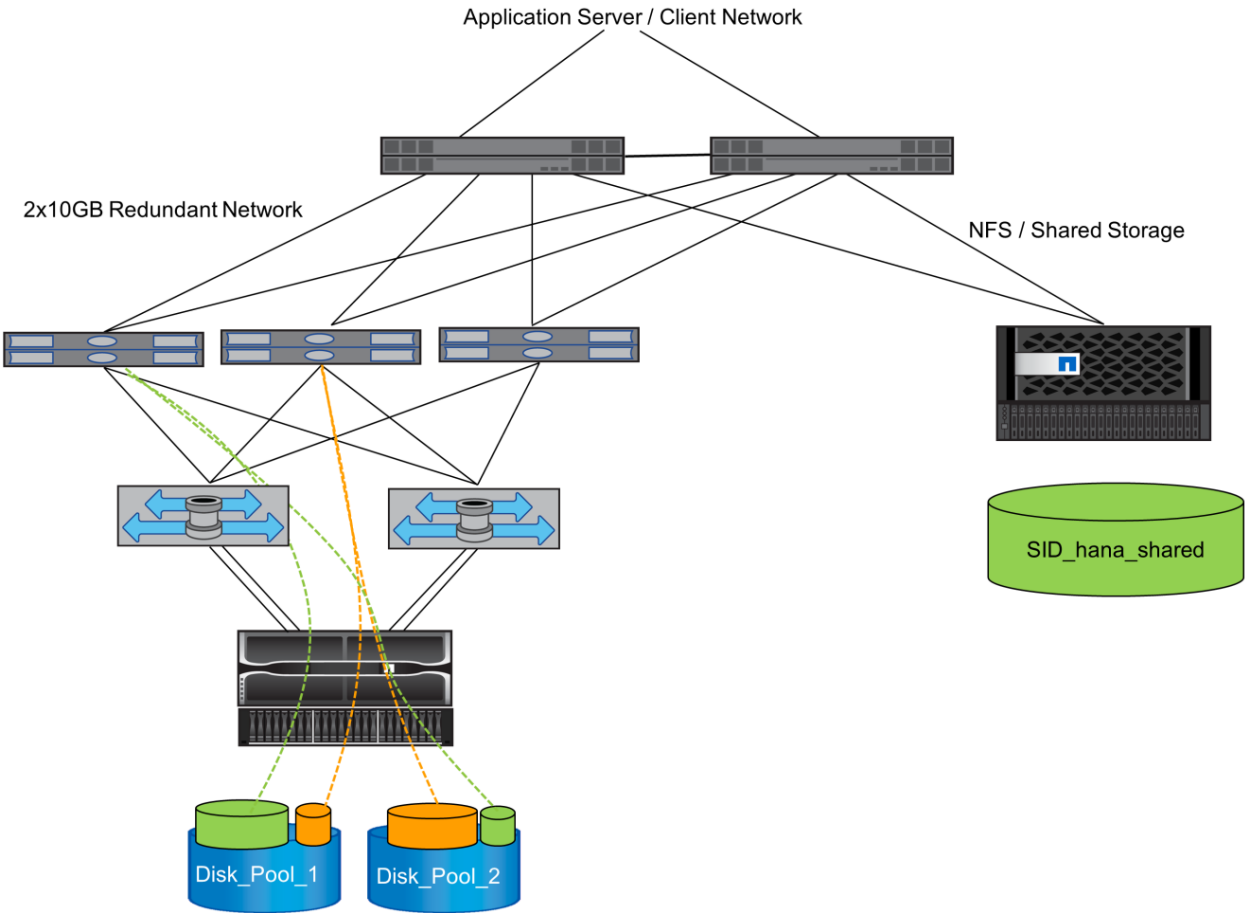
Figure 14 shows a sample screenshot of one of those disk pools displayed by the SANtricity management software. Within each pool two volumes/LUNs were created, one for data, one for log.

For convenience, meaningful aliases should be created to make handling of log and data of each HANA storage partition easier. The following nomenclature applies to this setup. The number of the HANA storage partitions and its usage type define the alias name as E55_{Log|Data}{storage_partition_nr}. The alias list looks like this:

```
lrwxrwxrwx 1 root root      7 Jan 17 03:23 E55_Data1
lrwxrwxrwx 1 root root      7 Jan 17 03:24 E55_Data2
lrwxrwxrwx 1 root root      7 Jan 17 03:21 E55_Log1
lrwxrwxrwx 1 root root      7 Jan 17 03:21 E55_Log2
```

All nodes share the same aliases to allow using the HANA storage API. Figure 15 shows the distribution of HANA storage partitions.

Figure 15) DDP and volume layout.



Volume Configuration

For each HANA node, a data volume and log volume need to be created, each equal to the size of the RAM (or larger).

The shared data of all HANA nodes is stored in a volume with the total RAM size of all active HANA nodes belonging to one HANA database. As a “shared” volume, it needs to be provisioned using an NFS server such as a FAS system in the data center. The default location for HANA backup is within the /hana/shared area, and thus /hana/shared is used also used as a backup volume in this case.

Table 4 shows a configuration for the SID=E55 and a system with two 1TB HANA worker nodes.

Table 4) Volume configuration.

Accessibility	Purpose	Disk Pool/Volume	Mount Point (\$size)
HANA node 1	Data and log partition for node 1	Disk_Pool_1/Data1 Disk_Pool_2/Log1	/hana/data/E55/mnt00001 (1TB) /hana/log/E55/mnt00001 (0.5TB)
HANA node 2	Data and log partition for node 2	Disk_Pool_2/Data2 Disk_Pool_1/Log1	/hana/data/E55/mnt00002 (1TB) /hana/log/E55/mnt00002 (0.5TB)

Accessibility	Purpose	Disk Pool/Volume	Mount Point (\$size)
All HANA nodes	Share for HANA	Shared volume: 4TB: /vol_HANA_E55_shared	/hana/shared (4TB; 2xRAM + backup)

Primary Controller/Preferred Path Assignment

Each of the LUNs for a HANA node is assigned to a different primary controller, and each LUN in a DDP is assigned to a different controller. This way the load is optimally balanced between DDPs and controllers.

Zoning

All nodes must have access to all volumes, which needs to be reflected in the zoning. Also, the storage controller must be configured so that all volumes are visible to all hosts.

Multipathing

Multipathing has to be turned on to aggregate traffic and make sure of proper failover behavior. The following multipath.conf can be used on the OS level.

```
# > cat /etc/multipath.conf
defaults {

    user_friendly_names no
    max_fds max
    flush_on_last_del yes

}

blacklist {

    devnode "^hd[a-z]"
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
    devnode "^cciss.*"

    device {
        vendor "*"
        product "Universal Xport"
    }

}

devices {

    device {
        vendor "NETAPP"
        product "INF-01-00"
        path_grouping_policy "group_by_prio"
        path_selector "service-time 0"
        path_checker "rdac"
        features "2 pg_init_retries 50"
        hardware_handler "1 rdac"
        prio "rdac"
        failback "immediate"
        no_path_retry 30
    }

}

multipaths {
    multipath {
        alias      E55_Data1
        wwid       360080e50002952680000029f52d62ab9
    }
    multipath {
```

```

        alias    E55_Data2
        wwid     360080e50002953f0000002a952d6172f
    }
    multipath {
        alias    E55_Log1
        wwid     360080e50002953f0000002ab52d6174d
    }
    multipath {
        alias    E55_Log2
        wwid     360080e50002952680000029952d61771
    }
}

```

The section `multipaths {}` is used to define aliases for the WWIDs to simplify the configuration and handling of the `global.ini`.

The device attribute features "2 pg_init_retries 50" and "no_path_retry 30" are used in combination with the `global.ini` settings `partition_*_*__prtype = 5` (see next paragraph).

Storage Connector API (global.ini)

The setup uses the SAP Storage Connector API for block devices. The preceding multipath configuration allows a simplified customization of the API in `global.ini`.

```

#> cat global.ini

[persistence]
basepath_datavolumes = /hana/data/E55
basepath_logvolumes = /hana/log/E55

[storage]
ha_provider = hdb_ha.fcClient
partition_*_*__prtype = 5
partition_1_data__wwid = E55_Data1
partition_1_log__wwid = E55_Log1
partition_2_data__wwid = E55_Data2
partition_2_log__wwid = E55_Log2

```

The configuration is using `prtype =5`, that is, write exclusive locks to make sure that only one single node is able to write to its data and log devices.

References

This report references the following documents and resources:

- SAP Certified Enterprise Storage Hardware for SAP HANA
<http://scn.sap.com/docs/DOC-48516>
- SAP HANA Storage Requirements
www.saphana.com/docs/DOC-4071
- SAP HANA Tailored Datacenter Integration Frequently Asked Questions
www.saphana.com/docs/DOC-3634
- SAP HANA Tailored Datacenter Integration Overview Presentation
www.saphana.com/docs/DOC-3633

Version History

Version	Date	Document Version History
Version 1.0	March 2014	Initial version

Version	Date	Document Version History
Version 2.0	July 2015	Including new model 5600/EF560

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2015 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the Web at <http://www.netapp.com/us/legal/netapptmlist.aspx>. TR-4296-0715

