

Storage Configuration Best Practices for SAP HANA TDI on Dell EMC PowerMax Arrays

PowerMax 2000 and 8000 Storage Arrays

August 2020

H17290.2

Validation Guide

Abstract

This validation guide describes storage configuration best practices for SAP HANA in Tailored Data Center (TDI) deployments on Dell EMC PowerMax storage arrays. The solution enables customers to use PowerMax arrays for SAP HANA TDI deployments in a fully supported environment with existing data center infrastructures.

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 © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA 08/20 Validation Guide H17290.2.

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

Contents

Chapter 1	Introduction	5
	Solution overview	6
	Key benefits	6
	Document purpose.....	7
	Audience.....	7
	We value your feedback.....	7
Chapter 2	Technology Overview	8
	SAP HANA deployment models	9
	SAP HANA database	10
	Dell EMC PowerMax storage	10
Chapter 3	Design Principles and Recommendations	14
	Introduction	15
	SAP HANA capacity requirements	15
	SAN network considerations	17
	SAP HANA I/O patterns	18
	SAP HANA shared file system on a PowerMax array.....	19
	PowerMax scalability for SAP HANA.....	19
	Masking views.....	20
	PowerMaxOS service levels and competing workloads.....	20
	SAP HANA NSE	23
Chapter 4	Configuration and Installation	25
	Configuration and installation overview	26
	Configuring the PowerMax array	26
	Preparing the SAP HANA nodes	32
	Installing the SAP HANA scale-out cluster	34
	Performing post-installation configuration.....	37
Chapter 5	Data Domain Backup and Recovery for SAP HANA	40
	Data Domain overview	41
	Configuring Data Domain for SAP HANA	41
	Backing up the database from SAP HANA studio	43
	Recovering the database with SAP HANA studio	44
Chapter 6	Conclusion	47
	Summary	48

Contents

Findings	48
Chapter 7 References	49
Dell Technologies documentation	50
SAP HANA documentation.....	50

Chapter 1 Introduction

This chapter presents the following topics:

Solution overview	6
Key benefits	6
Document purpose	7
Audience	7
We value your feedback	7

Solution overview

SAP HANA is an in-memory data platform that can be deployed on-premises or in the cloud. Organizations use the SAP HANA platform to analyze large volumes of data and develop and deploy applications in real time. The Dell EMC PowerMax enterprise storage family offers SAP enterprise customers trusted data services, along with simplicity and capacity, while delivering high performance and consistently low latency with guaranteed service levels.

This solution uses SAP HANA in a TDI deployment scenario on the PowerMax enterprise storage platform. The PowerMax family systems meet SAP HANA performance and functional requirements. They are certified by SAP, which means that customers can use PowerMax arrays for SAP HANA TDI deployments in a fully supported environment using their existing data center infrastructures.

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

Key benefits

The TDI solution increases server and network vendor flexibility while reducing hardware and operational costs. Customers using SAP HANA TDI on PowerMax arrays can:

- Integrate SAP HANA into an existing data center.
- Use Fibre Channel (FC) storage area network (SAN) connectivity from the SAP HANA hosts to the storage array.
- Avoid the significant risks and costs associated with operational change by using existing operational processes, skills, and tools.
- Transition easily from an appliance-based model to the PowerMax based TDI architecture while relying on Dell EMC Professional Services to minimize risk.
- Use PowerMax shared enterprise storage to rely on already-available, multisite concepts and benefit from established automation and operations processes.
- Use the significant performance and scale benefits of PowerMax arrays to obtain real-time insights across the business.
- Use flash and SCM drives for the SAP HANA persistence and benefit from reduced SAP HANA startup, host auto-failover, and backup times.
- Expect storage efficiencies and capacity savings with PowerMax advanced data reduction with inline compression and deduplication.

Document purpose

This validation guide describes a solution that uses the SAP HANA platform in a TDI deployment scenario on PowerMax enterprise storage arrays. The guide provides configuration recommendations based on SAP requirements for HA and KPIs for data throughput and latency for the TDI deployment. Topics that the guide addresses include:

- Introduction to the key technologies in the solution
- Best practices and storage design principles for PowerMax storage with SAP HANA
- Configuration steps for deploying the SAP HANA scale-out database on PowerMax storage systems
- Example of an SAP HANA scale-out installation using PowerMax storage
- Backup and recovery for SAP HANA with the Dell EMC PowerProtect DD series appliance

Note: This validation guide describes SAP HANA TDI deployments in physical environments. If you plan to use SAP HANA in VMware vSphere virtualized environments, see the *VMware Virtualized SAP HANA with EMC Storage Solution Guide*.

Audience

This document is for system integrators, system or storage administrators, customers, partners, and members of Dell EMC Professional Services who want to configure a PowerMax storage array for use in an SAP HANA TDI environment.

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the Dell Technologies Solutions team by [email](#) or provide your comments by completing our [documentation survey](#).

Authors: Donagh Keeshan, Dat Nguyen, Benjamin Simroth

Note: For links to additional documentation for this solution, see the [Dell Technologies Solutions Info Hub for SAP](#).

Chapter 2 Technology Overview

This chapter presents the following topics:

SAP HANA deployment models	9
SAP HANA database.....	10
Dell EMC PowerMax storage	10

SAP HANA deployment models

Introduction

SAP HANA is an in-memory data platform that can be deployed on-premises (locally) or in the cloud. Organizations use the SAP HANA platform to analyze large volumes of data and develop and deploy applications in real time. The SAP HANA database is at the core of this data platform.

The SAP HANA system combines SAP software components that are optimized on proven, SAP partner-provided hardware. Two models are available for on-premises deployment, as shown in the following figure:

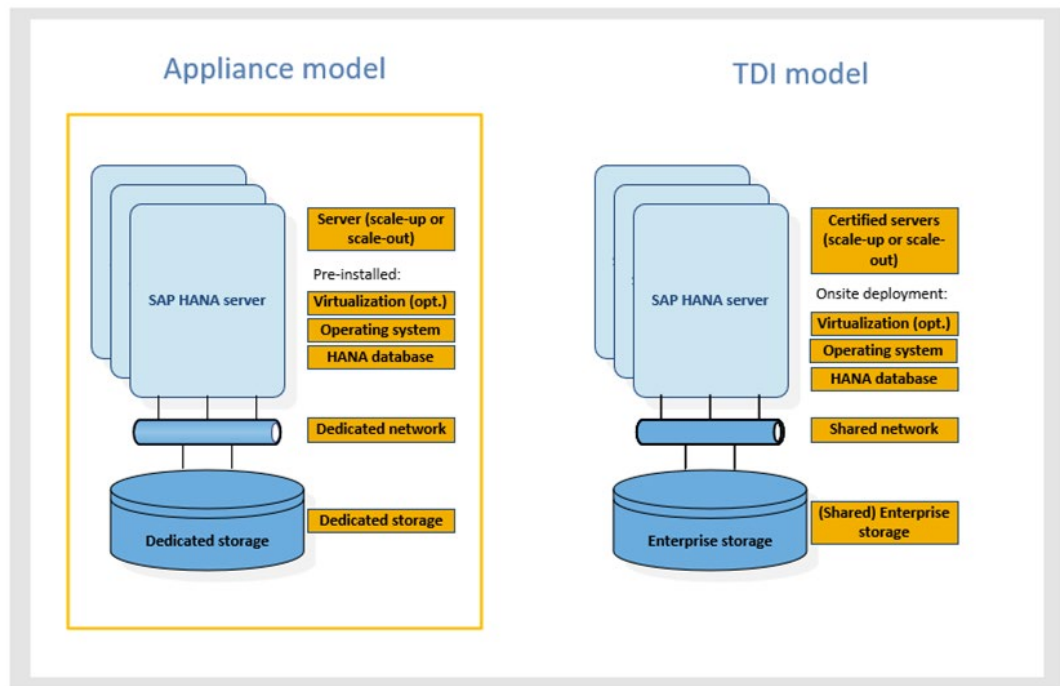


Figure 1. SAP HANA appliance model and TDI model comparison

Appliance model

An SAP HANA appliance includes integrated storage, compute, and network components by default. The appliance is certified by SAP, built by one of the SAP HANA hardware partners, and shipped to customers with all its software components preinstalled, including the operating systems and SAP HANA software.

The SAP HANA appliance model presents customers with the following limitations:

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

TDI model

The SAP HANA servers in a TDI model must be certified by SAP and meet the SAP HANA requirements, but the network and storage components, including arrays, can be shared in customer environments. Customers can integrate SAP HANA seamlessly into existing data center operations such as disaster recovery (DR), data protection, monitoring, and management, reducing the time-to-value, costs, and risk of an overall 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

Introduction

SAP HANA is an in-memory database: The data is kept in the RAM of one or more SAP HANA worker hosts, and 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 provides the ability to restore the SAP HANA database to its most recent committed state in case of failure. The log captures all changes by database transactions (redo logs), and data and undo log information is automatically saved to disk at regular savepoints.

Scale-up compared to scale-out deployments

As SAP-certified enterprise storage arrays for SAP HANA, PowerStore arrays 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 online transaction processing (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. A worker host is an active component that accepts and processes database requests, while a standby host is a passive component that has all database services running but no data in RAM. A standby host waits for a failure of a worker host so that it can take over its role, a process known as host auto-failover. Because the in-memory capacity in these deployments can be high, scale-out SAP HANA clusters are perfectly suited for online analytical processing (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, SAP requires a site-specific certification.

Dell EMC PowerMax storage

Introduction

The [PowerMax family](#) is the first Dell EMC data storage product to fully use NVMe technology for customer application data. The PowerMax array is built using an end-to-end NVMe storage architecture, allowing it to reach unprecedented I/O densities and

performance by eliminating the flash media choke points found with traditional SAS and SATA interfaces. With the PowerMax array, customers can deploy innovative applications in the areas of real-time analytics, machine learning, and big data, which demand lower latency and higher performance than legacy all-flash offerings provided.

The Dell EMC PowerMax family consists of two models—the PowerMax 2000 and the flagship PowerMax 8000—as shown in the following figure. The PowerMax 2000 is designed to provide Dell EMC customers with efficiency and maximum flexibility in a 20U footprint. The PowerMax 8000 is designed for massive scale, performance, and IOPS density, all within a one- or two-floor-tile footprint.



Figure 2. Dell EMC PowerMax product family

Both PowerMax models have at their foundation the Dynamic Virtual Matrix architecture and PowerMaxOS, which was created for the NVMe platform. PowerMax systems are true all-flash arrays—that is, they are targeted to meet the storage capacity and performance requirements of the all-flash enterprise data center. The feature-rich PowerMax products are designed to take advantage of ultrahigh-performing storage class memory (SCM) and high-capacity standard flash drives to create the densest storage configuration possible.

The PowerMax platforms offer customers these main benefits:

- I/O density and predictable performance. The PowerMax system is capable of predictably delivering approximately 187 K IOPS per rack unit (U) or up to 15M IOPS in a two-rack system (two floor tiles), regardless of workload and storage capacity utilization.
- Future-proof design. The PowerMax NVMe design can seamlessly implement current and future SCM and NAND flash drive capacities and allow for hosts to connect through current Gen 6 and future Gen 7 high-speed FC-NVMe SANs.
- Efficient workload consolidation for block, file, and mainframe workloads on a single platform.

- Ultrahigh reliability, availability, and serviceability for the enterprise.
- Excellent data services, including advanced data reduction using inline deduplication and compression.

Note: For more information about the Dell EMC PowerMax family, see the [Dell EMC PowerMax: Family Overview White Paper](#).

System specifications

The following table shows the system specifications for the PowerMax 2000 and 8000 arrays:

Table 1. PowerMax system specifications

PowerMax model	Bricks per model	Brick CPU	Brick cache	Maximum cache	Maximum number of front-end modules per Brick	Maximum number of drives	Maximum capacity
2000	1–2	12 core, 2.5 GHz Intel Broadwell (2 per director) Total: 48 cores per Brick	512 GB, 1 TB, and 2 TB (DDR4)	4 TB	8 (32 total front-end ports per Brick) Total: 64 front-end ports per system	96	1 PB
8000	1–8	18 core, 2.8 GHz Intel Broadwell (2 per director) Total: 72 cores per Brick	1 TB and 2 TB (DDR4)	16 TB	6 (24 total front-end ports per Brick—open systems/mixed) Total: 256 front-end ports per system	288	4 PB

Expandable modular architecture

PowerMax configurations consist of modular building blocks called PowerMax Bricks. The modular Brick architecture reduces complexity and allows for easier system configuration and deployment. This architecture also allows the system to scale while continuing to deliver predictable high performance.

Two types of Bricks are available for PowerMax configurations:

- The open systems Brick supports configurations with FC, FC-NVMe, and iSCSI connectivity with FBA device formatting. The open systems Brick can also be configured for file storage using embedded NAS.
- The mainframe zBrick supports configurations with FICON connectivity and CKD device formatting. The initial Brick system includes a single engine consisting of two directors, two system power supplies (SPS), and two 24-slot 2.5 in. NVMe disk

array enclosures (DAE24) that are preconfigured with an initial total usable capacity.

The Brick concept allows PowerMax platforms to scale up and scale out. Customers can scale up by incrementally adding Flash Capacity Packs. For the PowerMax 8000, each Flash Capacity Pack has 13 TB of usable storage. For the PowerMax 2000, each Flash Capacity Pack has either 11 TB or 13 TB of usable storage, depending on the RAID protection type selected. PowerMax systems scale out by aggregating up to two Bricks for the PowerMax 2000 and up to eight Bricks for the PowerMax 8000. Scaling out a PowerMax system by adding Bricks produces a predictable, linear performance improvement regardless of the workload.

Note: For detailed information about available PowerMax Brick configurations, see the [PowerMax Family Specification Sheet](#).

SCM drives on PowerMax systems

PowerMax 2000 and 8000 systems support storage class memory (SCM) drives. These drives are powered by dual-port Intel Optane technology, together with machine learning algorithms, for data placement when both NVMe and SCM drives are used. SCM drives are high-performance drives that have a significantly lower latency than NVMe flash drives, thus increasing response times.

A PowerMax model that contains both types of drives effectively has two storage tiers—the higher-performance SCM drives and the NVMe drives. Automated data placement takes advantage of the performance difference to optimize access to data that is frequently accessed. Automated data placement can also help to optimize access to storage groups that are used for workloads such as SAP HANA that have higher-priority service levels.

For optimum cost per performance, we recommend that the total usable capacity (TBu) of SCM Tier 0 is between 3 and 12 percent of the required effective capacity (TBe) of the system.

Advanced data reduction using inline compression and deduplication

PowerMax arrays use the PowerMax Adaptive Compression Engine (ACE) for inline hardware compression. The ACE data-reduction method provides the highest space-saving capability with a negligible performance impact. PowerMax technology uses inline hardware-based data deduplication, which identifies repeated data patterns on the array and stores the repeated data only once. This functionality prevents the consumption of critical core resources on the PowerMax system and limits performance impact.

Note: For information about PowerMax data services and how to configure and use SCM drives on PowerMax arrays, see the [Dell EMC PowerMax: Family Overview White Paper](#).

Chapter 3 Design Principles and Recommendations

This chapter presents the following topics:

- Introduction.....15**
- SAP HANA capacity requirements.....15**
- SAN network considerations17**
- SAP HANA I/O patterns18**
- SAP HANA shared file system on a PowerMax array19**
- PowerMax scalability for SAP HANA19**
- Masking views.....20**
- PowerMaxOS service levels and competing workloads.....20**
- SAP HANA NSE.....23**

Introduction

SAP HANA production systems in TDI environments must meet the SAP performance KPIs. The following topics describe system requirements, general considerations, and best-practice recommendations for connecting SAP HANA to PowerMax arrays:

- [SAP HANA capacity requirements](#)
- [SAN network considerations](#)
- [SAP HANA I/O patterns](#)
- [SAP HANA shared file system on a PowerMax array](#)
- [PowerMax scalability for SAP HANA](#)
- [Masking views](#)
- [PowerMaxOS service levels and competing workloads](#)
- [SAP HANA NSE](#)

SAP HANA capacity requirements

Introduction

Every SAP HANA node requires storage devices and capacity for the following purposes:

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

Operating system boot image

For the SAP HANA nodes to be able to start up from a volume on a PowerMax array (boot from the SAN), the overall capacity calculation for the SAP HANA installation must include the required operating system capacity. Every SAP HANA node requires approximately 100 GB capacity for the operating system, including the `/usr/sap/` directory. For information about best practices when booting from a SAN, see the [Dell EMC Host Connectivity Guide for Linux](#).

SAP HANA installation

Every SAP HANA node requires access to a file system mounted under the local mount point, `/hana/shared/`, for installation of the SAP HANA binary files, configuration files, traces, and logs. An SAP HANA scale-out cluster requires a single shared file system, which must be mounted on every node. Most SAP HANA installations use an NFS for this purpose. PowerMax arrays provide this file system with the embedded eNAS option.

You can calculate the size of the `/hana/shared/` file system by using the formulas provided in [SAP HANA Storage Requirements](#). Version 2.10 provides the following formulas:

Single node (scale-up):

$$\text{Size}_{\text{installation(single-node)}} = \text{MIN}(1 \times \text{RAM}; 1 \text{ TB})$$

Multinode (scale-out):

$$\text{Size}_{\text{installation(scale-out)}} = 1 \times \text{RAM_of_worker per 4 worker nodes}$$

SAP HANA persistence (data and log)

The SAP HANA in-memory database requires disk storage for the following purposes:

- Data:
 - To maintain the persistence of the in-memory data on disk to prevent a data loss resulting from a power outage
 - To allow a host auto-failover, where a standby SAP HANA host takes over the in-memory data of a failed worker host in scale-out installations
- Log: To log information about data changes (redo log)

Every SAP HANA scale-up node and scale-out (worker) node requires two disk volumes to save the in-memory database on disk (data) and keep a redo log (log). The size of these volumes depends on the anticipated total memory requirement of the database and the RAM size of the node. [SAP HANA Storage Requirements](#) provides references to help you with disk sizing. Version 2.10 of the document states that you can calculate the size of the data volume by using the following formula:

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

“Net disk space” is the anticipated total memory requirement of the database plus 20 percent free space.

If the database is distributed across multiple nodes in a scale-out cluster, divide the net disk space by the number of SAP HANA worker nodes in the cluster. For example, if the net disk space is 2 TB and the scale-out cluster consists of 4 worker nodes, every node must have a 616 GB data volume assigned to it (2 TB / 4 = 512 GB x 1.2 = 616 GB).

If the net disk space is unknown at the time of storage sizing, we recommend using the RAM size of the node plus 20 percent free space to calculate the capacity of the data file system.

The size of the log volume depends on the RAM size of the node. [SAP HANA Storage Requirements](#) provides the following formulas to calculate the minimum size of the log volume:

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

Backup

SAP HANA supports backup to a file system or the use of SAP-certified third-party tools. Dell Technologies supports data-protection strategies for SAP HANA backup using Dell EMC Data Domain and Dell EMC Networker. Even though you can back up an SAP HANA database to an NFS on a PowerMax array, we do not recommend backing up the SAP HANA database to the storage array on which the primary persistence resides. If you plan to back up SAP HANA to an NFS on a different PowerMax array, see [SAP HANA Storage Requirements](#) for information about sizing the backup file system. The capacity depends not only on the data size and the frequency of change operations in the database, but also on the number of backup generations kept on disk.

For information about using Dell EMC Data Domain data protection for SAP HANA, see [Data Domain Backup and Recovery for SAP HANA](#).

SAN network considerations

General SAN connectivity considerations

The SAN connectivity, which includes host bus adapters (HBAs), SAN ports, switches, and array front-end ports, requires careful planning. The SAP HANA KPIs for TDI deployments require a maximum bandwidth of 400 MB/s per SAP HANA node. For example, if 24 nodes in a SAN are connected to a PowerMax array, a total bandwidth of 9,600 MB/s is required. Assuming that a 16 Gbps front-end port provides approximately 3,000 MB/s of bandwidth, at least four dedicated 16 Gbps front-end ports are required to support 10 SAP HANA nodes ($4 \times 3,000 \text{ MB/s} = 12,000 \text{ MB/s}$).

While this maximum bandwidth requirement only arises in the unlikely event that all nodes require the maximum bandwidth simultaneously, the ability of storage arrays to sustain this peak workload is one of the SAP HANA certification criteria.

This maximum bandwidth requirement does not affect only the storage front-end configuration. In the example with 24 nodes, the complete path through the SAN network must be configured to support the maximum bandwidth. In a multihop SAN, where multiple switches are connected through Inter-Switch Links (ISLs), the bandwidth of the ISLs must also support the required bandwidth.

Storage ports

When you are planning storage connectivity for performance and availability, we recommend “going wide before going deep.” In other words, connecting storage ports across different engines and directors is preferable to using all the ports on a single director. Even if a component fails, the storage can continue to service host I/O.

Note: Each PowerMax engine has two redundant directors.

Dynamic core allocation, where each director provides services such as front-end connectivity, back-end connectivity, and data management, is available with the PowerMax arrays. Each service has its own set of cores on each director. The cores are pooled to provide CPU resources that can be allocated as necessary. For example, even if host I/O arrives through a single front-end port on the director, the front-end pool with all its CPU cores will be available to service that port. Because I/O arriving to other directors has its own core pools, for best performance and availability connect each host to ports on different directors before using additional ports on the same director.

SAP requires isolation of the SAP HANA workload from non SAP HANA applications. Thus, we recommend using dedicated front-end ports for SAP HANA and not sharing these ports with non SAP HANA applications.

To connect a host or cluster, we recommend the following best practices:

- Configure two to four front-end paths in the port group for masking and zones to the host (single-initiator zoning is recommended).
- For cabling options, connect all even-numbered ports to fabric A and all odd-numbered ports to fabric B.
 - In single-engine systems, select two I/O ports spanning both SAN fabrics on each director, with each port being on a separate I/O module—for example, ports 4 and 24 on both directors 1 and 2.

- In a multi-engine system, distribute the paths further across directors spanning different engines to spread the load for performance and ensure fabric redundancy—for example, port 4 in directors 1, 2, 3, and 4.

The following figure shows SAN connectivity in a single-engine environment:




Figure 3. SAN connectivity in a single-engine environment

HBA ports

Each HBA port (initiator) creates a path for I/O between the host and the SAN switch, which then continues to the PowerMax storage. Use two HBA ports, preferably on two separate HBAs. Two ports provide more connectivity and also enable the Linux native multipathing (DM-MPIO) to load-balance and fail over across HBA paths.

SAP HANA I/O patterns

Introduction

The SAP HANA persistent devices use different I/O patterns. For more information, see [SAP HANA Storage Requirements](#).

Data volume

Access to the data volume is primarily random, with blocks ranging in size from 4 KB to 64 MB. The data is written asynchronously with parallel I/O to the data file system. During normal operations, most I/O to the data file system consists of writes. Data is read from the file system only during database restarts, SAP HANA backups, host auto-failover, or a column store table load or reload operation.

Log volume

Access to the log volume is primarily sequential, with blocks ranging in size from 4 KB to 1 MB. SAP HANA keeps a 1 MB buffer in memory for the redo log. When the buffer is full, it is synchronously written to the log volume. When a database transaction is committed before the log buffer is full, a smaller block is written to the file system. Because data is written synchronously to the log volume, a low latency for the I/O to the storage device is important, especially for the smaller 4 KB and 16 KB block sizes. During normal database operations, most I/O to the log volume consists of writes, and data is read from the log volume only during database restart, HA failover, log backup, or database recovery.

SAP HANA I/O can be optimized for specific storage environments. For more information, see [Optimizing file I/O after the SAP HANA installation](#).

SAP HANA shared file system on a PowerMax array

In an SAP HANA scale-out implementation, install the SAP HANA database binary files under the `/hana/shared` mount point on a shared file system that is exposed to all hosts of a system. If a host must write a memory dump that can read up to 90 percent of the RAM size, the memory dump is stored in that file system. Based on your infrastructure and requirements, the following options are available:

- PowerMax eNAS or other NAS systems

NAS systems can provide an NFS share for the SAP HANA shared file system.

- NFS server-based shared file system
- PowerMax block storage

Block storage can create a shared file system using a cluster file system such as the General Parallel File System (GPFS) or the Oracle Cluster File System 2 (OCFS2) on top of the block LUNs. SUSE provides OCFS2 capabilities with the high availability (HA) package. The HA package is also part of the SUSE Linux Enterprise Server (SLES) for SAP applications distribution from SAP that most SAP HANA appliance vendors use.

Note: A SuSE license is required for HA.

PowerMax scalability for SAP HANA

We performed tests on a PowerMax 2000 dual (2 Brick) engine using the SAP HCMT tool for HANA-HWC-ES 1.1 certification. The following table provides guidelines that we derived from the test results for estimating the initial number of SAP HANA production hosts that can be connected to a PowerMax array:

Table 2. PowerMax all-flash scalability

PowerMax model	Number of Bricks	Number of SAP HANA worker hosts
2000	1	24
	2	40
8000	1	34
	2	54
	3	80
	4	106
	5	132
	6	158
	7	184
	8	210

Note: We determined the scalability of PowerMax 8000 systems and additional PowerMax Bricks by extrapolating the PowerMax 2000 test results and using the performance characteristics of the PowerMax 8000 systems.

Depending on the workload, the number of SAP HANA hosts that can be connected to a PowerMax array can be higher or lower than the number specified in Table 2. Use the SAP HCMT tool with scenario HANA-HWC-ES 1.1 to validate the SAP HANA performance and determine the maximum number of SAP HANA hosts on the storage array.

Masking views

Introduction

PowerMax systems use masking views to assign storage to a host. We recommend creating a single masking view for each SAP HANA system whether host (scale-up) or cluster (scale-out). A masking view consists of the following groups:

- Initiator
- Port
- Storage

Initiator group

The initiator group contains the initiators (WWNs) of the HBAs on the SAP HANA host. For redundancy, connect each SAP HANA host to the PowerMax array with at least two HBA ports.

Port group

The port group contains the front-end director ports to which the SAP HANA hosts are connected.

Storage group

An SAP HANA scale-out cluster uses the shared-nothing concept for the persistence of the database, where each SAP HANA worker host uses its own pair of data and log volumes and has exclusive access to these volumes during normal operations. If an SAP HANA worker host fails, the SAP HANA persistence of the failed host is used on a standby host. All persistent volumes must be visible to all SAP HANA hosts because every host can become a worker or a standby host.

The PowerMax storage group of an SAP HANA database must contain all persistent devices of the database cluster. The SAP HANA name server and the SAP HANA storage connector API handle persistence mounting and I/O fencing, which ensures that only one node at a time has access to a pair of data and log volumes.

PowerMaxOS service levels and competing workloads

Overview

Service levels for PowerMaxOS provide open-systems customers with the ability to separate applications based on performance requirements and business criticality. You can specify service levels to ensure that your highest-priority application response times are not affected by lower-priority applications.

Service levels address an organization's requirement to ensure that applications have a predictable and consistent level of performance while running on the array. The available service levels are defined in PowerMaxOS and can be applied to an application's storage group at any time, which enables the storage administrator to initially set and change the performance level of an application as needed. A service level is applied to a storage group using the PowerMax management tools—Unisphere for PowerMax, REST API, Solutions Enabler, and SMI-S.

You can use service levels with host I/O limits to make application performance more predictable while enforcing a specified service level. Setting host I/O limits enables you to define front-end port performance limits on a storage group. These front-end limits can be set by IOPS, host Mb per host, or a combination of both. You can use host I/O limits on a storage group that has a specified service level to throttle IOPS on applications that are exceeding the expected service-level performance.

PowerMaxOS provides six service levels to choose from, as shown in the following table:

Table 3. PowerMaxOS service levels

Service level	Expected average response time
Diamond (highest priority)	0.6 milliseconds (0.4 milliseconds with SCM drives)
Platinum	0.8 milliseconds (0.6 milliseconds with SCM drives)
Gold	1 millisecond
Silver	3.6 milliseconds
Bronze (lowest priority)	7.2 milliseconds
Optimized	N/A

Diamond, platinum, and gold service levels have a response-time upper limit but no lower limit, ensuring that I/O is serviced as fast as possible. Silver and bronze service levels have both an upper limit and a lower limit to allow higher-priority IOPS to be unaffected.

For systems configured with SCM, the diamond average response time drops to 0.4 milliseconds while platinum drops to 0.6 milliseconds.

Storage groups that are set to "Optimized" are throttled for higher-priority IOPS on all service levels aside from bronze.

Note: PowerMaxOS service levels and host I/O limits are available at no additional cost for both PowerMax and VMAX All Flash systems that are running PowerMaxOS 5978.

For more information, see the [Dell EMC PowerMax: Service Levels for PowerMaxOS White Paper](#).

SLO and workload type best practices for SAP HANA

The following table shows the service-level objective (SLO) configurations that we recommend for different SAP HANA installation types:

Table 4. Recommended SLO and workload configurations for SAP HANA

Installation type	SLO	Reason
SAP HANA persistence (data and log) for production installations	Diamond	<p>Diamond ensures that the PowerMax system tries to keep the latency below 1 millisecond, which is an SAP requirement for small (4 KB and 16 KB) block sizes on the log volume.</p> <p>Using the diamond SLO with all-flash devices provides the following benefits:</p> <ul style="list-style-type: none"> • Reduced SAP HANA startup times when data is read from the data volume into memory • Reduced SAP HANA host auto-failover times in scale-out deployments when a standby node takes over the data from a failed worker node • Reduced SAP HANA backup times when the backup process must read the data from the data volume • Submillisecond latencies for small block sizes on the log volume • Optimized performance for SAP HANA NSE
SAP HANA persistence for nonproduction installations	Gold	Although the SAP performance KPIs do not apply to SAP HANA nonproduction installations, those installations are still critical components in an overall SAP landscape.
SAP HANA installation	Bronze	Bronze is sufficient when you are using eNAS in a PowerMax array to provide the NFS share for the SAP HANA installation file system (/hana/shared).
Operating system boot image	Bronze	The operating system boot image can also reside on a bronze SLO.

SLO considerations for “noisy neighbors” and competing workloads

In highly consolidated environments, SAP HANA and other databases and applications compete for storage resources. PowerMax arrays can provide the appropriate performance for each application when the user specifies SLOs and workload types. By using different SLOs for each such application or group of applications, it is easy to manage a consolidated environment and modify the SLOs when business requirements change. See [Host I/O limits and multitenancy](#) for additional ways of controlling performance in a consolidated environment.

Service levels enable you to insulate specific storage groups from any performance impact from other “noisy neighbor” applications. You can assign critical applications to higher service levels such as diamond, platinum, or gold, which allow these storage groups to use all available resources at all times. These critical applications are not managed unless the system exhibits resource constraints that cause the applications to fail to maintain required performance levels.

Host I/O limits and multitenancy

The quality of service (QoS) feature that limits host I/O was introduced in the previous generation of VMAX arrays. The feature continues to offer PowerMax customers the

option to place specific IOPS or bandwidth limits on any storage group, regardless of the SLO assigned to that group. For example, assigning a host I/O limit for IOPS to a storage group of noisy SAP HANA neighbors with low performance requirements can ensure that a spike in I/O demand does not affect the SAP HANA workload and performance.

SAP HANA NSE

SAP HANA NSE, a native SAP HANA feature, is a “warm” data solution that provides an improved cost-to-performance ratio. It is comparable to other SAP HANA warm data-tiering solutions such as SAP HANA Extensions Node and SAP HANA Dynamic Tiering. Mission-critical (or “hot”) data is stored in memory on the SAP HANA database for real-time processing and analytics. Less frequently accessed (or “warm”) data is saved on disk, although it is still managed as a unified part of the SAP HANA database. Managing and retaining older data in this cost-effective way helps maintain data growth while minimizing the expense of hardware growth.

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

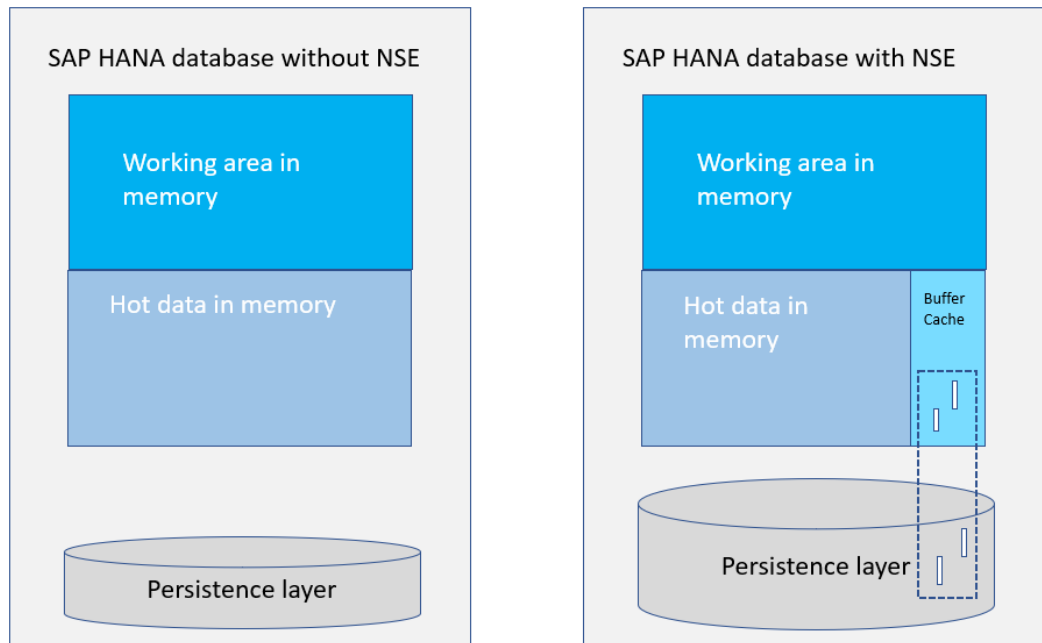


Figure 4. Comparison of standard SAP HANA database and SAP HANA database with NSE

The capacity of a standard SAP HANA database is limited by the amount of main memory. Using SAP HANA NSE, you can bypass these limits by storing warm data on disk 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.

Note: Since SAP HANA 2.0 SP05, scale-out SAP HANA systems are supported with NSE.

While hot data is “column loadable,” residing completely 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 page-loadable data is loaded into memory page by page as required for query processing. Unlike column-loadable data, page-loadable data does not have to reside completely in memory.

The following figure depicts the SAP HANA database with NSE:

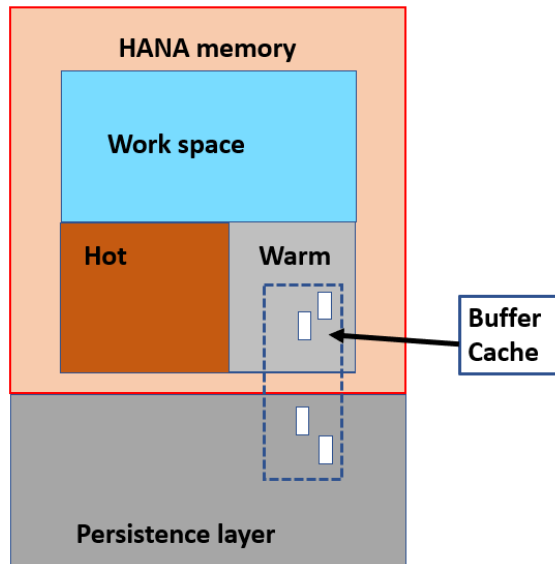


Figure 5. SAP HANA database with NSE

NSE reduces the memory footprint for page-loadable data. The database is partly in memory and partly on disk, as illustrated in Figure 5. You can use the PowerMax storage array, together with SAP HANA NSE, to substantially increase SAP HANA data capacity and reduce TCO.

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 the key features of NSE and related topics, see [SAP HANA Native Storage Extension](#).

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

Chapter 4 Configuration and Installation

This chapter presents the following topics:

Configuration and installation overview	26
Configuring the PowerMax array	26
Preparing the SAP HANA nodes	32
Installing the SAP HANA scale-out cluster	34
Performing post-installation configuration	37

Configuration and installation overview

The following sections describe how to configure and install an SAP HANA scale-out cluster on a PowerMax array. The procedures include:

- **Configuring the PowerMax array**—Creating and configuring the persistent storage (data and log) on a PowerMax array for an SAP HANA scale-out cluster with three worker nodes and one standby node (3+1)
- **Preparing the SAP HANA nodes**—Initializing the SAP HANA persistence
- **Installing the SAP HANA cluster**—Preparing the global.ini file, preparing the installation parameters file, and installing the cluster using the SAP life cycle management command-line tool hdblicm
- **Performing post-installation configuration**—Optimizing file I/O and configuring Native Storage Extensions

Configuring the PowerMax array

Create a parent storage group and volumes

We used the Unisphere for PowerMax UI to configure all storage devices, storage groups, port groups, host groups, and the masking view for the SAP HANA scale-out cluster.

To create a parent storage group and volumes:

1. Log in to Unisphere for PowerMax.

The following figure shows the Unisphere for PowerMax UI:

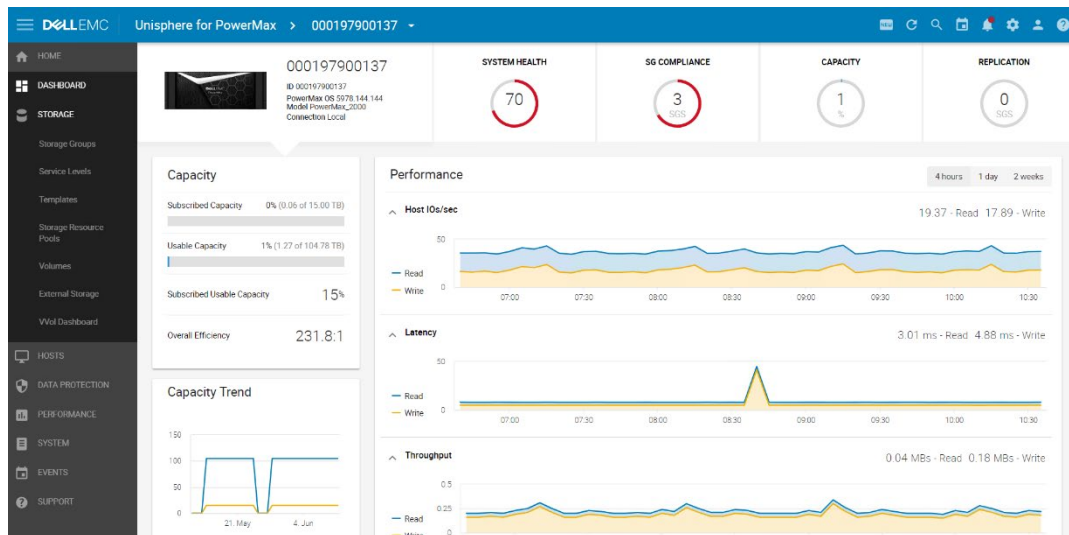


Figure 6. Unisphere for PowerMax

2. Select **Storage > Storage Groups** and click **Create**, as shown in the following figure:

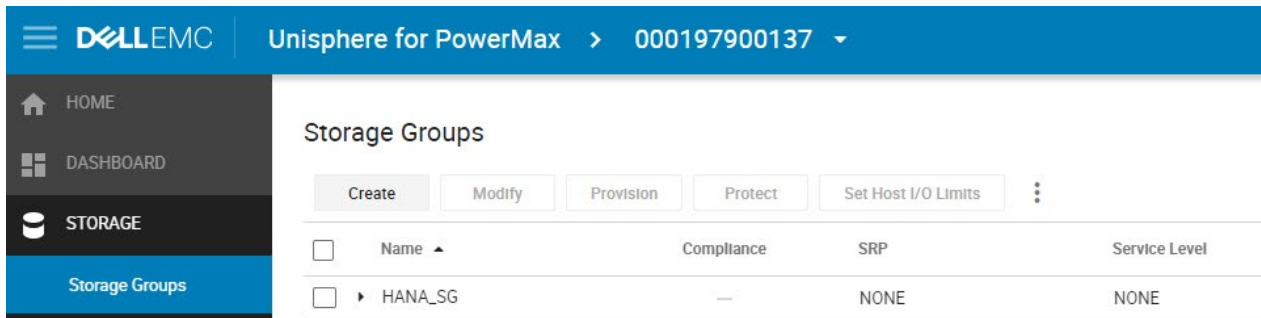


Figure 7. Storage Groups page

3. Enter a name for the parent storage group, and then hover your mouse over **Service Level** and click the + sign to create cascaded storage groups, as shown in the following figure:

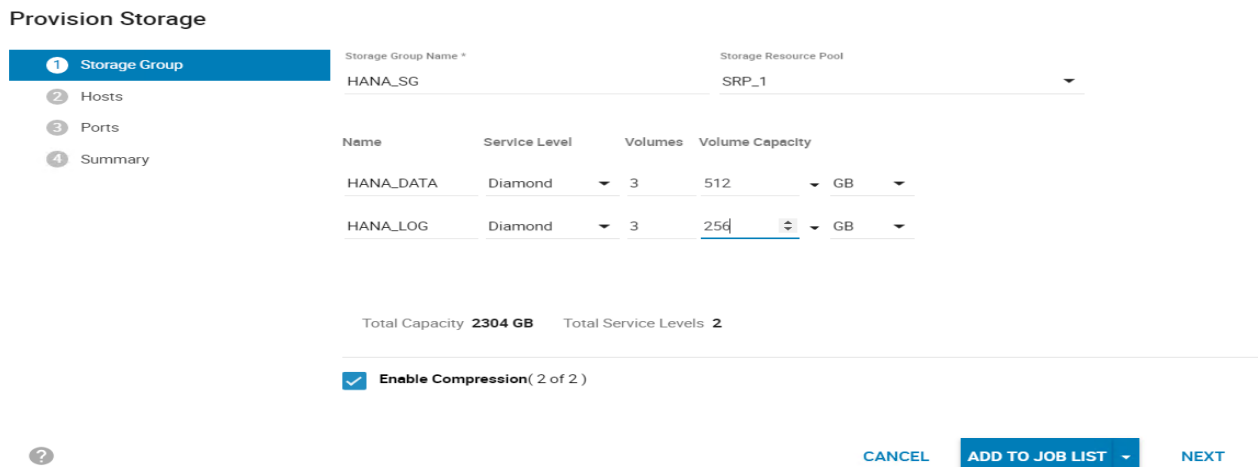


Figure 8. Provision Storage page

For our 3+1 SAP HANA cluster, we needed three data volumes of 512 GB capacity each and three log volumes of 256 GB capacity each:

- Top-level group: **HANA_SG**
 - Subgroup for all data volumes: **HANA_DATA**
 - Subgroup for all log volumes: **HANA_LOG**
4. Specify the number and size of the volumes to be created, and then click the down arrow at **Add to Job List** and select **Run Now**.

The new cascaded storage group is created.

- To view information about the volumes that you created, in the right pane, click the hyperlink on the number listed for **Volumes**, as shown in the following figure:

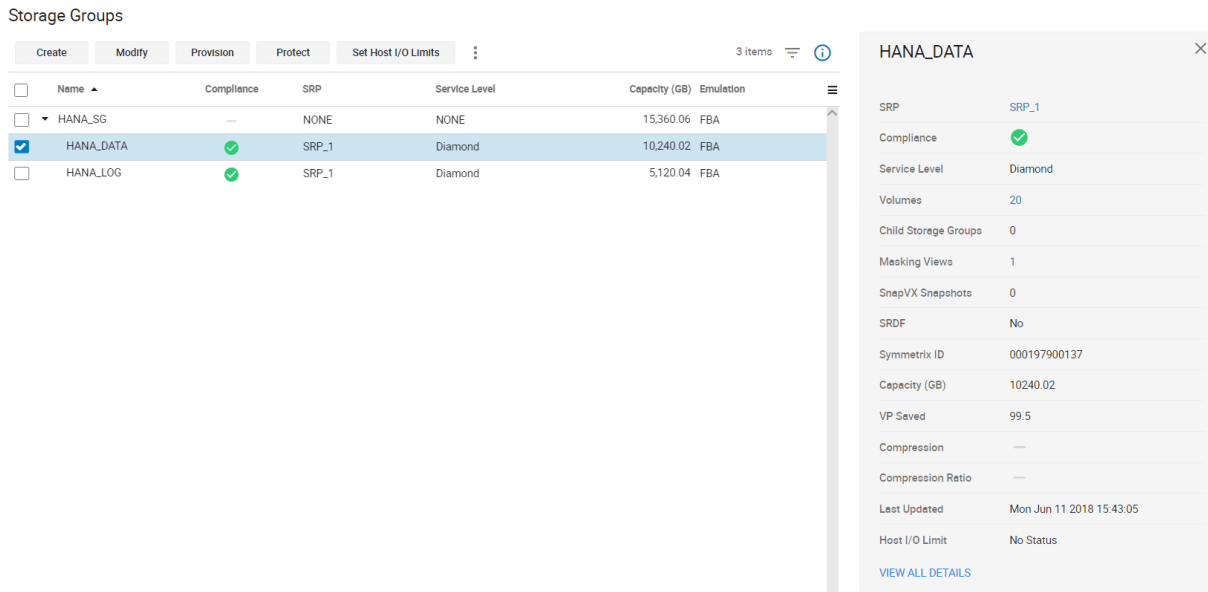


Figure 9. HANA_DATA storage groups

The Volumes page is displayed, as shown in the following figure:

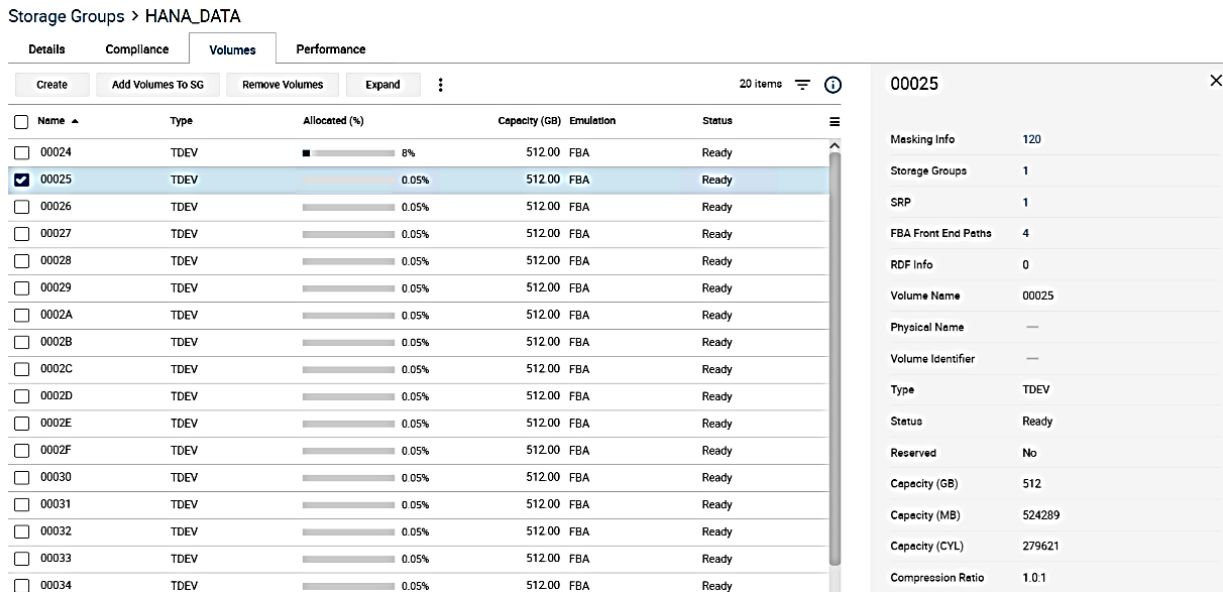


Figure 10. Volumes page

- Select a volume and note the WWN of the volume in the right panel. You might have to scroll to find the WWN of the volume.
- Repeat the preceding steps for all your data and log volumes.

The SAP HANA storage connector fcClient uses the WWN specified in the SAP HANA `global.ini` file to identify a storage LUN.

Create a host group

To create a host group:

1. Select **Hosts > Hosts > Create Host**.

The following page appears:

Create Host

Host Name *
C460-03U

Initiator Type

Fibre iSCSI

Select Initiators

Available Initiators

<input type="checkbox"/>	Name
<input type="checkbox"/>	10000090fab37196
<input type="checkbox"/>	10000090fab397da
<input type="checkbox"/>	10000090fab3981a
<input type="checkbox"/>	10000090fab39856
<input type="checkbox"/>	10000090fab39876

> Add

Initiators in Host

<input type="checkbox"/>	Name
<input type="checkbox"/>	10000090fab36738
<input type="checkbox"/>	10000090fab36eac

Set Host Flags

? CANCEL **ADD TO JOB LIST** ▾

Figure 11. Create Host page

2. Enter a hostname, and then select the HBA initiators for that host from the **Available Initiators** list and move them to the **Initiators Host** area.
3. Click the down arrow next to **Add to Job List** and select **Run Now**.

The new host is created.

4. Select **Create > Create Host Group** and select the hosts that belong to the SAP HANA cluster; then click the right arrow to move them to the **Hosts in Host Group** area, as shown in the following figure:

Create Host Group

Host Group Name *
HANA_IG

Host Type

Fibre iSCSI

Select Hosts

Available Hosts

<input type="checkbox"/>	Name	Consistent LUNS	Flags
<input type="checkbox"/>	C460-02U	—	—
<input type="checkbox"/>	C460-09Q	—	—
<input type="checkbox"/>	C460-0A8	—	—
<input type="checkbox"/>	C460-0AE	—	—
<input type="checkbox"/>	C460-0AG	—	—

> <

Hosts in Host Group

<input type="checkbox"/>	Name	Consistent LUNS	Flags
--------------------------	------	-----------------	-------

Set Host Group Flags

? CANCEL **ADD TO JOB LIST** ▾

Figure 12. Create Host Group page

- Click the down arrow next to **Add to Job List** and select **Run Now**.

The new host group is created.

Create a port group

To create a port group:

- Select **Hosts > Port Groups > Create**, and enter a name, such as `HANA_PG`, as shown in the following figure:

Create Port Group

Port Group Name *
HANA_PG

Select Port Group Type

Fibre iSCSI

Select one or more Ports to be added to the Port Group

<input type="checkbox"/> Dir:Port ▲	Identifier	Port Groups	Masking Views	Volumes ≡
<input checked="" type="checkbox"/> FA-1D:4	50000973b0022404	1	1	40 ^
<input type="checkbox"/> FA-1D:5	50000973b0022405	0	0	0
<input checked="" type="checkbox"/> FA-1D:6	50000973b0022406	1	1	40
<input type="checkbox"/> FA-1D:7	50000973b0022407	0	0	1
<input checked="" type="checkbox"/> FA-2D:4	50000973b0022444	1	1	40
<input type="checkbox"/> FA-2D:5	50000973b0022445	0	0	0
<input checked="" type="checkbox"/> FA-2D:6	50000973b0022446	1	1	40
<input type="checkbox"/> FA-2D:7	50000973b0022447	0	0	0

? CANCEL **ADD TO JOB LIST** ▼

Figure 13. Creating a port group

- Select the ports that your initiators are logged in to by holding down the Control key, and then click the down arrow next to **Add to Job List** and select **Run Now**.
- Click **OK** if you receive a warning message stating that the port group will have multiple ports from the same director—in this example, FA-1D:4, FA1D:6, FA-2D:4, and FA-2D:6.

Note: For a single host, we recommend a 1:1 relationship between a host HBA and a storage front-end port. Because we created a port group for an SAP HANA cluster, we required throughput and bandwidth for multiple hosts.

Create a storage group

To create and configure a storage group:

1. Select **Storage > Storage Groups > Create**.

The following page is displayed:

Provision Storage

1 Storage Group

2 Hosts

3 Ports

4 Summary

Storage Group Name *

HANA_SG

Storage Resource Pool

SRP_1

Name	Service Level	Volumes	Volume Capacity
HANA_Log	Diamond	1	512 GB
HANA_Data	Diamond	0	3.6 TB
HANA_Share	Bronze	0	3 TB

Total Capacity **512 GB** Total Service Levels **3**

Enable Compression (3 of 3)

?
CANCEL
ADD TO JOB LIST
NEXT

Figure 14. Storage Group page

2. Enter a storage group name such as HANA__SG, and then select the storage resource pool.
3. Enter or select the following information, as applicable:
 - **Name**
 - **Service level**
 - **Number of created volumes**
 - **Volume capacity**
4. If applicable, select **Enable Compression**.
5. Click the down arrow next to **Add to Job List** and select **Run Now**.

Create a masking view

A PowerMax masking view combines the storage group, port group, and host group, and enables access from the SAP HANA nodes to the storage volumes.

Follow these steps to create a masking view:

1. Select **Hosts > Masking View > Create Masking View**.
2. Enter a masking view name, such as HANA_Cluster, and select the host group, port group, and storage group that you created in the preceding steps, as shown in the following figure:

Create Masking View

Masking View Name *

HANA_Cluster

Host *

HANA_IG

Port Group *

HANA_PG

Storage Group *

HANA_SG

Set Dynamic LUNs

?

CANCEL OK

Figure 1. Create Masking View page

3. Click **OK**.

The masking view is created. The SAP HANA nodes now have access to the storage volumes.

We created an SAP HANA scale-out system with three worker nodes and one standby node (3+1). In an SAP HANA cluster with four nodes, each node must have access to every SAP HANA device. During SAP HANA startup, the SAP HANA nameserver along with the SAP HANA storage connector `fcClient` mounts the volumes to the correct SAP HANA node and I/O fencing.

Ensure that you correctly prepare the SAP HANA nodes and the SAP HANA `global.ini` file, as described in subsequent sections, before you perform the SAP HANA installation.

Note: When using `fcClient`, do not auto-mount the device using `/etc/fstab`.

Preparing the SAP HANA nodes

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 properly configured in accordance with the SAP recommendations.
- An SAP HANA shared file system, `hana/shared/`, has been created on a NAS system, for example, on a PowerMax array with eNAS capability, and is mounted on all SAP HANA nodes.
- Linux native multipathing (DM-MPIO) is installed on the SAP HANA nodes.

- All network settings and bandwidth requirements for internode communications are configured in accordance with the SAP requirements.
- SSH keys have been exchanged between all SAP HANA nodes.
- System time synchronization has been configured using 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 only be installed on certified server hardware. A certified SAP HANA expert must perform the installation.

Linux native multipathing (DM-MPIO)

For both Red Hat Enterprise Linux and SUSE, the native MPIO already contains default configuration parameters for PowerMAX arrays to provide optimal performance in most environments. You do not have to create a device stanza for these arrays unless you want to modify the default behavior. The `/etc/multipath.conf` file is installed by default when you install the `device-mapper-multipath` package.

Note: For the latest information about the operating system version and storage array MPIO settings for native multipathing, always refer to the [Dell EMC Host Connectivity Guide for Linux](#).

Initialize the SAP HANA persistence

The SAP HANA persistence must be visible to every node in the SAP HANA cluster. Run the `rescan-scsi-bus.sh` command or reboot each node, as described in this section.

To verify that the volumes are visible:

1. On one of the nodes, list all 512 GB data volumes by running the following command:

```
multipath -ll | grep -B1 -A5 512G
```

```
c460-2ww:~ # multipath -ll | grep -B1 -A5 512G
3600009700001979001375330303242 dm-39 EMC,SYMMETRIX
size=512G features='1 retain_attached_hw_handler' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
  |- 3:0:2:7 sdm 8:192 active ready running
  |- 4:0:2:7 sdb 67:128 active ready running
  |- 3:0:3:7 sddp 71:112 active ready running
  `-- 4:0:3:7 sdc 70:0 active ready running
--
3600009700001979001375330303239 dm-37 EMC,SYMMETRIX
size=512G features='1 retain_attached_hw_handler' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
  |- 3:0:2:5 sdk 8:160 active ready running
  |- 4:0:2:5 sdb 67:96 active ready running
  |- 3:0:3:5 sddm 71:64 active ready running
  `-- 4:0:3:5 sdcg 69:224 active ready running
3600009700001979001375330303241 dm-38 EMC,SYMMETRIX
size=512G features='1 retain_attached_hw_handler' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
  |- 3:0:2:6 sdl 8:176 active ready running
  |- 4:0:2:6 sdbd 67:112 active ready running
  |- 3:0:3:6 sddn 71:80 active ready running
  `-- 4:0:3:6 sdc 69:240 active ready running
```

Figure 2. List data volumes command

2. List all 256 GB log volumes by running:

```
multipath -ll | grep -B1 -A5 256G
```

```
c460-2ww:~ # multipath -ll | grep -B1 -A5 256G
360000970000197900137533030303342 dm-18 EMC,SYMMETRIX
size=256G features='1 retain_attached_hw_handler' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
  |- 3:0:2:23 sdac 65:192 active ready running
  |- 4:0:2:23 sdbu 68:128 active ready running
  |- 3:0:3:23 sdev 129:112 active ready running
  `-- 4:0:3:23 sddq 71:128 active ready running
--
360000970000197900137533030303442 dm-35 EMC,SYMMETRIX
size=256G features='1 retain_attached_hw_handler' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
  |- 3:0:2:39 sdas 66:192 active ready running
  |- 4:0:2:39 sdck 69:128 active ready running
  |- 3:0:3:39 sdfm 130:128 active ready running
  `-- 4:0:3:39 sdew 129:128 active ready running
360000970000197900137533030303339 dm-16 EMC,SYMMETRIX
size=256G features='1 retain_attached_hw_handler' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
  |- 3:0:2:21 sdaa 65:160 active ready running
  |- 4:0:2:21 sdbs 68:96 active ready running
  |- 3:0:3:21 sdeg 129:32 active ready running
  `-- 4:0:3:21 sddl 71:48 active ready running
```

Figure 3. List log volumes command

The unique device identifier of the multipath device must match the WWN of the volumes that you created in the Unisphere UI. (When viewed on a Linux host, the WWN of the volume from the PowerMax array is now preceded by a 3.)

3. Initialize the devices and create the Linux XFS file system on each of the devices by running a command such as:

```
mkfs.xfs /dev/mapper/360000970000197900137533030303234
```

After all the file systems are created, you are ready to install the SAP HANA scale-out cluster.

Installing the SAP HANA scale-out cluster

Introduction

Before you run the installation script, prepare the following two configuration files:

- A `global.ini` file with a storage section describing the SAP HANA storage partitions, mount options, and the storage connector to use
- An installation parameter file with customized parameters to be used by the `hdblcmm` command-line script

Prepare the `global.ini` file

The installation uses the `global.ini` file to describe the SAP HANA storage partitions and the storage connector. The SAP HANA certified Dell EMC storage platforms use `fcClient`, which is part of the SAP software distribution.

Follow these steps:

1. Ensure that the `global.ini` file has the content shown in the following figure:

```
[storage]
ha_provider = hdb_ha.fcClient
partition_*_*_prtype = 5
partition_*_data_mountoptions = -o relative,inode64
partition_*_log_mountoptions = -o relative,inode64,nobarrier
partition_1_data_wwid = 3600009700001979001375330303234
partition_1_log_wwid = 3600009700001979001375330303338
partition_2_data_wwid = 3600009700001979001375330303235
partition_2_log_wwid = 3600009700001979001375330303339
partition_3_data_wwid = 3600009700001979001375330303236
partition_3_log_wwid = 3600009700001979001375330303341
```

Figure 4. Global.ini file storage content

2. Ensure that the partition entries match the displayed unique device identifier when you run the `multipath -ll` command.
3. Place the `global.ini` file in a directory on the `/hana/shared/` file system—for example, `/hana/shared/PMX_cfg`.

Do not alter the name of the `global.ini` file.

Larger SAP HANA scale-out installations require additional partition entries.

Prepare the installation parameters file

SAP HANA SPS 07 introduced the SAP HANA Database Lifecycle Manager to provide the efficiency of installing all components at one time, while automating the installation and providing further flexibility to customers. Our example uses `hdblcm` to install our SAP HANA 3+1 scale-out cluster.

Follow these steps:

1. In the shared file system, go to the `HDB_LCM_LINUX_X86_64` directory into which the SAP HANA installation DVD ISO file has been extracted by running the following command:

```
cd /<installation media>/DATA_UNITS/HDB_LCM_LINUX_X86_64
```

2. Create a template installation parameters file by running the following command:

```
./hdblcm --action=install --
dump_configfile_template=PMX_install.cfg
```

3. Modify the following parameters in the template file to match your environment:

```
Directory root to search for components
component_root=/SAPShare/software/SAP_HANA_SPS11_IM/51050506/
```

```
Components ( Valid values: all | client | es | ets | lcapps
| server | smartda | streaming | rdsync | xs | studio | afl
| pos | sal | sca | sop | trd | udf )
components=server,client
```

```
Installation Path ( Default: /hana/shared )
```

```
sapmnt=/hana/shared

Local Host Name ( Default: server06 )
hostname=C460-2WW

Directory containing a storage configuration
storage_cfg=/hana/shared/PMX_cfg
```

Note: The `storage_cfg` parameter points to the directory where you have placed the customized `global.ini` file.

```
SAP HANA System ID
sid=PMX

Instance Number
number=00

# System Administrator User ID
userid=1001

ID of User Group (sapsys)
groupid=79

Action to be performed ( Default: exit; Valid values:
install | update | extract_components )
action=install

Additional Hosts
addhosts=C460-2V4:storage_partition=2:role=worker,C460-
22L:storage_partition=3:role=worker,C460-39A:role=standby
```

Note: The `Additional Hosts` parameter describes the additional hosts and their roles in the scale-out installation.

4. Specify passwords for the root user, SAP Host Agent User (`sapadm`), system administrator user (`<sid>adm`), and database user (`SYSTEM`) in the parameters file.

You can also use encrypted passwords. For more information, see the [SAP HANA Studio Installation and Update Guide](#).

The `hdblcm` installation procedure prompts you for any missing passwords or parameters.

5. Review the entire template file and specify any additional parameters that might be required for your environment.

Installing the scale-out cluster

After you have created and customized the `global.ini` file and the installation parameters file, start the installation by running the following command:

```
./hdblcm --action=install --configfile=PMX_Install.cfg
```

Performing post-installation configuration

Optimizing file I/O after the SAP HANA installation

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

- **Simple File**—Used for small, simple I/O requests on configuration files, traces, and so on. The simple file interface uses lightweight, platform-independent wrappers around system calls.
- **FileFactory and File**—Used for large, 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 with configuration parameters to optimize file I/O for a given storage array and file system. The Linux XFS file system is used on all Dell EMC storage LUNs for the HANA persistence.

After the SAP HANA persistence is installed on PowerMax LUNs, set the following file I/O layer parameters for optimal I/O processing:

- `max_parallel_io_requests=256`
- `async_read_submit=on`
- `async_write_submit_blocks=all`

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=256
hdbparam --paramset fileio [LOG].max_parallel_io_requests=256
hdbparam --paramset fileio [DATA].async_read_submit=on
hdbparam --paramset fileio [LOG].async_read_submit=on
hdbparam --paramset fileio [DATA].async_write_submit_blocks=all
hdbparam --paramset fileio [LOG].async_write_submit_blocks=all
```

SAP HANA 2.0

Starting with SAP HANA 2.0, the `hdbparam` command-line tool was deprecated. Instead, the parameters are defined in `global.ini > [fileio]`.

Using the SQL console in SAP HANA studio, set the `max_parallel_io_requests=256` parameter in the `global.ini` file by running the following commands:

```
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio', 'max_parallel_io_requests[DATA]') = '256';
```

```
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio', 'max_parallel_io_requests[LOG]') = '256' WITH
RECONFIGURE;
```

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

fileio		
<code>async_read_submit</code>	on	
<code>async_write_submit_active</code>	on	
<code>async_write_submit_blocks</code>	all	
<code>max_parallel_io_requests</code>	64	
<code>max_parallel_io_requests[data]</code>		● 256
<code>max_parallel_io_requests[log]</code>		● 256
<code>max_submit_batch_size</code>	64	
<code>min_submit_batch_size</code>	16	
<code>num_completion_queues</code>	1	
<code>num_submit_queues</code>	1	

Figure 5. fileio section of the global.ini file in SAP HANA studio

Both `async_read_submit=on` and `async_write_submit_blocks=all` are set by default during the installation. For more information, see SAP Note 2399079: [Elimination of hdbparam in SAP HANA 2](#) (access requires an SAP username and password).

Note: The instructions in this guide for tuning file I/O parameters are based on SAP HANA 1.0 and SAP HANA 2.0 SPS03. See the latest SAP HANA documentation for any updates.

Configure SAP HANA Native Storage Extensions (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 take into account 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 about the SAP HANA buffer cache, see the [SAP HANA Administration Guide](#).

SAP recommends not exceeding a ratio of 1:8 for the buffer cache and the total amount of warm data handled by SAP HANA NSE. When using SAP HANA NSE, you can store data in the warm tier in the following 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 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 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, go to the [SAP Help Portal](#).

Chapter 5 Data Domain Backup and Recovery for SAP HANA

This chapter presents the following topics:

Data Domain overview	41
Configuring Data Domain for SAP HANA	41
Backing up the database from SAP HANA studio	43
Recovering the database with SAP HANA studio.....	44

Data Domain overview

Introduction

Dell EMC Data Domain systems work seamlessly with a range of backup, archive, and enterprise applications. The new generation of midsize and large enterprise Data Domain systems are powered with flash solid-state drives (SSDs). For more information, see the [Dell EMC Data Domain Deduplication Storage Systems Spec Sheet](#). By consolidating backup and archive data on a Data Domain system, you can reduce storage requirements by 10 to 55 times, making disks cost-effective for onsite retention and highly efficient for network-based replication to DR sites.

Data Domain Boost connection to SAP HANA backup interface

The SAP HANA database provides a backup interface called Backint for SAP HANA. This backup interface enables third-party backup tools such as Data Domain Boost for Enterprise Applications (DDBEA) to connect to the backup and recovery capabilities of the SAP HANA database. Because Backint for SAP HANA is fully integrated into the SAP HANA database, you can individually configure data and log backups to be created and recovered using DDBEA.

A DDBEA backup to a Data Domain system uses the DD Boost feature as follows:

- The DD Boost library API enables the backup software to communicate with the Data Domain system.
- The DD Boost distributed segment processing (DSP) component reviews the data that is already stored on the Data Domain system and sends only unique data for storage. The DSP component enables the backup data to be deduplicated on the database or application host to reduce the amount of data that is transferred over the network. During the restore of a backup to the client, the Data Domain system converts the stored data to its original non-deduplicated state before sending the data over the network.

Configuring Data Domain for SAP HANA

To configure Data Domain for SAP HANA:

1. Install the DDBEA software on the operating system of the database host.
2. Manually create the following subdirectory if it does not already exist:

```
/usr/sap/<SID>/SYS/global/hdb/opt
```

3. Copy the `/opt/ddbda/bin/hdbbackint` file to the subdirectory that you specified in step 2.

Alternatively, as shown in the following figure, create a symbolic link that points to the executable file from the following directory:

```
/usr/sap/<SID>/SYS/global/hdb/opt/hdbbackint
```

```
x22adm@c460-09g:/usr/sap/x22/SYS/global/hdb/opt> ll
total 4
lrwxrwx--- 1 x22adm sapsys 38 Oct 20 09:45 hdbbackint -> /opt/dpsapps/dbappagent/bin/hdbbackint
```

Figure 6. Create symbolic link command

4. Modify the parameter settings of the SAP HANA template configuration file `/opt/ddbda/config/sap_hana_ddbda.utl`, as shown in the following figure:

```
# #####
# General Parameters
# #####
[GENERAL]

CLIENT = c460-09q.sse.sap.local
LOCKBOX_PATH = /opt/dpsapps/common/lockbox
PARALLELISM = 5

#####
# Primary Data Domain
# #####
[PRIMARY_SYSTEM]

DDBOOST_USER = ddx22hana
DEVICE_HOST = dd9300
DEVICE_PATH = /sap_x22hana
```

Figure 7. Modifying the configuration file parameter settings

5. Configure the DDBDA lockbox by using `ddbmadmin -P -Z <configuration_file>`, where `<configuration_file>` is the file that you used in step 4.

Note: This example of configuring SAP HANA for backup and recovery uses X22 as the SAP HANA `<SID>`. If you are configuring the multinode cluster, repeat steps 2–5 for all nodes.

6. In the SAP HANA studio UI, specify the location of the DDBEA configuration file for data and log backup, as shown in the following figure:

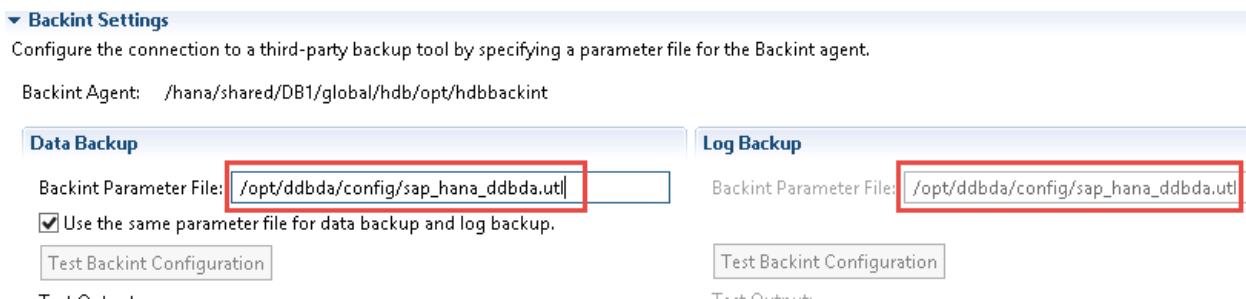


Figure 8. Specifying the DDBEA configuration file location

For more information, see the [Dell EMC Data Domain Boost for Enterprise Applications and ProtectPoint Database Application Agent Installation and Administration Guide](#).

Backing up the database from SAP HANA studio

To back up the SAP HANA database:

1. Log in to SAP HANA studio, and then select **Backup and Recovery**, as shown in the following figure:

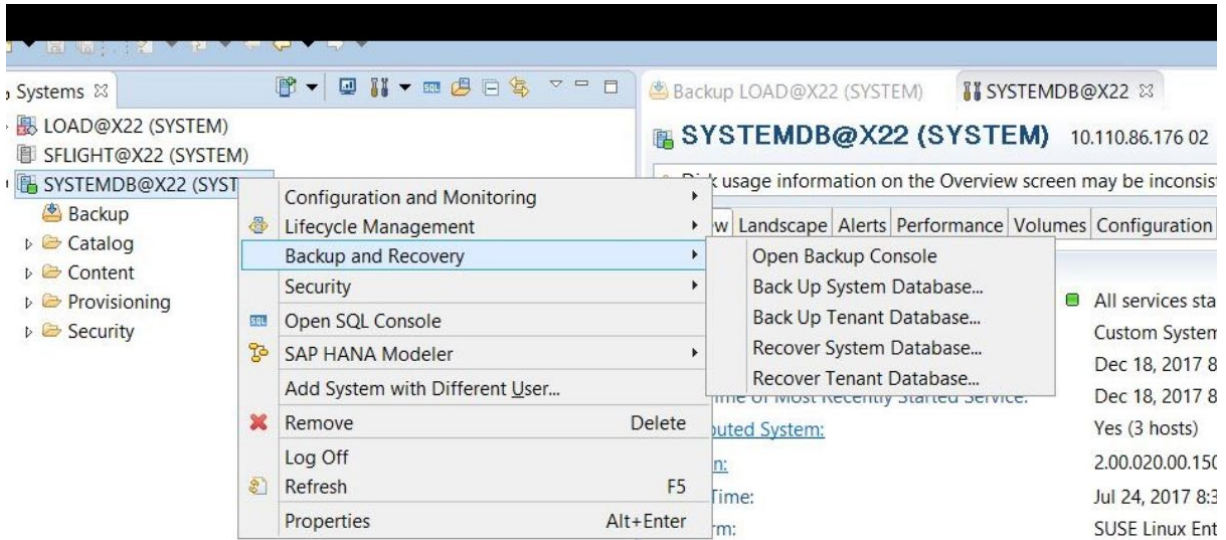


Figure 9. SAP HANA studio: Selecting Backup and Recovery

2. Select the tenant database, as shown in the following figure.
The **LOAD** database in the figure is for illustration purposes only.

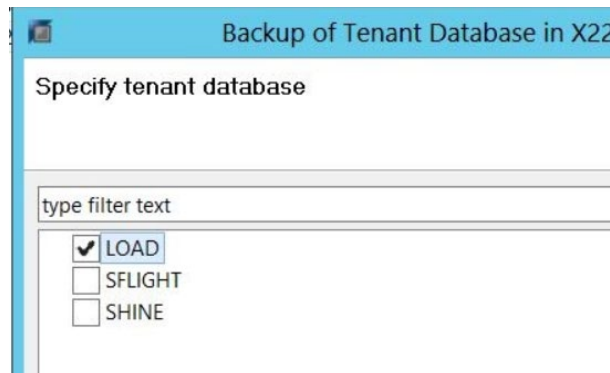


Figure 10. Selecting the tenant database

3. Select **Backint** as the destination type, as shown in the following figure:

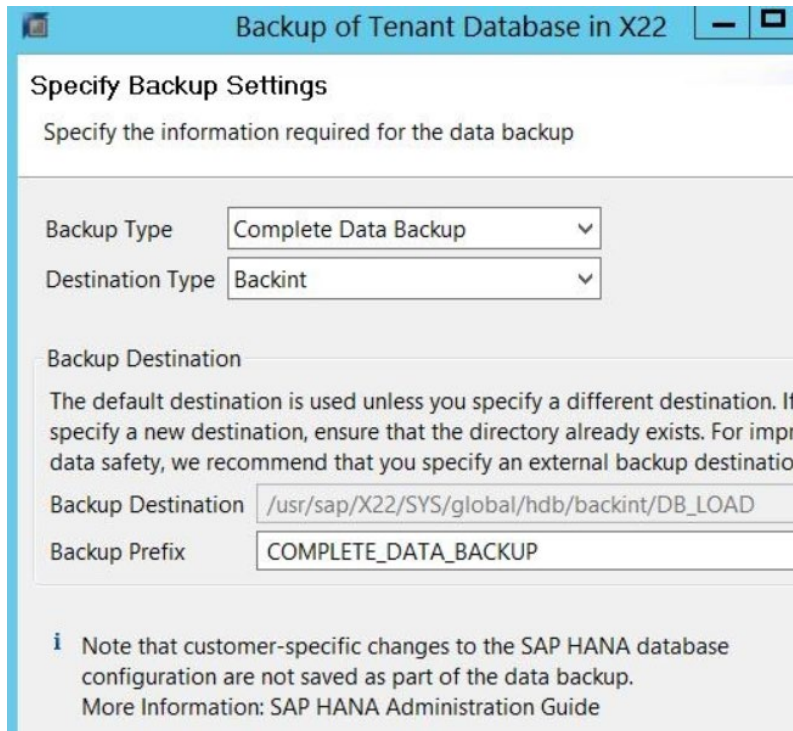


Figure 11. Specifying the backup destination type

4. Review the backup settings, as shown in the following figure:

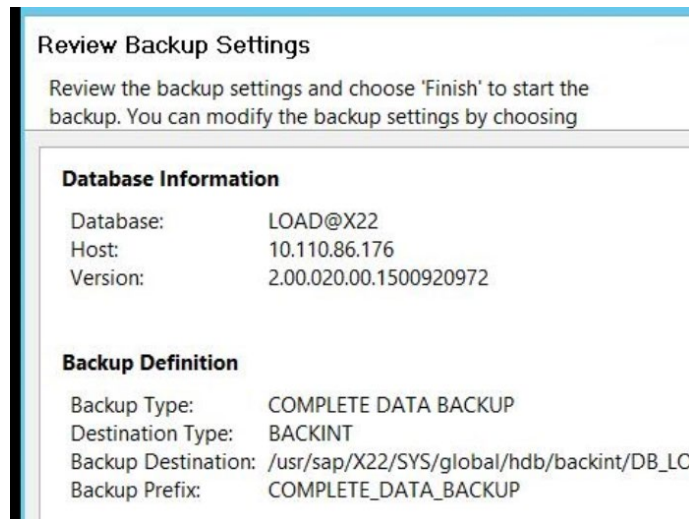


Figure 12. Review Backup Settings page

5. Click **Finish** to start the backup.

Recovering the database with SAP HANA studio

Restore the SAP HANA database to the point when you backed it up:

1. Log in to the SAP HANA studio, and then select **Backup and Recovery**.
2. Select the tenant database.

3. Select **Recover the database to the most recent state**, as shown in the following figure:

Recovery of Tenant Database in X22

Specify Recovery Type

Select a recovery type.

Recover the database to its most recent state [?]

Recover the database to the following point in time [?]

Date: Time:

Select Time Zone:

i System Time Used (GMT): 2018-01-31 05:26:29

Recover the database to a specific data backup [?]

Figure 13. Specifying the recovery type

4. Specify a location for the log backups, as shown in the following figure:

Recovery of Tenant Database in X22

Locate Log Backups

Specify location(s) of log backup files to be used to recover the database.

i Even if no log backups were created, a location is still needed to read data that will be used for recovery.

If the log backups were written to the file system and subsequently moved, you need to specify their current location. If you do not specify an alternative location for the log backups, the system uses the location where the log backups were first saved. The directory specified will be searched recursively.

Locations:

Figure 14. Specifying the log backup location

5. Select the point in time to which you want to restore the database, as shown in the following figure:

Select a Backup

Select a backup to recover the SAP HANA database

Selected Point in Time

Database will be recovered to its most recent state.

Backups

The overview shows backups that were recorded in the backup catalog as successful. The backup at the top is estimated to have the shortest recovery time.

Start Time	Location	Backup Pref...	A...
2018-01-08 09:53...	/usr/sap/X22/SYS/glo...	COMPLETE_...	◇
2017-12-20 15:05...	/usr/sap/X22/SYS/glo...	COMPLETE_...	◇
2017-12-19 10:22...	/usr/sap/X22/SYS/glo...	COMPLETE_...	◇
2017-12-18 16:45...	/usr/sap/X22/SYS/glo...	COMPLETE_...	◇
2017-12-05 18:29...	/usr/sap/X22/SYS/glo...	COMPLETE_...	◇
2017-12-05 17:44...	/usr/sap/X22/SYS/glo...	COMPLETE_...	◇

Refresh Show More

Details of Selected Item

Start Time: 2017-12-19 10:22:17 Destination Type: BACKINT Source System: X22
 Size: 206.72 GB Backup ID: 15136501 External Backup ID: 15
 Backup Name: /usr/sap/X22/SYS/global/hdb/backint/DB_LOAD/COMPLETE_DATA_BACKI

Figure 15. Specifying a recovery point in time

6. Select **Third-Party Backup Tool (Backint)** as the backup tool.
7. Review the recovery settings, and then click **Finish**.

Chapter 6 Conclusion

This chapter presents the following topics:

Summary	48
Findings	48

Summary

Using SAP HANA in TDI deployments with Dell EMC PowerMax enterprise storage arrays provides many benefits, including reducing hardware and operational costs, lowering risk, improving availability and performance, and increasing hardware vendor flexibility.

SAP has certified the PowerMax arrays for use in SAP HANA installations on production and nonproduction systems and on single-node (scale-up) and scale-out systems.

Findings

During our tests with SAP HANA on PowerMax arrays, we made the following observations:

- The SAP HANA-HWC-ES 1.1 certification scenario makes higher demands in terms of disk configuration.
- SAP HANA production installations on PowerMax systems require SSDs for the SAP HANA persistence.
- PowerMax arrays provide significant performance and scaling benefits.
- Using flash and SCM drives for the SAP HANA persistence provides benefits such as reduced time for SAP HANA startup, host auto-failover, and backup.
- PowerMax advanced data reduction with inline compression and deduplication provide storage efficiencies and capacity savings.
- Data Domain systems and DDBEA enable you to consolidate backup and archive SAP HANA systems with greater network efficiency, reducing your storage protection footprint through deduplication and compression.

Chapter 7 References

This chapter presents the following topics:

Dell Technologies documentation50
SAP HANA documentation.....50

Dell Technologies documentation

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

- [Dell EMC PowerMax: Family Overview Technical White Paper](#)
- [Dell EMC PowerMax Family Product Guide](#)
- [Dell EMC PowerMax: Service Levels for PowerMaxOS Technical White Paper](#)
- [Dell EMC PowerMax Reliability, Availability, and Serviceability Technical White Paper](#)
- [Dell EMC PowerMax: Data Reduction](#)
- [Dell EMC PowerMax and VMAX All Flash: eNAS Best Practices Technical White Paper](#)
- [VMware Virtualized SAP HANA with Dell EMC Storage Deployment Guide](#)
- [Dell EMC Host Connectivity Guide for Linux](#)
- [Dell EMC Data Domain Boost for Enterprise Applications and ProtectPoint Database Application Agent Installation and Administration Guide](#)

For additional documentation, see the [Dell Technologies Solutions Info Hub for SAP](#).

SAP HANA documentation

The following documentation on the [SAP website](#) provides additional relevant information:

- [SAP HANA Master Guide](#)
- [SAP HANA Server Installation and Update Guide](#)
- [SAP HANA Studio Installation and Update Guide](#)
- [SAP HANA Technical Operations Manual](#)
- [SAP HANA Administration Guide](#)
- [SAP HANA Native Storage Extension](#)
- [SAP HANA Storage Requirements](#)
- [SAP HANA Platform](#)
- [SAP HANA One](#)
- [SAP HANA Enterprise Cloud](#)
- [SAP HANA Tailored Data Center Integration](#)
- [SAP HANA Tailored Data Center Integration: Frequently Asked Questions](#)
- [Certified and supported SAP HANA hardware directory](#)
- [SAP Note 1943937—Hardware Configuration Check Tool - Central Note \(requires SAP username and password\)](#)
- [SAP Note 2399079—Elimination of hdbparam in HANA 2 \(requires SAP username and password\)](#)