

Storage Configuration Best Practices for SAP HANA TDI on Dell EMC VPLEX Systems

VPLEX VS6 systems

August 2018

H17323

Validation Guide

Abstract

This validation guide describes storage configuration best practices for SAP HANA TDI on VPLEX systems. The solution enables you to use VPLEX systems for TDI deployments in a supported environment with existing data center infrastructures.

Dell EMC 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 © 2018 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 08/18 Validation Guide H17323.

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 and considerations.....7**
- Design principles and recommendations for SAP HANA on VPLEX systems.....9**
- Configuring and installing an SAP HANA scale-out cluster on a VPLEX system.....19**
- Backing up and recovering the SAP HANA database.....31**
- Conclusion.....37**
- References.....38**

Executive summary

Introduction

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 SAP HANA database is at the core of this real-time data platform.

SAP HANA combines SAP software components that are optimized on proven hardware that is provided by SAP partners. Two models are available for on-premises deployment, as the following figure shows.

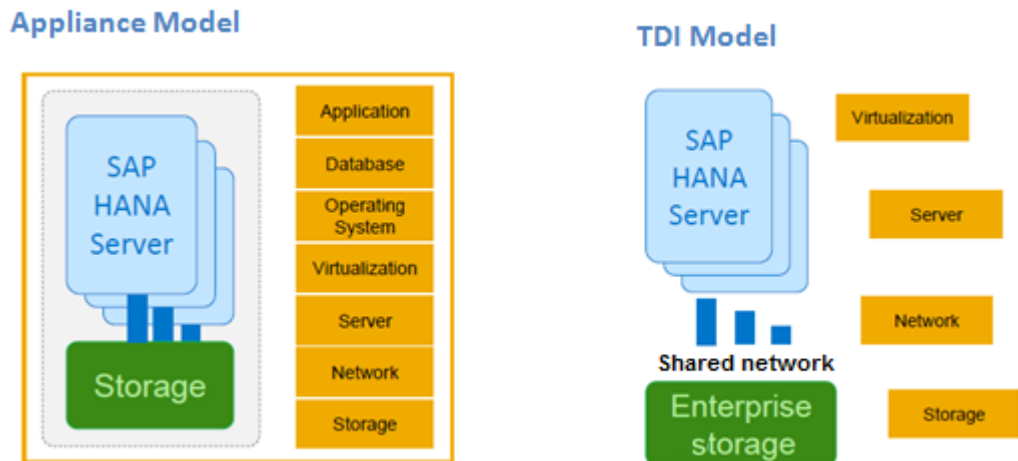


Figure 1. SAP HANA appliance model and TDI model comparison (© SAP SE)

Appliance model

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

Customers using the SAP HANA appliance model have experienced 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 tailored data center integration (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, and reduce the time-to-value, costs, and risk of an overall SAP HANA adoption.

Solution overview

SAP certifies Dell EMC VPLEX systems as meeting all performance and functional requirements for SAP HANA. That means that customers can use VPLEX systems for SAP HANA TDI deployments in a fully supported environment with their existing data center infrastructures. SAP allows storage virtualization with a system such as VPLEX only in combination with a certified SAP HANA TDI enterprise storage system. For more information, see the [Certified and Supported SAP HANA Hardware Directory](#).

Using the SAP HANA hardware configuration check tool (hwcct), Dell EMC engineering performed extensive testing on the VPLEX product family in accordance with the SAP HANA-HWC-ES-1.1 certification scenario. The storage configuration recommendations for the VPLEX systems that Dell EMC engineering derived from the tests 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 specific 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 VPLEX storage systems can:

- Integrate the SAP HANA platform into an existing data center.
- Use VPLEX storage to rely on already-available, multisite concepts and benefit from established automation and operations processes.
- Transition easily from an appliance-based model to the VPLEX-based TDI architecture while relying on Dell EMC Professional Services to minimize risk.
- Avoid the significant risks and costs that are associated with operational change by using existing operational processes, skills, and tools.
- Use the performance and scale benefits of VPLEX systems to obtain real-time insights across the business.

Document purpose

This guide describes a solution that uses SAP HANA in a TDI deployment scenario on VPLEX systems. The guide provides configuration recommendations for the TDI deployment based on SAP requirements for high availability (HA) and on the performance test results that are needed to meet the SAP key performance indicators (KPIs) for data throughput and latency. Topics include:

- Key technologies in the solution
- Configuration requirements and storage design principles for VPLEX systems with SAP HANA
- Best practices and tips for deploying the SAP HANA database on VPLEX systems
- Example of an SAP HANA scale-out installation using VPLEX systems

Note: This validation guide describes SAP HANA TDI deployments in physical environments. If you plan to use SAP HANA in virtualized environments on VMware vSphere, see the [VMware Virtualized SAP HANA with Dell EMC Storage Solution Guide](#).

Audience

This document is intended for system integrators; SAP system, database, and storage administrators; customers; partners; and members of Dell EMC Professional Services who need to configure a VPLEX system for use in a TDI environment for SAP HANA.

We value your feedback

Dell EMC and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the [Dell EMC Solutions team](#) with your comments.

Authors: Fergal Murphy, Donagh Keeshan, Aighne Kearney

Technology overview and considerations

VPLEX system overview

VPLEX systems enable connectivity to heterogeneous storage arrays with the ability to manage storage that is provisioned from multiple arrays from a single interface within a data center. The VPLEX system resides between the servers and the storage assets and uses a unique clustering architecture that gives servers at multiple data centers read/write access to shared block storage devices. The VPLEX VS6 architecture makes use of Dell EMC's 20-plus years of expertise in designing and implementing enterprise-class, intelligent-cache, and distributed data protection solutions, and represents the next-generation architecture for data mobility and continuous availability. VPLEX technology enables transparent movement of data between storage arrays for the connected hosts without disruption to operations and protects the data from failures of single components and disasters.

VPLEX systems address these distinct use cases:

- **Data mobility**—The ability to move applications and data across storage arrays within the same data center, between data centers across a campus, or within a geographical region
- **Continuous availability**—The ability to create a continuously available storage infrastructure with unmatched resiliency across the same data center or between data centers

VPLEX cluster configuration

A VPLEX VS6 cluster consists of one, two, or four engines to cater to different performance requirements. Each VS6 engine contains two directors and has dual integrated management module control station (MMCS) modules in engine 1 that replace the management server. Dual-engine and quad-engine clusters also contain a pair of InfiniBand switches for communication between directors within the cluster. Each engine is protected by integrated battery backup (BBU) modules on the VS6 engine, and each internal switch gets its power through an uninterruptible power supply (UPS).

The following table shows the VS6 engine cluster configuration:

Table 1. VPLEX VS6 engine cluster configuration

	Single-engine cluster	Dual-engine cluster	Quad-engine cluster
Number of engines	1	2	4
Front-end 16 Gb/s Fibre Channel (FC) ports	8	16	32
Back-end 16 Gb/s FC ports	8	16	32
WAN communication	MetroIP 10 GbE or MetroFC 16 Gb FC		
Inter-engine communication	InfiniBand 40 Gb/s		
Inter-engine communication switch	12-port 40 Gb/s InfiniBand		

If performance demands grow over time, the VPLEX system can scale from a single-engine to a dual-engine cluster or from a dual-engine to a quad-engine cluster. These upgrades are fully supported and are accomplished nondisruptively.

VPLEX deployment options

You can deploy your VPLEX system in one of the following ways:

- **VPLEX Local**—Uses a single VPLEX cluster within one data center, offering the following benefits:
 - Ability to combine Dell EMC and non-Dell EMC storage arrays by allowing transparent data mobility for fast and simple data movement and technology refreshes
 - Improved storage utilization by using pooling and capacity aggregation across multiple arrays
 - Increased data protection HA by mirroring data across arrays
 - Standardized LUN presentation and management
- **VPLEX Metro**—Uses two VPLEX clusters across synchronous distances over an intercluster link with less than 10 ms round-trip time, offering the following benefits:
 - Transparent data relocation over distance
 - Disaster protection
 - Efficient collaboration
 - Data mirroring to another site with full access at near-local speeds
 - Simplified management through one interface

Scale-out upgrades from a VPLEX Local to a VPLEX Metro without disruption to normal operations are supported.

Design principles and recommendations for SAP HANA on VPLEX systems

Overview

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

- Underlying certified storage arrays
- SAP HANA capacity requirements
- SAP HANA I/O patterns
- SAP HANA shared file systems
- SAN connectivity for VPLEX
- VPLEX scalability for SAP HANA
- Competing workloads

Underlying certified storage arrays

Storage virtualization with a VPLEX system is available for SAP HANA only in combination with a certified SAP HANA TDI enterprise storage system. Dell EMC provides a range of certified enterprise storage arrays that can be used as the underlying storage with a VPLEX system. For best-practice recommendations for configuring the underlying storage volumes, see the relevant guide in the following list:

- [*Storage Configuration Best Practices for SAP HANA Tailored Data Center Integration on Dell EMC VMAX and VMAX3 Storage Systems*](#)
- [*Dell EMC XtremIO Storage Configuration Best Practices for SAP HANA TDI Configuration Guide*](#)
- [*Storage Configuration Best Practices for SAP HANA TDI on Dell EMC PowerMax Arrays Validation Guide*](#)
- [*Storage Configuration Best Practices for SAP HANA TDI on Dell EMC Unity Storage Systems Solution Guide*](#)
- [*Storage Configuration Best Practices for SAP HANA Tailored Data Center Integration on Dell EMC Compellent Storage Systems*](#)

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)
- Backups

Operating system boot image

For the SAP HANA nodes to be able to boot from a volume on a storage array (that is, 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.

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

SAP HANA installation (/hana/shared/)

For installation of the SAP HANA binaries and the configuration files, traces, and logs, every SAP HANA node requires access to a file system that is mounted under the local mount point, `/hana/shared/`. An SAP HANA scale-out cluster requires a single shared file system, which must be mounted on every node. Most SAP HANA scale-out installations use an NFS server-based shared file system for this purpose. NAS systems such as Dell EMC PowerMax embedded NAS (eNAS) arrays, Dell EMC Unity file systems, and Dell EMC Isilon systems also provide this `/hana/shared/` file system. You can calculate the size of the `/hana/shared/` file system by using the formula in [SAP HANA Storage Requirements](#). Version 2.10 of the paper 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):

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

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 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 and worker (scale-out) node requires two disk volumes to save the in-memory database on disk (data) and to 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. To help you prepare the disk sizing, SAP provides references to tools and documents in [SAP HANA Storage Requirements](#). Version 2.10 of the paper provides the following formula to calculate the size of the data volume:

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

In this formula, `net disk space for data` 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 four worker nodes, every node must have a data volume of 616 GB 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, 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:

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

Backups

SAP HANA supports backup to a file system or backups to SAP-certified third-party tools. Dell EMC supports data protection strategies for SAP HANA backups using Dell EMC Data Domain systems and Dell EMC Networker software. Although it is possible to back up an SAP HANA database to an NFS file system on a Dell EMC array, Dell EMC does 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 Dell EMC storage array, see [SAP HANA Storage Requirements](#) for details about sizing the backup file system. The capacity depends on not only the data size and the frequency of change operations in the database but also on the number of backup generations kept on disk.

For more information, see [Backing up and recovering the SAP HANA database](#) on page 31.

SAP HANA I/O patterns

The SAP HANA persistent devices use multiple I/O patterns, as described in this section.

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/Os to the data file system. During normal operations, most of the I/Os to the data file system are 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 of the I/Os to the log volume are writes; data is read from the log volume only during database restart, HA failover, log backup, or database recovery.

SAP HANA shared file systems

In an SAP HANA scale-out implementation, install the SAP HANA database binaries 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. Depending on your infrastructure and requirements, the following options are available:

- Using an NFS server-based shared file system.
- Using a NAS system such as Dell EMC VMAX3 eNAS, PowerMax eNAS, or Unity to provide an NFS share for the SAP HANA shared file system.
- Using VMAX3 eNAS, PowerMax eNAS, or Unity block storage to create a shared file system using a cluster file system such as General Parallel File System (GPFS) or Oracle Cluster File System 2 (OCFS2) on top of the block LUNs. SUSE provides OCFS2 capabilities with the 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.

SAN connectivity for VPLEX systems

This section describes hardware connectivity best practices for connecting VPLEX systems to the SAN. These best practices are based on dual fabric SAN.

A single VPLEX VS6 engine consists of two directors, A and B, as the following figure shows:

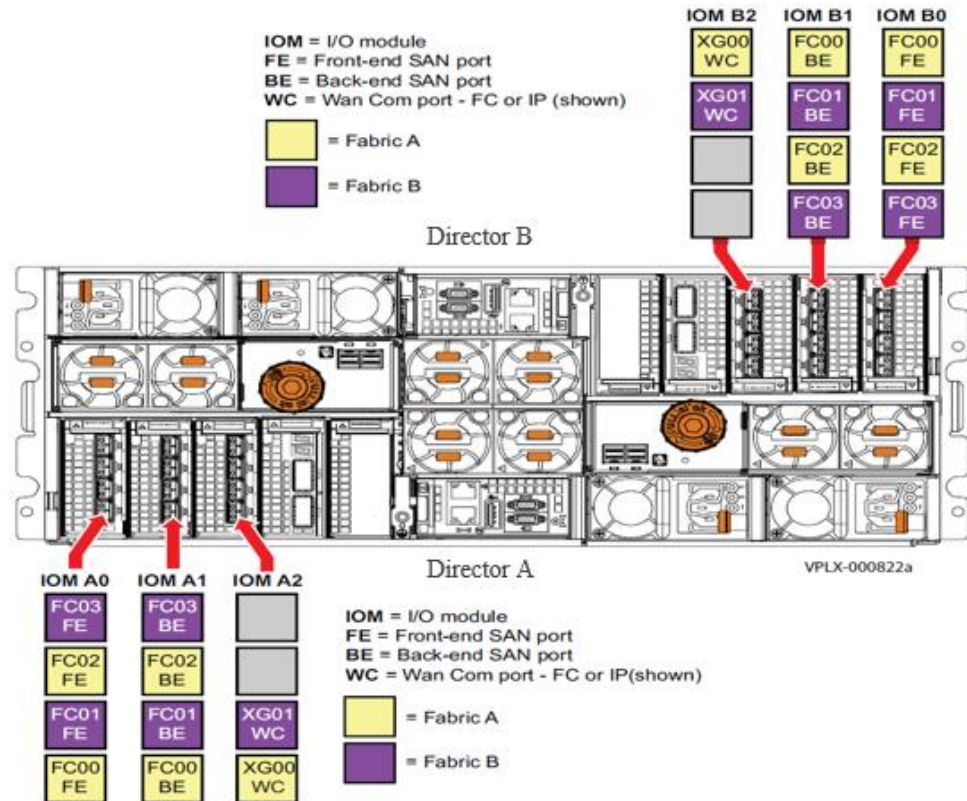


Figure 2. VPLEX engine with two directors

Director A and Director B have four I/O modules each. I/O modules A0 and B0 are configured for host connectivity and are identified as front-end, while the A1 and B1 modules are configured for array connectivity and identified as back-end. The front-end ports log in to the fabrics and present themselves as targets for zoning to the host initiators, and the back-end ports log in to the fabrics as initiators to be used for zoning to the array targets. Each director connects to both SAN fabrics with both front-end and back-end ports. VPLEX VS6 FC ports are 16 Gb. Therefore, they run at 16 Gb/s but can be configured for lower speeds. For SAP HANA, Dell EMC recommends using 16 Gb/s FC SAN networks or higher.

Front-end connectivity

Dell EMC recommends that you follow these best practices when connecting SAP HANA nodes to VPLEX systems:

- Use dual fabric design.
- Ensure that the front-end I/O modules on each director have a minimum of two physical connections, one to each fabric.
- Ensure that each host has at least one path to an A director and one path to a B director on each fabric, for a total of four logical paths.
- Follow best practices for the SAN topology, with all redundant components and links (dual fabric environment).

- Do not connect a single host bus adapter (HBA) to both ports of the same director of one engine.
- Use multipathing or path failover software at the host, which is required for access across the dual fabrics.
- Ensure that each host has fabric zoning that provides redundant access to each LUN from, at a minimum, an A and B director from each fabric.
- Use spanning engines, which are required by dual- and quad-engine VPLEX clusters for host connectivity on each fabric.
- Avoid an incorrect FC port speed between the fabric and VPLEX system. Use the highest possible bandwidth to match the VPLEX maximum port speed, and use dedicated port speeds—that is, do not use oversubscribed ports on SAN switches.

The following figure shows dual HBAs connected to two fabrics, with each HBA connecting to two VPLEX directors on the same engine in a single-engine cluster. This is the minimum configuration required for an SAP HANA node.




Figure 3. Host connectivity for HA performance across a single VPLEX engine

Back-end connectivity

Follow these best practices when connecting your VPLEX system to a storage array:

- Each director in a VPLEX cluster has a minimum of two I/O paths to every local back-end storage array and to every storage volume presented to that cluster.

- Each director has redundant physical connections to the back-end storage across dual fabrics. Each director is required to have redundant paths to every back-end storage array across both fabrics.
- Each storage array has redundant controllers connected to dual fabrics, with each VPLEX director having a minimum of two ports connected to the back-end storage arrays through the dual fabrics.
- Each director has no more than four active paths to a given LUN. Each director will load-balance across the four active paths to the storage volume.
- Connect all VPLEX engines that are used by the SAP HANA hosts to the storage array.
- Connect 1:n VPLEX back-end ports to front-end ports without reusing any front-end ports, if permitted by the number of front-end ports on the array. Such connections ensure resiliency in the event of component failure.

Note: For more information about VPLEX front-end and back-end array connectivity, see the [Dell EMC VPLEX SAN Connectivity Implementation Planning and Best Practices Technical Notes](#).

VPLEX scalability for SAP HANA

The Dell EMC VPLEX system is an additional layer between the SAP HANA hosts and the certified storage arrays, providing new functionality while also affecting throughput and latency.

This section describes VPLEX configuration with LUNs that are exported 1:1 from the array through the VPLEX system to the hosts. Note that adding local mirrors, splitting volumes across multiple arrays, and using other mobility, availability, and collaboration features might have an additional impact on performance.

Note: For SAP HANA TDI enterprise storage certification, the SAP HANA performance KPIs do not take into account DR support. The KPIs were fulfilled in this configuration without DR support.

The maximum number of hosts that can be connected depends on the capabilities of the VPLEX cluster and the capabilities of the underlying storage array. A VPLEX system can serve about 80 percent of the maximum number of hosts of the underlying arrays. If, for example, the storage array is certified for up to ten hosts, the VPLEX system can cater to eight hosts. In addition, the VPLEX system itself has a maximum number of hosts that can be connected in SAP HANA production systems.

We tested a VPLEX VS6 single engine using the SAP hwcct tool for HANA-HWC-ES 1.1 certification. From the test results, we derived guidelines for estimating the initial number of SAP HANA production hosts that can be connected to the number of VPLEX engines that are used in a VPLEX cluster. The following table shows the guidelines.

Table 2. VPLEX VS6 scalability

Number of VPLEX engines	Maximum number of SAP HANA workers
1	18
2	36
4	72

Note: We determined the scalability of additional engines by extrapolating the VPLEX VS6 single-engine test results using the performance characteristics of the dual and quad engines.

The number of SAP HANA hosts that can be connected to a VPLEX VS6 system can be higher or lower than the number listed in Table 2, depending on the actual workload. Use the SAP HANA hwcct tool with scenario HANA-HWC-ES 1.1 to validate the SAP HANA performance and determine the maximum possible number of SAP HANA hosts on a given VPLEX storage array configuration. Although test results might vary in different environments, Dell EMC recommends not exceeding the number of hosts specified in this guide.

To determine the maximum number of SAP HANA hosts that can be supported, we used the `min()` function to compare two numeric values: the maximum number of SAP HANA hosts that the VPLEX VS6 system supports compared with 80 percent of the maximum number of SAP HANA hosts that the certified physical storage array supports:

$$\text{max hosts} = \min(\langle \text{max VPLEX hosts} \rangle, \langle \text{max array(s) hosts} \rangle * 0.8)$$

The smaller of the two values is the number of SAP HANA worker nodes that can be supported.

Example 1

This example demonstrates how to physically connect an XtremIO two-brick system to SAP HANA servers using a single-engine VPLEX Local configuration. An XtremIO X2 two-brick cluster can support up to 14 SAP HANA nodes. For more information, see the [Dell EMC XtremIO Storage Configuration Best Practices for SAP HANA TDI Validation Guide](#).

Using one VPLEX engine with this array gives the following calculation:

$$\text{max hosts} = \min(18, 14 * 0.8) = 11.2$$

Eight front-end ports are required (four per X-Brick block) to support 11 SAP HANA hosts on an XtremIO X2 system.

Each VPLEX engine uses the front-end ports FC00 and FC02 of each director for fabrics A and B. The following table shows an example of back-end zoning.

Table 3. VPLEX back-end zoning example

VPLEX back-end port	XtremIO port
Director A – FC00	X1-SC1-FC2

VPLEX back-end port	XtremIO port
Director A – FC01	X2-SC1-FC1
Director A – FC02	X1-SC2-FC2
Director A – FC03	X2-SC2-FC1
Director B – FC00	X2-SC1-FC2
Director B – FC01	X1-SC1-FC1
Director B – FC02	X2-SC1-FC2
Director B – FC03	X1-SC2-FC1

On the VPLEX front-end, each HBA is connected to one port of each VPLEX director. This leads to a connectivity of four paths from the host to the VPLEX system.

The following table shows an example of front-end zoning.

Table 4. VPLEX front-end zoning example

VPLEX front-end port	HBA
Director A – FC00	HBA1 of each server
Director B – FC00	HBA1 of each server
Director A – FC02	HBA2 of each server
Director B – FC02	HBA2 of each server

Example 2

This example shows how to connect multiple arrays to a VPLEX single-engine cluster to set up an SAP HANA system, using two XtremIO X2 two-brick clusters with a maximum support of 14 SAP HANA nodes each. Calculate the maximum number of hosts as follows:

$$\text{max hosts} = \min[18, (14+14)*0.8 = 22] = 18$$

Previous-generation VPLEX scalability and certification

The SAP HANA enterprise certification for the Dell EMC VPLEX VS2 release expired on July 7, 2018, but you can still use VPLEX VS2 systems for SAP HANA that were purchased before the certification expiry date until the end of maintenance by Dell EMC. For more information, see the [Certified and Supported SAP HANA Hardware Directory](#).

The following table shows our recommendations for the VS2 system.

Table 5. VPLEX VS2 scalability

Number of VPLEX engines	Maximum number of SAP HANA workers
1	12
2	24
4	48

To determine the maximum number of SAP HANA hosts that can be supported, we used the `min()` function to compare two numeric values: the maximum number of SAP HANA hosts that the VPLEX VS6 system supports compared with 80 percent of the maximum number of SAP HANA hosts that the certified physical storage array supports:

$$\text{max hosts} = \min(\langle \text{max VPLEX VS2 hosts} \rangle, \langle \text{max array(s) hosts} \rangle * 0.8)$$

The smaller of the two values is the number of SAP HANA worker nodes that can be supported.

Competing workloads

In highly consolidated environments, SAP HANA and other databases and applications compete for storage resources. When creating and presenting storage volumes to VPLEX systems, consider Quality of Service (QoS) features such as host I/O limits and service levels for the underlying back-end storage arrays.

The host I/O limits feature controls noisy neighbors and protects the performance of the SAP HANA system. Host I/O limits can be effective in consolidated environments if an overload on the storage resources affects the performance of more critical applications such as production installations. To protect the performance of SAP HANA production systems, you can limit the total IOPS or bandwidth, which might otherwise consume a large portion of the system’s resources. To do so, configure host I/O limits on the underlying back-end storage array for the storage volume LUNs and the datastores of non-SAP HANA applications or SAP HANA non-production systems that are presented to the VPLEX system.

PowerMax, VMAX3, and VMAX All Flash arrays can provide the required performance by means of service levels. Service levels enable users to insulate specific storage groups from any performance impact by other “noisy neighbor” applications. The user 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.

For more information, see [References](#) for links to the configuration best practices guide for your certified Dell EMC underlying storage.

Configuring and installing an SAP HANA scale-out cluster on a VPLEX system

This section of the guide describes how to:

- Provision VPLEX storage from the already-created data and log volumes on the underlying XtremIO X2 two-brick block for an SAP HANA scale-out cluster with three worker nodes and one standby node (3+1).
- Prepare the SAP HANA hosts.
- Install the SAP HANA cluster using `hdblcm`, the SAP lifecycle management command-line tool.

Provisioning VPLEX storage

The following figure depicts basic VPLEX storage provisioning.

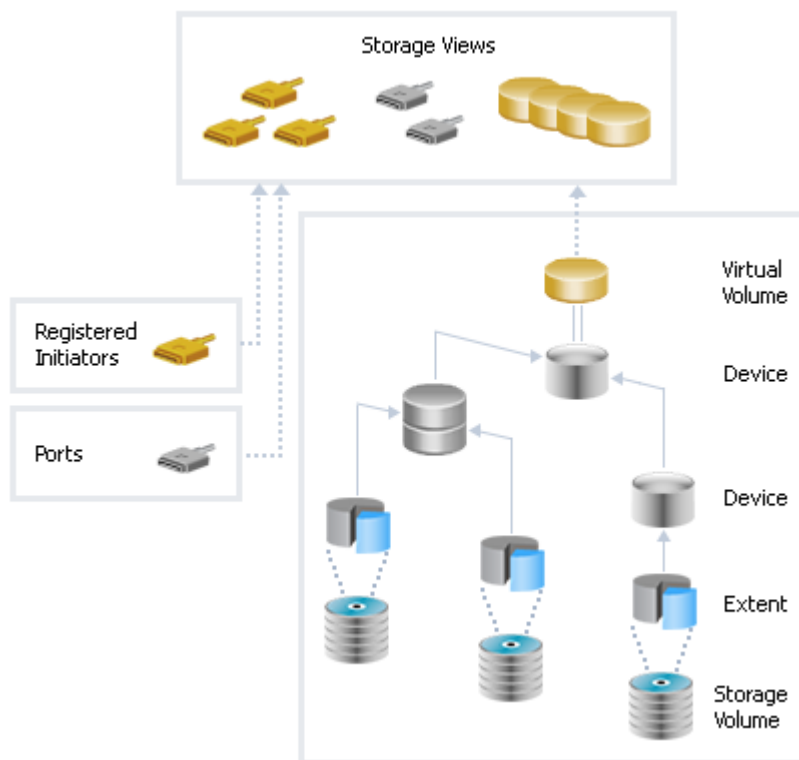


Figure 4. VPLEX storage provisioning

At the bottom of the figure are the storage volumes, the LUNs of the underlying storage arrays. Volumes are split into extents, which can be used for device partitioning. One or more extents are combined into devices using a certain RAID technique. VPLEX storage offers RAID-0 (striped), RAID-1 (mirrored), and RAID-C (concatenated) devices, while 1:1 mapping creates one device for every extent. Virtual volumes are created from the top-level devices and exported to the hosts.

Note: Dell EMC recommends only a 1:1 mapping of storage volumes to extents and a 1:1 mapping of extents to devices with SAP HANA because performance is crucial.

The following figure shows an example of a VPLEX virtual volume configuration.

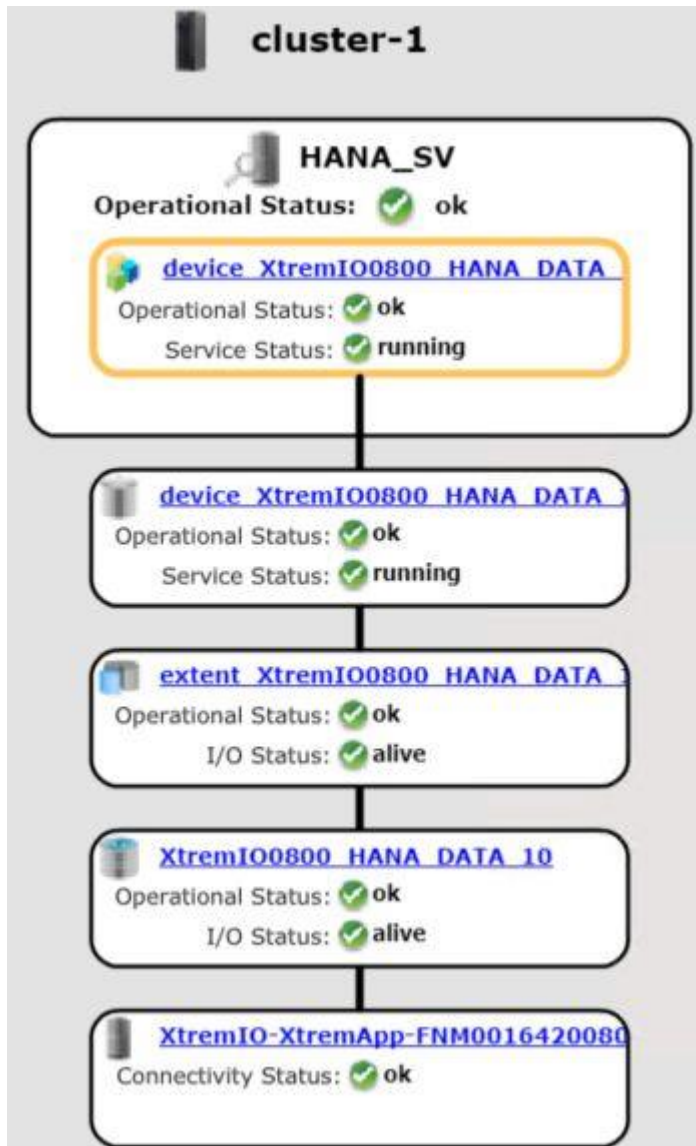


Figure 5. VPLEX virtual volume configuration example

To successfully make virtual volumes available to the hosts, the ports of the VPLEX front-end are used with the registered host initiators for building storage views. Storage views provide full control over the LUN visibility.

Registering initiators

When all of the HBAs have been made visible to VPLEX using zoning, follow these steps:

1. In the navigation pane, select **Initiators**.
2. If the initiators are not already present:
 - a. Click **Rediscover Initiators**, select all your hosts' unregistered initiators, and then click **Register**.
 - b. Type a meaningful name for the initiator, such as `<servername>_HBA<X>`, select **default** as the **Host Type**, and then click **OK**.

Claiming storage volumes

To claim storage volumes for VPLEX systems, start by exporting the SAP HANA volumes from the underlying storage array to the VPLEX cluster.

For specific configuration and sizing recommendations, see the configuration guide for the underlying storage array. Then follow these steps:

1. Log in to the VPLEX UI and select **Provision Storage > cluster-1 > Storage Arrays**.
2. Select the required storage arrays and click **Rediscover Array**.

The Storage Arrays page appears, as shown in the following figure.



Figure 6. Storage Arrays page

3. Click **OK**, wait until the operation is completed, and then click **Close**.
4. Click **Claim Storage**.
5. In the dialog box that appears, click **Next**.
VPLEX discovers all unclaimed LUNs.
6. Based on the recommendations for your storage array for SAP HANA, select **Use Thin Rebuilds** or, if thick provisioning is recommended, select **Do not use Thin Rebuilds**.
7. Click **Next**.
8. Select all LUNs to be claimed, click **Add** or **Add All**, and then click **Next**.
9. Optionally, rename the LUNs by specifying a base name for the set of LUNs in the upper-right corner.

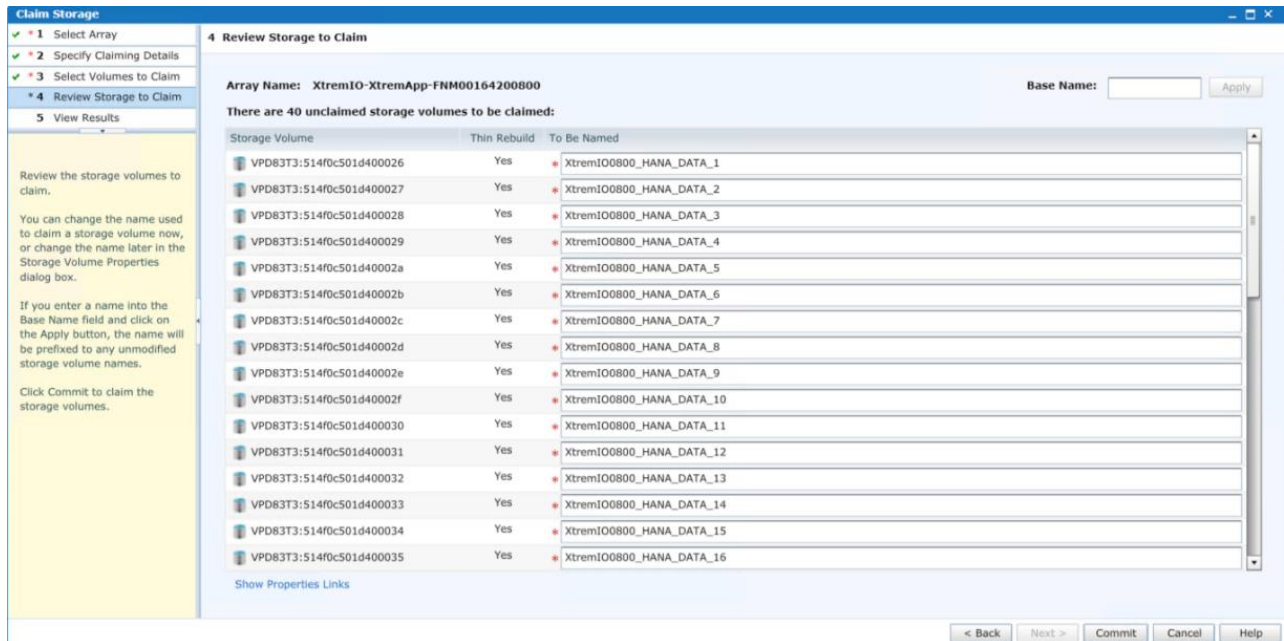


Figure 7. Renaming LUNs

Dell EMC recommends using a meaningful name such as *<basename>_HANA_Data_<X>* and *<basename>_HANA_Log_<X>* to distinguish between the LUNs. You can rename them later if necessary.

10. Click **Commit**.
All LUNs are now claimed—that is, available to the VPLEX system.
11. Click **Close**.
12. To review the results, in the navigation pane, select **Storage Volumes**.

Creating extents

Create extents from the previously claimed LUNs as follows:

1. In the navigation pane, select **Storage Volumes**, select one or more LUNs, and then click **Create Extents**.
2. If you need to adjust the selection, select all the LUNs that should be available to SAP HANA, click **Add** or **Add All**, and then click **Next**.

The Review & Finish page appears.

3. Accept the default values (the maximum capacity), review the results, and then click **Commit**.

For each LUN, an extent with the total capacity of the LUN is created.

4. Click **Close**.

Creating devices

To create devices from the extents:

1. Select **Extents** in the navigation pane, and then click **Create Devices**.
2. Select **1:1 Mapping of Extents to Devices** and click **Next**.

The **Select Device Type** page appears.

3. Select all the extents to use for creating the devices, click **Add** or **Add All**, and then click **Next**.
4. When prompted to create virtual volumes, select **Yes** and click **Next**.
5. Review the configuration and click **Commit**.

A summary of the created devices appears.

6. Review the summary and click **Close**.
7. In the navigation pane, click **Devices**.

The **Devices** page appears.

8. To view a summary of the created virtual volumes, as shown in the following figure, in the navigation pane, click **Virtual Volumes**.

Name	Capacity	Health	Operational Status	Service Status	Thin Enabled	Expandable By	Expansion Method	Expansion Status	Supporting Device
device_XtremIO0800_HANA_DATA_10_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_10_1
device_XtremIO0800_HANA_DATA_11_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_11_1
device_XtremIO0800_HANA_DATA_12_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_12_1
device_XtremIO0800_HANA_DATA_13_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_13_1
device_XtremIO0800_HANA_DATA_14_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_14_1
device_XtremIO0800_HANA_DATA_15_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_15_1
device_XtremIO0800_HANA_DATA_16_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_16_1
device_XtremIO0800_HANA_DATA_17_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_17_1
device_XtremIO0800_HANA_DATA_18_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_18_1
device_XtremIO0800_HANA_DATA_19_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_19_1
device_XtremIO0800_HANA_DATA_1_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_1_1
device_XtremIO0800_HANA_DATA_20_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_20_1
device_XtremIO0800_HANA_DATA_2_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_2_1
device_XtremIO0800_HANA_DATA_3_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_3_1
device_XtremIO0800_HANA_DATA_4_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_4_1
device_XtremIO0800_HANA_DATA_5_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_5_1
device_XtremIO0800_HANA_DATA_6_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_6_1
device_XtremIO0800_HANA_DATA_7_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_7_1
device_XtremIO0800_HANA_DATA_8_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_8_1
device_XtremIO0800_HANA_DATA_9_1_vol	512.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_DATA_9_1
device_XtremIO0800_HANA_LOG_10_1_vol	256.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_LOG_10_1
device_XtremIO0800_HANA_LOG_11_1_vol	256.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_LOG_11_1
device_XtremIO0800_HANA_LOG_12_1_vol	256.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_LOG_12_1
device_XtremIO0800_HANA_LOG_13_1_vol	256.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_LOG_13_1
device_XtremIO0800_HANA_LOG_14_1_vol	256.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_LOG_14_1
device_XtremIO0800_HANA_LOG_15_1_vol	256.00 GB	OK	OK	Unexported	disabled	0.00 MB	storage-volume	-	device_XtremIO0800_HANA_LOG_15_1

Figure 8. Summary of virtual volumes

The volumes are now ready to use with SAP HANA.

Creating storage views

The combination of virtual volumes, ports, and initiators controls the visibility of the LUNs to the hosts. To create storage views, follow these steps:

1. In the navigation pane, select **Storage Views**, and then click **Create**.

2. In the **Storage View Name** field, enter the name of the host to which you are connecting.
3. Select all the initiators that belong to the host, click **Add**, and then click **Next**.
4. Select all the required ports of the VPLEX engine to which you are connecting the host and click **Add**, and then click **Next**.

At a minimum, Director A ports FC00 and FC01 and Director B ports FC00 and FC01 are required. Do not connect one host to multiple engines. For more information, see [Front-end connectivity](#) on page 13.

5. Select all SAP HANA LUNs to be in this storage view, and then click **Next**.
The SAP HANA HA features require all LUNs to be exported to all hosts, including the standby host.
6. At the top of the page, select **Show virtual volumes already exported to other storage views**, select all SAP HANA LUNs, and click **Yes** in the safety question dialog box.
7. When configuring the second, third, and every succeeding host, confirm that already-exported LUNs are made visible to another host by selecting **Auto-assign LUN numbers starting with a value of 1**.

The LUN number 0 is usually preserved for boot LUNs and will not be detected when the host SCSI bus is scanned on reboot.

8. Click **Apply**.
9. Click **Next**.
10. Review the information that appears, and then click **Commit**.
The **Review Results** page appears.
11. Review the summary and click **Close**.
12. Repeat the preceding steps for all other hosts belonging to the SAP HANA system.

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, together with the SAP HANA storage connector fcClient, mounts the volumes to the correct SAP HANA node and I/O fencing.¹ Correct preparation of the SAP HANA nodes and the SAP HANA `global.ini` file is required before the actual SAP HANA installation. The following section describes how to prepare the SAP HANA nodes.

Preparing the SAP HANA nodes

This example assumes that the following basic installation and configuration operations are complete on the SAP HANA nodes:

- The operating system is installed and properly configured using the SAP recommendations—in this example, we used SUSE Linux 12 SP3 for SAP applications.

¹ When using the SAP HANA storage connector fcClient do not auto-mount the device using `/etc/fstab`.

- An SAP HANA shared file system, `/hana/shared/`, has been created on a NAS system such as an Isilon, and 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 according to the SAP requirements.
- SSH keys have been exchanged between all SAP HANA nodes.
- System time synchronization has been configured using a NTP server.
- The SAP HANA installation DVD ISO file has been 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.

The remainder of this section describes how to configure Linux native multipathing and initialize the SAP HANA persistence.

Configure Linux native multipathing (DM-MPIO)

Configure Linux native multipathing as follows:

1. Edit the `/etc/multipath.conf` file as shown in the following figure.

```
device {
    vendor "EMC"
    product "Invista"
    product_blacklist LUNZ
    path_grouping_policy multibus
    path_selector "round-robin 0"
    path_checker tur
    features "0"
    hardware_handler "0"
    rr_weight uniform
    no_path_retry 5
    rr_min_io 1000
}
```

Figure 9. Entries for the `/etc/multipath.conf` file

2. Restart multipathing by typing the following command:
`service multipath restart`

Note: For the operating system version and storage array MPIO configuration 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. Either use the `rescan-scsi-bus.sh` command or reboot each node.

To verify that the volumes are visible, run the following commands on one of the nodes:

1. To list all 512 GB data volumes, type:

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

The following figure shows the command results.

```
3600014400000010f010268c81e76088 dm-25 EMC,Invista
size=512G features='2 queue_if_no_path retain_attached_hw_handler' hwhandler='0' wp=rw
`--+ policy='round-robin 0' prio=1 status=active
  |-- 3:0:0:31 sdag 66:0 active ready running
  |-- 4:0:0:31 sdgk 132:0 active ready running
  |-- 3:0:1:31 sdbu 68:128 active ready running
  `-- 4:0:1:31 sdhy 134:128 active ready running
--
3600014400000010f010268c81e7607f dm-22 EMC,Invista
size=512G features='2 queue_if_no_path retain_attached_hw_handler' hwhandler='0' wp=rw
`--+ policy='round-robin 0' prio=1 status=active
  |-- 3:0:0:29 sdae 65:224 active ready running
  |-- 4:0:0:29 sdgi 131:224 active ready running
  |-- 3:0:1:29 sdbb 68:96 active ready running
  `-- 4:0:1:29 sdhw 134:96 active ready running
```

Figure 10. Listing all 512 GB data volumes

- To list all 256 GB log volumes, type:

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

The following figure shows the command results.

```
3600014400000010f010268c81e76099 dm-11 EMC,Invista
size=256G features='2 queue_if_no_path retain_attached_hw_handler' hwhandler='0' wp=rw
`--+ policy='round-robin 0' prio=1 status=active
  |-- 3:0:0:19 sdu 65:64 active ready running
  |-- 4:0:0:19 sdfy 131:64 active ready running
  |-- 3:0:1:19 sdbi 67:192 active ready running
  `-- 4:0:1:19 sdhm 133:192 active ready running
3600014400000010f010268c81e760a4 dm-38 EMC,Invista
size=256G features='2 queue_if_no_path retain_attached_hw_handler' hwhandler='0' wp=rw
`--+ policy='round-robin 0' prio=1 status=active
  |-- 3:0:0:8 sdj 8:144 active ready running
  |-- 4:0:0:8 sdfn 130:144 active ready running
  |-- 3:0:1:8 sdax 67:16 active ready running
  `-- 4:0:1:8 sdhb 133:16 active ready running
```

Figure 11. Listing all 256 GB data volumes

The unique device identifier of the multipath device must match the WWN of the volumes that you created in Dell EMC Unisphere for VPLEX. When viewed on a Linux host, the WWN of the volume from the VPLEX is now preceded by a 3.

- Initialize the devices and create the Linux XFS file system on each of the devices by following this command example:

```
mkfs.xfs /dev/mapper/ 3600014400000010f010268c81e76081
```

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

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 installation 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 all use `fcClient`, which is part of the SAP software distribution.

To view the universally unique identifiers (UUIDs) that SAP HANA uses to identify the correct storage partitions for the `global.ini` file, follow these steps:

1. In the navigation pane, select **Storage Views**, and then select the previously created SAP HANA storage view.

The right pane displays the details.

2. Select **Virtual Volumes**, and then, at the bottom of the page, select **VPD ID**.

The **VPD ID** column displays the UUIDs.

The UUIDs are used to describe the SAP HANA storage partitions within the `global.ini` file.

3. Ensure that the `global.ini` file has the content that is 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 = 360001440000000010f010268c81e76081
partition_1_log_wwid = 360001440000000010f010268c81e7609b
partition_2_data_wwid = 360001440000000010f010268c81e76082
partition_2_log_wwid = 360001440000000010f010268c81e7609c
partition_3_data_wwid = 360001440000000010f010268c81e76083
partition_3_log_wwid = 360001440000000010f010268c81e7609d
```

Figure 12. Global.ini file storage section

4. Ensure that the partition entries match the UUIDs that are displayed by running the `multipath -ll` command with a preceding 3.
5. Place the `global.ini` file in a directory on the `/hana/shared/` file system—for example, `/hana/shared/VS6_cfg`.

The `global.ini` name is mandatory. Note that 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 offer the efficiency of installing all components at one time while automating the installation and providing further flexibility to customers. The following example uses the `hdblcm` CLI to install our SAP HANA 3+1 scale-out cluster:

1. Open the shared file system and go to the `HDB_LCM_LINUX_X86_64` directory into which you extracted the SAP HANA installation DVD ISO file by typing the following command:

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

2. Create a template installation parameter file by typing the following command:

```
./hdblcm --action=install --  
dump_configfile_template=VS6_install.cfg
```

3. After the template is created, modify the following parameters in the file to match the environment:

```
Directory root to search for components  
component_root=/SAPShare/software/SAP_HANA_SPS03_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/VS6_cfg
```

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

```
SAP HANA System ID  
sid=VS6
```

```
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 (unencrypted or encrypted) for the root user, SAP Host Agent User (`sapadm`), system administrator user (`<sid>adm`), and database user (`SYSTEM`) in the parameter file.

For information, see the [SAP HANA Server Installation and Update Guide](#). The `hdblcm` installation procedure prompts you for any missing passwords or parameters.

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

Install the SAP HANA scale-out cluster

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

```
./hdblcm --action=install --configfile=VS6_Install.cfg
```

Optimizing file I/Os after the 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. This 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/Os for a given storage array and file system. The Linux XFS file system is used on all Dell EMC storage LUNs for the SAP HANA persistence.

After the SAP HANA persistence is installed on VMAX 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`

Note: The following instructions for tuning file I/O parameters are based on SAP HANA 1.0 and SAP HANA 2.0 SPS01. See the latest SAP HANA documentation for later versions and updates.

SAP HANA 1.0

After the initial SAP HANA installation is complete, set the SAP HANA 1.0 parameters by running the `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 has been deprecated. Instead, the parameters are defined as normal parameters 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 how the `fileio` section of `global.ini` looks after you set the parameters.

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 13. 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 installation.

For more information, see [SAP Note 2399079](#)—Elimination of `hdbparam` in HANA 2 (access requires an SAP username and password).

Backing up and recovering the SAP HANA database

Dell EMC Data Domain deduplication storage systems enable you to redefine SAP backup, archive, and availability with deduplication and consolidated data protection. 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 is powered with flash SSD. 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 disaster recovery 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 Dell EMC 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 Dell EMC Data Domain Boost (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 your Data Domain system for SAP HANA

Configure the Data Domain system as follows:

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


```
/usr/sap/<SID>/SYS/global/hdb/opt
```
3. Copy the `/opt/ddbda/bin/hdbbackint` file to the subdirectory that is specified in the preceding step.

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

```
/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 14. Creating a symbolic link

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 15. Modifying the configuration file parameter settings

5. Configure the DDBEA lockbox by running the command `ddbadmin -P -Z <configuration_file>`, where `<configuration_file>` is the file that you used in step 4.

Note: This configuration example uses x22 as the SAP HANA backup and recovery SID. If you are configuring the multinode cluster, repeat steps 2–5 for all nodes.

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

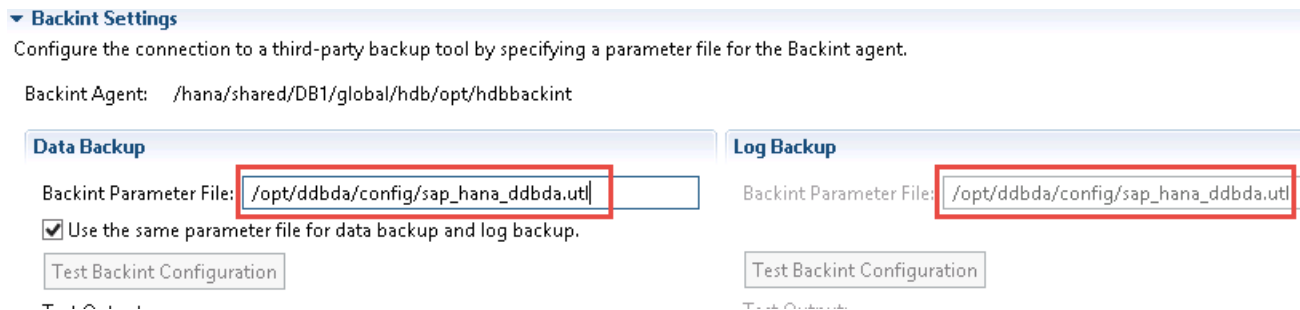


Figure 16. Specifying the DDBEA configuration file location

For more information about configuring DD Boost software with SAP HANA, 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

Back up the database from SAP HANA Studio as follows:

1. Log in to SAP HANA Studio, and then select **Backup and Recovery > Back up Tenant Database**, as shown in the following figure.

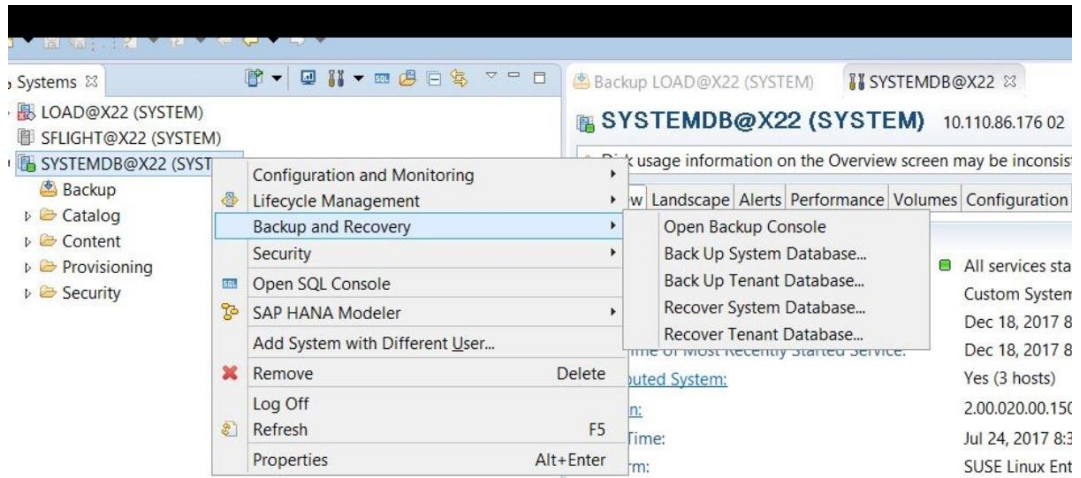


Figure 17. SAP HANA Studio Backup and Recovery menu

2. Select the tenant database, as shown in the following figure.

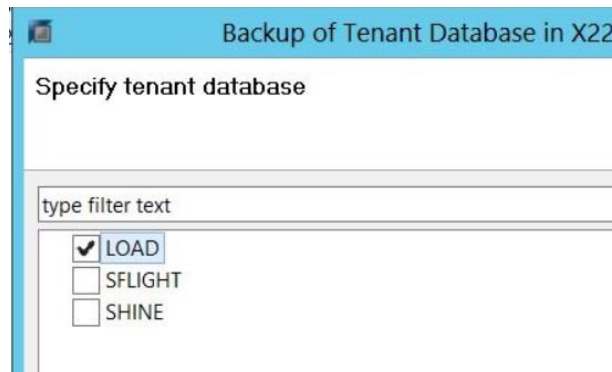


Figure 18. Specifying the tenant database

Note: The LOAD database is for illustration purposes only.

3. When specifying settings such as **Backup Type**, **Backup Destination**, and **Backup Prefix**, ensure that you select **Backint** in the **Destination Type** field, as shown in the following figure.

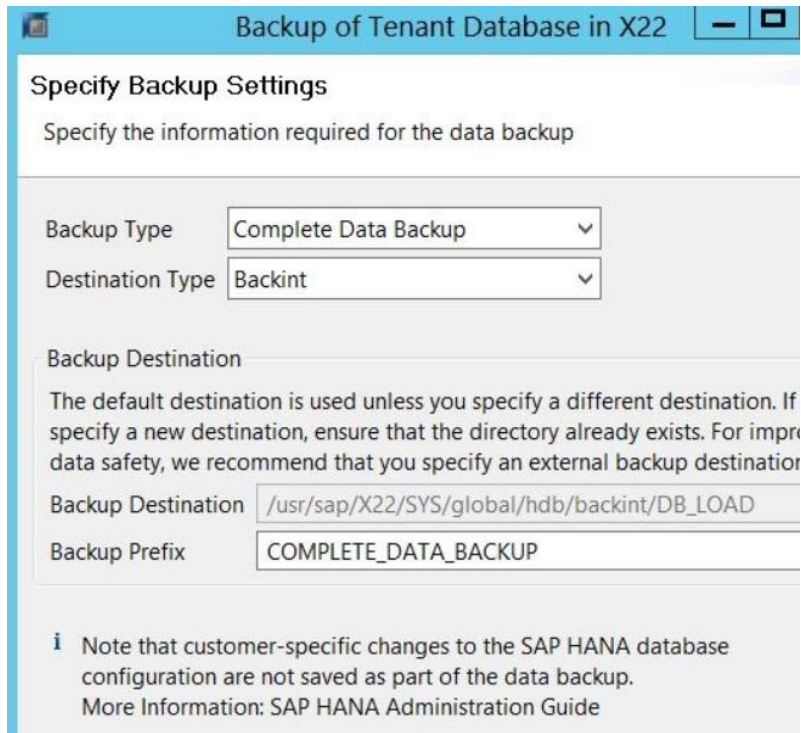


Figure 19. Specifying backup settings

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

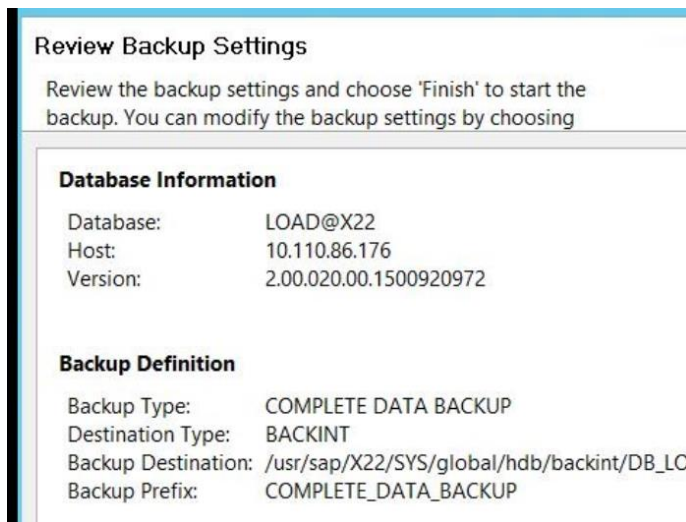


Figure 20. Review Backup Settings page

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

Recovering the database with SAP HANA Studio

To restore the SAP HANA database to the point when you backed it up, follow these steps:

- Log in to the SAP HANA Studio, and then select **Backup and Recovery > Recover Tenant Database**.
- Select the tenant database.

3. Select **Recover the database to its most recent state**.
4. Specify a location for the log backups, as shown in the following figure.

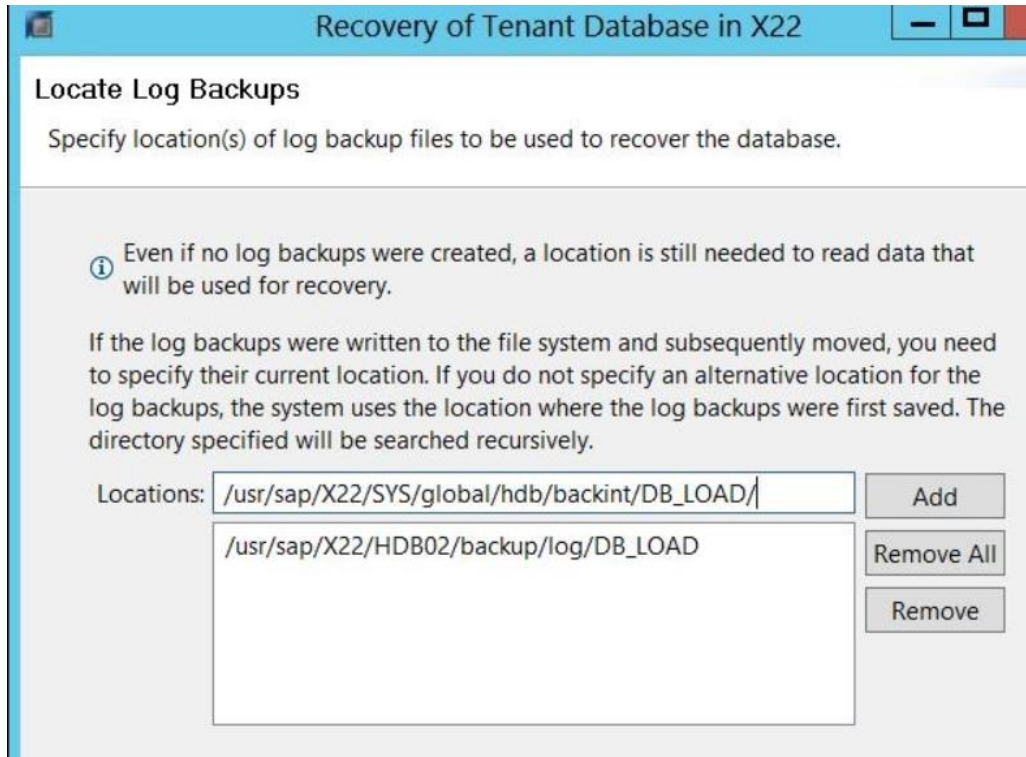


Figure 21. Specifying the log backup location

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

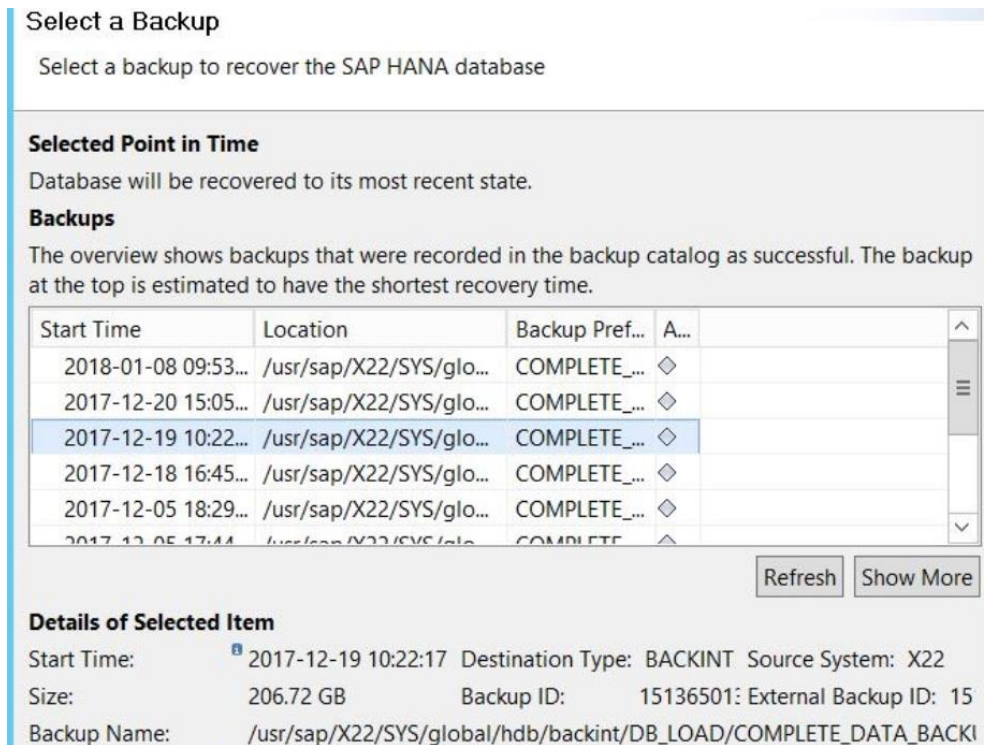


Figure 22. Specifying a recovery point in time

6. Select **Third-Party Backup Tool (Backint)** as the backup tool.
7. Review the recovery settings, as shown in the following figure.



Figure 23. Tenant database recovery settings

8. Click **Finish** to complete the recovery operation.

Conclusion

Summary

Using SAP HANA in TDI deployments with Dell EMC VPLEX systems provides many benefits, including lowering hardware and operational costs because storage and network components can be shared, reducing risk, improving availability and performance, and increasing hardware vendor choice.

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

VPLEX systems combine data that is located on heterogeneous certified storage arrays to create dynamic, distributed, highly available data centers. The VPLEX VS6 model represents the next-generation architecture for data mobility and continuous availability.

Findings

During our tests with SAP HANA on VPLEX systems, 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 VPLEX systems require SSDs for the SAP HANA persistence.
- Using SSDs for the SAP HANA persistence provides significant benefits, including:
 - Reduced SAP HANA startup and host auto-failover times
 - Reduced SAP HANA backup times
 - No need to consider spindle count because initial array and disk configuration can be performed based on capacity
- Data Domain systems with DDBEA enable you to consolidate SAP HANA backups and archives with greater network efficiency, reducing your storage protection footprint through deduplication and compression.

References

Dell EMC documentation

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

- [*Dell EMC Host Connectivity Guide for Linux*](#)
- [*Dell EMC VPLEX Family Specification Sheet*](#)
- [*Dell EMC XtremIO Storage Configuration Best Practices for SAP HANA TDI Configuration Guide*](#)
- [*EMC VPLEX Data Mobility and Migrations*](#)
- [*EMC VPLEX Host Multipathing*](#)
- [*EMC VPLEX Metro Cross-Cluster Host Connectivity*](#)
- [*EMC VPLEX Overview and General Best Practices*](#)
- [*EMC VPLEX SAN Connectivity*](#)
- [*EMC VPLEX with XtremIO 3.0*](#)
- [*Storage Configuration Best Practices for SAP HANA Tailored Data Center Integration on Dell EMC VMAX Solution Guide*](#)
- [*Storage Configuration Best Practices for SAP HANA Tailored Data Center Integration on Dell EMC PowerMax Arrays Validation Guide*](#)
- [*Storage Configuration Best Practices for SAP HANA TDI on Dell EMC Unity Storage Systems Solution Guide*](#)
- [*Storage Configuration Best Practices for SAP HANA Tailored Data Center Integration on Dell EMC Compellent Storage Systems*](#)
- [*VMware Virtualized SAP HANA with Dell EMC Storage Solution Guide*](#)
- [*VPLEX ensures uptime for business critical applications—Data Sheet*](#)

For SAP Ready Solution documentation, see [Solutions for SAP Server Info Hub: Technical Documentation](#) in the Everything SAP at Dell EMC Community forum.

SAP HANA documentation

The following SAP documentation and web pages provide additional relevant information:

- [*Certified and Supported SAP HANA Hardware Directory*](#)
- [*SAP HANA Administration Guide*](#)
- [*SAP HANA In-Memory Data Platform*](#)
- [*SAP HANA Master Guide*](#)
- [*SAP HANA One*](#)
- [*SAP HANA Platform*](#)
- [*SAP HANA Server Installation and Update Guide*](#)
- [*SAP HANA Storage Requirements*](#)
- [*SAP HANA Studio Installation and Update Guide*](#)
- [*SAP HANA Tailored Data Center Integration - Frequently Asked Questions*](#)

- [SAP HANA Tailored Datacenter Integration \(TDI\)](#)
- [SAP HANA Technical Operations Manual](#)

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](#)