

# STORAGE CONFIGURATION BEST PRACTICES FOR SAP HANA TAILORED DATA CENTER INTEGRATION ON DELL EMC VMAX

- VMAX 10K, 20K, 40K
- VMAX3 100K, 200K, 400K
- VMAX All Flash 250F/FX, 450F/FX, 850F/FX, 950F/FX

June 2017

## ABSTRACT

This solution guide describes a concept that overcomes limitations of the SAP HANA model. Using Tailored Data Center Integration (TDI) on Dell EMC VMAX, VMAX3, and VMAX All Flash storage systems, customers can integrate SAP HANA into an existing, well-established data center infrastructure, providing multiple benefits.

H12342.7

**This document is not intended for audiences in China, Hong Kong, Taiwan, and Macao.**

## 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 © 2017 Dell Inc. or its subsidiaries. All Rights Reserved. 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 the property of their respective owners. Published in the USA 6/2017. Solution Guide, H12342.7

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

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect




# Contents

Executive summary .....	4
Using Dell EMC VMAX arrays for SAP HANA: General considerations .....	7
Storage design principles for SAP HANA on VMAX arrays .....	12
Storage design principles for SAP HANA on VMAX3 and VMAX All Flash arrays .....	17
VMAX All Flash inline compression .....	21
Storage configuration and installation of an SAP HANA scale-out cluster on a VMAX All Flash array: Example .....	24
Conclusion .....	39
References .....	40

## Executive summary

**Business case** SAP HANA is an in-memory data platform that can be deployed on-premises (locally) or in the cloud. Customers can use the SAP HANA platform to analyze large volumes of data and to 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 provided by SAP hardware partners. It can be deployed on-premises in two different models, as shown in Figure 1.



**Figure 1. SAP HANA appliance model compared with the TDI model (picture ©SAP SE)**

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

The TDI model provides greater flexibility than the appliance model. The SAP HANA servers must still be certified SAP HANA servers and meet the SAP HANA requirements, but the network and storage components can be shared in customer environments. This allows customers to use their existing enterprise storage arrays for SAP HANA and integrate SAP HANA seamlessly into existing data center operations (such as disaster recovery (DR), data protection, monitoring, and management). This reduces the time-to-value, costs, and risk of an overall SAP HANA adoption.

## Solution overview

SAP certifies the enterprise storage arrays used in SAP HANA TDI deployments to ensure that they meet the SAP HANA performance and functional requirements<sup>1</sup>. Using the SAP HANA Hardware Configuration Check Tool (hwccct), Dell EMC performed extensive testing on the Dell EMC™ VMAX™ family and on VMAX All Flash storage systems using the following SAP certification scenarios:

- **VMAX**—HANA-HWC-ES 1.0
- **VMAX3™ and VMAX All Flash**—HANA-HWC-ES-1.1

Based on the test results, this solution guide provides storage configuration recommendations for the arrays. The recommendations meet SAP performance requirements and ensure the highest availability for database persistence on disk.

---

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

---

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

## Key benefits

Customers using SAP HANA TDI on VMAX, VMAX3, and VMAX All Flash arrays can:

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

---

<sup>1</sup> Dell EMC VMAX, VMAX3, and VMAX All Flash arrays are certified by SAP.

## Document purpose

Before SAP introduced the SAP HANA TDI deployment model, customers using the SAP HANA appliance model 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 due to lack of capacity and inability to respond rapidly to unexpected growth demands

This guide describes a solution that uses SAP HANA in a TDI deployment scenario on VMAX, VMAX3, and VMAX All Flash enterprise storage. The solution reduces hardware and operational costs, lowers risks, and increases server and network vendor flexibility.

This guide provides configuration recommendations based on SAP requirements for high availability and on the performance tests and results that are needed to meet the SAP key performance indicators (KPIs) for data throughput and latency for the SAP HANA TDI. Specific topics include:

- Best practices and tips for deploying the SAP HANA database on VMAX, VMAX3, and VMAX All Flash storage systems
- Introduction to the key technologies in the SAP HANA TDI on VMAX solution
- Description of the configuration requirements and storage design principles for VMAX, VMAX3, and VMAX All Flash storage with SAP HANA
- Example of a HANA scale-out installation using VMAX All Flash storage devices

## Audience

This document is intended for system integrators, system or storage administrators, customers, partners, and members of Dell EMC Professional Services who must configure a VMAX, VMAX3, or VMAX All Flash storage array to 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 [EMC.Solution.Feedback@emc.com](mailto:EMC.Solution.Feedback@emc.com) with your comments.

**Authors:** Werner Katzenberger, Donagh Keeshan, Aighne Kearney

## Using Dell EMC VMAX arrays for SAP HANA: General considerations

This section describes general considerations for connecting HANA to VMAX arrays. VMAX systems have different disk, pool, and LUN configurations to VMAX3 or VMAX All Flash systems. The following considerations are common to all platforms:

- SAP HANA capacity requirements
- SAN network
- SAP HANA I/O patterns
- Symmetrix Remote Data Facility (SRDF™) storage replication

The SAP HANA enterprise storage certification does not use the same certification scenarios for the VMAX as for the VMAX3 and VMAX All Flash arrays.

### SAP HANA capacity requirements

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

- Operating system (OS) 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 boot from a volume on a VMAX array (that is, boot from the storage area network, or SAN), the overall capacity calculation for the HANA installation must include the required OS capacity. Every SAP HANA node requires approximately 100 GB capacity for the OS. This includes the `/usr/sap/` directory.

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

#### SAP HANA installation (/hana/shared/)

Every SAP HANA node requires access to a file system mounted under the local mount point, `/hana/shared/`, for installation of the SAP HANA binaries and the configuration files, traces, and logs. An SAP HANA scale-out cluster requires a single shared file system. The file system must be mounted on every node. Most SAP HANA installations use an NFS file system for this purpose. VMAX3 and VMAX All Flash arrays can provide this file system with the embedded NAS (eNAS) option. You can calculate the size of the `/hana/shared/` file system by using the formula in the SAP white paper [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):

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

## 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 so as 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 HANA node (scale-up) and worker node (scale-out) 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 their [SAP HANA Storage Requirements](#) white paper. Version 2.10 of the white paper states that you can calculate the size of the data volume by using the following formula:

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

In the formula, *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 four worker nodes, then every node must have a 616 GB data volume assigned to it (2 TB / 4 = 512 GB x 1.2 = 616 GB).

If the net disk space is unknown at the time of storage sizing, 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. The [SAP HANA Storage Requirements](#) white paper 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 HANA backup using Dell EMC Data Domain™ and Dell EMC NetWorker™. Although an SAP HANA backup to an NFS file system on a VMAX3 or VMAX All Flash array is possible, 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 VMAX3 or VMAX All Flash array, see the [SAP HANA Storage Requirements](#) white paper for details 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.



## 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 HANA node. If, for example, ten nodes are connected in a SAN to a VMAX array, a total bandwidth of 4000 MB/s is required. Assuming that an 8 Gb/s front-end port provides ~750 MB/s bandwidth, then at least six dedicated 8 Gb/s front-end ports are required to support 10 HANA nodes ( $6 \times 750 \text{ MB/s} = 4500 \text{ MB/s}$ ). Three ports are required with 16 Gbps front-end ports.

While this maximum bandwidth requirement will only arise 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 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<sup>2</sup> than to use all the ports on a single director. This way, even if a component fails, the storage can continue to service host I/Os.

Dynamic core allocation is new to VMAX3 and to VMAX All Flash arrays. Each VMAX3 director provides services such as front-end connectivity, back-end connectivity, and data management. Each such service has its own set of cores on each director. The cores are pooled together to provide CPU resources that can be allocated as necessary. For example, even if host I/Os arrive 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/Os arriving to other directors will have their 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 HANA workload from non-HANA applications. Dell EMC therefore recommends using dedicated front-end ports for HANA and not sharing these ports with non-HANA applications.

### HBA ports

Each HBA port (initiator) creates a path for I/Os between the host and the SAN switch, which then continues to the VMAX storage. You must 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.

---

<sup>2</sup> Each VMAX engine has two redundant directors.

## SAP HANA I/O patterns

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

### 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/Os to the data file system. During normal operations, most of the I/Os to the data file system are 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 up 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 of the I/Os to the log volume are writes, and data is read from the log volume only during database restart, high availability (HA) failover, log backup, or database recovery.

SAP HANA I/Os can be optimized for specific storage environments. [Optimizing file I/Os after the SAP HANA installation](#) describes the optimization for the VMAX arrays.

## Symmetrix Remote Data Facility storage replication

SAP HANA supports two replication technologies: replication of the storage volumes with storage replication, and application-based replication with SAP HANA system replication.

Dell EMC validated the storage replication using Dell EMC Symmetrix Remote Data Facility (SRDF), as described in the following solution guides:

- [Business Continuity and Disaster Recovery with EMC VMAX3 for SAP HANA TDI Deployments](#)
- [Business Continuity Best Practices for SAP HANA TDI with EMC Symmetrix VMAX](#)

Even though synchronous storage replication using SRDF/S has been validated and is supported in SAP HANA environments, the impact on the latency of SAP HANA log I/Os requires careful analysis. Depending on the distance and the network infrastructure between the SRDF/S sites, the latency of these I/Os might exceed the acceptable threshold value and lead to significantly slower transaction processing, which might lead to freezes or out-of-memory situations in some cases.

While SAP does not provide a latency threshold value for storage replications for small block (4 KB and 16 KB) I/Os, Dell EMC recommends that these latencies do not significantly exceed the SAP-defined KPIs for nonreplicated environments.

If the latency using SRDF/S in a customer environment is higher and users are experiencing SAP HANA performance problems, Dell EMC recommends switching to asynchronous replication using SRDF/A if a non-zero recovery point objective (RPO) is acceptable.

## SAP HANA shared file system on VMAX3 and VMAX All Flash

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 the customer's infrastructure and requirements, the following options are available:

- NFS server-based shared file system.
- NAS systems such as VMAX3 eNAS can be used to provide an NFS share for the SAP HANA shared file system.
- VMAX block storage can create a shared file system using a cluster file system such as 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 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.

---

## SAP HANA certification scenarios

Two different certification scenarios are used to test the performance of the SAP HANA persistence (data and log) in TDI environments and validate that the storage array meets the SAP KPIs for bandwidth (MB/s) and latency (microseconds).

### HANA-HWC-ES 1.0

The HANA-HWC-ES 1.0 certification scenario was the first scenario that SAP provided. It used the underlying `fstest` tool to perform and validate I/O operations on the file systems. From a file system perspective, `fstest` performed sequential and random I/O operations, while most of the random file system I/Os were treated as sequential I/Os by the storage array. Therefore, it was possible to achieve the SAP KPIs using a smaller number of disks.

The VMAX 10K, 20K, and 40K arrays were certified using the HANA-HWC-ES 1.0 scenario. The configuration and scalability recommendations for these models that this guide provides are based on the HANA-HWC-ES 1.0 scenario. When revalidating a VMAX 10K, 20K, or 40K at a customer site, see [SAP Note 1943937 - Hardware Configuration Check Tool - Central Note](#) (access requires an SAP username and password).

### HANA-HWC-ES 1.1

With SAP HANA 1.0 SPS 10 and higher, SAP introduced a new scenario for enterprise storage certifications. Version 1.1 uses the underlying `fsperf` tool for file system performance validation. With `fsperf`, random I/O operations required additional resources and the KPIs could be achieved only with more disks (HDDs) or with flash disks.

VMAX3 arrays were originally certified using HANA-HWC-ES 1.0, but have been recertified with the VMAX All Flash arrays using the HANA-HWC-ES 1.1 scenario. The configuration and scalability recommendations in this solution guide are based on the HANA-HWC-ES 1.1 scenario. When revalidating VMAX3 or VMAX All Flash arrays at a customer site, see [SAP Note 1943937 - Hardware Configuration Check Tool - Central Note](#) (access requires an SAP username and password).

## Storage design principles for SAP HANA on VMAX arrays

This section describes configuration recommendations that apply to SAP HANA production systems deployed on VMAX 10K, 20K, and 40K enterprise storage arrays. SAP HANA production systems in TDI environments must meet the SAP performance requirements and you must apply the specific configuration requirements described here.

---

**Note:** [Preparing the HANA nodes](#) and [Installing the SAP HANA scale-out cluster](#) provide instructions for installing SAP HANA on VMAX 10K, 20K and 40K arrays. The instructions describe an SAP HANA installation on a VMAX All Flash array, but they can also be used as an example of an installation on VMAX 10K, 20K and 40K arrays.

---

### Scalability

Table 1 provides guidelines for estimating the initial number of production SAP HANA hosts that can be connected, based on the performance tests we carried out on a VMAX 10K single engine using the SAP hwcct tool (for scenario HANA-HWC-ES 1.0).

---

**Note:** We determined the scalability of higher models and additional engines by extrapolating the VMAX 10K test results using the performance characteristics of the higher models.

---

The actual number of SAP HANA hosts that can be connected to a VMAX array in a customer environment might be higher or lower than the number listed in Table 1, depending on the actual workload. Dell EMC recommends using the SAP HANA hwcct tool from scenario HANA-HWC-ES 1.0 in customer environments to validate the SAP HANA performance and determine the maximum possible number of SAP HANA hosts on a given storage array.

**Table 1. VMAX 10K, 20K, and 40K scalability**

VMAX model	Number of available engines	Number of SAP HANA worker hosts
10K	1	12
	2	18
	3	24
	4	30
20K	1	12
	2	20
	3	28
	4	36
	5	44
	6	52
	7	60
	8	68
40K	1	12
	2	22

VMAX model	Number of available engines	Number of SAP HANA worker hosts
	3	32
	4	42
	5	52
	6	62
	7	72
	8	82

### FA director and port requirements

In addition to the information provided in the [Storage ports](#) section of this guide, special requirements apply when you connect SAP HANA nodes to the front-end director ports (FA-ports) of a VMAX 10K, 20K, or 40K array.

On a VMAX director, two FA-ports share a dedicated CPU core. Assume for illustration purposes that FA-1E:0 and FA-1E:1 share the same core. To achieve full I/O performance for SAP HANA deployments, use only one FA-port per CPU core on the I/O module. For example, use FA-1E:0 and leave FA-1E:1 unused. Do not use the adjacent port for non-SAP HANA applications.

Figure 2 and Figure 3 show the rear view of the VMAX engines with 4-port FC I/O modules (8 Gbps) for host connectivity. The four ports are named 0, 1, 2, and 3. Dell EMC recommends using the I/O ports marked with a yellow box (port 0 and port 2) for SAP HANA connectivity because ports 0 and 1 share one CPU core and ports 2 and 3 share another CPU core. Leave the adjacent ports unused.




Figure 2. Rear view of the VMAX 10K engine




Figure 3. Rear view of the VMAX 20K and 40K engine



## Virtual provisioning considerations

VMAX 10K, 20K, and 40K arrays use Dell EMC Virtual Provisioning™ to provide storage capacity to an application. The capacity is allocated using thin data devices (TDATs) and provided in thin pools based on the disk technology and RAID type. Thin devices (TDEVs) are host-accessible devices bound to thin pools and natively striped across the pool to provide the highest performance.

### Fully Automated Storage Tiering for Virtual Pools (FAST VP)

Dell EMC Fully Automated Storage Tiering (FAST™) for Virtual Pools (VP) moves data in a VMAX from one storage tier to another to optimize performance and economy. Most-used data is stored on the fastest (and most expensive) storage tier, and least-used data is stored on the slowest (and least expensive) storage tier. In SAP HANA environments, these data movement algorithms may not lead to the expected results. SAP HANA changes data in the memory of the servers. When a savepoint is written to the file system, it might be written to a different location in the file system even if data has not been changed in memory.

Therefore, using FAST VP does not provide any advantage with the SAP HANA persistence, and a single storage tier based on a 10000 rpm or 15000 rpm hard disk is preferred on VMAX 10K, 20K and 40K arrays. Using SSDs for the SAP HANA storage tier is optional—it is not needed to meet the performance requirements for the HANA-HWC-ES 1.0 certification scenario.

### Disks and disk groups

Disk groups in a VMAX array contain disks of the same technology, that is, FC, SAS, or SSDs. Disk groups are preconfigured with the array or created when drives are added to the array to increase available capacity. Contact Dell EMC Customer Services to create disk groups.

To further isolate the SAP HANA workload from non-SAP HANA applications on a shared array, consider a dedicated disk group for SAP HANA. Carefully consider the choice between workload isolation and performance and avoid dedicated SAP HANA disk groups with fewer than 40 disks (HDDs).

Each SAP HANA worker node requires at least ten disks (10000 rpm or 15000 rpm HDDs) to meet the IOPS requirements based on the SAP HANA-HWC-ES 1.0 certification scenario. For best performance configurations for SAP HANA, the number of disks in a disk group should be divisible by eight when you are using HDDs.

### RAID considerations

To provide the best write performance for the SAP HANA persistence, Dell EMC recommends a RAID1 mirrored configuration for the TDATs on 10 k rpm or 15 k rpm disks. When you use SSDs, RAID5 3+1 is recommended.

### Thin pools

Dell EMC recommends creating one thin pool for all SAP HANA data volumes in the VMAX and a second thin pool for the SAP HANA log volumes. However, if a limited number of disks are available in smaller SAP HANA environments, you could improve performance by using a single thin pool for both types of volumes. Thin pools consist of TDATs. The number and size of the TDATs in a thin pool depends on the SAP HANA

capacity requirements and must be configured according to VMAX configuration best practices.

When creating TDATs, ensure that eight hypervolumes (splits) are allocated on each disk. Adjust the TDAT size according to the usable disk capacity. Also ensure that TDATs are created on all available disks.

Example:

- 64 disks are available → 512 hypervolumes (64 x 8 = 512)
- A RAID1 TDAT has two hypervolumes → create 256 TDATs (512 / 2 = 256)
- The size of the TDATs should be the usable capacity of a disk divided by 8.
- Create thin pools for data and log files using the number of TDATs that meet the capacity requirements.

### Meta volumes for data and log

Each SAP HANA worker host requires one data volume and one log volume for the persistent file systems. The volume sizes depend on the [SAP HANA capacity requirements](#).

- Create a TDEV Meta volume for HANA data using 32 members.
- Create a TDEV Meta volume for HANA log using 8 members.

## Masking view

VMAX 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:

- Initiator group
- Port group
- Storage group

### Initiator group

The initiator group contains the initiators (WWNs) of the HBAs on the SAP HANA host. Connect each SAP HANA host to the VMAX array with at least two HBA ports for redundancy.

### Port group

The port group contains the front-end director ports to which the SAP HANA hosts are connected. See [Storage ports](#) and [FA director and port requirements](#) to determine the number of ports required for the SAP HANA installation.

### Storage group

An SAP HANA scale-out cluster uses the shared-nothing concept for the database persistence, 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.

## Storage design principles for SAP HANA on VMAX arrays

The VMAX 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.



## Storage design principles for SAP HANA on VMAX3 and VMAX All Flash arrays

The configuration recommendations described in this section apply to SAP HANA systems deployed on VMAX3 and VMAX All Flash enterprise storage arrays.

- **VMAX3**—100K, 200K, 400K arrays
- **VMAX All Flash**—250F/FX, 450F/FX, 850F/FX, 950F/FX arrays

SAP HANA production systems in TDI environments must meet the SAP KPIs and the following configuration requirements.

### Scalability

Based on our performance tests on a VMAX 100K single engine and both the VMAX250F and VMAX450F with a single V-Brick using the SAP hwcct tool (for HANA-HWC-ES 1.1 certification), Table 2 and Table 3 provide guidelines for estimating the initial number of SAP HANA production hosts that can be connected.

---

**Note:** We determined the scalability of higher models and additional engines and V-Bricks by extrapolating the VMAX 100K, 250F, and 450F test results using the performance characteristics of the higher models.

---

The actual number of SAP HANA hosts that can be connected to a VMAX array in a customer environment can be higher or lower than the number of SAP HANA hosts listed in Table 2 and Table 3, depending on the actual workload. 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.

**Table 2. VMAX3 scalability**

VMAX3 model	Engines	Number of SAPHANA worker hosts
100K	1	12
	2	20
200K	1	16
	2	28
	3	40
	4	52
400K	1	20
	2	32
	3	44
	4	56
	5	68
	6	80
	7	92

VMAX3 model	Engines	Number of SAPHANA worker hosts
	8	104

**Table 3. VMAX All Flash scalability**

VMAX All Flash model	V-Bricks	Number of SAP HANA worker hosts
250F	1	14
	2	22
450F	1	16
	2	28
	3	42
	4	56
850F	1	20
	2	32
	3	48
	4	64
	5	80
	6	96
	7	112
	8	128
950F	1	24
	2	36
	3	54
	4	72
	5	90
	6	108
	7	126
	8	144

**Service level objective-based provisioning**

**FAST in VMAX3**

With VMAX3, FAST is enhanced to include both intelligent storage provisioning and performance management using service level objectives (SLOs). SLOs automate the allocation and distribution of application data to the correct data pool and storage tier without manual intervention. In contrast to FAST VP on the VMAX 10K, 20K, or 40K, where data movement is triggered by data aging, VMAX3 SLOs are tied to an expected average I/O latency for both reads and writes. Therefore, both the initial provisioning and ongoing application performance are automatically measured and managed based on

compliance with storage tiers and performance goals. FAST samples the storage activity every 10 minutes and, when necessary, moves data at FAST's sub-LUN granularity, which is 5.25 MB (42 extents of 128 KB). SLOs can be dynamically changed (promoted or demoted) at any time, and FAST continuously monitors and adjusts data location at sub-LUN granularity across the available storage tiers to match the performance goals provided. All this is done automatically within the VMAX3 storage array, without the need to deploy a complex application ILM<sup>3</sup> strategy or use host resources for migrating data due to performance needs.

A storage resource pool (SRP) is a collection of data pools that provides FAST with a domain for capacity and performance management. By default, a single default SRP is preconfigured at the factory. The FAST data movements are performed within the boundaries of the SRP. SAP HANA does not require a separate SRP and can co-exist with non-SAP HANA applications in the same SRP when proper SLOs are used, as described in [SLO and workload type best practices for SAP HANA](#).

## SLOs and workload types

Five SLOs are available, varying in expected average response time targets, as shown in Table 4. An additional optimized SLO is available that has no explicit response time target associated with it.

**Table 4. SLO information**

SLO	Expected average response time
Diamond	0.8 – 2.3 ms
Platinum	3.0 – 4.4 ms
Gold	5.0 – 6.5 ms
Silver	8.0 – 9.5 ms
Bronze	14.0 – 15.5 ms
Optimized (default)	N/A

When you select any SLO other than an optimized SLO, you can further qualify it by workload type: online transaction processing (OLTP) or decision-support system (DSS). The OLTP workload is focused on optimizing performance for small-block I/Os, while the DSS workload is focused on optimizing performance for large-block I/Os. The workload type can also specify whether to account for any overhead associated with replication (local or remote). The workload type qualifiers for replication overhead are OLTP\_Rep and DSS\_Rep, where Rep denotes replicated. Table 5 shows the available workload types.

**Table 5. VMAX3 workload types**

Workload	Description
OLTP	Small block I/O workload
OLTP with replication	Small block I/O workload with local or remote replication

<sup>3</sup> Information Lifecycle Management (ILM) refers to a strategy of managing application data based on policies. It involves complex data analysis, mapping, and tracking practices.

Workload	Description
DSS	Large block I/O workload
DSS with replication	Large block I/O workload with local or remote replication

**SLO and workload type best practices for SAP HANA**

Dell EMC recommends the following SLO configurations for SAP HANA installations. Note that VMAX All Flash arrays deliver a single top-tier Diamond service level only.

**SAP HANA persistence for SAP HANA production installations**

For SAP HANA production installations, Dell EMC recommends using the Diamond SLO for the SAP HANA persistence (the data and log volumes). Although SAP HANA uses various block sizes, we recommend specifying the OLTP workload type. This ensures that the VMAX3 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 for production installations:

- 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

**SAP HANA persistence for nonproduction installations**

Although the SAP performance KPIs do not apply to SAP HANA nonproduction installations, those installations are still critical components in an overall SAP landscape. Dell EMC recommends using the Gold SLO for all nonproduction installations.

**SAP HANA installation (/hana/shared/)**

A Bronze SLO is sufficient when you are using eNAS in a VMAX3 array to provide the NFS share for the SAP HANA installation file system.

**Operating system boot image**

The OS boot image can also reside on a Bronze SLO.

**What if an existing VMAX3 does not support Diamond SLO?**

When a new VMAX3 array is configured for an SAP HANA environment, SSDs are automatically added with the requested capacity to enable the Diamond SLO for SAP HANA production installations. However, customers may want to use existing VMAX3 arrays without SSDs for SAP HANA. In these environments, SAP HANA performance and response times might be impacted and operations such as restarts, table reloads, or backups might perform slower than expected. Also, the number of SAP HANA production nodes that can be connected to the array (as shown in [Table 1](#)) will be lower. In such cases, use the highest available SLO for SAP HANA.

## SLO considerations for “noisy neighbors” and competing workloads

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

## 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. It offers VMAX3 and VMAX All Flash 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 neighbor with low performance requirements can ensure that a spike in I/O demand does not affect the SAP HANA workload and performance.

# VMAX All Flash inline compression

VMAX All Flash inline compression significantly increases the effective capacity of the VMAX All Flash array by reducing the physical footprint of the dataset and therefore the number of flash disks required. Inline compression also intelligently optimizes system resources to ensure the system is delivering the best balance of performance and efficiency at all times.

## Adaptive Compression Engine

VMAX Hypermax OS 5977 and later provides an Adaptive Compression Engine (ACE). ACE performs the following functions to minimize any impact on the array performance:

- Offloads compression onto the hardware compression I/O module (Asteroid SLIC)
- Dynamically adjusts the backend compression pools
- Uses activity-based compression to tag the most active data in the SRP to skip the compression workflow, ensuring an optimal response time

---

**Note:** ACE is available for all VMAX All Flash storage arrays (250/450/850/950 F/X) running Hypermax OS 5977 or later. Dell EMC supports all data services offered on the VMAX All Flash array with compression enabled, including SnapVX and SRDF.

---

## Compression-ready

Compression-ready describes the system state when the default SRP is capable of storing compressed data. For a system to be able to compress data, it must have:

- One compression I/O module per director (this is standard on All Flash arrays)
- Compression enabled on the storage group (this is the default)
- A system compression ratio (CR) set within the system—the expected average of the compressed data

Compression is managed at the storage group level. You can easily enable or disable compression, allowing customers to target the workloads that would benefit most. You can also apply compression to existing data. Each storage group reports the effective CR related to the data specific to that storage group. The storage group CR can therefore differ from the overall system CR.

For more information, see [VMAX All Flash with the Adaptive Compression Engine](#)

### SAP HANA and VMAX All Flash inline compression

The expected compression for the SAP HANA database is typically lower than for traditional databases. This is because the SAP HANA column store compresses data automatically within the database and optimizes SAP HANA compression after any changes. Therefore, the achievable CRs with VMAX inline compression can vary depending on the SAP HANA dataset.

To take a closer look at the VMAX inline compression effect on SAP HANA, Dell EMC performed laboratory tests with the VMAX 250F. The array has 11 TB of usable capacity and a system CR setting of 2:1.

Dell EMC engineers installed SAP ERP IDES<sup>4</sup> on SAP HANA and generated activity on test datasets. We used the Mitrend tool to determine the expected compressibility of the SAP HANA persistence. The expected compressibility we observed was between 1.5:1 and 2:1. This is the optimal achievable CR for the sample SAP HANA datasets used.

---

**Note:** Mitrend is a pre-sales tool that enables field specialists to work with customers to determine the expected compressibility of their data. For more information, see the *Data Analysis* and *Selecting Compression Reduction* topics on the [Mitrend website](#) (login required).

---

You can use Dell EMC Unisphere<sup>TM</sup> for VMAX to observe the individual storage group CR efficiencies, as Figure 4 shows. Each storage group contains SAP HANA databases with different datasets. Compression is not enabled on the NoCOMP\_SG storage group. The other two storage groups with compression enabled have achieved CR efficiencies of 1.4:1 and 1.6:1, resulting in physical storage capacity savings of between 40 and 60 percent.

> [Storage](#) > [Storage Groups Dashboard](#) > [Storage Groups](#)

Name	1 ▼ SRP	Service Level	Compression Ratio	Capacity (GB)
NoCOMP_SG	SRP_1	Diamond	1.0:1	2560.01
HANA_SG	SRP_1	Diamond	1.4:1	4352.01
HANA_NonProd_SG	SRP_1	Diamond	1.6:1	4608.01

**Figure 4. Storage group compression ratios on VMAX All Flash**

With VMAX inline compression, activity-based compression algorithms tag 20 percent of the busiest data in the SRP\_1 storage resource pool as active data. To ensure optimal performance, the active data skips the compression workflow and remains uncompressed until it becomes less active. The active data affects the compression efficiency shown for a particular storage group. More than 20 percent of the SRP is allocated, so a higher

---

<sup>4</sup> Enterprise resource planning internet demonstration and evaluation system (ERP IDES)

effective CR is achieved because more of the data in the SRP is tagged as not active and therefore goes through the compression workflow.

We used the SAP hwcct tool (for HANA-HWC-ES 1.1) to examine the throughput and latency from provisioned storage with compression enabled. The KPI results showed no material impact on performance compared with KPI test results with compression disabled.

A VMAX All Flash array typically consists of a variety of storage groups, SAP HANA production and nonproduction databases, and non-SAP HANA workloads, each with its own CR. The overall system CR is therefore a mix of the various underlying storage group ratios. With a normal mix of workloads, you can expect to see an approximately 2:1 system CR. This ratio could be higher or lower depending on the workload mix. When inline compression is combined with other VMAX All Flash space-saving capabilities (such as virtual provisioning, zero space reclaim, and space-efficient snapshots), an overall efficiency rate of 4:1 is achievable.

# Storage configuration and installation of an SAP HANA scale-out cluster on a VMAX All Flash array: Example

This section shows how to:

- Create and configure the persistent storage (data and log) on a VMAX All Flash array 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 the SAP lifecycle management command-line tool hdblcm

The storage configuration on VMAX3 arrays follows the same steps. The only difference is the service level selection when the storage group is created. VMAX3 arrays offer additional service level options for SAP HANA installations. Select a Diamond service level for production installations and a Gold service level for nonproduction installations.

## Configuring the VMAX All Flash array

We used the Unisphere for VMAX GUI to configure all storage devices, storage groups, port groups, host groups, and the masking view for the SAP HANA scale-out cluster. Log in to Unisphere and follow these steps:

1. Go to the Storage Group dashboard, as shown in Figure 5.




Figure 5. Storage Group dashboard

2. Click the **Total** tile to view the existing storage groups, as shown in Figure 6.




Figure 6. Existing storage groups



3. Click **Create Storage** to create a new storage group for our SAP HANA cluster, as shown in Figure 7.

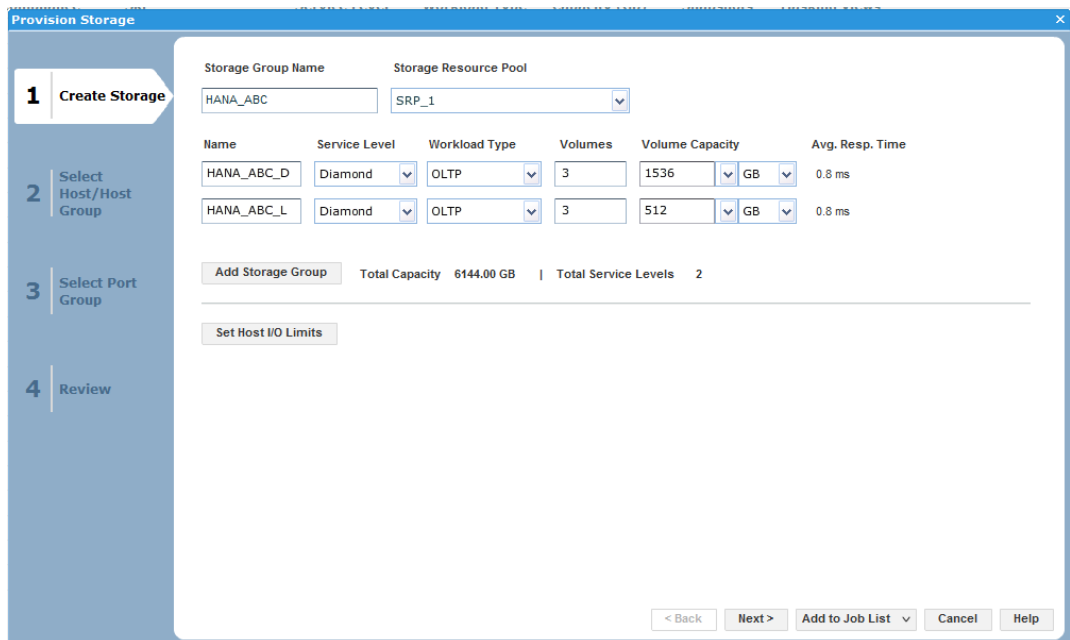


Figure 7. Creating a new storage group

For our 3+1 SAP HANA cluster we needed three data volumes, each of 1.5 TB capacity, and three log volumes, each of 512 GB capacity. We therefore created a cascaded storage group with one top-level group (**HANA\_ABC**), one sub-group for all data volumes (**HANA\_ABC\_D**), and a second sub-group for all log volumes (**HANA\_ABC\_L**). We then had to specify the number and size of the volumes to be created. On a VMAX All Flash array, the only available service level is **Diamond**. On VMAX3 hybrid arrays, we would select **Diamond** service level for SAP HANA production instances and **Gold** for all SAP HANA nonproduction instances.

4. Click the down arrow on **Add to Job List** and select **Run Now**. The new cascaded storage group is created, as shown in Figure 8.

Name	Compliance	SRP	Service Level	Workload Type	Capacity (GB)	Snapshots	Masking Views
EMBEDDED_NAS_DM_SG	●	None	Diamond	OLTP	105.82	0	1
Y HANA_ABC	—	None	None	None	6144	0	0
HANA_ABC_D	●	SRP_1	Diamond	OLTP	4608	0	0
HANA_ABC_L	●	SRP_1	Diamond	OLTP	1536	0	0
HANA_AFA	—	None	None	None	6144.01	0	1

Figure 8. Cascaded storage group

5. Select the **HANA\_ABC\_D** storage group and click **View Details** to view information about the volumes created, as shown in Figure 9.

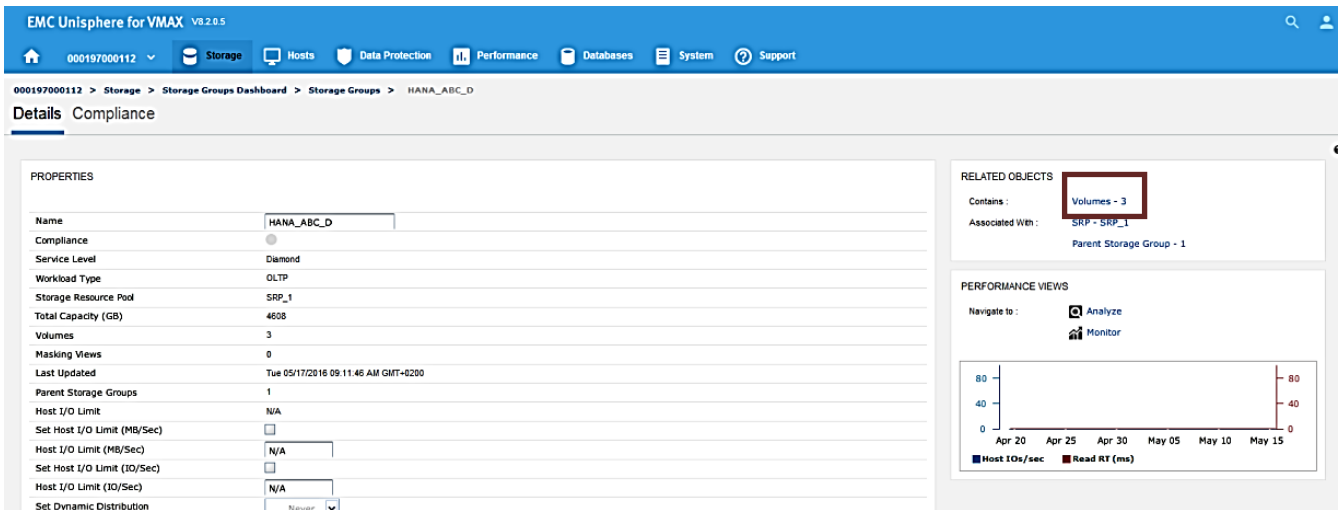


Figure 9. Created volumes

6. Click **Volumes – 3** in the **RELATED OBJECTS** area to view the list of data volumes. Figure 10 shows an example.

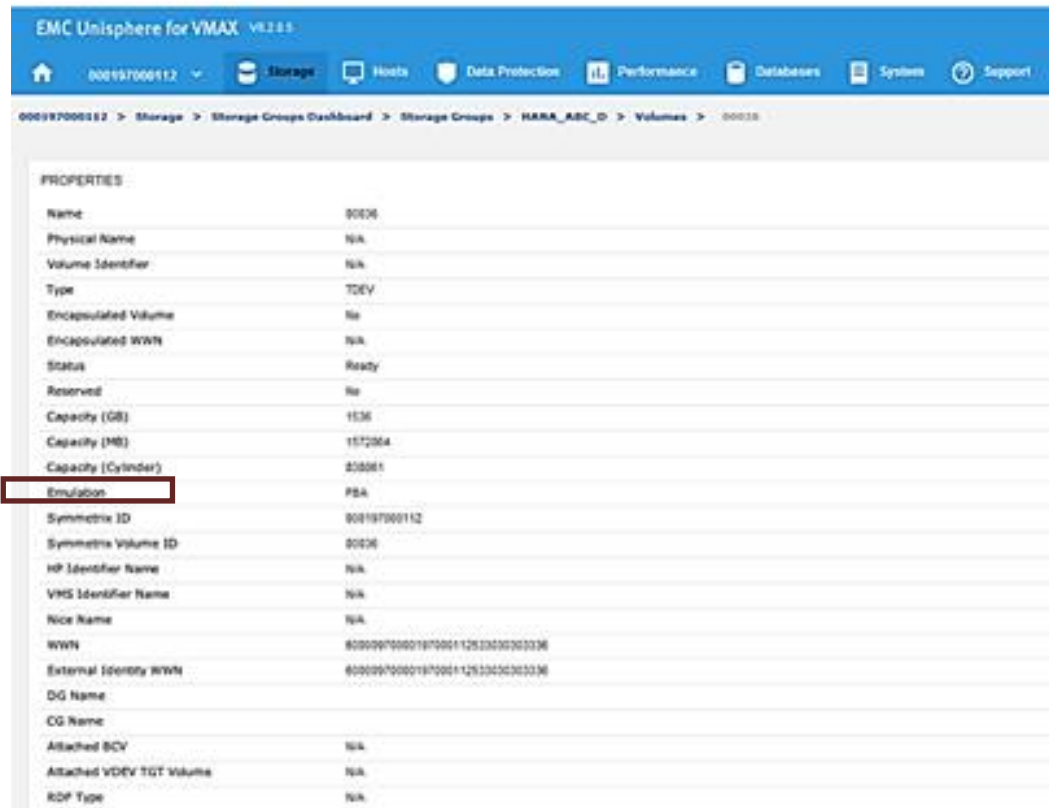


Figure 10. Data volumes list

7. Note the WWN of the volume.
8. Repeat the previous steps for all your data and log volumes. The SAP HANA storage connector fcClient uses the WWN to identify a storage LUN. The WWN is specified in the SAP HANA `global.ini` file.

## Set initiator aliases (host names and HBAs)

Click **Host > Initiators**. Right-click on the initiator, select **Rename Alias**, and specify an alias for each SAP HANA host and HBA port (initiator), as shown in Figure 11. Do this just once for each initiator, even though an initiator may be connected to multiple storage ports.

Initiator	Dir:Port	Alias	Logged In	On Fabric	Port Flag Overrides	Hosts
10000090fa53fd1c	FA-2D:28	c240-02c/hba2	Yes	Yes	No	2
10000090fa53fd1c	FA-1D:28	c240-02c/hba1	Yes	Yes	No	2
10000090fa53fd1c	FA-2D:30	c240-02c/hba1			No	2
10000090fa53fd1c	FA-1D:30	c240-02c/hba1			No	2
10000090fa53fd1d	FA-2D:28	c240-02c/hba2			No	2
10000090fa53fd1d	FA-1D:28	c240-02c/hba1			No	2
10000090fa53fd1d	FA-2D:30	c240-02c/hba2	Yes	Yes	No	2
10000090fa53fd1d	FA-1D:30	c240-02c/hba2	Yes	Yes	No	2

**Figure 11. Setting initiator aliases**

In this example, each HBA is connected to one port on each Director (1D and 2D) of the VMAX All Flash single V-Brick. Make a note of the ports on which your initiators are logged in.

If you do not know the initiator WWNs of your SAP HANA nodes, run the following Linux command on the node to find the information below it:

```
# systool -c fc_host -v | grep -i port_name
```

The output is:

```
port_name      = "0x10000090fa53fd1c"
port_name      = "0x10000090fa53fd1d"
```

## Create a Host Group

To create a host group:

1. Select **Host > Create Host Group**.
2. Enter a name in the format `HANA_ABC`, select the hosts that belong to the SAP HANA cluster, and click **Add**, as shown in Figure 12.

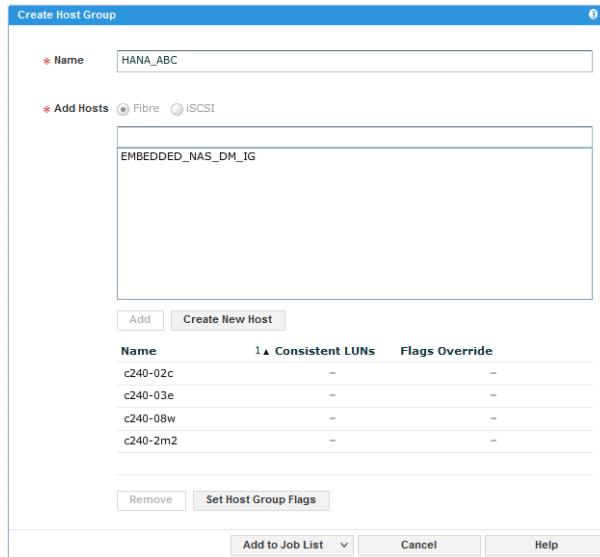


Figure 12. Creating host groups

3. Click the down arrow on **Add to Job List** and select **Run Now**. Your host group is created, as shown in Figure 13.

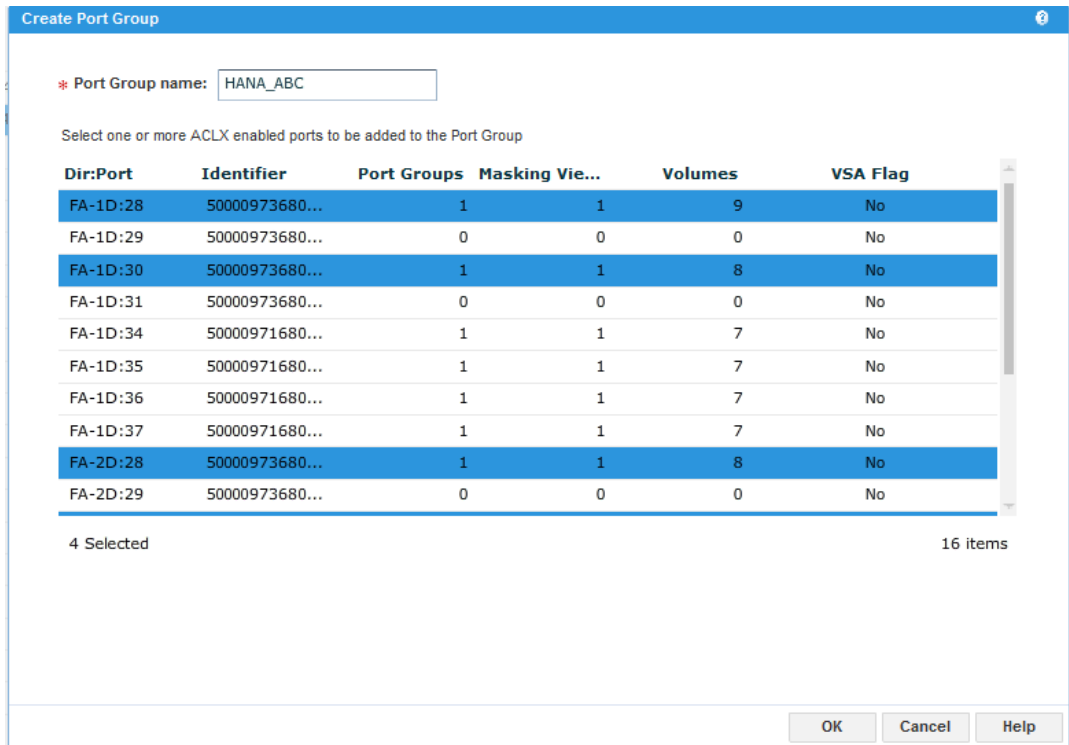
Name	Masking Views	Initiators	Consistent LUNs	Port Flag Overrides	Last Update
EMBEDDED_NAS_DM_IG	1	8	--	--	Wed 04/13/2016 07:44:21 PM GMT+0200
▼ HANA_ABC	0	8	--	--	Tue 05/17/2016 09:35:58 AM GMT+0200
c240-02c	1	2	--	--	Sat 04/16/2016 08:42:55 AM GMT+0200
c240-03e	1	2	--	--	Sat 04/16/2016 08:42:55 AM GMT+0200
c240-08w	1	2	--	--	Sat 04/16/2016 08:42:55 AM GMT+0200
c240-2m2	1	2	--	--	Sat 04/16/2016 08:42:55 AM GMT+0200

Figure 13. Created host group

### Create a port group


To create a port group:

1. Select **Hosts > Port Groups > Create Port Group**. Enter a name such as **HANA\_ABC**, as shown in Figure 14, and mark the ports your initiators are logged into by holding down the Control key.



**Figure 14. Creating a port group**

2. Click **OK**. The warning message shown in Figure 15 may appear.



**Figure 15. Port group warning message**

3. Click **OK** to confirm that you will have multiple ports from the same director in the port group—in this example, FA-1D:28, FA1D:30, FA-2D:28, and FA-2D:30.

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 masking view

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

1. Select **Hosts > Masking View > Create Masking View**.
2. Enter a masking view name (`HANA_ABC`) and select the host group, port group, and storage group you created in the previous steps, as shown in Figure 16.




Figure 16. Creating a Masking View

3. Click **OK**. The Masking View is created, as shown in Figure 17.




Figure 17. Masking View

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, together with the SAP HANA storage connector fcClient, mounts<sup>5</sup> the volumes to the proper SAP HANA node and I/O fencing. Correct preparation of the SAP HANA nodes and the SAP HANA `global.ini` file is required before the actual SAP HANA installation. [Preparing the HANA nodes](#) provides these steps.

## Preparing the HANA nodes

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

- The OS is installed and properly configured using the SAP recommendations (in this example we used SUSE Linux 12 SP2 for SAP applications).

<sup>5</sup> When using the 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, for example, on a VMAX3 or VMAX All Flash with eNAS capability, 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 an 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 following sections of this guide provide the steps for:

- Configuring storage on the SAP HANA nodes
- Preparing the SAP HANA `global.ini` file
- Installing an SAP HANA scale-out instance with the SAP HANA database lifecycle management command-line tool `hdblcm`, using the storage volumes created in the previous steps.

### Linux native multipathing (DM-MPIO)

Use the following entries in the `/etc/multipath.conf` file:

```
defaults {
    user_friendly_names    no
}
devices {
    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           "1"
        path_checker        "directio"
        prio                "const"
        prio_args           ""
        flush_on_last_del   yes
        fast_io_fail_tmo    off
        dev_loss_tmo        120
    }
}
```

Restart multipathing after changing the `/etc/multipath.conf` file by using the following command:

```
# service multipath restart
```

---

**Note:** For the OS version and storage array MPIO configuration settings for native multipathing, see the [Dell EMC Host Connectivity Guide for Linux](#).

---

### Initializing the HANA persistence

The HANA persistence must be visible to every node in the HANA cluster. To achieve this, either use the `rescan-scsi-bus.sh` command or reboot each node.

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

#### List all 1.5 TB data volumes:

```
# multipath -ll | grep -B1 -A5 1.5T

360000970000197000112533030303338 dm-6 EMC,SYMMETRIX
size=1.5T features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
  |- 0:0:0:11 sdm 8:192 active ready running
  |- 0:0:4:11 sdcs 70:0 active ready running
  |- 1:0:7:11 sdlk 68:288 active ready running
  `-- 1:0:0:11 sdfs 130:224 active ready running
--
360000970000197000112533030303337 dm-5 EMC,SYMMETRIX
size=1.5T features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
  |- 0:0:0:10 sdl 8:176 active ready running
  |- 0:0:4:10 sdcr 69:240 active ready running
  |- 1:0:7:10 sdlj 68:272 active ready running
  `-- 1:0:0:10 sdfr 130:208 active ready running
--
360000970000197000112533030303336 dm-18 EMC,SYMMETRIX
size=1.5T features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
  |- 0:0:0:9 sdk 8:160 active ready running
  |- 0:0:4:9 sdcq 69:224 active ready running
  |- 1:0:7:9 sdli 68:256 active ready running
  `-- 1:0:0:9 sdfq 130:192 active ready running
```

#### List all 512GB log volumes:

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

360000970000197000112533030303339 dm-4 EMC,SYMMETRIX
size=512G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
  |- 0:0:0:12 sdn 8:208 active ready running
  |- 0:0:4:12 sdct 70:16 active ready running
  |- 1:0:7:12 sdll 68:304 active ready running
  `-- 1:0:0:12 sdft 130:240 active ready running
360000970000196701016533030313841 dm-47 EMC,SYMMETRIX
size=512G features='1 queue_if_no_path' hwhandler='0' wp=rw
```



```

`-+- policy='service-time 0' prio=1 status=active
  |- 0:0:1:15 sdaf 65:240 active ready running
  |- 0:0:5:15 sddi 71:0 active ready running
  |- 1:0:5:15 sdjj 8:464 active ready running
  ` 1:0:2:15 sdgu 132:160 active ready running
360000970000197000112533030303342 dm-7 EMC,SYMMETRIX
size=512G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
  |- 0:0:0:13 sdo 8:224 active ready running
  |- 0:0:4:13 sdcu 70:32 active ready running
  |- 1:0:7:13 sdml 68:320 active ready running
  ` 1:0:0:13 sdfu 131:0 active ready running
    
```

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

Next, initialize the devices and create the Linux XFS file system on each of the devices. Follow this example command:

```
# mkfs.xfs /dev/mapper/360000970000197000112533030303341
```

After all the file systems have been created, 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 be used
- An installation parameter file used by the `hdblcm` command-line script with customized installation parameters

### Prepare the global.ini file

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

Ensure that the `global.ini` file has the following content:

```

[storage]
ha_provider = hdb_ha.fcClient
partition_*_*_prtype = 5
partition_*_data__mountoptions = -o inode64
partition_*_log__mountoptions = -o inode64,nobarrier
partition_1_data__wwid = 360000970000197000112533030303336
partition_1_log__wwid = 360000970000197000112533030303339
partition_2_data__wwid = 360000970000197000112533030303337
partition_2_log__wwid = 360000970000197000112533030303341
partition_3_data__wwid = 360000970000197000112533030303338
partition_3_log__wwid = 360000970000197000112533030303342
    
```

Ensure that the partition entries match the unique device identifier displayed. Do this by using the `multipath -ll` command with a preceding **3**.

Place the `global.ini` file (this name is mandatory) in a directory on the `/hana/shared/` file system, for example, `/hana/shared/ABC_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 offer the efficiency of installing all components at one time, while automating the installation and providing further flexibility to customers. In the following example, we used the `hdblcm` command-line interface to install our HANA 3+1 scale-out cluster.

1. In the shared file system, go to the `HDB_LCM_LINUX_X86_64` directory into which the HANA installation DVD ISO file has been extracted:
2. Create a template installation parameter file by using the following command:
3. After the template has been created, modify the following parameters in the file to match our environment:

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

```
# ./hdblcm --action=install --  
dump_configfile_template=ABC_install.cfg
```

```
# Directory root 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=C240-08W
```

```
# Directory containing a storage configuration  
storage_cfg=/hana/shared/ABC_cfg
```

---

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

---

```
# SAP HANA System ID  
sid=ABC
```

```
# 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=C240-03E:storage_partition=2:role=worker,C240-02C:storage_partition=3:role=worker,C240-2M2: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. See the [SAP HANA Installation and Upgrade Guide](#) for more information. 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 specific 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 using the following command:

```
# ./hdblcm --action=install --configfile=ABC_Install.cfg
```

```
SAP HANA Lifecycle Management - SAP HANA 1.00.110.00.1447753075
*****
```

```
Scanning Software Locations...
```

```
Detected components:
```

```
  SAP HANA Database (1.00.110.00.1447753075) in
/SAPShare/software/SAP_HANA_SPS11_IM/51050506/DATA_UNITS/HDB_SERVER_LINUX_X86_64/server
  SAP HANA AFL (incl.PAL,BFL,OFL,HIE) (1.00.110.00.1447766426)
in
/SAPShare/software/SAP_HANA_SPS11_IM/51050506/DATA_UNITS/HDB_AFL_LINUX_X86_64/packages
  SAP TRD AFL FOR HANA (1.00.110.00.1447766426) in
/SAPShare/software/SAP_HANA_SPS11_IM/51050506/DATA_UNITS/HDB_TRD_AFL_LINUX_X86_64/packages
  SAP HANA Database Client (1.00.110.00.1447753075) in
/SAPShare/software/SAP_HANA_SPS11_IM/51050506/DATA_UNITS/HDB_CLIENT_LINUX_X86_64/client
  SAP HANA Studio (2.2.8.000000) in
/SAPShare/software/SAP_HANA_SPS11_IM/51050506/DATA_UNITS/HDB_STUDIO_LINUX_X86_64/studio
  SAP HANA Smart Data Access (1.00.6.001.0) in
/SAPShare/software/SAP_HANA_SPS11_IM/51050506/DATA_UNITS/SAP_HANA_SDA_10_LINUX_X86_64/packages
  SAP HANA XS Advanced Runtime (1.0.9.258635) in
/SAPShare/software/SAP_HANA_SPS11_IM/51050506/DATA_UNITS/XSA_RT_10_LINUX_X86_64/packages
```

```

    XS Monitoring 1 (1.001.1) in
    /SAPShare/software/SAP_HANA_SPS11_IM/51050506/DATA_UNITS/XSA_CONTE
NT_10/XSAC_MONITORING-1.1.1.zip
    XS Services 1 (1.001.0) in
    /SAPShare/software/SAP_HANA_SPS11_IM/51050506/DATA_UNITS/XSA_CONTE
NT_10/XSAC_SERVICES-1.1.0.zip
    SAP Hana Demo Model for XS Advanced 1.0 (1.001.4) in
    /SAPShare/software/SAP_HANA_SPS11_IM/51050506/DATA_UNITS/XSA_CONTE
NT_10/XSAC_SHINE-1.1.4.zip

```

```

Collecting information from host 'c240-03e'...
Collecting information from host 'c240-02c'...
Collecting information from host 'c240-2m2'...
Information collected from host 'c240-03e'.
Information collected from host 'c240-2m2'.
Information collected from host 'c240-02c'.
Restrict maximum memory allocation? [n]: n
Enter Certificate Host Name For Host 'c240-08w' [c240-08w]:
Enter Certificate Host Name For Host 'c240-03e' [c240-03e]:
Enter Certificate Host Name For Host 'c240-02c' [c240-02c]:
Enter Certificate Host Name For Host 'c240-2m2' [c240-2m2]:
Enter System Administrator (abccadm) Password:
Confirm System Administrator (abccadm) Password:
Enter Database User (SYSTEM) Password:
Confirm Database User (SYSTEM) Password:

```

Summary before execution:

=====

```

SAP HANA Components Installation
  Installation Parameters
    Remote Execution: ssh
    Installation Path: /hana/shared
    Local Host Name: c240-08w
    Root User Name: root
    Directory containing a storage configuration:
    /hana/shared/ABC_cfg
    SAP HANA System ID: ABC
    Instance Number: 09
    Database Mode: single_container
    System Usage: custom
    Location of Data Volumes: /hana/data/ABC
    Location of Log Volumes: /hana/log/ABC
    Certificate Host Names: c240-08w -> c240-08w, c240-2m2 ->
c240-2m2, c240-03e -> c240-03e, c240-02c -> c240-02c
    System Administrator Home Directory: /usr/sap/ABC/home
    System Administrator Login Shell: /bin/sh
    System Administrator User ID: 1010
    ID of User Group (sapsys): 79
  Software Components
    SAP HANA Database
      Install version 1.00.110.00.1447753075
      Location:
    /SAPShare/software/SAP_HANA_SPS11_IM/51050506/DATA_UNITS/HDB_SERVE
R_LINUX_X86_64/server
      SAP HANA AFL (incl.PAL,BFL,OFL,HIE)
      Do not install
      SAP TRD AFL FOR HANA

```

```
Do not install
SAP HANA Database Client
Do not install
SAP HANA Studio
Do not install
SAP HANA Smart Data Access
Do not install
SAP HANA XS Advanced Runtime
Do not install
Additional Hosts
c240-2m2
  Role: Database Standby (standby)
  Storage Partition: N/A
c240-03e
  Role: Database Worker (worker)
  Storage Partition: 2
c240-02c
  Role: Database Worker (worker)
  Storage Partition: 3

Do you want to continue? (y/n): y

[...]

SAP HANA system installed
You can send feedback to SAP with this form: https://c240-
08w:1129/lms1/HDBLCM/ABC/feedback/feedback.html
Log file written to '/var/tmp/hdb_ABC_hdblcmm_install_2016-04-
17_17.09.22/hdblcmm.log' on host 'C240-08W'.
```

## Optimizing file I/Os 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. It uses lightweight, platform-independent wrappers around system calls.
- **FileFactory & File**—Used for large, complex streams of I/O requests on the data and log volumes and for backup and recovery. It 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`

## SAP HANA 1.0

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

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

```
# hdbparam --paramset fileio.max_parallel_io_requests=256
# hdbparam --paramset fileio.async_read_submit=on
# hdbparam --paramset fileio.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]*.

Set the parameter as follows in the `global.ini` file:

- `max_parallel_io_requests=256`

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).

**Note:** The instructions in this section 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.

## Conclusion

### Summary

Using SAP HANA in TDI deployments with Dell EMC VMAX, VMAX3, and VMAX All Flash enterprise storage arrays provides many benefits, including reducing hardware and operational costs, lowering risk, improving availability and performance, and increasing flexibility in hardware vendor selection.

All VMAX arrays are certified by SAP and can be used for SAP HANA installations on production and nonproduction systems and on single-node (scale-up) and scale-out systems.

### Findings

During our tests with HANA on VMAX arrays, we observed the following:

- The SAP HANA-HWC-ES 1.0 certification scenario requires fewer disk resources, as described in [Storage design principles for SAP HANA on VMAX arrays](#).
- The SAP HANA-HWC-ES 1.1 certification scenario makes higher demands in relation to disk configuration.
- SAP HANA production installations on VMAX3 and VMAX All Flash systems require SSDs for the SAP HANA persistence.
- VMAX All Flash inline compression delivers space-efficiency savings with expected compression ratios of between 1.4:1 and 1.6:1 for SAP HANA environments, with no material impact on SAP HANA storage performance.
- 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

## References

### Dell EMC documentation

The following documentation on [Dell EMC.com](#) or [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.

- [EMC Symmetrix VMAX Family with Enginuity Product Guide](#)
- [EMC VMAX3 Family Product Guide](#)
- [EMC VMAX All Flash Product Guide](#)
- [Unisphere for VMAX 8.0.3 Documentation Set](#)
- [Business Continuity and Disaster Recovery with EMC VMAX3 for SAP HANA TDI Deployments Solution Guide](#)
- [Business Continuity Best Practices for SAP HANA TDI with EMC Symmetrix VMAX White Paper](#)
- [Dell EMC Host Connectivity Guide for Linux](#)
- [VMAX All Flash Compression FAQ](#)
- [VMAX All Flash with the Adaptive Compression Engine](#)
- [VMware Virtualized SAP HANA with EMC Storage Solution Guide](#)

### SAP HANA documentation

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

- [SAP HANA Administration Guide](#)
- [SAP HANA Master Guide](#)
- [SAP HANA Server Installation and Update Guide](#)
- [SAP HANA Studio Installation and Update Guide](#)
- [SAP HANA Technical Operations Manual](#)

#### Web resources

- [SAP HANA Enterprise Cloud](#)
- [SAP HANA One](#)
- [SAP HANA Platform](#)
- [SAP HANA Storage Requirements](#)
- [SAP HANA Tailored Data Center Integration](#)

---

**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](#)