

Dell Validated Design for SAP HANA TDI with PowerStore Storage Systems

September 2022

H18336.3

Design Guide

Abstract

This design guide describes storage best practices for SAP HANA in Tailored Data Center Integration (TDI) deployments on SAP-certified Dell PowerStore enterprise storage systems.

Dell Technologies Solutions



Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2022 Dell Inc. or its subsidiaries. Published in the USA 09/22 Design Guide H18336.3.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Revision history

Publication date	Part number	Description
June 2020	H18336	Initial release: PowerStoreOS 1.0 Fibre Channel (FC) certification
March 2021	H18336.1	Updated for PowerStoreOS 1.0 NAS certification
October 2021	H18336.2	Updated for PowerStoreOS 2.0 recertification to include model 500T
September 2022	H18336.3	Updated for PowerStoreOS 3.0, NAS certification for new models

Contents

Introduction	4
Technology overview.....	6
Network design considerations	14
PowerStore storage configuration	15
SAP HANA host configuration and setup	24
SAP HANA installation with FC connectivity.....	40
SAP HANA installation with NAS/NFS connectivity.....	42
Post-installation configuration	52
Conclusion.....	55
References.....	56

Introduction

Solution overview

SAP HANA is an in-memory data platform that can be deployed on-premises or in the cloud. SAP HANA uses in-memory computing to store compressed data in RAM rather than in a relational database on disk. Organizations use the SAP HANA platform for Enterprise Resource Planning (ERP) and other business applications, enabling them to process and analyze live data and make real-time business decisions.

Dell PowerStore is a midrange enterprise storage system that is optimized for simplicity across the storage life cycle. The PowerStore platform is simple to acquire, deploy, manage, and service.

The solution that this guide describes uses SAP HANA in a tailored data center integration (TDI) deployment scenario on PowerStore enterprise storage systems. These storage systems meet the SAP HANA performance and functional requirements and have SAP certification. SAP certification means that customers can use PowerStore systems for SAP HANA TDI deployments in a fully supported environment using their existing data center infrastructures.

The configuration recommendations in this guide are based on SAP's high availability (HA) requirements and the performance test results that are required to meet SAP's key performance indicators (KPIs) for SAP HANA TDI.

Key benefits

This Dell solution for SAP HANA TDI deployments on PowerStore systems reduces hardware and operational costs and lower risks by integrating SAP HANA into existing data center operational processes and increases server and network vendor flexibility for customers. Customers can:

- Integrate SAP HANA into an existing data center.
- Choose between the following storage protocols and connectivity options for the SAP HANA nodes:
 - **NAS:** Shared file system (NAS/NFS)
 - **SAN:** Fibre Channel (FC SAN)
- Use their existing operational processes, skills, and tools, thus avoiding the significant risks and costs that are associated with operational change.
- Use the performance and scale benefits of PowerStore systems to obtain real-time insights across the business.
- Expect significant benefits from using NVMe storage class memory (SCM) and solid-state device (SSD) drives for the SAP HANA persistence by reducing SAP HANA startup, host autofailover, and backup times.
- Expect storage efficiencies and capacity savings with PowerStore always-on inline data reductions such as zero detect, compression, and deduplication.
- Transition easily from an appliance-based model to the PowerStore-based TDI architecture while relying on Dell Technologies Professional Services to minimize risk.

Document purpose

This guide provides:

- Best practices and tips for deploying the SAP HANA database on PowerStore enterprise storage systems.
- Configuration requirements and storage design recommendations for PowerStore systems that are used with SAP HANA.
- Network requirements for the FC SAN and NAS/NFS protocols and connectivity.
- Steps for configuring PowerStore storage for SAP HANA using PowerStore Manager for FC SAN and NAS/NFS.
- Detailed instructions for accessing PowerStore storage from the SAP HANA nodes.

Note: This guide describes SAP HANA TDI deployments in physical environments. If you plan to use SAP HANA in virtualized environments on VMware vSphere, see [SAP HANA on VMware vSphere](#).

Audience

This guide is intended for Dell Technologies customers, partners, and Professional Services personnel who must configure a PowerStore system for an SAP HANA TDI environment.

Term table

The following table defines abbreviations that are used in this guide:

Term	Definition
iDRAC	integrated Dell Remote Access Controller
LACP	Link Aggregation Control Protocol
LDAP	Lightweight Directory Access Protocol
MLAG	Multi-Chassis Link Aggregation
NAS	Network-attached storage
NDMP	Network Data Management Protocol
NFS	Network file system
NIS	Network Information Service
OLTP	Online transaction processing
OLAP	Online analytical processing
PCIe	Peripheral Component Interconnect express
SAN	Storage area network

We value your feedback

Dell Technologies and the authors of this guide welcome your feedback on the solution and the solution documentation. Contact the Dell Technologies Solutions team by [email](#).

Author: Donagh Keeshan

Contributors: Aighne Kearney, Stephen Wright

Technology overview

SAP HANA deployment models

The SAP HANA platform combines SAP software components that are optimized on proven and certified SAP partner-provided hardware. Two models are available for on-premises deployment:

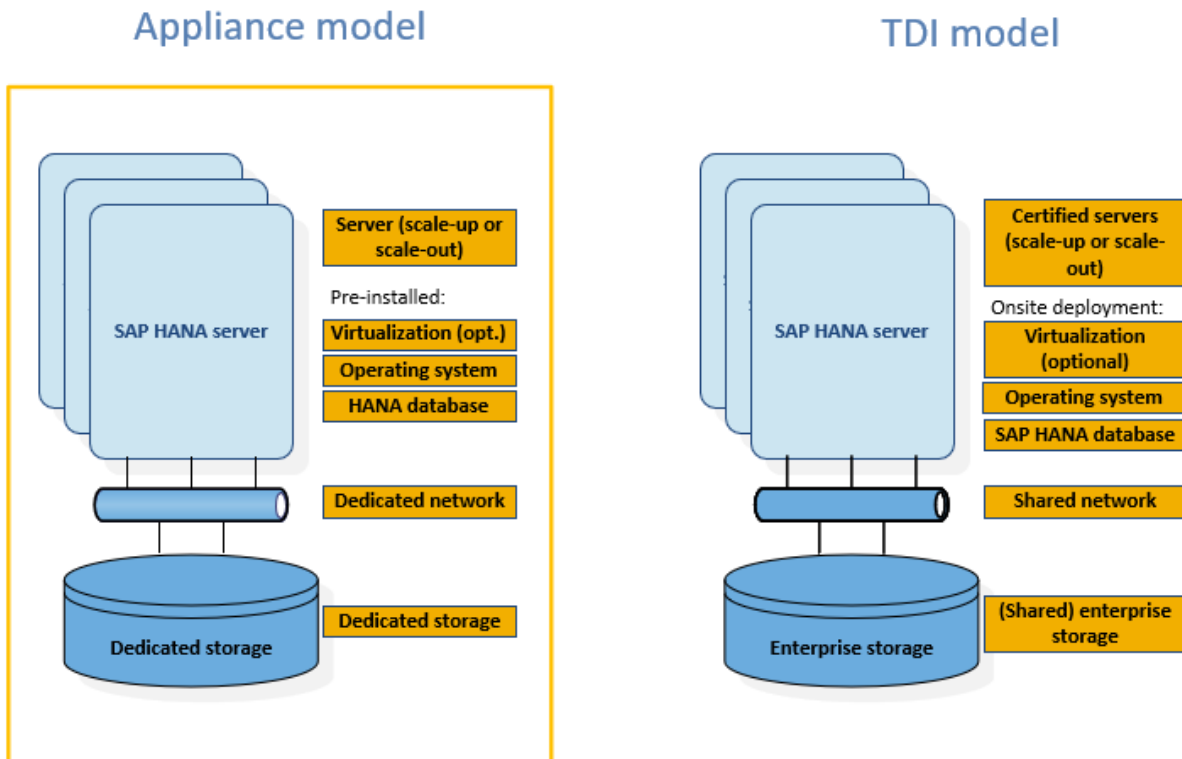


Figure 1. Comparison of SAP HANA deployment models

Appliance model

By default, an SAP HANA appliance includes integrated storage, compute, and network components. The appliance is certified by SAP, built by an SAP HANA hardware partner, and shipped to customers with all its software components preinstalled, including operating systems.

Dell Technologies provides preinstalled SAP HANA appliance solutions for a faster time-to-market and easy integration into an SAP landscape. However, the SAP HANA appliance model presents the following limitations for customers:

- Limited choice of servers, networks, and storage
- Inability to use existing data center infrastructure and operational processes
- Little control of the critical components in the appliance
- Fixed sizes for SAP HANA server and storage capacities, leading to higher costs from a capacity shortfall and inability to respond rapidly to unexpected growth demands

TDI model

The TDI deployment model enables customers to choose from a broad portfolio of SAP HANA-certified servers that can be combined with SAP-certified network and storage components. Different workloads can share the storage and network components to optimize the total cost of ownership (TCO). Customers can seamlessly integrate SAP HANA into existing data center operations such as disaster recovery, data protection, monitoring, and management, reducing the cost, time-to-value, and risk of an SAP HANA adoption.

For more information, see the following SAP documents:

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

SAP HANA database

An SAP HANA database is an in-memory database. The data is kept in the RAM of one or more SAP HANA worker hosts (active components that accept and process database requests) while all database operations (reads, inserts, updates, and deletions) are performed in the main memory of the host. This feature differentiates the SAP HANA database from traditional databases, where only a part of the data is cached in RAM and the remaining data resides on disk.

Persistent storage enables you to restore the SAP HANA database to its most recent committed state if there is a failure. The log captures all changes by database transactions (redo logs). Data and undo log information are automatically saved to disk at regular savepoints (the default is five minutes).

Scale-up and scale-out deployments

As certified enterprise storage systems for SAP HANA, PowerStore systems can be used for both single-host (scale-up) and multihost (scale-out) systems in TDI deployments.

In single-host environments, the database must fit into the RAM of a single server. Single-host environments are preferred for OLTP-type workloads such as S/4HANA and SAP Business Suite on SAP HANA.

In multihost environments, the database tables are distributed across the RAM of multiple servers. These environments use worker and standby hosts. The worker hosts accept and process database requests. By contrast, standby hosts are passive components that have the database services running but no data in RAM. A standby host waits for a worker host to fail and then takes over its role, a process known as host autofailover. Because the in-memory capacity in these deployments can be high, scale-out SAP HANA clusters are ideally suited for online OLAP-type workloads with large datasets, such as SAP Business Warehouse on SAP HANA and BW/4HANA. By default, SAP supports scale-out deployments of up to 16 worker hosts. If more than 16 worker hosts are needed, a site-specific SAP certification is required.

Dell PowerStore

The PowerStore system is a midrange storage product with a container-based active/active architecture. The PowerStore system supports the latest storage media, such as NVMe SCM, NVMe SSDs, and NVMe NVRAM drives. This system is a highly available 2U two-node-based appliance that incorporates a flexible consumption model. The PowerStore system can scale out to up to four appliances (eight nodes) in a cluster and scale up to three expansion enclosures per appliance.



Figure 2. PowerStore base enclosure

The PowerStore platform consists of ten different models, from the PowerStore 500 model through the PowerStore 9200 model. PowerStoreOS 1.0 was launched with the PowerStore 1000T/X, 3000T/X, 5000T/X, 7000T/X, and 9000T/X models, while PowerStoreOS 2.0 introduced the entry-level 500 model. PowerStoreOS 3.0 introduced a platform refresh with the 1200T, 3200T, 5200T, and 9200T models. All PowerStore models use a common base enclosure and I/O modules. The models differ by CPU core count and speed, memory size, and number of NVMe NVRAM drives. These hardware differences give each model a unique performance profile, scaling from entry-level to high-end.

Appliances in the PowerStore series are available in one of the following configurations:

- **PowerStore T appliances:** These models are storage-centric. Use them to manage and provision block and file storage to external hosts. During the initial configuration, you can configure an appliance for block-only storage or for unified (block and file) storage.

In SAP environments, you can configure the appliance for block-only storage if a shared NFS file system is already available for the provision of the required SAP shared file systems such as `/hana/shared/` and `/sapmnt/SID/`. If a shared NFS file system is not available, you can configure the appliance for unified storage, where either the NAS/NFS file storage or block storage is used for the SAP HANA persistence or the PowerStore NFS share file components are used for the provision of the `/hana/shared` file system. In a PowerStore cluster, only the primary appliance would be configured as unified; the subsequent appliances in the cluster would be configured as block-only.

Note: PowerStore T models are certified for running SAP HANA.

- **PowerStore X appliances:** These models are application-centric and storage-centric. PowerStore X appliances provide a hypervisor layer in addition to block storage. This additional capability enables you to deploy and manage virtual machines and applications.

Note: PowerStore X models are not certified for running SAP HANA in production at the time of publication, but can be used for non-production SAP HANA systems.

Each of these configurations is available in different models with varying or expanding memory and CPU capabilities, as described in Table 1.

PowerStore clusters

A PowerStore cluster combines multiple PowerStore appliances into a single grouping, which is managed as a single storage system. The cluster delivers aggregate performance from all appliances in the cluster, but a single volume is serviced by only one appliance at any given time. Volumes can be migrated between appliances in a cluster. A cluster must contain only PowerStore T appliances or only PowerStore X appliances; these models cannot be mixed within a single cluster.

Dell Technologies recommends that:

- All appliances in a cluster are of the same model and have similar physical capacities to provide consistent performance across the cluster.
- Any host that is connected to a PowerStore cluster has equivalent connectivity to all appliances in the cluster.
- All appliances in a cluster are physically located in the same data center and are connected to the same LAN.

Clustering is applicable to block storage resources only. While a PowerStore T model in unified mode can serve as the cluster's primary appliance, the file resources cannot migrate to a different appliance. When deploying multiple appliances for file access, plan to have multiple clusters. For more information about unified deployment mode for PowerStore T models, see [Dell PowerStore: Introduction to the Platform](#).

PowerStore management tools

The HTML5-based PowerStore Manager UI is the primary management tool for PowerStore configuration and administration.

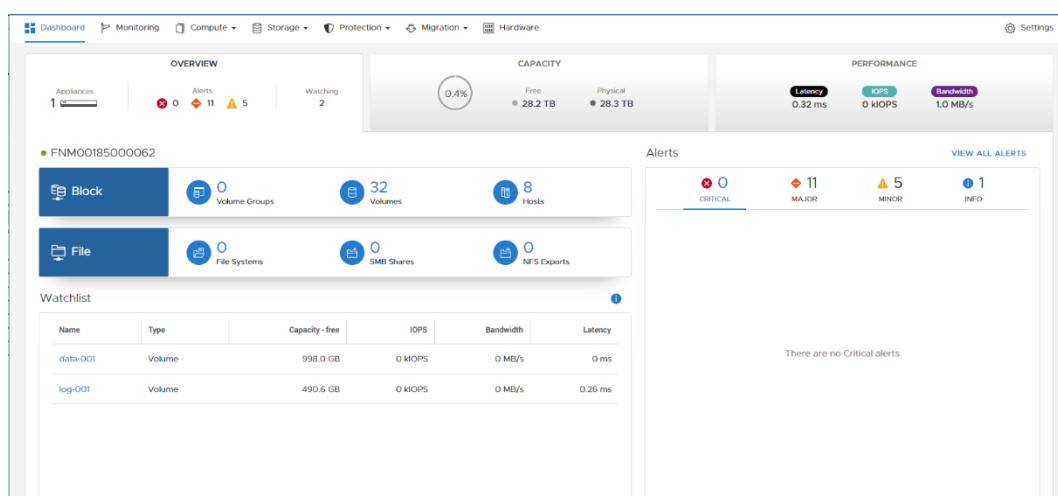


Figure 3. PowerStore Manager UI

The PowerStore system also supports management, configuration, and monitoring using the PowerStore command-line interface (PSTCLI) and a REST API. For relevant documentation, see the [PowerStore Info Hub](#).

PowerStore T models certified for SAP HANA

Dell Technologies has certified the following PowerStore models for running SAP HANA: PowerStore 500T, 1000T, 1200T, 3000T, 3200T, 5000T, 5200T, 7000T, 9000T, and 9200T: Table 1 shows high-level specifications for the new platform refresh models with PowerStoreOS 3.0, and Table 2 shows high-level specifications for the PowerStoreOS 1.0 models.

For more information, see the following specification sheets:

- [PowerStore 3.0 models](#)
- [PowerStore 1.0 models](#)

Note: Dell Technologies performed SAP HANA certification tests for both FC SAN and NAS/NFS on the 500T, 1000T, 5000T, and 9200T models with a base enclosure configuration of 19 x NVMe SSD drives and Dynamic Resiliency Engine (DRE) double tolerance protection (R6). The PowerStore X models are not certified for SAP HANA TDI production deployments but can be used for nonproduction deployments.

Table 1. Storage specifications for PowerStore 3.0 models

Specification	500T	1200T	3200T	5200T	9200T
CPU per appliance	2 x 12c Intel CPUs @ 2.2GHz	4 x 10c Intel CPUs @ 2.4GHz	4 x 16c Intel CPUs @ 2.1 GHz	4 x 24c Intel CPUs @ 2.2 GHz	4 x 28c Intel CPUs @ 2.2GHz
Memory	192 GB	384 GB	768 GB	1152 GB	2,560 GB
NVRAM drives per appliance	N/A	2		4	
Maximum number of storage drives per appliance	25	93			
Supported drives	NVMe Flash/SCM	NVMe SSD, NVMe SCM			
Embedded ports	4 x 25/10 GbE optical	4 x 25/10 GbE optical or 4 x 10 GbE Base T			
Supported I/O modules (two slots per node)	4 x 32/16/8 GB FC 4 x 25/10 GbE optical 4 x 10/1 GbE base T	2 x 100 GbE Optical, 4 x 32/16/8Gb Fibre Channel (FC) 4 x 25/10 GbE Optical, 4 x 10 GbE-T			
Supported expansion shelves	N/A (cluster only)	24-drive NVMe SSD			
Clustering	Up to four appliances (mix and match any model)				

Table 2. Storage specifications for PowerStore 1.0 models

Specification	1000T	3000T	5000T	7000T	9000T
CPUs per appliance	4 x Intel CPUs, 32 cores, 1.8GHz	4 x Intel CPUs, 48 cores, 2.1GHz	4 x Intel CPUs, 64 cores, 2.1GHz	4 x Intel CPUs, 80 cores, 2.4GHz	4 x Intel CPUs, 80 cores, 2.4GHz
System cache/Memory	384 GB	768 GB	1152 GB	1536 GB	2560 GB
NVRAM drives per appliance	2			4	
Maximum number of storage drives per appliance	96				
Supported drives	NVMe flash/SCM, SAS SSD				
Embedded ports	4 x 25/10 GbE optical or 4 x 10 GbE Base T				
Supported I/O modules (two slots per node)	4 x 32/16/8 GB FC 4 x 25/10 GbE optical 4 x 10/1 GbE base T				
Supported expansion shelves	25-drive SAS SSD				
Clustering	Up to four appliances (mix and match any model)				

SAP HANA scalability on PowerStore

PowerStore storage systems are certified for SAP HANA with both:

- SAN: Fibre Channel (FC)
- NAS: Shared file system

Note: PowerStore 3.0 1200T, /3200T, 7200T, and /9200T models are certified for NAS connectivity. FC SAN certification is planned and pending soon for these new models.

SAP HANA FC SAN certification for PowerStore storage

SAP has certified PowerStore storage for SAP HANA for using the FC protocol and the SAP fcClient storage connector. For information about the host and network requirements for the FC connectivity protocol, see [Network design considerations](#).

To determine the proper storage system model, number, and type of front-end I/O modules and disk configuration for an SAP HANA TDI deployment, you will need the following information:

- Number of SAP HANA nodes to be deployed on the storage system
- Capacity requirements of the SAP HANA nodes

An SAP HANA node can be either a single (scale-up) SAP HANA server or an SAP HANA worker node that is part of an SAP HANA multihost cluster. Multiple SAP HANA nodes can be connected to a PowerStore T storage system up to a recommended maximum number. For more information, see [SAP HANA capacity requirements](#).

For the SAP HANA certification, Dell Technologies performed extensive testing using the SAP HANA-HWC-ES-1.1 certification scenario to determine the scalability of the PowerStore product family. The following table shows the number of SAP HANA nodes that are supported on a specified PowerStore storage system model or cluster. These numbers represent the recommended maximum number of SAP HANA production nodes that can be connected to the PowerStore system in block-only deployments with either single parity or dual parity tolerance levels while still meeting the SAP performance KPIs. For PowerStore unified deployments, some additional system resources are used to provide the FILE capabilities. The overhead could reduce the SAP HANA scalability by up to a maximum of two SAP HANA nodes per single appliance.

Scalability refers to the supported number of worker nodes. The following table shows the supported number of nodes per appliance and cluster in production deployments:

Table 3. SAP HANA FC scalability per appliance and cluster in production

Appliance/appliances per cluster	Number of SAP HANA nodes per appliance, two-appliance cluster, three-appliance cluster, or four-appliance cluster
500T/up to 4	6, 12, 18, 24
1000T/up to 4	8, 16, 24, 32
3000T/up to 4	10, 20, 30, 40
5000T/up to 4	16, 32, 48, 64
7000T/up to 4	18, 36, 54, 72
9000T/up to 4	20, 40, 60, 80

The **Appliance/Appliances per cluster** column in the table shows the PowerStore T appliance models that can be clustered with up to four appliances. Each appliance has two active/active storage nodes, and each node has two four-port front-end (FE) I/O (I/O) modules.

The second column shows the number of production SAP HANA nodes that can be deployed on the specified single appliance model and in two-appliance, three-appliance, and four-appliance cluster configurations using FC SAN connectivity. A PowerStore cluster of appliances has linear scalability. PowerStore supports clustering for block-only for up to four appliances. In a four-appliance cluster, only the first or primary appliance can be unified while the other three must be block-only. For example, in a 4 x 9000T appliance cluster, the primary 9000T appliance could be a unified configuration supporting 18 SAP HANA nodes while the remaining three 9000T appliances would be configured as block-only in the cluster, with each appliance supporting 20 SAP HANA nodes. In this example, a total of 78 SAP HANA nodes is therefore supported.

Note: SAP HANA standby nodes in an SAP HANA scale-out cluster are not counted for scalability because they do not have a storage persistence.

Achieving these numbers in a customer environment requires:

- FC network with 16 or 32 Gb/s link speed on all components
- A four-port I/O module (16 or 32 Gb/s) per PowerStore node

Note: An I/O module with four FE ports per PowerStore node (eight FE ports per appliance) are required for the maximum number of SAP HANA nodes.

- Eight dedicated FE ports for SAP HANA per appliance, with SAP HANA nodes equally balanced across the FE ports
- Capacity for the SAP HANA persistence allocated on PowerStore volumes or file systems with the “Volume Performance Policy” set to high
- Host connections configured as described in [SAP HANA host configuration and setup](#)

SAP HANA NAS certification for PowerStore storage

SAP has certified PowerStore storage for SAP HANA for use with NAS and the NFS protocol and with the Dell custom HA/DR STONITH provider. For more information about the host and network requirements for the NAS connectivity protocol, see [Network design considerations](#).

To determine the proper storage system model and the disk configuration for an SAP HANA TDI deployment, you will need the following information:

- Number of SAP HANA nodes to be deployed on the storage system
- Capacity requirements of the SAP HANA nodes

An SAP HANA node can be either a single (scale-up) SAP HANA server or an SAP HANA worker node that is part of an SAP HANA multihost cluster.

Multiple SAP HANA nodes can be connected to a PowerStore T storage system up to a recommended maximum number, as specified in the following table. For more information, see [SAP HANA capacity requirements](#).

As part of the SAP HANA certification, Dell Technologies performed extensive testing using the SAP HANA-HWC-ES-1.1 certification scenario to determine the scalability of the PowerStore product family. The following table shows the recommended maximum number of SAP HANA production nodes that can be connected to a specified PowerStore model while still meeting the SAP performance KPIs:

Table 4. SAP HANA NAS scalability per appliance

Appliance	Number of SAP HANA nodes
500T	4
1000T	6
1200T	10
3000T	10
3200T	12
5000T	16
5200T	18
7000T	20
9200T	22

Table 4 shows the number of production SAP HANA nodes that can be deployed on the specified configuration using NAS/NFS connectivity. SAP HANA standby nodes in a HANA scale-out cluster are not counted for scalability because they do not have a storage persistence.

Note: PowerStore appliance clustering for NAS storage is not yet available.

Achieving these numbers in a customer environment requires:

- Storage IP network with 25 Gb/s link speed on all components
- 2 x 25 Gb/s interfaces on the first two ports of the embedded four-port card per PowerStore node (4 x 25 Gb/s per appliance)
- SAP HANA nodes equally balanced across two NAS servers, one on each PowerStore node.
- NAS servers configured as described in [Creating NAS file systems for the SAP HANA persistence](#)
- Host connections configured as described in [SAP HANA host configuration and setup](#)

Network design considerations

Introduction

SAP HANA production systems in TDI environments must meet the SAP KPIs for storage performance. These systems also demand HA and redundancy in storage and network components such as network cards, (HBAs), network switches, and storage FE components. HA and redundant systems protect SAP HANA database operations from component failure.

The following sections of this guide provide configuration recommendations and best practices for the networks and storage systems depending on the connectivity option that you choose for the SAP HANA production systems.

Network configuration

This section describes the considerations that apply when you connect SAP HANA to PowerStore systems over an FC SAN network or an Ethernet network for NAS/NFS.

FC SAN

For FC connectivity, configure dual redundant fabrics with each PowerStore node and each external host having connectivity on each of the fabrics. Minimize the number of hops between host and PowerStore. For performance, load balancing, and redundancy, each host must have at least two paths to each PowerStore node (four ports per PowerStore appliance). Dell Technologies recommends that a host has no more than eight paths per volume.

The FC SAN network must support 16 Gb/s or 32 Gb/s speed on all components: HBAs, SAN ports, and PowerStore FE I/O modules. When SAN switches are connected using Inter-Switch Links (ISLs), the aggregated bandwidth of the ISLs must be able to support the aggregated total bandwidth of all the SAP HANA nodes that are using these ISLs. An SAP HANA node requires the maximum bandwidth of 400 MB/s.

Ethernet networks: IP network for NAS/NFS

For Ethernet connectivity, use multiple switches that are connected with Virtual Link Trunking interconnect (VLTi) and Link Aggregation Control Protocol (LACP) or equivalent technologies. Each PowerStore node must have connectivity to all linked switches. For PowerStore T models, the first two ports of the embedded module 4-port card on each PowerStore node are bonded together within the PowerStoreOS. For highest performance and availability from these ports, Dell Technologies also recommends configuring link aggregation across the corresponding switch ports. LACP is not applicable to PowerStore X models.

With the introduction of PowerStoreOS 3.0, user-defined link aggregations are supported for Ethernet ports. Link aggregations can include 2 to 4 different Ethernet ports. The ports must be on the same node and operate at the same speed. A mirror link aggregation will automatically be created on the peer node. User-defined link aggregations only support NAS server interfaces.

To connect the SAP HANA nodes using NAS/NFS, you need a dedicated storage IP network that has redundant switches with 25 Gb/s link speed on all components, including host network interface cards (NICs) and switch ports. Jumbo frames with maximum transmission unit (MTU) size 9,000 are also required and must be set on all network components. When IP switches are connected using uplinks, the aggregated bandwidth of the uplinks must be able to support the aggregated total bandwidth of all the SAP HANA nodes that are using these uplinks. An SAP HANA node requires the maximum bandwidth of 400 MB/s.

Direct attachments

PowerStore systems support direct attachment of the SAP HANA hosts to the appliance through FC if the hosts connect to both the PowerStore nodes and have the required native Linux multipathing software (DM-MPIO). Although direct-attached FC connections are supported with PowerStore, Dell Technologies recommends using switches for production environments to provide HA and redundancy and protect SAP HANA against any component failures.

PowerStore storage configuration

Overview

This section provides configuration recommendations and best practices for connecting the PowerStore storage system to SAP HANA nodes.

Module connectivity options

All PowerStore models have the same connectivity options. The following figure shows these connectivity options as seen from the rear of the unit:




Figure 4. PowerStore T rear view

PowerStore node A is the bottom node, while PowerStore node B is the top node. Each node has one embedded module (red rectangles) with management and expansion ports and a four-port Mezzanine card for connections such as cluster interconnects. Each node has two additional slots for four-port I/O modules (green rectangles), Slot 0 and Slot 1.

PowerStore front-end ports

PowerStore supports Ethernet connectivity through ports on the embedded module and on optional I/O modules. PowerStore supports FC connectivity through ports on optional I/O modules. On PowerStore 1000 to PowerStore 9200 models, the I/O module in slot 0 is 16-lane PCIe Gen3, while the I/O module in slot 1 is 8-lane. For this reason, if the system is configured with a 100 GbE I/O module, the module must be installed in slot 0. When installing an FC I/O module, always use I/O module slot 0 first unless the system contains a 100 GbE I/O module. Both I/O modules on PowerStore 500 models are 8-lane PCIe, and the 100 GbE I/O module is not supported on this platform.

PowerStore FC ports

PowerStore FC ports support speeds for 32 Gb/s, 16 Gb/s, 8 Gb/s, and 4 Gb/s. This speed depends on the SFP that is used and the switchport or HBA that is connected. Because higher speeds allow for greater MBPS and IOPS capabilities, use the highest speed that is supported by the environment. FC ports are available on I/O modules that are inserted into I/O module slots on the nodes. The FC I/O module is 16-lane PCIe Gen3. On PowerStore 1000 to 9200 models, I/O module slot 0 is also 16-lane, while I/O module slot 1 is 8-lane. If you install FC I/O modules in both I/O module slots, cable the ports in I/O module slot 0 first because of the PCIe difference. The PCIe lanes in I/O module slot 1 are a limiting factor only for total MBPS, and only when all four ports on the FC I/O module are operating at 32 Gb/s. Both I/O module slots on PowerStore 500 are 8-lane PCIe, and therefore there is no slot preference.

PowerStore Ethernet ports

PowerStore optical Ethernet ports support speeds of up to 25 Gb/s based on the SFP that is used. Copper Ethernet ports support speeds of up to 10 Gb/s. Because higher speeds allow for greater MBPS and IOPS capabilities, use the highest speed that is supported by your environment. With PowerStoreOS 3.0, a new two-port Ethernet card is introduced that supports speeds of up to 100 Gb/s. This 100 GbE card is supported on PowerStore 1000 to 9200 models in I/O module slot 0. Jumbo frames (MTU 9000) are recommended for increased network efficiency and must be supported on all parts of the network between PowerStore and the host.

To improve network performance and redundancy, use bonded ports for NAS connectivity. Prior to PowerStoreOS 3.0, NAS servers automatically created their interfaces on the two bonded ports on the embedded module four-port card. With PowerStoreOS 3.0, you can use user-defined link aggregations to reserve different physical ports for file access only. For highest performance and availability from any aggregated ports, configure link aggregation across the corresponding switch ports.

PowerStore drive configuration

PowerStore can be configured with NVMe solid-state devices (SSDs) or NVMe storage class memory (SCM) drives for user data. You can also expand SSD-based systems with additional drives to increase the amount of available storage capacity. PowerStore 1000, 3000, 5000, 7000, and 9000 series models can be expanded with SAS SSD-based expansion shelves.

Starting from PowerStoreOS 3.0, all PowerStore models can be expanded with NVMe SSD-based expansion shelves if they meet the hardware prerequisites that are defined in the [Dell PowerStore: Introduction to the Platform white paper](#). Dell Technologies recommends that all drives in a PowerStore system be the same size, which can maximize the usable capacity from each drive.

PowerStore Dynamic Resiliency Engine (DRE) is used to manage the drives in the system. All drives are automatically used to provide storage capacity. DRE groups the drives into resiliency sets to protect against drive failure. User configuration of the drives is not necessary, and dedicated hot spare drives are not required in PowerStore. Spare space for rebuilds is automatically distributed across all drives within each resiliency set. This configuration provides better resource utilization and enables faster rebuilds if there is a drive failure.

Except for PowerStore 500, PowerStore systems use NVMe NVRAM drives to provide persistent storage for cached write data. PowerStore 1000 to 3200 arrays have two NVRAM drives per system, while PowerStore 5000 to 9200 arrays have four NVRAM drives per system. The extra drives mean that these systems can provide higher MBPS for large-block write workloads.

The following figure shows how the NVRAM drives are configured in the base enclosure:



Figure 5. PowerStore base enclosure with drives (front view)

NVMe SSD-based systems (except for the 500T model) can also be expanded with an expansion enclosure using up to 24 SAS SSD drives to increase the amount of available storage capacity, as shown in the following figure:




Figure 6. PowerStore with base and expansion enclosure with NVMe SSD drives

SCM drives

NVMe storage class media (SCM) drives are dual-ported, extreme-high-performance, nonvolatile drives that are designed with Intel Optane technology. NVMe SCM drives have lower latency and improved performance compared to other SSD drives, and the PowerStore operating system accesses them with NVMe.

NVMe SCM drives serve as a storage tier for PowerStore, are used for user data or metadata, and come at a capacity point of 750 GB. PowerStore can use SCM drives either by having only SCM drives installed in the system, or, with PowerStoreOS 2.0 or higher, by mixing SCM and SSD drives. Systems with all SCM drives are recommended for small-block workloads that require the absolute lowest latencies. A system with all SCM drives place both data and metadata on the SCM drives. Systems with mixed SSD and SCM drives use the SCM drives for metadata acceleration; and the SCM drives store metadata for faster lookups. This configuration can reduce latency on read operations in systems with large physical capacities. When mixing SCM and SSDs, Dell Technologies recommends that at least five percent of the system's physical capacity be SCM.

PowerStore disks for SAP HANA

For SAP HANA, we recommend using NVMe SSDs with a disk capacity size that allows the system capacity to fit within the base enclosure. Use a minimum of ten drives for single-drive failure tolerance, or nineteen drives for double-drive failure tolerance. Dell Technologies has tested SAP HANA performance and scalability with single parity (single tolerance R5) and dual parity (double tolerance R6) and found that the performance levels were the same for both. The SAP HANA scalability in Table 3 (FC SAN) and Table 4 (NAS/NFS) is therefore valid for both single and double failure tolerance levels.

Data reduction

PowerStore systems provide data-reduction capabilities such as zero-detect, compression, and deduplication. Data reduction is integrated into the PowerStore architecture and is always active.

All storage resources in PowerStore systems are thin-provisioned and space-efficient, including snapshots and thin clones. Creating a snapshot or thin clone requires only a quick duplication of pointers, after which the resources behave as independent storage resources and do not affect the performance of the source resource.

NAS servers

The PowerStore file system uses virtualized file servers that are known as NAS servers. A NAS server contains the configuration, interfaces, and environmental information that is used to facilitate access to the file systems. This information includes services such as DNS, LDAP, NIS, NDMP, anti-virus, and more.

New NAS servers are automatically assigned on a round-robin basis across the available PowerStore nodes. The preferred node acts as a marker to indicate the node that the NAS server runs based on this algorithm. After it is provisioned, the preferred node for a NAS server never changes. The current node indicates the node on which the NAS server is running. Changing the current node moves the NAS server to a different node, which can be used for load-balancing purposes. When a NAS server is moved to a new node, all file systems on the NAS server are moved along with it.

Volumes and file systems

As a unified storage system, the PowerStore system allocates block LUNs (volumes) and file systems on the same storage resources.

Performance policy

All block storage resources in a PowerStore system have a defined performance policy. By default, this policy is set to medium. The policy does not affect system behavior unless some volumes are set to low performance while other volumes are set to medium or high performance. During times of system resource contention, the PowerStore system devotes fewer compute resources to volumes with a low-performance policy. Reserve the low-performance setting for volumes that have fewer critical performance needs and use the high-performance setting for all SAP HANA persistence volumes.

Appliance balance for SAP HANA block storage workloads

PowerStore block storage resources are accessed using ALUA¹ optimized/non-optimized paths between the hosts and the two nodes within the PowerStore appliance. The PowerStore system automatically chooses one of the nodes for the optimized path when the volume is mapped to the host. This node affinity feature maintains a balanced workload across the PowerStore nodes. You can view and modify node assignment by using the PSTCLI command line. For the SAP HANA persistence (data and log) of a single SAP HANA server, it is a good practice to keep data and log assigned to alternate PowerStore nodes.

To view the current node assignment, run the following command:

```
pstcli -d <IP address> -u <username> -p <password> volume -name <volume name> show -output nvp | grep "ffin"
```

Here is an example of the view node affinity command and output:

```
pstcli -d 10.xxx.xx.xx -u admin -p ***** volume -name HANA-
data-tdi-hana-storage-test01-vol show -output nvp | grep "ffin"
node_affinity                = System_Selected_Node_A
node_affinity_l10n            = System Selected Node A
```

¹ ALUA is an industry standard protocol for identifying optimized paths between a storage system and a host.

To change the node affinity for a volume, run the following command:

```
pstcli -d <IP address> -u <username> -p <password> volume -name <volume name> set -node_affinity Preferred_Node_A (or Preferred_Node_B)
```

Appliance balance for SAP HANA file system workloads

For best performance, it is important to balance the file system workload across the PowerStore nodes. For the SAP HANA persistence (data and log), it is a good practice to always use two NAS servers (one on each PowerStore node) and, when creating the file systems, to distribute them evenly across the two NAS servers.

SAP HANA capacity requirements

Every SAP HANA node requires storage devices and capacity for:

- Operating system boot image
- SAP HANA installation (/hana/shared)
- SAP HANA persistence (data and log)
- Backup

Operating system boot image

When the SAP HANA nodes boot from a volume on PowerStore (boot from SAN), the required capacity for the operating system must be included in the overall capacity calculation for the SAP HANA installation. Every SAP HANA node requires approximately 100 GB capacity for the operating system, including space for the /usr/sap/ directory.

When booting from a SAN, follow the best practices that are described in the [Dell EMC Host Connectivity Guide for Linux](#).

SAP HANA persistence (data and log)

The SAP HANA in-memory database requires disk storage to:

- Maintain the persistence of the in-memory data on disk to prevent a data loss caused by a power outage. Disk storage must allow a host autofailover, where a standby SAP HANA host takes over the in-memory data of a failed worker host in scale-out installations.
- Log information about data changes (redo log).

Every SAP HANA scale-up or scale-out (worker) node requires two disk volumes or file systems to save the in-memory database on disk (data) and to keep a redo log (log). The size of these volumes or file systems depends on the anticipated total memory requirement of the database and the RAM size of the node. To help you prepare the disk sizing, the SAP HANA Storage Requirements white paper that is attached to SAP Note 1900823: [SAP HANA Storage Connector API](#) provides references to tools and documents. (Access to the note requires SAP login credentials.) The size of the data and log devices depends on the database size. Version 2.10 of the white paper provides the following formula to calculate the size of the data volume:

Option 1: If an application-specific sizing program can be used:

$$\text{Size}_{data} = 1.2 \times \text{anticipated net disk space for data}$$

where *anticipated net disk space* is the anticipated total memory requirement of the database plus an additional 20 percent free space. If the database is distributed across multiple nodes in a scale-out cluster, divide the *anticipated net disk space* by the number of SAP HANA worker nodes in the cluster. For example, if the anticipated net disk space is 2 TB and the scale-out cluster consists of four worker nodes, every node must be assigned a 616 GB data volume: (2 TB / 4 = 512 GB x 1.2 = 616 GB).

If the anticipated net disk space is unknown at the time of the storage sizing, use the RAM size of the node plus 20 percent free space for a capacity calculation of the datafile system.

- Option 2: If no application-specific sizing program is available, the recommended size of the data volume of an SAP HANA system is equal to the total memory required for the system:

$$\mathbf{Size_{data} = 1 \times RAM}$$

Note: The preceding calculations are based on in-memory data (hot data). If you are deploying SAP HANA NSE, the size of the data volume must increase with the size of the buffer cache and anticipated warm data.

The size of the log volume depends on the RAM size of the node. The SAP HANA Storage Requirements white paper that is attached to SAP Note 1900823: [SAP HANA Storage Connector API](#) provides the following formulas to calculate the minimum size of the log volume:

$$\begin{aligned} [\text{systems} \leq 512 \text{ GB}] \text{ Sizeredolog} &= 1/2 \times \text{RAM} \\ [\text{systems} > 512 \text{ GB}] \text{ Sizeredolog} &= (\text{min}) 512 \text{ GB} \end{aligned}$$

SAP HANA installation (/hana/shared/)

Before you install the SAP HANA binary, configuration, trace, and log files, every SAP HANA node must have access to a file system that is mounted under the local `/hana/shared/` mount point. In an SAP HANA scale-out cluster, a single shared file system is required and must be mounted on every node. Most SAP HANA installations use a Network File System (NFS) for this purpose. PowerStore unified storage systems can provide this file system with the NAS option. If there is a NAS or NFS option available to provide the `/hana/shared` mount point, PowerStore can be configured as block-only.

You can calculate the size of the `/hana/shared/` file system by using the formula in the SAP HANA Storage Requirements white paper, which is included as an attachment in SAP Note 1900823: [SAP HANA Storage Connector API](#). Version 2.10 of the white paper provides the following formulas.

- Single node (scale-up):
Size_{installation(single-node)} = 1 x RAM but max 1 TB
- Multinode (scale-out):
Size_{installation(scale-out)} = 1 x RAM_of_worker per four worker nodes

Backup

SAP HANA supports backup to a file system or the use of third-party tools that SAP has certified. Dell Technologies supports data protection strategies for SAP HANA backup using Dell Data Domain systems and Dell NetWorker software. While an SAP HANA backup to an NFS file system on a PowerStore all-flash or hybrid storage system is possible, backing up the SAP HANA database to the storage system where the primary persistence resides is not recommended. If you plan to back up SAP HANA to an NFS file system on a different Dell storage array, see the SAP HANA Storage Requirements white paper that is attached to SAP Note 1900823: [SAP HANA Storage Connector API](#) for details about sizing. The capacity depends on not only the data size and the frequency of change operations in the database, but also the number of backup generations that are kept on disk.

SAP HANA NSE

SAP HANA Native Storage Extension (NSE) is a general-purpose, onboard warm data store in SAP HANA that lets you manage less-frequently accessed data without fully loading it into memory. NSE integrates flash-drive based database technology with the SAP HANA in-memory database for an improved price-to-performance ratio. Using SAP HANA NSE enables the customer to achieve dramatic savings in SAP licensing costs.

- Hot data is used to store mission-critical data for real-time processing and analytics. It is retained continuously in SAP HANA memory for fast performance and is persisted to storage.
- Warm data is primarily used to store mostly read-only data that need not be accessed frequently. The data need not reside continuously in SAP HANA memory and can be unloaded to disk. It is still managed as a unified part of the SAP HANA database, transactionally consistent with hot data, participates in SAP HANA backup and system replication operations, and is stored in lower-cost stores within SAP HANA.

Note: NSE is subject to certain functional restrictions. For more information, see SAP Note 2771956: [SAP HANA NSE Functional Restrictions with HANA 2.0 SPS 04](#) and SAP Note 2927591: [SAP HANA NSE Functional Restrictions with HANA 2.0 SPS 05](#) (access requires an SAP username and password).

The following figure shows the difference between standard SAP HANA in-memory storage and the storage that is offered with NSE:

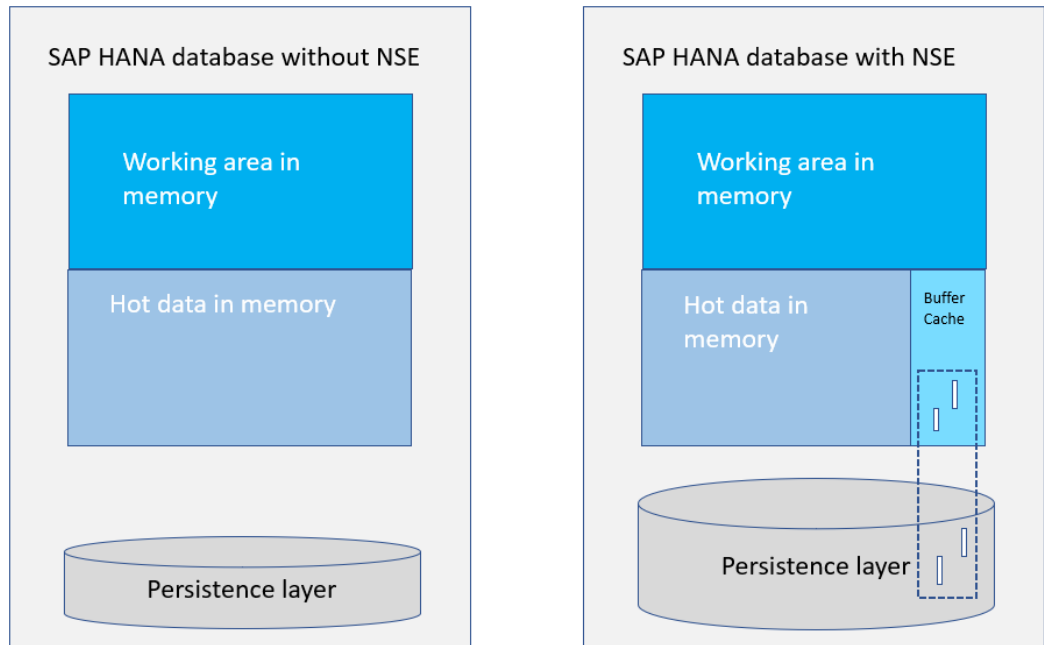


Figure 7. Comparison: standard SAP HANA database with SAP HANA database with NSE

The capacity of a standard SAP HANA database is limited by the amount of main memory. Using SAP HANA NSE, customers can bypass these limits by storing warm data on PowerStore storage. Paging operations require a relatively small amount of SAP HANA memory for the NSE buffer cache because the buffer cache can handle up to eight times the size of warm data on disk. For example, a 2 TB SAP HANA system without NSE equates to a 1 TB database in memory. With NSE and the addition of a 500 GB buffer cache, you can expand your 1 TB database to a 5 TB database consisting of 1 TB of hot data, 4 TB of warm data, and a 500 GB buffer cache to page data between memory and disk. The default buffer cache is set to 10 percent of SAP HANA memory.

While hot data is 'column loadable', residing in-memory for fast processing and loaded from disk into SAP HANA memory in columns, with SAP HANA NSE you can specify certain warm data as 'page loadable.' This data is loaded into memory page by page as required for query processing. Unlike column-loadable data, page-loadable data need not reside completely in memory.

The following figure depicts the SAP HANA database with NSE:




Figure 8. SAP HANA database with NSE

NSE reduces the memory footprint for page-loadable data. The database is partly in memory and partly on disk, as illustrated in Figure 8. The PowerStore storage system together with SAP HANA NSE can be used to substantially increase SAP HANA data capacity and reduce TCO for customers.

NSE is integrated with other SAP HANA functional layers, such as the query optimizer, query execution engine, column store, and persistence layers. For more information about SAP HANA NSE and related topics, see [SAP HANA Native Storage Extension Whitepaper](#).

SAP HANA host configuration and setup

Host configuration

Because PowerStore is a new storage type, host operating systems might not natively recognize PowerStore volumes. For optimal performance, ensure that the appropriate configuration changes have been applied to all hosts that are connected to a PowerStore. For information about the appropriate settings for a host type, see the [Dell EMC PowerStore Host Configuration Guide](#). The guide provides recommendations for:

- **MPIO settings:** Path checker and timeout values
- **iSCSI settings:** Timeout and queue depth values; disabling delayed ACK

- **FC settings:** Queue depth values
- **Network settings:** Jumbo frames and flow control
- **Unmap operations**
- **VMware ESXi claim rules**

Host connection and setup using FC SAN

When you connect an SAP HANA host to a PowerStore storage system using the FC protocol, you must connect two host bus adapter (HBA) ports supporting 16 Gb/s or 32 Gb/s link speed to the PowerStore system. Even though it is possible to use a single dual-port HBA, Dell Technologies highly recommends using two HBAs for HA and redundancy.

Connect and zone each port over the FC SAN network to two FC FE ports per storage processor on the PowerStore storage system. This setup produces two active and two standby paths per LUN, as shown in the following figure:




Figure 9. Recommended redundant FC connection

After the HBAs of the SAP HANA server have been zoned to the PowerStore FC FE ports, use the PowerStore Manager UI to create a host entry in the PowerStore system and add volumes to the host.

1. From the PowerStore Manager UI, select **Compute > Host & Host Groups** and click **+ADD HOST**.

The Add Host dialog is displayed, as shown in the following figure:

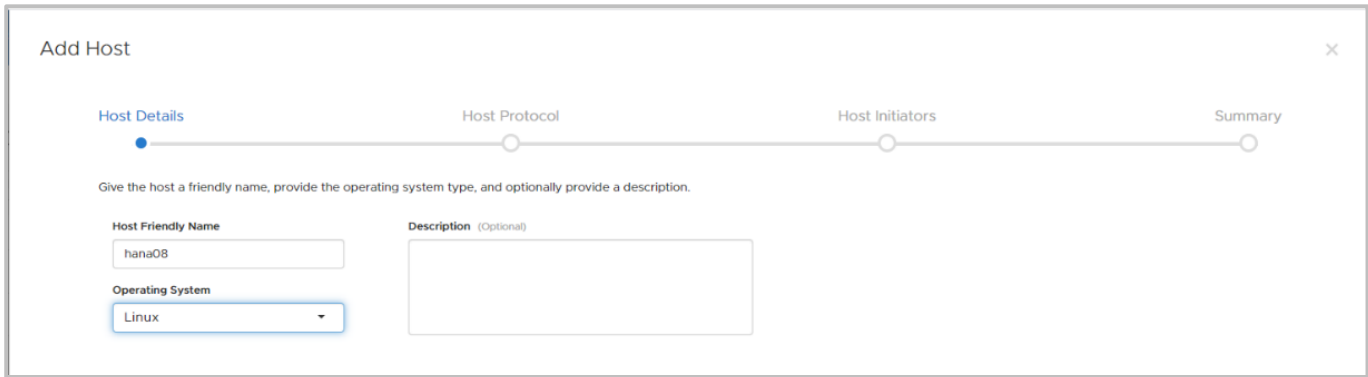


Figure 10. Adding the SAP HANA host to PowerStore

2. Add the following host details and click **Next**:

- User-friendly name
- Host operating system: Linux
- Optionally, a server description

The following page is displayed:

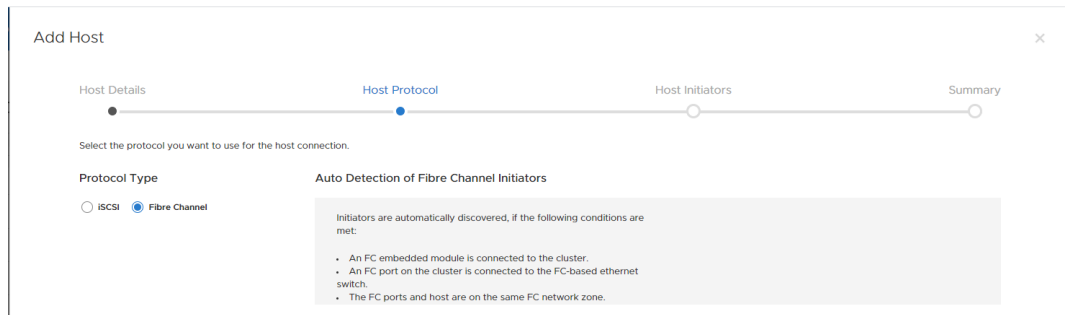


Figure 11. Selecting the host protocol

3. Select **Fibre Channel** as the host protocol type and click **Next**.

A list of automatically discovered initiators is displayed, as shown in the following figure:

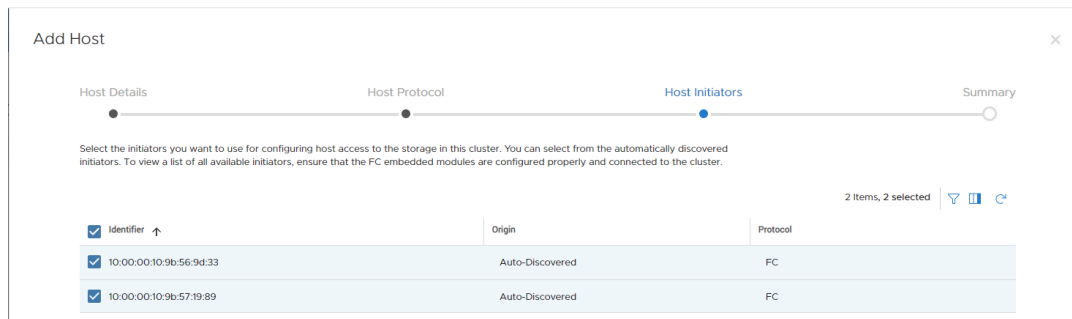


Figure 12. Selecting host initiators

4. Select the two FC initiators (WWNs) of the SAP HANA host HBAs from the list, and then click **Next**.

A Host Summary page opens, as shown in the following figure:

The screenshot shows the 'Add Host' wizard in the Summary step. A progress bar at the top indicates the steps: Host Details, Host Protocol, Host Initiators, and Summary (current step). Below the progress bar, a message reads: 'Review host details about the new host and when ready, click ADD HOST button.' The main content area is divided into three sections: Properties, Protocol, and Initiators.

Properties		Protocol		Initiators
Name	hana08	Selected Protocol	Fibre Channel	10:00:00:10:9b:56:9d:33
OS	Linux			10:00:00:10:9b:57:19:89
Description				

At the bottom right, there are three buttons: CANCEL, BACK, and ADD HOST.

Figure 13. Host summary page

5. Click **ADD HOST** to complete the process.

Volume creation and mapping for FC SAN

After the host is created, you can create and add (map) the volumes. Each host requires a data and a log volume for the SAP HANA persistence.

1. From the PowerStore Manager UI, select **Storage > Volumes > + Create**.

The Create Volumes page opens, as shown in the following figure:

The screenshot shows the 'Create Volumes' wizard in the Properties step. A progress bar at the top indicates the steps: Properties (current step), Host Mappings, and Summary. The main content area is divided into two sections: General and Additional Properties.

General

Name (Or Prefix):

Description (Optional):

Quantity: Size:

Additional Properties

Associated Volume Group (Optional):

Volume Protection Policy (Optional):

Volume Performance Policy:

At the bottom right, there are three buttons: CANCEL, FINISH, and NEXT.

Figure 14. Creating a volume

2. Provide the following information for the new volume:
 - **Name (or prefix)**

- Optionally, a description
- **Quantity:** You can create multiple volumes. When you specify a quantity, PowerStore automatically adds a number suffix to the volume name.
- **Required size**
- **Volume performance policy:** Select **High**.

Create volume groups

PowerStore can group volumes together into a volume group. A volume group is useful for operations that affect multiple volumes simultaneously, for example, when the same volumes are added (mapped) to multiple hosts.

To create a volume group:

1. On the Create Volumes page, as shown in the following figure, select the hosts on which to add the new volume and click **NEXT**.

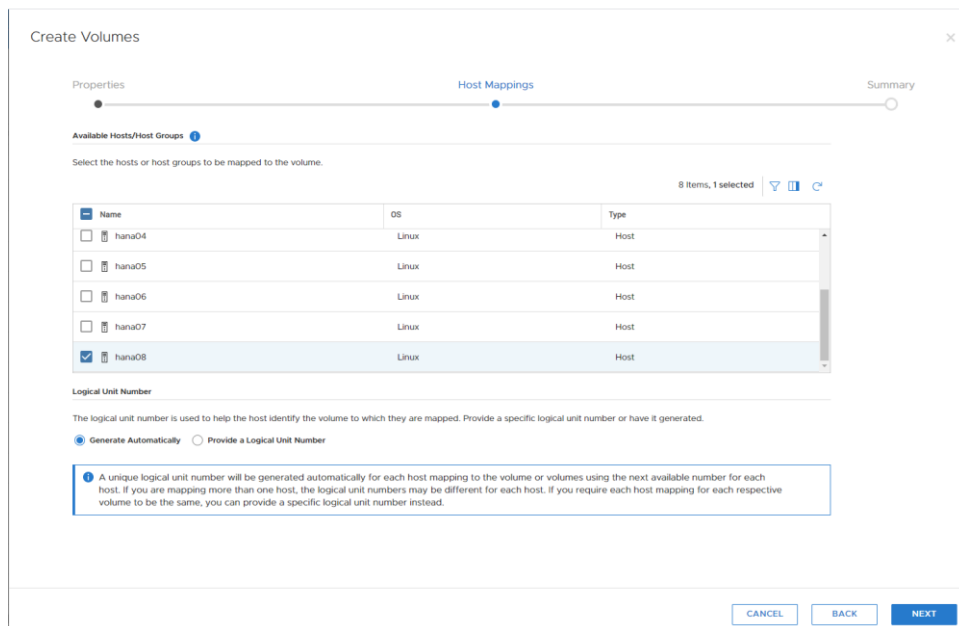


Figure 15. Mapping a volume to a host

2. Check the host and let PowerStore automatically generate the logical unit number under which the host discovers the volume. PowerStore selects the next available number.
3. Click **NEXT** to access a summary page, as shown in the following figure:

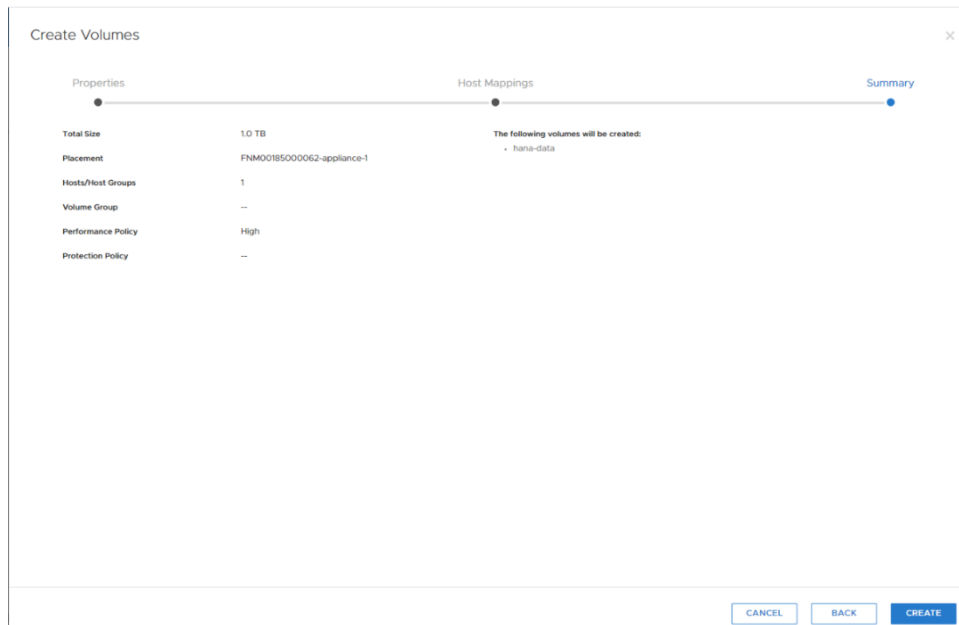


Figure 16. Create volumes summary page

4. Click **CREATE** to complete the volume creation.
5. Repeat these steps for all SAP HANA persistence volumes (data and log volumes).

Using FC SAN volumes on the SAP HANA host

After the SAP HANA data and log volumes have been created using the PowerStore Manager UI, format them on the SAP HANA host.

Scan LUNs on the host and create file systems

Follow these steps:

1. To scan the new volumes on the SAP HANA host, run the following command:

```
rescan-scsi-bus.sh
```
2. To verify that the host can see the new volumes, run:

```
fdisk -l | grep -B1 -A4 PowerStore
```

The output should look like this example:

```
hana01:~ # fdisk -l | grep -B1 -A3 Power
Disk /dev/sdr: 512 GiB, 549755813888 bytes, 1073741824 sectors
Disk model: PowerStore
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 4096 bytes / 4096 bytes
--
Disk /dev/sdq: 1 TiB, 1099511627776 bytes, 2147483648 sectors
Disk model: PowerStore
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
```

I/O size (minimum/optimal): 4096 bytes / 4096 bytes

3. After the new LUNs have been rescanned and are visible to the host, add them to multipathing. Native Linux multipathing (DM-MPIO) must be enabled on the Linux host. When multipathing is enabled, the host accesses the block LUNs over multiple paths, providing redundancy if there is a component failure.

The `/etc/multipath.conf` file controls multipathing. PowerStore default multipath values are distributed with both SLES and Red Hat Enterprise Linux and should not need to be configured. This file does not exist by default, however. If you want to view or edit the `multipath.conf` settings, create the file by running the following command:

```
multipath -t > /etc/multipath.conf
```

4. If the entry for PowerStore is not there, add the following content to the `devices` section of `/etc/multipath.conf`:

```
device {
    vendor "DellEMC"
    product "PowerStore"
    path_selector "queue-length 0"
    path_grouping_policy "group_by_prio"
    path_checker "tur"
    detect_prio "yes"
    failback "immediate"
    no_path_retry 3
    rr_min_io_rq 1
    fast_io_fail_tmo 15
    max_sectors_kb 1024
}
```

Note: For the latest PowerStore MPIO configuration settings, see the [Dell EMC PowerStore Host Configuration Guide](#).

5. Restart DM-MPIO after modifying `/etc/multipath.conf` by running:

```
systemctl restart multipathd
```

6. After the block devices are under multipath control, display the devices by running:

```
multipath -ll
```

The following output shows two devices:

```
368ccf09800dce784009f298f053f1e03 dm-4 DellEMC,PowerStore
size=1.0T features='1 queue_if_no_path' hwhandler='1 alua'
wp=rw
|+- policy='queue-length 0' prio=50 status=active
| |- 18:0:0:3 sdp 8:240 active ready running
| `-- 20:0:1:3 sdgh 131:208 active ready running
`+- policy='queue-length 0' prio=10 status=enabled
  |- 18:0:1:3 sdt 65:48 active ready running
  `-- 20:0:0:3 sdgd 131:144 active ready running
```

```
--
368ccf0980058a0131a024b9cfd171c03 dm-2 DelleMC,PowerStore
size=512G features='1 queue_if_no_path' hwhandler='1 alua'
wp=rw
|+- policy='queue-length 0' prio=50 status=active
| |- 18:0:1:1 sds 65:32 active ready running
| `-- 20:0:0:1 sdgc 131:128 active ready running
`+- policy='queue-length 0' prio=10 status=enabled
  |- 18:0:0:1 sdo 8:224 active ready running
  `-- 20:0:1:1 sdgg 131:192 active ready running
```

When the block devices are under multipath control, you can format them by using XFS and then mount them as required.

7. To format the block device, run the following (sample) command:

```
mkfs.xfs /dev/mapper/368ccf09800dce784009f298f053f1e03
```

8. When installing an SAP HANA single-node instance, automate the mounting of the SAP HANA persistent devices by adding entries such as the following to the `/etc/fstab` file:

```
/dev/mapper/368ccf09800dce784009f298f053f1e03 /hana/data xfs
rw,relatime,noquota 0 0
/dev/mapper/368ccf0980058a0131a024b9cfd171c03 /hana/log xfs
rw,relatime,noquota 0 0
```

When you install an SAP HANA multimode scale-out cluster, the SAP HANA storage connector (fcClient) mounts the devices during SAP HANA startup.

Adjust queue depth for HBAs and LUNs (optional)

The default queue depth settings for HBAs and LUNs work well with SAP HANA devices in most situations. For recommendations for queue depth to achieve optimal performance and maximum scalability, see the latest [Dell PowerStore Host Configuration Guide](#).

To display the current queue depth values, run the commands for your HBA type:

Emulex HBAs:

```
cat /sys/class/scsi_host/host20/lpfc_hba_queue_depth
cat /sys/class/scsi_host/host20/lpfc_lun_queue_depth
```

Determine the host number (in this example, `host20`) by running the `multipath -ll` command, as shown in step 6 of [Scan LUNs on the host and create file systems](#). Linux assigns a specific number to every connected HBA port. The example uses `host18` and `host20`.

QLogic HBAs:

```
cat /sys/module/qla2xxx/parameters/ql2xmaxqdepth
```

To change the LUN queue depth to a new value:

1. Create or edit the `/etc/modprobe.d/lpfc.conf` file.
2. Depending on your HBA type, insert one of the following:

Emulex HBAs: `options lpfc lpfc_lun_queue_depth=nn`

QLogic HBAs: `options qla2xxx ql2xmaxqdepth=nn`

setting `nn` to a numeric value.

3. Rebuild your Linux RAM disk or kernel image by following the instructions in the Administration Guide for your Linux distribution. For example, for SUSE Linux Enterprise Server 15, run the following commands and then reboot your system:

```
cd /boot
sync
dracut -f
```

Using NAS/NFS for SAP HANA shared directory with FC SAN deployments

In an SAP HANA scale-out cluster, a single shared file system (`/hana/shared`) is required and must be mounted on every node. Most SAP HANA installations use a Network File System (NFS) for this purpose. PowerStore unified storage systems can provide this file system with the NAS option. To use NAS for the shared file system, see [Creating the SAP HANA shared file system](#).

Host connection and setup using NAS/NFS

Each SAP HANA host requires a network connection with 25 Gb/s link speed to the dedicated storage network switches. For redundancy, the connection must have at least two NICs.

With two 25 Gb/s NICs on the SAP HANA hosts, you can optionally configure active/active interface groups on the SAP HANA hosts to use network bonding (bonds), where multiple network interfaces are aggregated into a single logical bonded interface. Hosts that are configured with bonds require MLAG and LACP on the switches. The following figure shows a sample network topology for eight SAP HANA nodes with two 25 Gb/s NICs, each connected through redundant switches to all four 25 Gb/s IP interface ports on the PowerStore system:

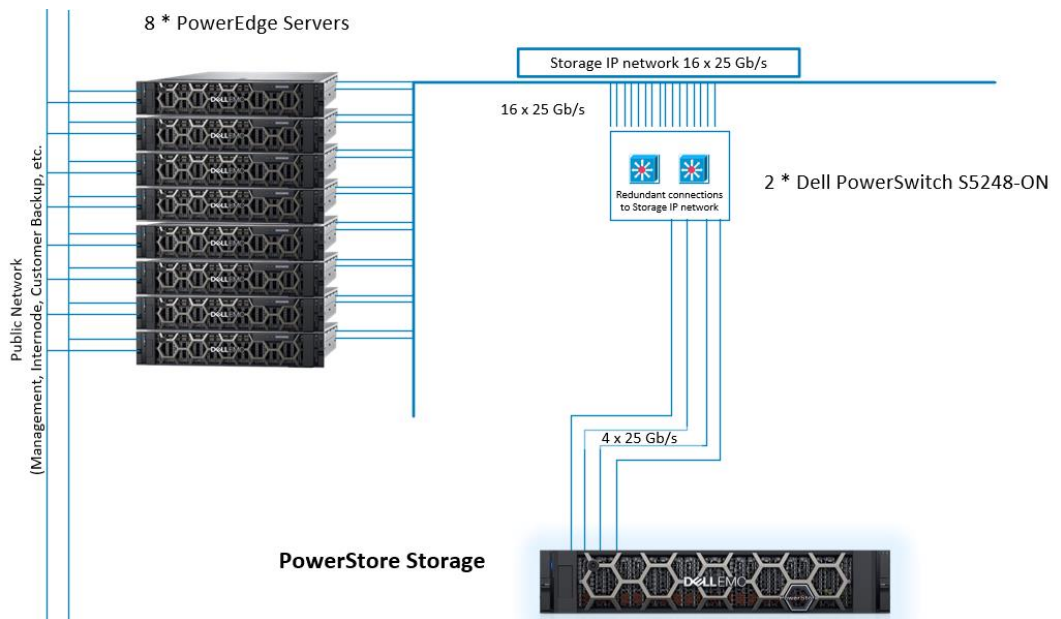


Figure 17. Storage network connections example for NAS/NFS to eight nodes

The file systems for the SAP HANA persistence are created on the PowerStore appliance, as described in [Creating NAS file systems for the SAP HANA persistence](#). Add the `_netdev` mount option when mounting NAS devices. This mount option prevents the system from attempting to mount these file systems until the network has been enabled on the system.

When you install an SAP HANA system, either as a single-node instance or a multimode scale-out cluster, you must automate the mounting of the SAP HANA persistent devices by using `/etc/fstab`. The mounts must be evenly distributed across the PowerStore IP address of the NAS servers on both PowerStore nodes within the appliance.

Creating the SAP HANA shared file system

This section is applicable for both FC SAN and NAS/NFS deployments. The SAP HANA shared file system is used for the SAP HANA binary, trace, and some configuration files. The shared file system must be mounted under the `/hana/shared` mount point. In SAP HANA single node (scale-up) systems, the file system can reside on a block volume using the XFS file system. For SAP HANA scale-out systems, this file system must be shared across all scale-out nodes, including worker and standby nodes.

Note: The `/hana/shared` file system must be shared on all hosts on SAP HANA scale-out systems in both FC SAN and NAS/NFS deployments.

Create a NAS server

To create a NAS/NFS shared file system on a PowerStore system, follow these steps in the PowerStore Manager UI. If NAS servers already exist in the UI, skip these steps.

1. Select **Storage > NAS Servers** and click **+ Create**.

The following page opens:

Figure 18. Creating a NAS server

2. Enter the required details and click **NEXT** to specify the sharing protocol, as shown in the following figure:

Figure 19. Specifying a sharing protocol

3. Choose NFSv3, NFSv4, or both depending on your requirements, and then define the UNIX directory servers and DNS. Click **NEXT** to view a summary, as shown in the following figure:

Name	Alerts	NFS Server	SMB Server	Preferred IPv4 Interface	Preferred IPv6 Interface	Current Node	Preferred Node
NAS_Server1	--	True	False	192.168.1.140	--	FNMO0185000062-appliance-1-node-A	FNMO0185000062-appliance-1-node-A
NAS_Server2	--	True	False	192.168.1.142	--	FNMO0185000062-appliance-1-node-B	FNMO0185000062-appliance-1-node-B

Figure 20. NAS servers summary page

4. Click **CREATE NAS SERVER**. For HA and load-balancing, Dell Technologies highly recommends creating one NAS server on node A and a second NAS server on node B.

Create a shared file system

To create an SAP HANA shared file system:

1. In the PowerStore Manager UI, select **Storage > File Systems** and click **+ CREATE**.

The Create File System page opens, as shown in the following figure:

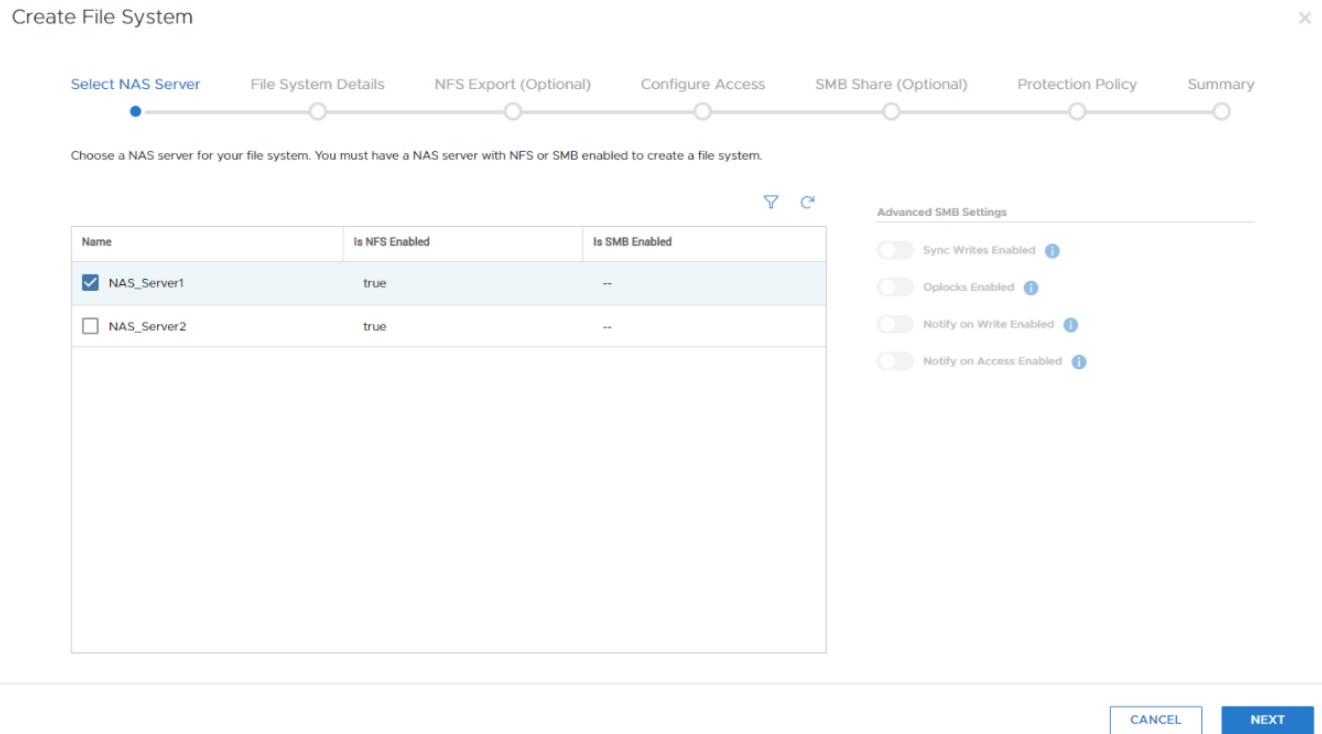


Figure 21. Selecting a NAS server

2. Select a NAS server and click **NEXT**:
3. Provide the following details for the file system:
 - Name
 - (Optional) description
 - File system size (see [SAP HANA capacity requirements](#) for information about capacity sizing)

The Create a File System Details page opens, as shown in the following figure:




Figure 22. Creating file system details

4. Click **NEXT** and then provide an NFS export name and (optionally) a description for the export, as shown in the following figure:




Figure 23. NFS Export details

- Click **NEXT** to configure access to the file system based on your security requirements, as shown in the following figure:




Figure 24. Configuring access

- Click **NEXT** to provide an optional protection policy. Select **CREATE FILE SYSTEM** to complete this step.

Mounting the HANA shared file system on the SAP HANA host

On the SAP HANA host:

- Create the `/hana/shared` mount point by running the following command:

```
mkdir -p /hana/shared
```

- Add the following entry to `/etc/fstab`:

```
<NAS Server>:/hanashared /hana/shared nfs4
rw,bg,hard,timeo=600,intr,noatime,vers=4.0,rsize=1048576,wsiz
e=1048576,lock,_netdev 0 0
```

where `<NAS Server>` is the IP address or hostname of your NAS server on the PowerStore system.

If you are using NFSv3, adjust the NFS version.

Creating NAS file systems for the SAP HANA persistence

When creating file systems for the SAP HANA persistence, follow these configuration guidelines:

1. Create two NAS servers that are dedicated to SAP HANA persistence: one NAS server on PowerStore node A and a second NAS server on PowerStore node B, as shown in the following figure. See [Create a NAS server](#) for the steps.

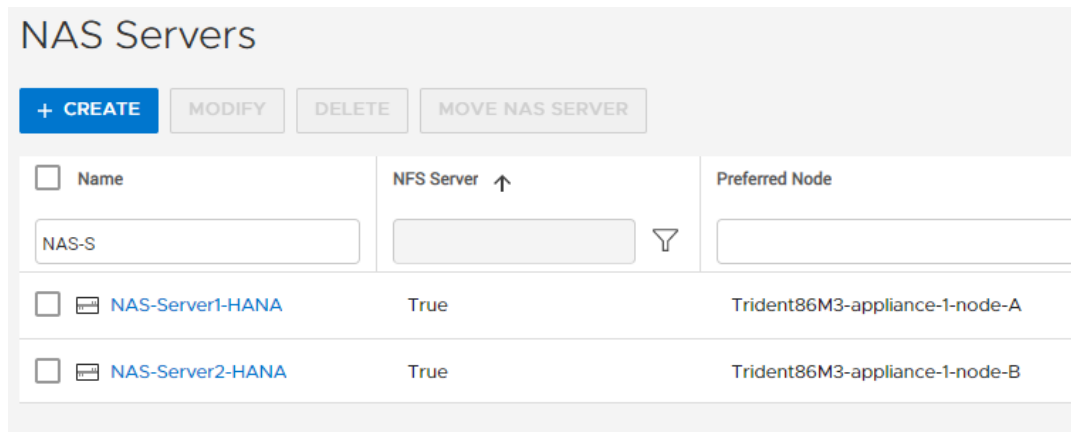


Figure 25. Dedicated NAS servers, one on node A and one on node B

2. Create two file systems for every SAP HANA node, one for **data** and the other for **log**. See [Create a shared file system](#) for the steps.
3. Distribute the file systems evenly across the two NAS servers.

For example, place **data** and **log** for the first SAP HANA node on NAS server A and place **data** and **log** for the second SAP HANA node on NAS server B. This step is important for performance and load-balancing. Continue to balance subsequent nodes across the two NAS servers. The following figure shows an example:

File system Name	NAS Server Name
hanadata1	NAS-Server1-HANA
hanalog1	NAS-Server1-HANA
hanadata2	NAS-Server2-HANA
hanalog2	NAS-Server2-HANA
hanadata3	NAS-Server1-HANA
hanalog3	NAS-Server1-HANA
hanadata4	NAS-Server2-HANA
hanalog4	NAS-Server2-HANA
hanadata5	NAS-Server1-HANA
hanalog5	NAS-Server1-HANA
hanadata6	NAS-Server2-HANA
hanalog6	NAS-Server2-HANA
hanadata7	NAS-Server1-HANA
hanalog7	NAS-Server1-HANA
hanadata8	NAS-Server2-HANA
hanalog8	NAS-Server2-HANA
hanashared	NAS-Server1-HANA

Figure 26. Data and log volumes balanced across two NAS servers

Mount the SAP HANA file system for data and log

On the SAP HANA hosts, create the mount points by running the following commands, where $x = 1, 2, 3$, and so on up to the number of hosts in the scale-out system:

```
mkdir -p /hana/data/SID/mnt0000x
mkdir -p /hana/data/SID/mnt0000x
chmod -R 777 /hana/data/SID/
chmod -R 777 /hana/log/SID/
```

Note: For SAP HANA scale-out systems, all the mount points must be created on each host.

Prepare the `/etc/fstab` file

Add the entries to `/etc/fstab`:

```
<NAS Server>:<filesystem> /hana/shared nfs4
rw,bg,hard,timeo=600,intr,noatime,vers=4.0,rsize=1048576,wsiz=104
8576,lock,_netdev 0 0
```

where `<NAS Server>` is the IP address or hostname of your NAS server on the PowerStore system. If you are using NFSv3, adjust the NFS version.

To achieve optimal performance, ensure that the SAP HANA mounts are evenly distributed across all available PowerStore front-end IP networks of the NAS servers in the `/etc/fstab` file. For more information, see [Host connection and setup using NAS/NFS](#).

In a 2+1 SAP HANA scale-out system with SID = NAS, the mount points are included in the `/etc/fstab` file on each of the SAP HANA clients including any standby hosts. The following code extract shows an example:

```
#PowerStore NAS nfs4
192.168.10.10:/hanadata /hana/data/NAS/mnt00001 nfs4
rw,bg,hard,timeo=600,intr,noatime,vers=4.0,rsize=1048576,wsiz=1048576,lock,_netdev 0 0
192.168.10.10:/hanalog /hana/log/NAS/mnt00001 nfs4
rw,bg,hard,timeo=600,intr,noatime,vers=4.0,rsize=1048576,wsiz=1048576,lock,_netdev 0 0
192.168.10.11:/hanadata2 /hana/data/NAS/mnt00002 nfs4
rw,bg,hard,timeo=600,intr,noatime,vers=4.0,rsize=1048576,wsiz=1048576,lock,_netdev 0 0
192.168.10.11:/hanalog2 /hana/log/NAS/mnt00002 nfs4
rw,bg,hard,timeo=600,intr,noatime,vers=4.0,rsize=1048576,wsiz=1048576,lock,_netdev 0 0
#PowerStore HANA shared
192.168.10.10:/hanashared /hana/shared nfs4
rw,bg,hard,timeo=600,intr,noatime,vers=4.0,rsize=1048576,wsiz=1048576,lock,_netdev 0 0
```

Figure 27. Example using NFS4 mount parameters

If preferred, use the following example for `/etc/fstab` with NFSv3 parameters.

```
# PowerStore NAS nfs3
192.168.10.10:/hanadata /hana/data/NAS/mnt00001 nfs
rw,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
192.168.10.10:/hanalog /hana/log/NAS/mnt00001 nfs
rw,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
192.168.10.11:/hanadata2 /hana/data/NAS/mnt00002 nfs
rw,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
192.168.10.11:/hanalog2 /hana/log/NAS/mnt00002 nfs
rw,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
#PowerStore HANA shared
192.168.10.10:/hanashared /hana/shared nfs
rw,vers=3,hard,timeo=600,rsize=1048576,wsiz=1048576,intr,noatime,nolock 0 0
```

Figure 28. Example using NFS3 mount parameters

The operating system command `mount -a` mounts all the directories from the PowerStore system. Run this command on each SAP HANA client.

SAP HANA installation with FC connectivity

Overview

This section describes the required settings in the SAP HANA `global.ini` file for an SAP HANA scale-out installation on PowerStore systems.

SAP HANA storage connector

Deploying an SAP HANA multinode scale-out cluster on FC-connected systems requires the SAP HANA storage connector (`fcClient` or `fcClientMpath`). The storage connector is responsible for mounting and unmounting the persistence (data and log) to the SAP HANA worker nodes.

The storage connector also writes SCSI-3 persistent reservations (PRs) to the devices. The Linux `sg_persist` or `mpathpersist` command initiates an operation that is known as I/O fencing, which ensures that only one SAP HANA worker host has simultaneous access to a set of data and log devices.

mpathpersist and fcClientMpath

With SAP HANA SP05 Rev 53, SAP released a newer version of the storage connector `fcClientMpath`, which uses the command `mpathpersist` rather than `sg_persist`. The main difference is that `mpathpersist` puts the reservations on the devices, enabling device paths to be added while the mount process is ongoing, whereas `sg_persist` puts the reservations on the device paths when mounting the devices. Paths cannot be added during the mount process, and a remount is necessary.

To use the `fcClientMpath` connector:

1. Add “reservation_key file” to the default section of the `multipath.conf` file.
The keys are added during the mount process to `/etc/multipath/prkeys`.
2. Set `ha_provider = hdb_ha.fcClientMpath` in the `global.ini` section.

**SAP HANA
global.ini file**

The storage connector API is controlled in the storage section of the SAP HANA `global.ini` file, as shown in the following code sample. The storage section of the file contains entries for the block devices. You can run the `multipath -ll` command on the SAP HANA hosts to determine the worldwide identifiers (WWIDs) of the partition entries.

```
[communication]
listeninterface = .global

[multidb]
mode = multidb
database_isolation = low
singletenant = yes

[persistence]
basepath_datavolumes = /hana/data/ANA
basepath_logvolumes = /hana/log/ANA
use_mountpoints = yes

[storage]
ha_provider = hdb_ha.fcClient
partition_*_*__prtype = 5
partition_1_data__wwid = 3600601604f804a00a79d565d90f63337
partition_1_log__wwid = 3600601604f804a00f99d565d224bd861
partition_2_data__wwid = 3600601604f804a00a89d565d2fd702d9
partition_2_log__wwid = 3600601604f804a00fa9d565d72e37054
partition_3_data__wwid = 3600601604f804a00a99d565d2e9c1d5b
partition_3_log__wwid = 3600601604f804a00fa9d565d2593cbc1

[trace]
ha_fcclient = info
```

For more information about the SAP HANA scale-out installation, the storage connector, and configuring the `global.ini` file, see the *SAP HANA Administration Guide* and the *SAP HANA Server Installation and Update Guide* on the [SAP Help Portal](#).

SAP HANA installation with NAS/NFS connectivity

Installing an SAP HANA scale-out cluster

Installing an SAP HANA scale-out cluster on PowerStore NAS systems involves:

- Configuring the `/etc/fstab` file on each SAP HANA client and mounting all the SAP HANA data, log, and shared directories from the PowerStore NAS storage to the SAP HANA nodes
- Installing an SAP HANA scale-out instance with the SAP HANA `hdblcm` command-line tool by using the NAS storage directories that you previously created

Prerequisites

The configuration example in this guide assumes that the following basic installation and configuration operations are complete on the SAP HANA nodes:

- The operating system is installed and configured according to the SAP recommendations. Our example uses SUSE Linux 15 SP1 for SAP applications.
- The directories to be used for mount points have been created for `/hana/shared` and for the data directories (`/hana/data/SID/mnt0000x`) and log directories (`/hana/log/SID/mnt0000x`) on each of the SAP HANA nodes with 777 permissions.
- All network settings and bandwidth requirements for internode communications are configured according to the SAP requirements.
- SSH keys have been exchanged between all SAP HANA nodes.
- System time synchronization has been configured through an NTP server.
- The SAP HANA installation DVD ISO file has been downloaded from the SAP website and is available on a shared file system.

Note: SAP HANA can be installed only on certified server hardware. A certified SAP HANA expert must perform the installation.

After all the file systems are mounted, you are ready to install the SAP HANA scale-out cluster. Our example uses the `hdblcm` tool to install the SAP HANA 2+1 scale-out cluster. For more information, see the [SAP HANA Studio Installation and Update Guide](#).

After the installation DVD ISO file has been extracted to a shared software-repository file system that is mounted on all hosts, begin the installation by running the following command from the extracted installation folder:

```
#!/SAPShareNew/software/SAP/HANA/hana2_rev2053/SAP_HANA_DATABASE/hdblcm
```

```
SAP HANA Lifecycle Management - SAP HANA Database
2.00.053.00.1605092543
*****
Scanning software locations...
Detected components:
```

SAP HANA Database (2.00.053.00.1605092543) in
/SAPShareNew/software/SAP/HANA/hana2_rev2053/SAP_HANA_DATABASE/ser
ver

SAP HANA Database Client (2.7.21.1611351107) in
/SAPShareNew/software/SAP/HANA/hana2_rev2053/SAP_HANA_CLIENT/clien
t

Choose an action

Index	Action	Description
1	install	Install new system
2	extract_components	Extract components
3	Exit (do nothing)	

Enter selected action index [3]: 1

SAP HANA Database version '2.00.053.00.1605092543' 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 2.7.21.1611351107

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

Enter Installation Path [/hana/shared]:

Enter Local Host Name [hana01]:

Do you want to add hosts to the system? (y/n) [n]: y

Enter comma-separated host names to add: hana02,hana03

Enter Root User Name [root]:

Collecting information from host 'hana02'...

Collecting information from host 'hana03'...

Information collected from host 'hana03'.

Information collected from host 'hana02'.

Select roles for host 'hana02':

Index	Host Role	Description
1	worker	Database Worker
2	standby	Database Standby

...

Enter comma-separated list of selected indices [1]: 1

Enter Host Failover Group for host 'hana02' [default]:

Enter Storage Partition Number for host 'hana02' [<<assign
automatically>>]:

Enter Worker Group for host 'hana02' [default]:

```

Select roles for host 'hana03':
  Index | Host Role          | Description
-----|-----|-----
---
  1     | worker            | Database Worker
  2     | standby           | Database Standby
...
Enter comma-separated list of selected indices [1]: 2
Enter Host Failover Group for host 'hana03' [default]:
Enter Worker Group for host 'hana03' [default]:
Enter SAP HANA System ID: NAS
Enter Instance Number [00]:
Enter Local Host Worker Group [default]:
  Index | System Usage | Description
-----|-----|-----
-----
  1     | production    | System is used in a production
environment
  2     | test          | System is used for testing, not
production
  3     | development   | System is used for development, not
production
  4     | custom        | System usage is neither production, test
nor development
Select System Usage / Enter Index [4]:
Enter Location of Data Volumes [/hana/data/NAS]:
Enter Location of Log Volumes [/hana/log/NAS]:
Restrict maximum memory allocation? [n]:
Enter Certificate Host Name For Host 'hana01' [hana01]:
Enter Certificate Host Name For Host 'hana02' [hana02]:
Enter Certificate Host Name For Host 'hana03' [hana03]:
Enter System Administrator (nasadm) Password:
Confirm System Administrator (nasadm) Password:
Enter System Administrator Home Directory [/usr/sap/NAS/home]:
Enter System Administrator Login Shell [/bin/sh]:
Enter System Administrator User ID [1001]:
Enter System Database User (SYSTEM) Password:
Confirm System Database User (SYSTEM) Password:
Restart system after machine reboot? [n]:
.....
Do you want to continue? (y/n): y
Installing components...
Installing SAP HANA Database...
  Preparing package ....
  Creating System...
  Extracting software...
  Installing package
  Starting SAP HANA Database system...
    All server processes started on host 'hana01' (worker).
  Importing delivery units...

```

```

Adding 2 additional hosts in parallel
Adding host 'hana03'...
Adding host 'hana02'...
  hana02: Adding host 'hana02' to instance '00'...
  hana03: Adding host 'hana03' to instance '00'...
  hana02: Starting SAP HANA Database...
  hana03: Starting SAP HANA Database...
  hana03: All server processes started on host 'hana03'
(standby).
  hana02: All server processes started on host 'hana02'
(worker).
Installing Resident hdblcm...
Installing SAP HANA Database Client...
...
Registering SAP HANA Database Components on Local Host...
  Regenerating SSL certificates...
  Deploying SAP Host Agent configurations...
Updating SAP HANA Database Instance Integration on Remote Hosts...
  Updating SAP HANA Database instance integration on host
'hana02'...
  Updating SAP HANA Database instance integration on host
'hana03'...
Creating Component List...
SAP HANA Database System installed

```

Implementing STONITH with the HA/DR provider for SAP HANA

Note: This section applies only to multihost SAP HANA scale-out instances on NAS and the host autofailover.

On failover, the database on the standby host must have read-access and write-access to the files of the failed active host. If the failed host can still write to these files, the files might become corrupted. Preventing this corruption is called fencing.

When you use shared file systems such as PowerStore NAS storage and NFSv3 or NFSv4, the STONITH method is implemented to achieve proper fencing capabilities and ensure that locks are always freed.

Note: For multihost SAP HANA scale-out instances and the host autofailover with NFSv3, the STONITH (SAP HANA HA/DR provider) implementation is mandatory. With NFSv4, a locking mechanism based on lease-time is available. The locking mechanism can be used for I/O fencing and STONITH is not required. However, STONITH can be used to speed up failover and ensure that locks are always released.

In such a setup, the storage connector API can be used for invoking the STONITH calls. During failover, the SAP HANA leading host calls the STONITH method of the custom storage connector with the hostname of the failed host as the input value.

A mapping of hostnames to management network addresses is maintained, which is used to send a reboot signal to the server through the management network. When the host restarts, it automatically starts in standby host role. The STONITH example uses the

Intelligent Platform Management Interface (IPMI) in bare-metal deployments with Dell PowerEdge servers.

Enable IPMI over LAN

For PowerEdge servers, you must configure IPMI over LAN for iDRAC to enable or disable IPMI commands over LAN channels to any external systems. If IPMI over LAN is not configured, external systems cannot communicate with the iDRAC server by using IPMI commands.

To enable IPMI over LAN:

1. In the iDRAC Settings Utility, select **Connectivity**.

The iDRAC Settings page opens, as shown in the following figure:

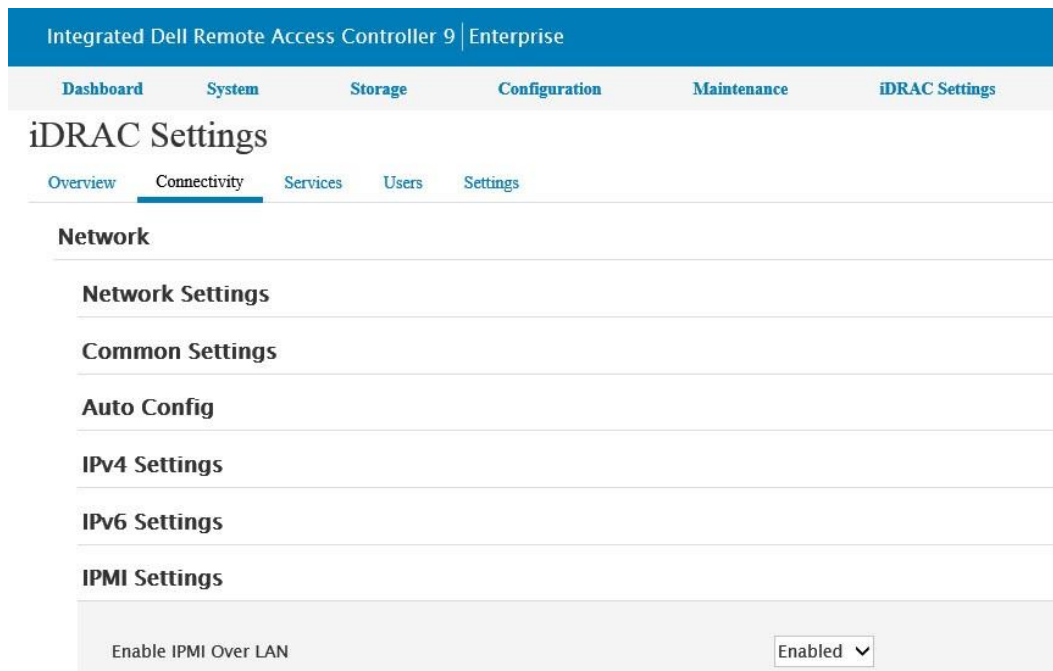


Figure 29. iDRAC Settings page

2. Enable **IPMI Over LAN**, and then click **Apply**.
3. Repeat the preceding steps for each host that is used in the SAP HANA scale-out instance.

The `/etc/hosts` file maintains a mapping of hostnames to IPMI IP addresses in STONITH using a standard naming convention:

```
# #
# IPMI mapping
#
10.230.79.85 hana01-ipmi
10.230.79.86 hana02-ipmi
10.230.79.87 hana03-ipmi
```

4. Verify that the IPMI tool is working on each SAP HANA host by running the following command as root:

```
ipmitool power status -H hana01-ipmi -U root -P xxxx
```

If IPMI is working successfully, *Chassis Power is on* is returned.

5. Set the `set-user-ID` bit for IPMItool to enable `sidadm` execution permissions by running:

```
chmod u+s /usr/bin/ipmitool
```

Create a custom HA/DR STONITH provider

To create your own HA/DR provider, perform the following steps and then add the hook method that you want to use. For more information, see the [SAP HANA Administration Guide](#).

The example here uses STONITH, which was written, tested, and validated for use with certified Dell PowerEdge servers in SAP HANA scale-out scenarios. To create the Dell recommended custom HA/DR STONITH provider, perform the following steps:

1. As `sidadm`, create a directory for the HA/DR provider. The directory must be within the `/hana/shared` storage of the SAP HANA installation but outside the `<SID>` directory structure. Our example uses the location `/hana/shared/HANA_Hooks`.
2. Copy `exe/python_support/hdb_ha_dr/HADRDummy.py` from an installed SAP HANA system to the new location. For example, copy the file and rename it to: `/hana/shared/HANA_Hooks/HA_STONITH_Hook.py`
3. Customize the contents of the new file by renaming the Python class to the name of the `HA_STONITH_Hook` file. Within the `HA_STONITH_Hook.py` we also customized `def __init__()`, `def about`, and the STONITH hook `def stonith`, as shown in the following code sample:

```
"""
```

Sample for a HA/DR hook provider.

When using your own code in here, please copy this file to location on `/hana/shared` outside the HANA installation.

This file will be overwritten with each `hdbupd` call! To configure your own changed version of this file, please add to your `global.ini` lines similar to this:

```
[ha_dr_provider_<HA_STONITH_Hook>]
provider = <HA_STONITH_Hook>
path = /hana/shared/HANA_Hooks
execution_order = 1
```

For all hooks, 0 must be returned in case of success.

```
"""
```

```
from hdb_ha_dr.client import HADRBase, Helper
import os, time
```

```

class HA_STONITH_Hook(HADRBase):

    def __init__(self, *args, **kwargs):
        # delegate construction to base class
        super(HA_STONITH_Hook, self).__init__(*args, **kwargs)

    def about(self):
        return {"provider_company" :    "Dell",
                "provider_name" :      "HA_STONITH_Hook", # provider name = class
name
                "provider_description" : "Dell STONITH HOOK for SAP HANA",
                "provider_version" :    "2.0"}

    def startup(self, hostname, storage_partition, system_replication_mode, **kwargs):
        self.tracer.debug("enter startup hook; %s" % locals())
        self.tracer.debug(self.config.toString())

        self.tracer.info("leave startup hook")
        return 0

    def shutdown(self, hostname, storage_partition, system_replication_mode,
**kwargs):
        self.tracer.debug("enter shutdown hook; %s" % locals())
        self.tracer.debug(self.config.toString())

        self.tracer.info("leave shutdown hook")
        return 0

    def failover(self, hostname, storage_partition, system_replication_mode, **kwargs):
        self.tracer.debug("enter failover hook; %s" % locals())
        self.tracer.debug(self.config.toString())

        self.tracer.info("leave failover hook")
        return 0

    def stonith(self, failingHost, **kwargs):
        self.tracer.debug("enter HANA HA stonith hook; %s" % locals())
        self.tracer.debug(self.config.toString())

        self.tracer.info( "Stonith - rebooting failing host %s" % failingHost)
        ipmi_host = "%s-ipmi" % failingHost

        power_cycle = "ipmitool power cycle -l lanplus -H %s -U root -P XXXXXXXX " %
ipmi_host
        power_on = "ipmitool power on -l lanplus -H %s -U root -P XXXXXXXX " %
ipmi_host
        rc = os.system(power_cycle)
        time.sleep(10)

```



```

    if rc == 0:
        msg = "Power cycle successfully executed to the failed host %s" %
failingHost
        self.tracer.info(msg)
        rc = 0
    elif rc !=0:
        msg = "failed to power cycle %s, will try again" % failingHost
        self.tracer.info(msg)
        rc = os.system(power_on)
        time.sleep(10)
        if rc == 0:
            msg = "Successfully powered on %s" % failingHost
            self.tracer.info(msg)
            rc = 0
        elif rc !=0:
            msg = "unable to power cycle %s - Please CHECK" % failingHost
            self.tracer.info(msg)
            return 1
    self.tracer.info("leaving HANA HA stonith hook")
    return rc

def preTakeover(self, isForce, **kwargs):
    """Pre takeover hook."""
    self.tracer.info("%s.preTakeover method called with isForce=%s" %
(self.__class__.__name__, isForce))

    if not isForce:
        # run pre takeover code
        # run pre-check, return != 0 in case of error => will abort takeover
        return 0
    else:
        # possible force-takeover only code
        # usually nothing to do here
        return 0

def postTakeover(self, rc, **kwargs):
    """Post takeover hook."""
    self.tracer.info("%s.postTakeover method called with rc=%s" %
(self.__class__.__name__, rc))

    if rc == 0:
        # normal takeover succeeded
        return 0
    elif rc == 1:
        # waiting for force takeover
        return 0
    elif rc == 2:

```

```

    # error, something went wrong
    return 0

def srConnectionChanged(self, parameters, **kwargs):
    self.tracer.debug("enter srConnectionChanged hook; %s" % locals())

    # Access to parameters dictionary
    hostname = parameters['hostname']
    port = parameters['port']
    volume = parameters['volume']
    serviceName = parameters['service_name']
    database = parameters['database']
    status = parameters['status']
    databaseStatus = parameters['database_status']
    systemStatus = parameters['system_status']
    timestamp = parameters['timestamp']
    isInSync = parameters['is_in_sync']
    reason = parameters['reason']
    siteName = parameters['siteName']

    self.tracer.info("leave srConnectionChanged hook")
    return 0

def srReadAccessInitialized(self, parameters, **kwargs):
    self.tracer.debug("enter srReadAccessInitialized hook; %s" % locals())

    # Access to parameters dictionary
    database = parameters['last_initialized_database']
    databasesNoReadAccess =
parameters['databases_without_read_access_initialized']
    databasesReadAccess = parameters['databases_with_read_access_initialized']
    timestamp = parameters['timestamp']
    allDatabasesInitialized = parameters['all_databases_initialized']

    self.tracer.info("leave srReadAccessInitialized hook")
    return 0

def srServiceStateChanged(self, parameters, **kwargs):
    self.tracer.debug("enter srServiceStateChanged hook; %s" % locals())

    # Access to parameters dictionary
    hostname = parameters['hostname']
    service = parameters['service_name']
    port = parameters['service_port']
    status = parameters['service_status']
    previousStatus = parameters['service_previous_status']
    timestamp = parameters['timestamp']

```

```

daemonStatus = parameters['daemon_status']
databaseId = parameters['database_id']
databaseName = parameters['database_name']
databaseStatus = parameters['database_status']

self.tracer.info("leave srServiceStateChanged hook")
return 0

def srSecondaryUnregistered(self, parameters, **kwargs):
    self.tracer.debug("enter srSecondaryUnregistered hook; %s" % locals())

    # Access to parameters dictionary
    siteName = parameters['site_name']
    siteId = parameters['site_id']
    reason = parameters['reason']

    self.tracer.info("leave srSecondaryUnregistered hook")
    return 0

```

Install the HA/DR provider script

You can add, configure, and monitor your custom provider scripts in the SAP HANA Cockpit.

1. After the HA/DR provider script is created, install the script on an SAP HANA system by adding an `ha_dr_provider_<classname>` section with the following parameters to the `global.ini` file:

`provider`: Class name

`path`: Location of the script

`execution_order`: Ordering of the HA/DR provider (if there is more than one; this is a number from 1 through 99) For example, add the following details to the `global.ini` file:

```

[ha_dr_provider_<HA_STONITH_Hook>]
provider = HA_STONITH_Hook
path = /hana/shared/HANA_Hooks
execution_order = 50

```

2. Using the SAP HANA Cockpit, in your SAP HANA database, select **Database Administration > Manage system configuration**. You can add, configure, and monitor the HA/DR provider information, as shown in the following figure:

Section	Parameter	Layer	Specific Value
<code>global.ini</code>			
[ha_dr_provider_<H A_STONITH_Hook>	+	execution_order	50
		path	/hana/shared/HANA_Hooks
		provider	HA_STONITH_Hook

Figure 30. SAP HANA Cockpit: `ha_dr_provider` section in `global.ini`

Verify the installation of the HA/DR provider script HA_STONITH.Hook.py

All scripts are loaded during the startup phase of the name server. You can monitor the name server trace file while general information is collected about the `ha_dr_provider` and return codes.

Perform host autofailovers to ensure that the failovers work as expected and that STONITH has been implemented on the failed host. The following figure shows an example of output from the name server trace file following a host autofailover and successful implementation of STONITH:

```
[4183][-1][-1] 2022-09-04 23:27:22.124705 i failover      TREXNameServer.cpp(05680) : failover from failed host hana02 to host hana03 started
[2494][-1][-1] 2022-09-04 23:27:22.124985 i ha_dr_provider PythonProxyImpl.cpp(00970) : calling HA/DR provider HA_STONITH_Hook.stonith(failing_host=hana02)
[2494][-1][-1] 2022-09-04 23:27:22.125179 i ha_dr_HA_STONITH HA_STONITH_Hook.py(00057) : stonith - rebooting failing host hana02
[2494][-1][-1] 2022-09-04 23:27:32.459448 i ha_dr_HA_STONITH HA_STONITH_Hook.py(00068) : Power cycle successfully executed to the failed host hana02
[2494][-1][-1] 2022-09-04 23:27:32.459478 i ha_dr_HA_STONITH HA_STONITH_Hook.py(00083) : leaving HANA HA stonith hook
[4183][-1][-1] 2022-09-04 23:27:32.614814 i failover      TREXNameServer.cpp(05737) : failover from failed host hana02 to host hana03 succeeded
```

Figure 31. Ha_dr_provider output from the leading name server trace file

Post-installation configuration

File I/O optimization after the SAP HANA installation

The base layer of SAP HANA provides two file I/O interfaces:

- **SimpleFile:** Used for small, simple I/O requests on configuration files, traces, and so on. This interface uses lightweight, platform-independent wrappers around system calls.
- **FileFactory & File:** Used for huge, complex streams of I/O requests on the data and log volumes and for backup and recovery. This interface uses synchronous and asynchronous I/O operations.

You can configure the SAP HANA file I/O layer to optimize file I/O for a specific file system and storage system.

Note: The Linux XFS file system is used on all Dell storage volumes for the SAP HANA persistence.

SAP HANA 1.0

After the initial SAP HANA installation is complete, set the parameters by running the SAP HANA `hdbparam` command as `<sid>adm` in the Linux shell:

```
su - <sid>adm
hdbparam -p          # lists current parameter setting

hdbparam --paramset fileio [DATA].max_parallel_io_requests=128
hdbparam --paramset fileio [LOG].max_parallel_io_requests=128
hdbparam --paramset fileio [LOG].num_completion_queues=4
hdbparam --paramset fileio [LOG]. num_submit_queues=8
```

SAP HANA 2.0

Effective from SAP HANA 2.0, the `hdbparam` command-line tool has been deprecated. Instead, the parameters are defined as normal parameters in

`global.ini` > [*fileio*]. Using the SQL console in the SAP HANA Cockpit, set the parameters in the `global.ini` file by running the following commands:

```
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio','max_parallel_io_requests[DATA]') = '128';
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio','max_parallel_io_requests[LOG]') = '128';
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio','num_completion_queues[LOG]') = '4';
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio','num_submit_queues[LOG]') = '8' WITH RECONFIGURE
```

The following figure shows what the `fileio` section of `global.ini` looks like in the Cockpit after the parameters are set:

fileio		
async_read_submit	on	
async_write_submit_active	on	
async_write_submit_blocks	all	
max_parallel_io_requests	64	
max_parallel_io_requests[data]		● 128
max_parallel_io_requests[log]		● 128
max_submit_batch_size	64	
min_submit_batch_size	16	
num_completion_queues	1	
num_completion_queues[log]		● 4
num_submit_queues	1	
num_submit_queues[log]		● 8
size_kernel_io_queue	512	

Figure 32. SAP HANA parameters in `global.ini`

All other parameters are set by default during installation. For more information, see SAP Note [2399079](#)—*Elimination of `hdbparam` in SAP HANA 2* (access requires SAP user credentials).

Configuring SAP HANA NSE

SAP HANA NSE uses the data volume that the main database is using. While no special configuration steps are necessary for the NSE persistence layer, sizing must consider the additional capacity for the feature.

When you use the SAP HANA NSE feature, a portion of DRAM is used as a buffer cache to dynamically load paged data from the persistence (data volume). By default, the size of this buffer cache is 10 percent of the total main memory of the system. For more information, see “SAP HANA Buffer Cache” in the [SAP HANA Administration Guide](#).

SAP recommends not exceeding a ratio of 1:8 for the total amount of hot data: warm data that is handled by SAP HANA NSE and the buffer cache. When using SAP HANA NSE, you can store data in the warm tier in the following specific granularities:

- Tables
- Columns
- Partitions

Data location handling is built into the SAP HANA Data Definition Language (DDL). Manage the configuration by using the SAP HANA CLI SQL client `hdbsql` or the SQL editor in SAP HANA Studio or SAP HANA Cockpit.

To create a table using SAP HANA NSE (the warm tier), run the following DDL command:

```
CREATE T_DELL_NSE (id INT, name VARCHAR(100)) PAGE LOADABLE;
```

To create a specific column in the warm tier, run:

```
CREATE T_DELL_NSE (id INT PAGE LOADABLE, name VARCHAR(100));
```

To move an entire existing table to the warm tier, run:

```
ALTER TABLE T_DELL_NSE PAGE LOADABLE CASCADE
```

To move a specific column to the warm tier, run:

```
ALTER TABLE T_DELL_NSE ALTER (name VARCHAR(100) PAGE LOADABLE);
```

[SAP HANA NSE Advisor](#), which is based on real-time statistics from an existing SAP HANA database, provides recommendations for which data to move from the hot (in-memory) tier to the warm tier (SAP HANA NSE). Use the SAP HANA NSE Advisor information for guidance regarding the amount of data to move to the warm tier on an existing system.

For more information about SAP HANA NSE, SAP HANA NSE Data Sizing, and related topics, see the [SAP Help Portal](#).

Conclusion

Summary

Using SAP HANA in TDI deployments with DellPowerStore enterprise storage systems provides many benefits, including reduced hardware and operational costs, lower risks, and greater hardware vendor flexibility.

You can easily transition to this new architecture and rely on Dell Technologies Professional Services to minimize risk.

SAP certifies PowerStore storage systems for both the FC SAN and NAS/NFS protocols, giving the customer the choice of block or file for SAP HANA. You can use these storage systems for SAP HANA production and nonproduction installations and for single-node and scale-out systems.

Findings

The Dell engineering tests with SAP HANA on PowerStore storage systems showed that:

- SAP HANA production installations on PowerStore meet the SAP HANA storage performance KPIs when the configuration and scalability rules that are described in this guide are applied.
- PowerStore can provide unified (block and file) or block-only certified storage for FC SAN and NAS/NFS deployments, including shared files systems such as `/hana/shared` or `/sapmnt/SID`, even for FC SAN deployments.
- SAP HANA performance and scalability are the same irrespective of whether PowerStore is configured with single parity (single tolerance R5) or dual parity (double tolerance R6).
- Using PowerStore clusters for block storage with up to four appliances enables you to linearly scale the number of FC SAN-connected SAP HANA nodes up to 80 nodes.
- PowerStore always-on inline data reduction provides significant storage efficiency and reduces the capacity requirement. Specific savings might vary depending on use case scenarios.

References

Dell Technologies documentation

The following documentation provides additional information. Access depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- [PowerStore Info Hub](#)
- [Dell PowerStore: Best Practices Guide](#)
- [Host Connectivity Guide for Linux](#)

For additional documentation that is relevant to this solution, see the [Dell Technologies Solutions Info Hub for SAP](#).

SAP HANA documentation

The following documentation provides relevant information:

- [SAP HANA Tailored Data Center Integration – Overview](#)[SAP HANA Tailored Data Center Integration Frequently Asked Questions](#)
- [SAP HANA Native Storage Extension Whitepaper](#)

For additional documentation that is relevant to this solution, see the [SAP Help Portal](#).

SAP Notes

The following SAP Notes are available at the [SAP Knowledge Base](#). Access requires an SAP username and password.

- SAP Note 1943937 - Hardware Configuration Check Tool – Central Note
- SAP Note 1969700: SQL statement collection for SAP HANA
- SAP Note 1999930: SAP HANA I/O Analysis
- SAP Note 2399079: Elimination of hdbparam in HANA 2.0