

Dell Validated Design for SAP HANA TDI with Dell PowerMax Arrays

PowerMax 2000 and 8000 storage arrays

September 2022

H17290.4

Design Guide

Abstract

This design guide describes storage configuration best practices for SAP HANA in Tailored Data Center (TDI) deployments on Dell PowerMax storage arrays.

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. All Rights Reserved. Dell Technologies, Dell, 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 09/22 Design Guide H17290.4.

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

Contents

- Executive summary 4
- Technology overview..... 7
- Configuring and installing an SAP HANA scale-out cluster on a PowerMax array 22
- Backup and recovery for SAP HANA with Dell EMC Data Domain..... 37
- Conclusion..... 43
- References..... 44

Executive summary

SAP HANA

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.

TDI deployment models

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:

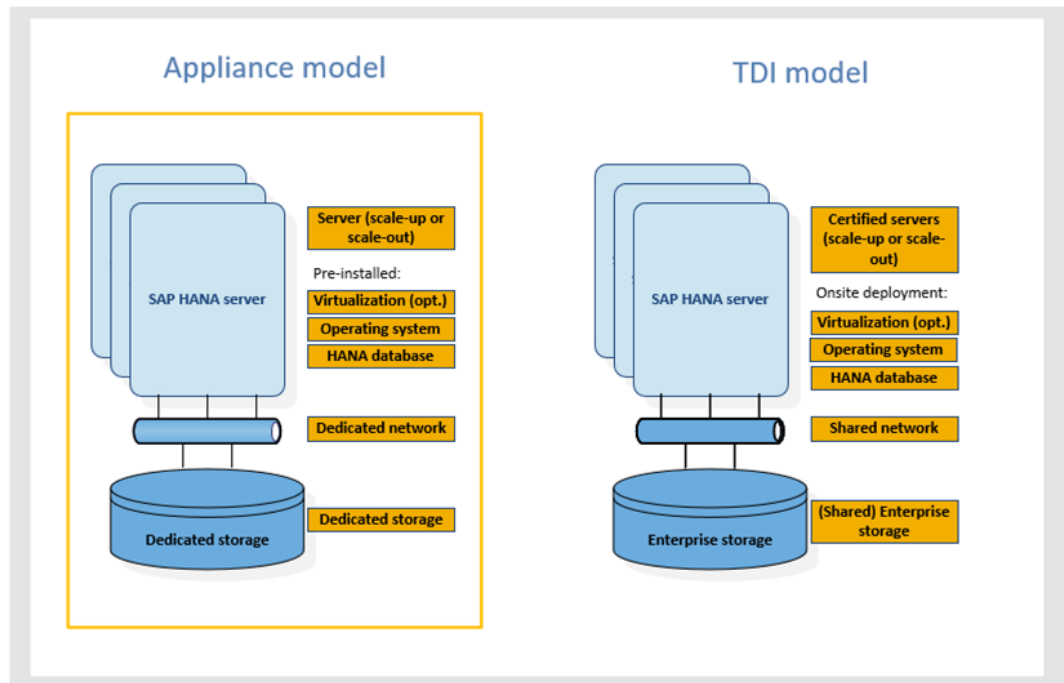


Figure 1. SAP HANA appliance model and TDI model comparison

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 HANA 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, data protection, monitoring, and management, reducing the time-to-value, costs, and risk of an overall SAP HANA adoption.

Solution overview

SAP has certified Dell EMC PowerMax storage arrays as meeting all performance and functional requirements for SAP HANA TDI deployments. This means that customers can use PowerMax arrays for SAP HANA TDI deployments in a fully supported environment with their existing data center infrastructures.

Note: For a complete list of the Dell EMC servers that SAP has certified, see the [SAP-Certified Dell Solutions website](#).

Using the SAP HANA Hardware Configuration Check Tool (hwcct), Dell EMC engineers performed extensive testing on the PowerMax family of products in accordance with the SAP HANA-HWC-ES-1.1 certification scenario. From these tests, we derived storage configuration recommendations for the PowerMax arrays that meet SAP performance requirements and ensure the highest availability for database persistence on disk.

Note: SAP recommends that TDI customers run the hwcct tool in their environment to ensure that their SAP HANA TDI implementation meets the SAP performance criteria.

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
- Avoid the significant risks and costs associated with operational change by using their 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 performance and scale benefits of PowerMax arrays to obtain real-time insights across the business
- Use flash drives for the SAP HANA persistence and benefit from reduced SAP HANA startup, host auto-failover, and backup times

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 high availability (HA) and key performance indicators (KPIs) for data throughput and latency for the TDI deployment. Topics that the guide addresses include:

- Introduction to the key technologies in the solution

Executive summary

- Configuration requirements and storage design principles for PowerMax storage with SAP HANA
- Best practices for deploying the SAP HANA database on PowerMax storage systems
- Example of an SAP HANA scale-out installation using PowerMax storage

Note: This validation guide describes SAP HANA TDI deployments in physical environments. If you plan to use SAP HANA in VMware virtualized environments on vSphere, 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 must configure a PowerMax storage array to 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: Dat Nguyen, Benjamin Simroth, Donagh Keeshan

Contributor: Aighne Kearney

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

Technology overview

PowerMax product family

The Dell EMC PowerMax product family consists of two models: the PowerMax 2000 array and the PowerMax 8000 array. The PowerMax 2000 array provides Dell EMC customers with efficiency and maximum flexibility in a 20U footprint, while the PowerMax 8000 array is designed for massive scale, performance, and IOPS density within a two-floor-tile footprint.

Both models are based on the Dell EMC Dynamic Virtual Matrix architecture and a version of HYPERMAX OS that was rewritten for the NVMe platform—PowerMaxOS 5978. PowerMaxOS can run natively on PowerMax systems and as an upgrade on legacy VMAX All Flash systems.

Like the VMAX All Flash arrays, PowerMax arrays are specifically designed to meet the storage capacity and performance requirements of an all-flash enterprise data center. PowerMax products take advantage of the higher-capacity NVMe flash drives used in the densest configuration possible. The PowerMax technology offers enterprise customers trusted data services along with the simplicity, capacity, and performance that their highly virtualized environments demand, while still meeting the economic needs of traditional storage workloads. In addition, the PowerMax platform enables you to deploy applications such as real-time analytics, machine learning, and big data, which require lower storage latency and higher IOPS densities than what was attainable with previous all-flash offerings.

The following figure shows the essential features of the PowerMax 2000 and 8000 arrays:



Figure 2. Dell EMC PowerMax product family

The main benefits that PowerMax platforms offer customers are:

- An end-to-end NVMe storage design providing industry-leading IOPS density in a system with a small footprint
- Readiness for next-generation data storage media, such as Storage Class Memory (SCM) and NVMe over Fabric (NVMe-oF) infrastructure

- Efficient workload consolidation for block, file, and mainframe workloads on a single platform
- Ultra-high reliability, availability, and serviceability for the enterprise
- Excellent data services, including advanced data reduction using inline deduplication and compression

System specifications

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

Table 1. PowerMax system specifications

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

Expandable modular architecture

PowerMax configurations consist of modular building blocks called PowerBricks. The modular PowerBrick 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 PowerBrick blocks are available for PowerMax configurations:

- The open systems PowerBrick block supports configurations with Fibre Channel (FC) or iSCSI connectivity with fixed-block architecture (FBA) device formatting. The PowerBrick block is also configurable for file storage using embedded NAS.
- The mainframe zPowerBrick block supports configurations with fiber connectivity (FICON) and count key data (CKD) device formatting.

The initial PowerBrick system includes a single engine consisting of two directors, two system power supplies (SPS), and two 24-slot 2.5" NVMe Drive Array Enclosures (DAE24) pre-configured with an initial total usable capacity. The PowerBrick 2000 array comes with an initial capacity of 11 or 13 TBu depending on the RAID configuration. The

PowerMax 8000 array comes with an initial capacity of 53 TBu for open systems, 13 TBu for mainframe systems, and 66 TBu for mixed systems. Each PowerBrick block comes preloaded with PowerMaxOS.

The PowerBrick concept makes it possible for the PowerMax platform to scale up and out. Customers can scale up by adding flash capacity packs, which include the NVMe flash drive capacity that can be added to a PowerMax array. Each flash capacity pack for the PowerMax 8000 system has 13 TBu of usable storage. For the PowerMax 2000 model each pack has 11 TBu or 13 TBu, depending on the RAID protection type selected. Scaling out a PowerMax system is done by aggregating up to two PowerBrick blocks for the PowerMax 2000 and up to eight blocks for the PowerMax 8000 in a single system with fully shared connectivity, processing, and capacity resources. Scaling out a PowerMax system by adding additional PowerBrick blocks produces a predictable linear performance improvement regardless of the workload.

Non-Volatile Memory Express (NVMe)

PowerMax arrays provide a full NVMe flash storage back-end for storing customer data. The PowerMax NVMe architecture provides:

- **I/O density with predictable performance**—PowerMax arrays are designed to deliver industry-leading I/O density. They are capable of delivering over 10 million IOPS in a two-rack system (two floor tiles), regardless of workload and storage capacity utilization.
- **NVMe storage density**—PowerMax delivers industry-leading NVMe TB per floor tile. While other all-flash alternatives use a proprietary flash drive design, PowerMax supports high-capacity, commercially available, dual-ported enterprise NVMe flash drives. PowerMax can therefore leverage increases in flash drive densities, economies of scale, and decreased time-to market.
- **Future-proof design**—The PowerMax NVMe design is ready for SCM flash and future NVMe-oF SAN connectivity options, which include 32 Gb FC and high-bandwidth converged Ethernet (RoCEv2 over 25 GbE/50GbE/100 GbE).

SCM drives on PowerMax systems

The Q3 2019 release of PowerMaxOS added support to the PowerMax 2000 and 8000 systems for storage class memory (SCM) drives. These drives are powered by dual port Intel Optane technology, together with Machine Learning (ML) algorithms, for data placement when both NVMe and SCM drives are used. SCM drives are new, high-performance drives that have a significantly lower latency than the NVMe flash drives, thus increasing response times.

A PowerMax model that contains both types of drive 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. The feature can also help to optimize access to storage groups that are used for workloads such as SAP HANA, which have higher priority service levels.

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

For information about how to configure and use SCM drives on PowerMax arrays, see [The Dell EMC PowerMax Family Overview](#).

Smart RAID

PowerMax uses a new active/active RAID group accessing scheme called smart RAID. Smart RAID allows the sharing of RAID groups across directors, giving each director active access to all drives on the PowerBrick or zPowerBrick block. Both directors on an engine can drive I/O to all the flash drives, creating balanced configurations in the system regardless of the number of RAID groups and providing performance benefits to customers.

Smart RAID also offers increased flexibility and efficiency. Customers can order PowerMax systems with a single RAID group, allowing for at least nine drives per engine with RAID 5 (7+1) or RAID 6 (6+2 and 1 spare) and as few as five drives per system for a PowerMax 2000 with RAID 5 (3+1 and 1 spare), leaving more drive slots available for capacity upgrades in the future. When the system is scaled up, customers have more flexibility because Flash Capacity Pack increments can be a single RAID group.

Data services

Remote replication with SRDF

The Symmetrix Remote Data Facility (SRDF) is considered a gold standard for remote replication in the enterprise data center. Up to 70 percent of Fortune 500 companies use SRDF to replicate their critical data to geographically dispersed data centers throughout the world. SRDF offers customers the ability to replicate tens of thousands of volumes, with each volume being replicated to a maximum of four different locations globally.

SRDF is available in three types:

- **SRDF Synchronous (SRDF/S)**—SRDF/S delivers zero data loss remote mirroring between data centers that are separated by up to 60 miles (100 km).
- **SRDF Asynchronous (SRDF/A)**—SRDF/A delivers asynchronous remote data replication between data centers that are up to 8,000 miles (12,875 km) apart. SRDF/S and SRDF/A can be used together to support three or four site topologies, as required by the world's most mission-critical applications.
- **SRDF/Metro**—SRDF/Metro delivers active-active HA for non-stop data access and workload mobility within a data center, or between data centers separated by up to 60 miles (100 km).

Local replication with TimeFinder SnapVX

Every PowerMax array includes the local replication data service TimeFinder SnapVX, which creates low-impact snapshots. Local replication with SnapVX starts as efficiently as possible by creating a snapshot, a pointer-based structure that preserves a point-in-time (PIT) view of a source volume. Snapshots do not require target volumes. Instead, they share back-end allocations with the source volume and other snapshots of the source volume, and only consume additional space when the source volume is changed. A single source volume can have up to 256 snapshots.

You can access a PIT snapshot from a host by linking it to a host-accessible volume known as a target. Target volumes are standard thin volumes. Up to 1,024 target volumes can be linked to the snapshots of a single source volume. By default, targets are linked in no-copy mode. This functionality significantly reduces the number of writes to the back-end flash drives by eliminating the requirement to perform a full-volume copy of the

source volume during the unlink operation to continue using the target volume for host I/O.

Advanced data reduction using inline compression and deduplication

PowerMax arrays use the 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 those that are repeated once only, preventing the consumption of critical core resources on the PowerMax system and limiting performance impact.

Embedded NAS

The embedded NAS (eNAS) data service extends the value of PowerMax to file storage by enabling customers to use vital enterprise features such as flash-level performance for both block and file storage, as well as to simplify management and reduce deployment costs.

The eNAS data service uses the hypervisor in PowerMaxOS to create and run a set of virtual machines (VMs) within the PowerMax array. These VMs host two major elements of eNAS software—data movers and control stations. The embedded data movers and control stations have access to shared system resource pools so that they can evenly consume PowerMax resources for both performance and capacity.

With eNAS, PowerMax becomes a unified block-and-file platform that uses a multicontroller, transactional NAS solution.

Design principles and recommendations for SAP HANA on PowerMax arrays

Overview

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 PowerMax
- PowerMax scalability for SAP HANA
- Masking views
- PowerMaxOS service levels and competing workloads
- SAP HANA NSE

SAP HANA capacity requirements

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 storage area network, or 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 file system 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):

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

$$\text{Size}_{\text{installation}(\text{scale-out})} = 1 \times \text{RAM_of_worker} \text{ per } 4 \text{ worker nodes}$$
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 in order 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 prepare the 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 ($2 \text{ TB} / 4 = 512 \text{ GB} \times 1.2 = 616 \text{ GB}$).

If the net disk space is unknown at the time of storage sizing, Dell EMC recommends 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:

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

Backup

SAP HANA supports backup to a file system or the use of SAP-certified third-party tools. Dell EMC 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 file system 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 file system on a different PowerMax array, see the [SAP HANA Storage Requirements](#) 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 [Backup and recovery for SAP HANA with Dell EMC Data Domain](#).

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. If, for example, 10 nodes are connected in a SAN to a PowerMax array, a total bandwidth of 4,000 MB/s is required. Assuming that a 16 Gb/s front-end port provides approximately 1,500 MB/s bandwidth, at least three dedicated 16 Gb/s front-end ports are required to support 10 SAP HANA nodes ($3 \times 1,500 \text{ MB/s} = 4,500 \text{ MB/s}$).

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 just affect the storage front-end configuration. In the example with 10 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 maximum required bandwidth.

Storage ports

When you are planning storage connectivity for performance and availability, Dell EMC recommends “going wide before going deep.” In other words, it is better to connect storage ports across different engines and directors than to use 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 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 Dell EMC recommends connecting 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. We therefore 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, port 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.

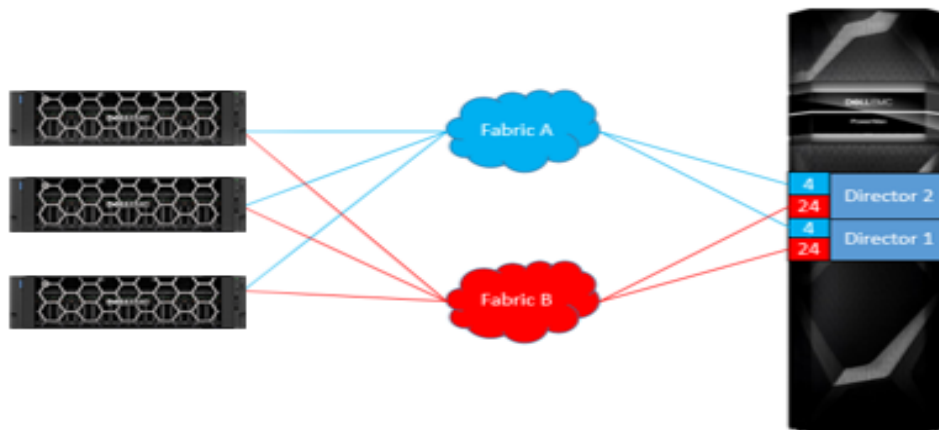


Figure 3. SAN connectivity in a single engine environment

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

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 from 4 KB to 64 MB in size. 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 and 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 from 4 KB to 1 MB in size. SAP HANA keeps a 1 MB buffer for the redo log in memory. 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 PowerMax

In an SAP HANA scale-out implementation, install the SAP HANA database binary files on a shared file system that is exposed to all hosts of a system under the `/hana/shared` mount point. If a host must write a memory dump which can read up to 90 percent of the RAM size, the memory dump is stored in this file system. Based on the customer's 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 single engine using the SAP hwcct 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 given 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 higher models and additional PowerBrick blocks by extrapolating the PowerMax 2000 test results using the performance characteristics of the higher models.

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

Masking view

PowerMax uses masking views to assign storage to a host. Dell EMC recommends creating a single masking view for each SAP HANA 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 given pair of data and log volumes.

PowerMaxOS service levels and competing workloads

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 ms (0.4 ms with SCM drives)
Platinum	0.8 ms (0.6 ms with SCM drives)
Gold	1 ms
Silver	3.6 ms
Bronze (lowest priority)	7.2 ms
Optimized	N/A

Diamond, platinum, and gold service levels have an 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 ms while platinum drops to 0.6 ms.

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 Dell EMC recommends for different SAP HANA installation types:

Table 4. Recommended SLO and workload configurations for SAP HANA

Installation type	SLO	Reason	Benefits
SAP HANA persistence (data and log) for production installations	Diamond	Ensures that the PowerMax tries to keep the latency below 1 ms, which is a SAP requirement for small (4 KB and 16 KB) block sizes on the log volume.	Using the diamond SLO with all-flash devices provides the following benefits: <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 needs to read the data from the data volume • Sub-millisecond 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.	
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 of them 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 utilize 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 desired 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 a 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:

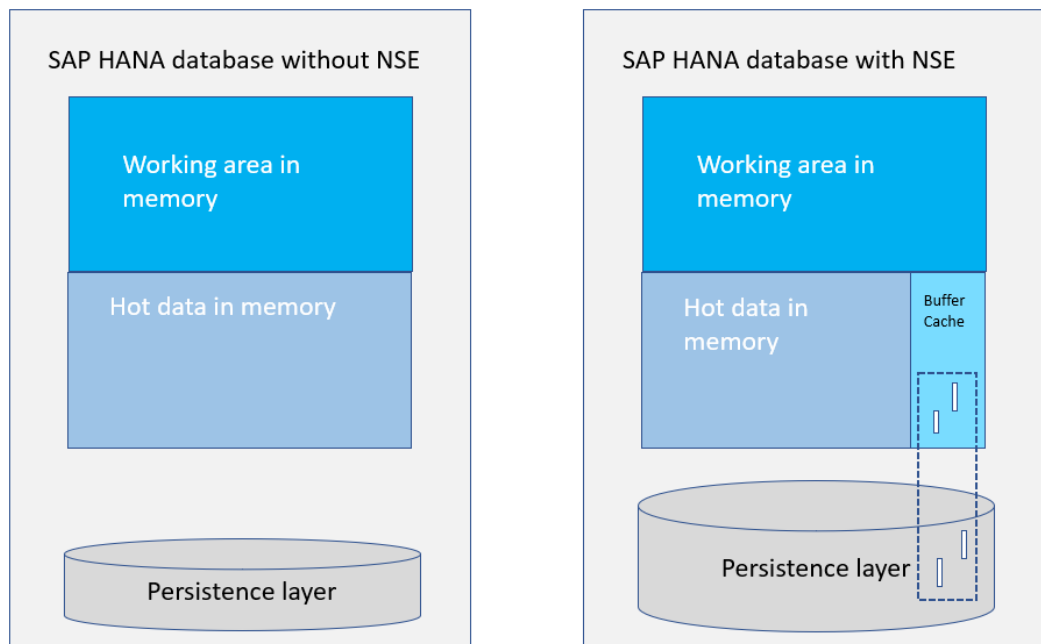


Figure 4. Standard SAP HANA database and SAP HANA database with NSE comparison

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 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: Only scale-up SAP HANA systems are currently 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 data is loaded into memory page by page as required for query processing. Unlike column-loadable data, page-loadable data does not need 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. The PowerMax storage array 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 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).

Configuring and installing an SAP HANA scale-out cluster on a PowerMax array

Overview

This section describes how to configure and install SAP HANA scale-out cluster on a PowerMax array. The procedure involves:

- 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 hosts
- Installing the SAP HANA cluster using the SAP lifecycle management command-line tool hdblcm

Configuring the PowerMax array

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 configure the PowerMax array:

1. Log in to Unisphere for PowerMax, as shown in the following figure.

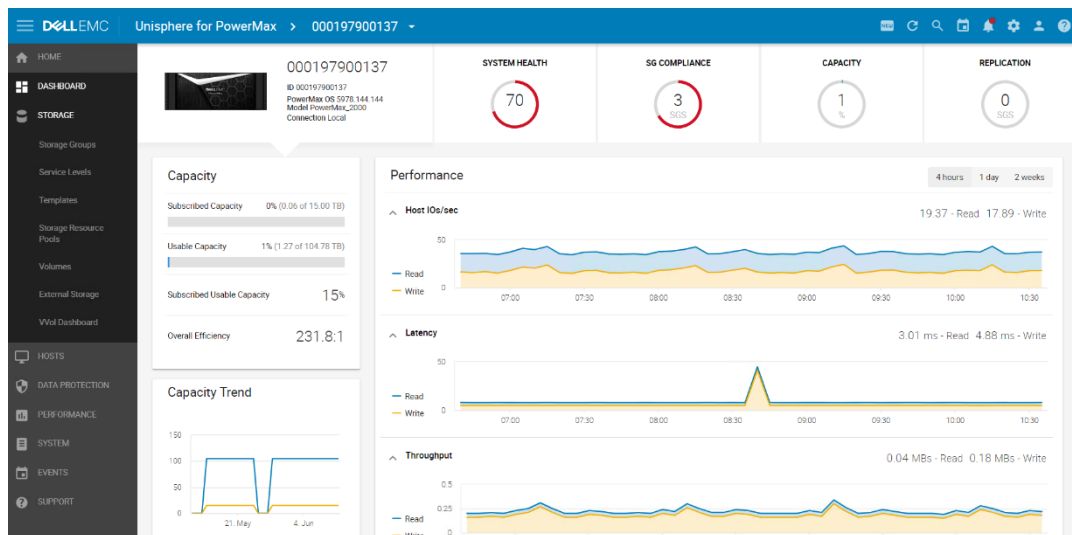


Figure 6. Unisphere for PowerMax

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

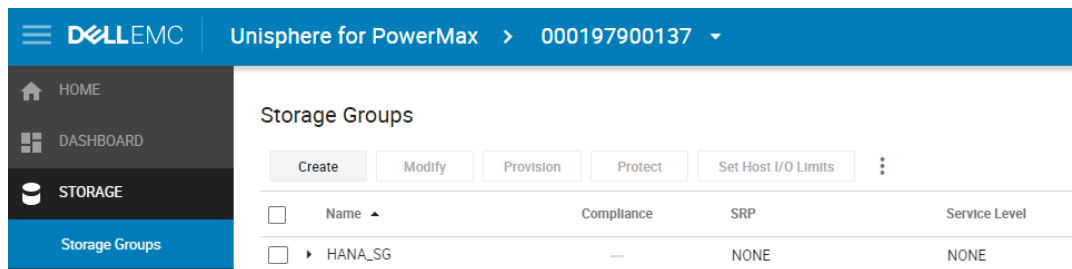


Figure 7. Create storage groups

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

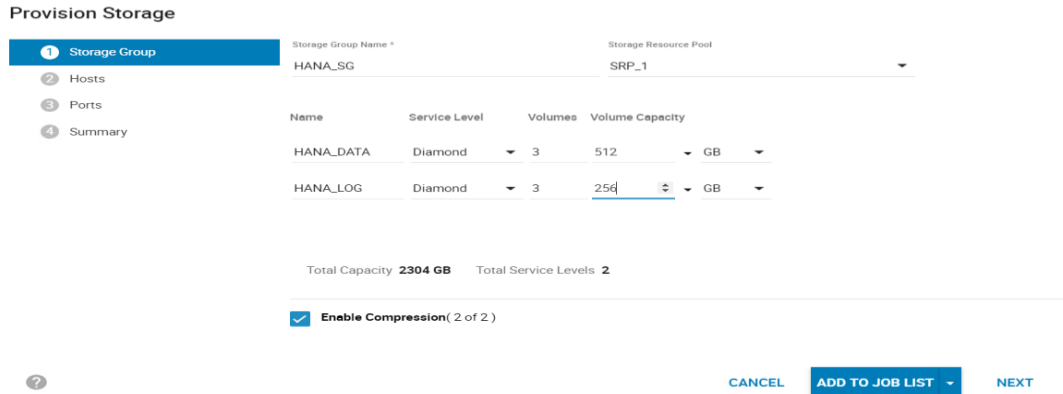


Figure 8. Provision Storage screen

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**
 - Sub-group for all data volumes: **HANA_DATA**
 - Sub-group for all log volumes: **HANA_LOG**
4. Specify the number and size of the volumes to be created. Then click the down arrow on **Add to Job List** and select **Run Now**.
The new cascaded storage group is created.
 5. To view information about the volumes that you created, click the hyperlink in the right-hand pane next to **Volumes**, as shown in the following figure:

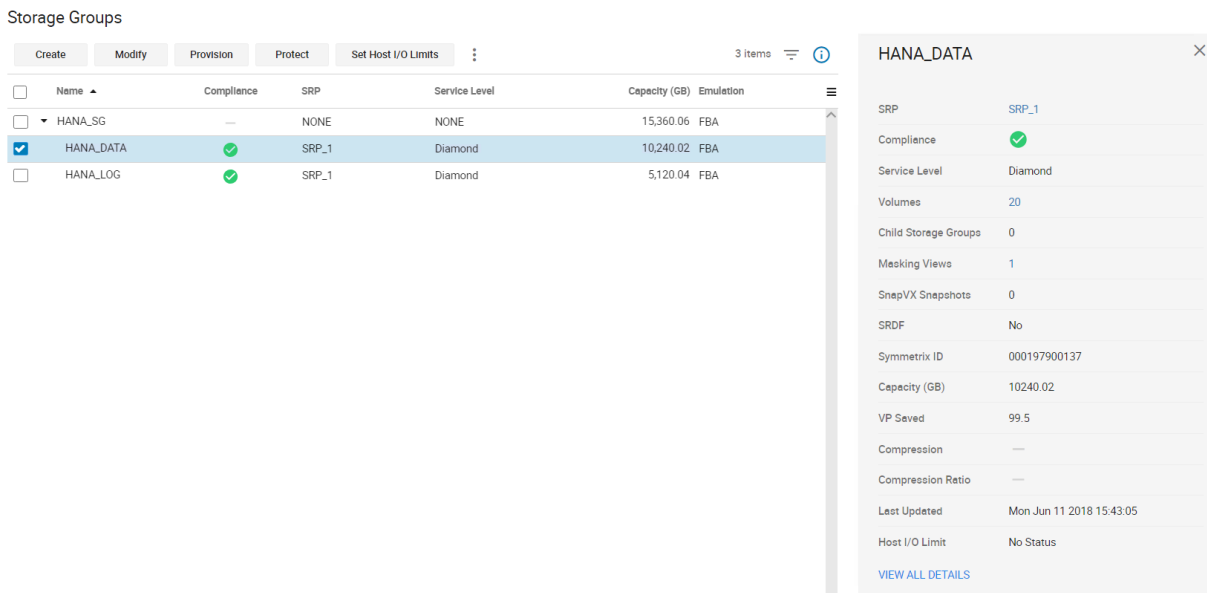


Figure 9. HANA_DATA storage groups

The Volumes page opens, as shown in the following figure.

Name	Type	Allocated (%)	Capacity (GB)	Emulation	Status
<input type="checkbox"/> 00024	TDEV	8%	512.00	FBA	Ready
<input checked="" type="checkbox"/> 00025	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 00026	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 00027	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 00028	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 00029	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 0002A	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 0002B	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 0002C	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 0002D	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 0002E	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 0002F	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 00030	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 00031	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 00032	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 00033	TDEV	0.05%	512.00	FBA	Ready
<input type="checkbox"/> 00034	TDEV	0.05%	512.00	FBA	Ready

Masking Info	120
Storage Groups	1
SRP	1
FBA Front End Paths	4
RDF Info	0
Volume Name	00025
Physical Name	---
Volume Identifier	---
Type	TDEV
Status	Ready
Reserved	No
Capacity (GB)	512
Capacity (MB)	524289
Capacity (CYL)	279621
Compression Ratio	1.0:1

Figure 10. Volumes page

6. Select a volume and note the WWN of the volume in the right-hand panel. You may need to scroll to find the WWN of the volume.
7. Repeat the previous 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	Initiators in Host
<input type="checkbox"/> 10000090fab37196	<input type="checkbox"/> 10000090fab36738
<input type="checkbox"/> 10000090fab397da	<input type="checkbox"/> 10000090fab36eac
<input type="checkbox"/> 10000090fab3981a	
<input type="checkbox"/> 10000090fab39856	
<input type="checkbox"/> 10000090fab39876	

Set Host Flags

CANCEL **ADD TO JOB LIST**

Figure 11. Create Host page

2. Enter a host name, 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 on **Add to Job List** and select **Run Now**.
The new host is created.
4. To create a host group, select **Create > Create Host Group** and select the hosts that belong to the SAP HANA cluster. Then 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

5. Click the down arrow on **Add to Job List** and select **Run Now**.
The new host group is created.

Create a port group

To create a port group:

1. Select **Hosts > Port Groups > Create**. 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

Figure 13. Creating a port group

2. Mark the ports your initiators are logged into by holding down the Control key. Then click the down arrow on **Add to Job List** and select **Run Now**.
3. Click **OK** if a warning message appears 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, Dell EMC recommends 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

Storage Group Name * HANA_SG Storage Resource Pool SRP_1

Name	Service Level	Volumes	Volume Capacity
HAN_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 on **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. Creating a Masking View

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. Make sure that you correctly prepare the SAP HANA nodes and the SAP HANA `global.ini` file 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

The following sections of this guide describe how to:

- Configure storage on the SAP HANA nodes
- Prepare the SAP HANA `global.ini` file
- Install an SAP HANA scale-out instance with the SAP HANA lifecycle management tool `hdblcm` by using the storage volumes you created in the preceding steps.

Prerequisites

The configuration example in this guide assumes 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 (in this example we used SUSE Linux 12 SP3 for SAP applications).
- An SAP HANA shared file system, `hana/shared/`, has been created on a NAS system, for example, on a PowerMax 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 is downloaded from the SAP website and made 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)

1. Use the entries shown in the following figure in the `/etc/multipath.conf` file.

```

device {
    vendor "EMC"
    product "SYMMETRIX"
    getuid_callout "/lib/udev/scsi_id -g -u -d /dev/%n"
    features "0"
    hardware_handler "0"
    path_selector "round-robin 0"
    path_grouping_policy "multibus"
    rr_weight "uniform"
    no_path_retry "fail"
    rr_min_io "100"
    path_checker "directio"
    prio "const"
    prio_args ""
    flush_on_last_del yes
    fast_io_fail_tmo off
    dev_loss_tmo 120
}

```

Figure 2. Linux native multipathing command

2. After changing the `/etc/multipath.conf` file, restart multipathing by running the following command:

```
service multipath restart
```

Note: For information about the operating system version and storage array MPIIO settings for native multipathing, see 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. To achieve this, you can 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 sdbe 67:128 active ready running
  |- 3:0:3:7 sddp 71:112 active ready running
  `-- 4:0:3:7 sdcx 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 sdbc 67:96 active ready running
  |- 3:0:3:5 sddm 71:64 active ready running
  `-- 4:0:3:5 sdcq 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 sdcr 69:240 active ready running
```

Figure 3. 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 sdeq 129:32 active ready running
  `-- 4:0:3:21 sddl 71:48 active ready running

```

Figure 4. List log volumes command

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

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

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

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

Configure 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 take into account 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 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 specific granularities:

- Tables
- Columns

- Partitions

Data location handling is built into SAP HANA's 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, go to the [SAP Help Portal](#).

Installing the SAP HANA scale-out cluster

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 `hdblcm` 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 relatime,inode64
partition_*_log__mountoptions = -o relatime,inode64,nobarrier
partition_1_data__wwid = 360000970000197900137533030303234
partition_1_log__wwid = 360000970000197900137533030303338
partition_2_data__wwid = 360000970000197900137533030303235
partition_2_log__wwid = 360000970000197900137533030303339
partition_3_data__wwid = 360000970000197900137533030303236
partition_3_log__wwid = 360000970000197900137533030303341
```

Figure 5. Global.ini file storage content

2. Ensure that the partition entries match the unique device identifier displayed. Do this by running the `multipath -ll` command with a preceding **3**.
3. Place the `global.ini` file (this name is mandatory) in a directory on the `/hana/shared/` file system—for example, `/hana/shared/PMX_cfg`.

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

Prepare the installation parameter 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 parameter 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:

```
Directoryroot to search for components
component_root=/SAPShare/software/SAP_HANA_SPS11_IM/51
050506/
```

```
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. You can specify passwords for the root user, SAP Host Agent User (`sapadm`), system administrator user (`<sid>adm`), and database user (`SYSTEM`) in the parameter file. You can also use encrypted passwords. For more information, see the [SAP HANA Installation and Upgrade 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 specific environment.

Installing the SAP HANA scale-out cluster

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

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

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		
async_read_submit	on	
async_write_submit_active	on	
async_write_submit_blocks	all	
max_parallel_io_requests	64	
max_parallel_io_requests[data]		● 256
max_parallel_io_requests[log]		● 256
max_submit_batch_size	64	
min_submit_batch_size	16	
num_completion_queues	1	
num_submit_queues	1	

Figure 6. 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 provided 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.

Backup and recovery for SAP HANA with Dell EMC Data Domain

Overview

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 specification 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@c460-09q:/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 7. 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 8. Modifying the configuration file parameter settings

5. Configure the DDBDA lockbox by using `ddbadmin -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.

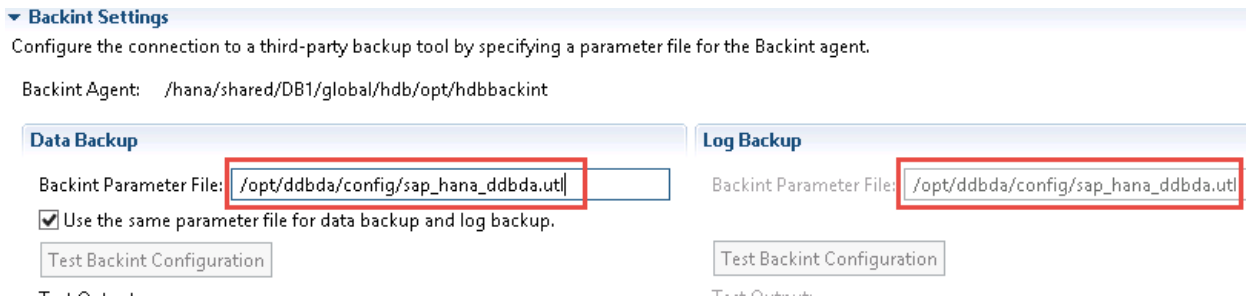


Figure 9. Specifying the DDBEA configuration file location

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.

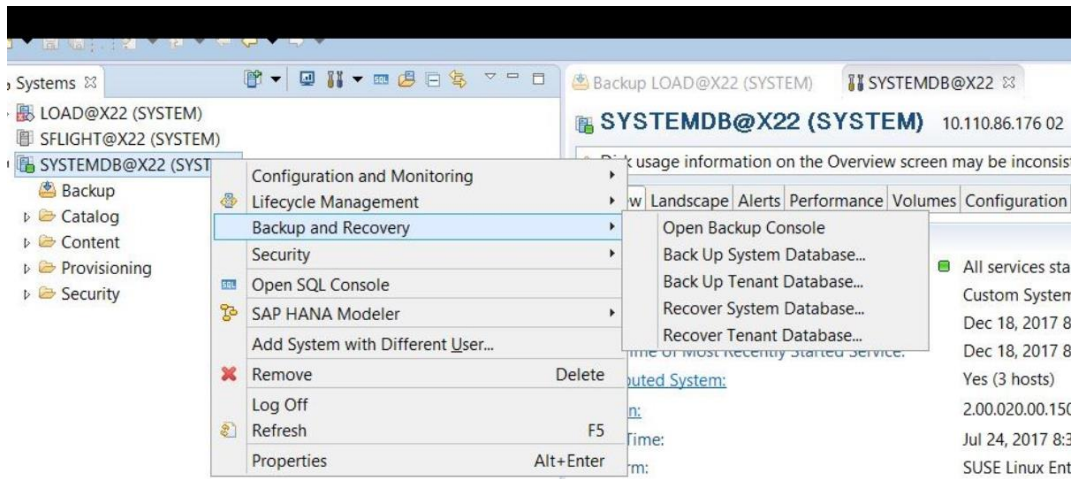


Figure 10. SAP HANA Studio Backup and Recovery screen

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

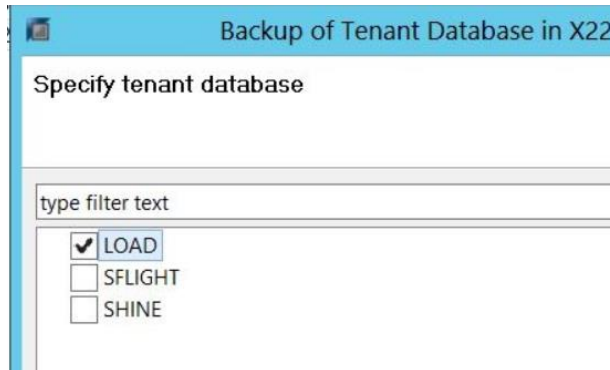


Figure 11. Specifying the tenant database

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

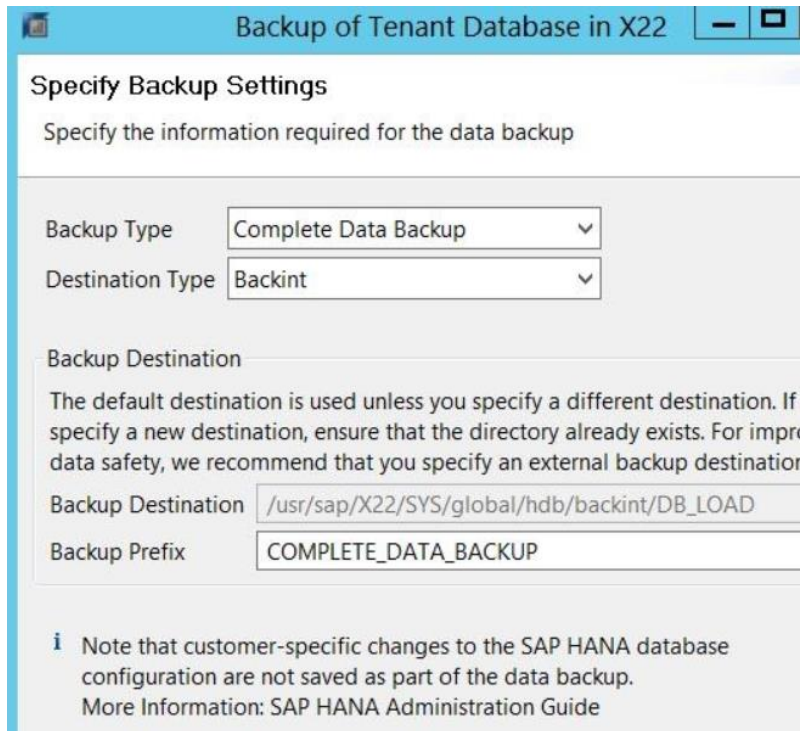


Figure 12. Specifying the backup destination

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

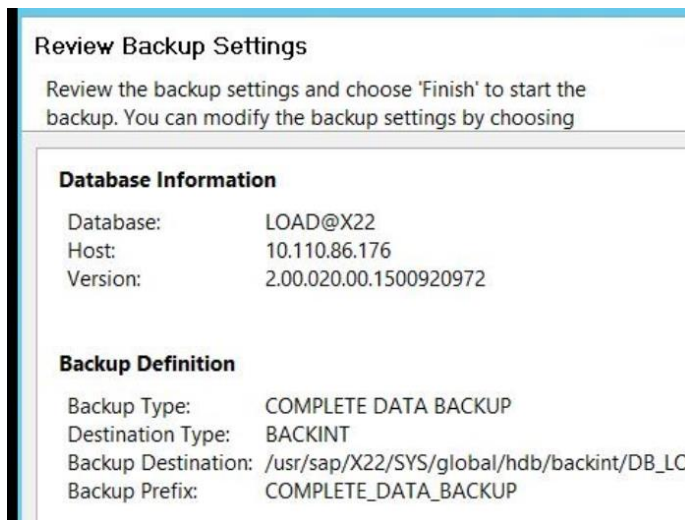


Figure 13. Review Backup Settings

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.

The screenshot shows a window titled "Recovery of Tenant Database in X22". The main heading is "Specify Recovery Type" with the instruction "Select a recovery type." Below this, there are three radio button options:

- Recover the database to its most recent state
- Recover the database to the following point in time
- Recover the database to a specific data backup

 The second option is expanded to show input fields:

- Date: 2018-01-31
- Time: 13:26:29
- Select Time Zone: (GMT+08:00) China Standard Time
- System Time Used (GMT): 2018-01-31 05:26:29

Figure 14. Specifying the recovery type

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

The screenshot shows a window titled "Recovery of Tenant Database in X22". The main heading is "Locate Log Backups" with the instruction "Specify location(s) of log backup files to be used to recover the database." Below this, there is an information icon and text:

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.

 Below the text is a list of locations:

- Locations: /usr/sap/X22/SYS/global/hdb/backint/DB_LOAD/
- /usr/sap/X22/HDB02/backup/log/DB_LOAD

 To the right of the list are three buttons: "Add", "Remove All", and "Remove".

Figure 15. 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 16. 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**.

Conclusion

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 observed the following:

- 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.
- 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.
- Using SSDs for the SAP HANA persistence provides significant benefits, including:
 - Reduced SAP HANA startup and host auto-failover times
 - Reduced SAP HANA backup and recovery times
 - No need to consider spindle count because initial array and disk configuration can be performed based on capacity

References

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.

- [The Dell EMC PowerMax Family Overview](#)
- [Dell EMC PowerMax Family Product Guide](#)
- [Dell EMC PowerMax: Service Levels for PowerMaxOS White Paper](#)
- [Dell EMC PowerMax Reliability, Availability and Serviceability Technical White Paper](#)
- [Dell EMC PowerMax: Data Reduction](#)
- [Dell EMC VMAX3 and VMAX All Flash ENAS Best Practices](#)
- [VMware Virtualized SAP HANA with Dell EMC Storage Solution 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 EMC 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](#)

Web resources

- [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](#)

Note: The following documentation requires an SAP username and password.

- [SAP Note 1943937 - Hardware Configuration Check Tool - Central Note](#)
- [SAP Note 2399079—Elimination of hdbparam in HANA 2](#)