Configuration and Deployment Best Practices for SAP HANA on Dell EMC PowerFlex Hyperconverged Infrastructure

August 2020

H17751.2

Validation Guide

Abstract

This validation guide provides best-practice configuration guidelines for SAP HANA deployments on Dell EMC PowerFlex hyperconverged systems. The solution incorporates Dell EMC PowerEdge servers, the VMware vSphere hypervisor, and PowerFlex software-defined storage.

Dell Technologies Solutions

D&LL Technologies

Contents

Introduction

Executive summary

SAP HANA and S/4HANA are disrupting IT

SAP HANA is foundational for next-generation data-driven intelligent enterprise applications and processes. Consider the momentum of SAP HANA:

- Since the 2011 release of the SAP HANA platform for real-time computing, 28,000 customers have adopted the platform.

- With the 2025 end of maintenance and support for SAP Business Suite 7 applications, IT departments must begin planning a strategy and architecture to run on the SAP HANA platform.

IT departments require a modern infrastructure to deliver the performance, scale, protection, and availability to power new mission-critical, data-driven intelligent applications with SAP S/4 HANA and SAP Leonardo.

Dell Technologies has invested in the SAP HANA certification program for hardware to help customers accelerate implementations with reduced risk. Dell first partnered with SAP in 2011, when they announced the SAP HANA appliance model for targeted workloads. Since then, the Dell Technologies portfolio has expanded to include SAP's tailored data center integration (TDI), converged infrastructure (CI), and hyperconverged infrastructure (HCI) models, creating a path to run SAP HANA like any other database on shared infrastructure.

In 2018, SAP extended their SAP HANA platform support to include a certification and validation process for hyperconverged platforms, including a software-defined storage platform and a software stack for the hypervisor. SAP has certified the PowerFlex HCI deployment model, where compute and storage reside in the same server, for SAP HANA deployments.

Dell EMC PowerFlex is a rack-scale hyperconverged system that delivers an exceptional combination of performance, resiliency, and flexibility to meet the needs of the modern enterprise data center. PowerFlex HCI includes the VMware vSphere hypervisor and the PowerFlex platform. This validation guide describes PowerFlex configurations for running SAP HANA workloads.

Solution overview

This hyperconverged solution for SAP HANA that is based on PowerFlex infrastructure incorporates Dell EMC PowerEdge servers with vSphere, integrated networking, and the PowerFlex platform. The solution encompasses design configurations and deployment options. Customers use SAP sizing tools to size SAP systems on VMware virtualized infrastructure, and then work with Dell Technologies representatives to determine the PowerFlex HCI requirements and configure and deploy the PowerFlex solution.

Key benefits

This PowerFlex HCI solution enables organizations to innovate faster and accelerate their IT operations. Customers implementing the solution can expect the following benefits:

- **Agility**—A modern SAP landscape management experience that delivers PowerFlex infrastructure for SAP applications including SAP HANA. The solution provides a faster time-to-value and the ability to rapidly respond to changing business requirements.

- **Engineering**—Compute, networking, storage, hypervisor, and management components that are integrated, manufactured, supported, and sustained as a single product to deliver a seamless solution experience.

- **Performance**—Supreme I/O performance, high bandwidth, and submillisecond latencies with PowerFlex software. Throughput and IOPS scale in direct proportion to the number of nodes that are added to the system.

- **Simplicity**—Simplified and unified PowerFlex management with management and operations software, making it easier to implement, consume, and manage resources.

- **Flexibility**—The ability to design and implement software-defined data centers that are based on your applications, budget, and growth demands. Add units incrementally and adjust compute and storage resources linearly or independently based on your needs.

Document purpose

This guide provides validated best practices for the design and configuration of virtualized SAP HANA deployments on PowerFlex integrated systems that are based on vSphere and PowerFlex storage virtualization.

The guide does not replace the requirement for PowerFlex implementation services by Dell Technologies Professional Services. Detailed PowerFlex installation and implementation documentation is available for Dell Technologies service personnel only.

Audience

This guide is for SAP Basis, system, and storage administrators and presales architects who design mission-critical SAP HANA systems for deployment on PowerFlex hyperconverged systems. You should have some knowledge of PowerFlex systems, PowerEdge servers, and VMware vSphere hypervisor.

We value your feedback

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

Authors: Werner Katzenberger, Donagh Keeshan

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

Dell EMC PowerFlex integrated systems

Introduction   The PowerFlex system creates a server-based storage area network (SAN) by combining PowerFlex virtualization with Dell EMC PowerEdge servers. Local storage resources are combined to create a virtual pool of block storage with varying performance tiers. The system architecture enables you to scale from as few as four nodes to over a thousand nodes. In addition, the PowerFlex system provides enterprise-grade data protection, multitenant capabilities, and add-on enterprise features such as quality of service (QoS), thin provisioning, and snapshots.

Note: Dell EMC PowerFlex was formely known as the VxFlex family and VxFlex OS. Some of the figures in this document and some of the cross-referenced documentation might still refer to "VxFlex."

Dell Technologies offers the PowerFlex system in the following consumption models:

- **Ready Nodes**—PowerFlex Ready Nodes are stand-alone servers that are preconfigured to run the PowerFlex platform with qualified hypervisors or operating systems.

    Because PowerFlex Ready Nodes are not included in the SAP HANA HCI certification, they are outside the scope of this guide.

- **Integrated rack**—The PowerFlex integrated rack comes preinstalled with integrated top-of-rack (ToR), aggregation, and out-of-band (OOB) management switches to provide optimal network traffic flow. The PowerFlex integrated rack is designed using Cisco Nexus aggregation switches that provide 10/40 GbE IP uplink connectivity to the external network for superior performance. Often, network bottlenecks can limit the scale of an HCI system. The 10/25 GbE ToR switches in the PowerFlex integrated rack eliminate these restrictions and provide a path for future growth.

    The PowerFlex integrated rack is included in the SAP HANA HCI certification for the PowerFlex HCI solution using VMware vSphere hypervisor and the PowerFlex platform.

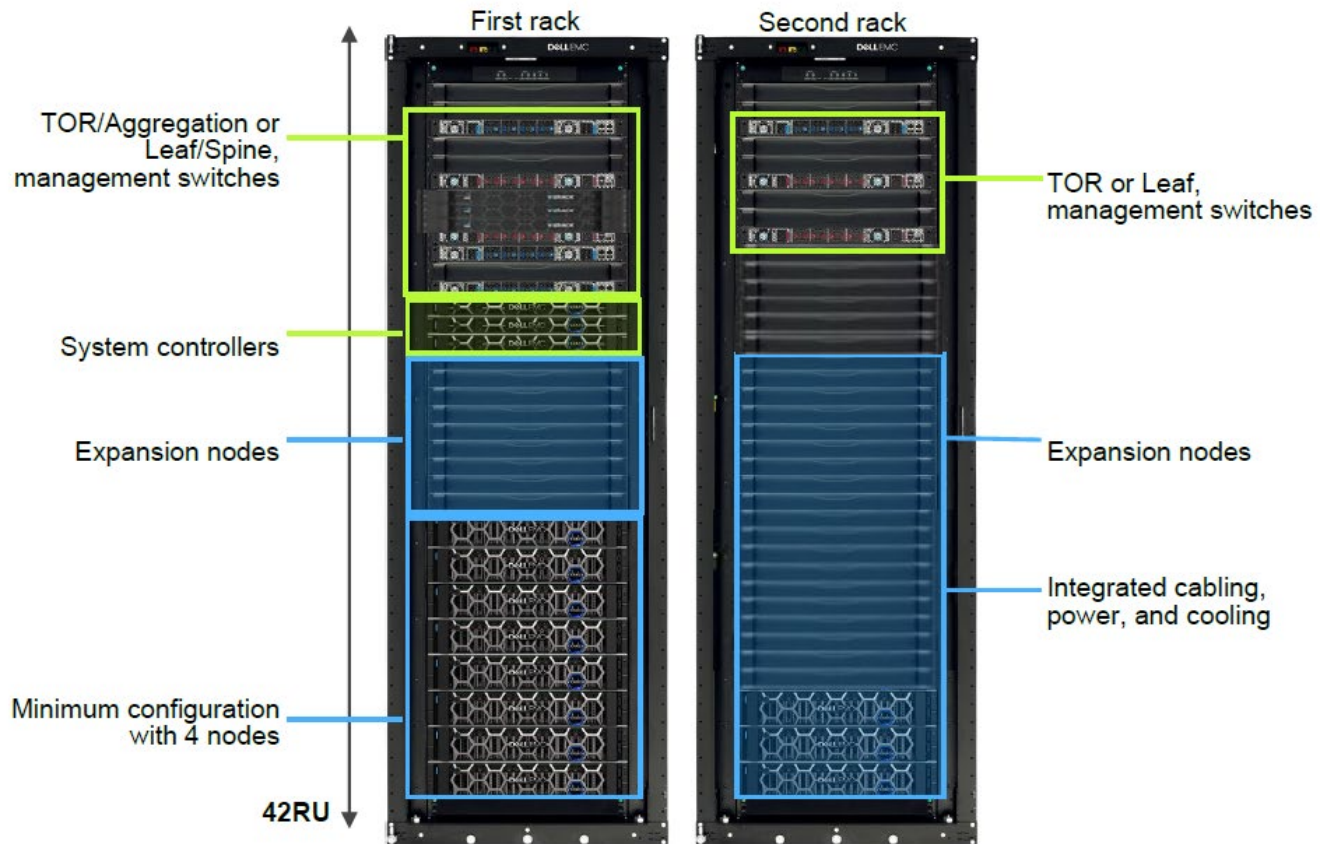   The following figure shows the integrated network:

Figure 1.   PowerFlex integrated rack, an engineered system with integrated rack-scale network fabric

- **PowerFlex appliance**—The PowerFlex appliance is an HCI system for high-performing, multi-hypervisor environments that is easy to acquire, own, and scale. Because its components are validated and tested prior to shipping, the PowerFlex appliance offers a turnkey experience and gives customers "bring your own" networking flexibility.

  The PowerFlex appliance offers enhanced lifecycle management and automation using PowerFlex Manager. PowerFlex Manager provides software alerts, monitoring, reporting, and lifecycle management.

  The PowerFlex appliance is available with smaller starting configurations than the PowerFlex integrated rack and provides greater network flexibility. The network switch must be in the support list for PowerFlex Manager.

SAP HANA-certified PowerFlex system components

The following table shows the PowerFlex system components that are certified for SAP HANA:

Table 1.   SAP HANA-certified PowerFlex components

Components	Configuration
Consumption model	Dell EMC PowerFlex integrated rack Dell EMC PowerFlex appliance
Compute	Dell EMC PowerEdge servers: PowerFlex-R640, PowerFlex-R740xd, PowerFlex-R840
CPU architecture	Intel Xeon Scalable Processors (Skylake and Cascade Lake): Silver, Gold, Platinum
Storage	Local SSD/NVMe drives
Storage virtualization	PowerFlex 2.6.1 and 3.0
Networking	Cisco Nexus switches
Hypervisor support	VMware vSphere 6.5 and 6.7
Bare-metal support	Not supported; SAP certification applies only to PowerFlex HCI nodes running VMware vSphere
Management infrastructure	PowerFlex Manager, PowerFlex GUI, VMware vCenter

PowerFlex deployment models

Two deployment options are available with PowerFlex systems:

- **HCI/single-layer architecture**—A model where compute and storage reside in the same server, creating a single-layer architecture.

- **Two-layer model**—A model resembling a traditional SAN architecture, with compute-only and storage-only servers.

The SAP rules for SAP HANA HCI certification enable deployment of SAP HANA only in an HCI/single-layer architecture. The two-layer model is not supported.

PowerFlex system hardware requirements for SAP HANA

Introduction

You can deploy SAP HANA on an existing or new PowerFlex hyperconverged system if your environment meets the hardware requirements that are described in this section.

SAP HANA HCI certification

SAP has certified SAP HANA on PowerFlex HCI systems. PowerFlex PowerEdge systems with Intel Skylake and Cascade Lake are supported and listed separately on the Certified and Supported SAP HANA hardware directory. Both PowerFlex certifications encompass:

- HCI/single-layer architecture with a vSphere hypervisor and the PowerFlex system. The SAP certification includes support for PowerFlex PowerEdge systems and both dual-socket R640 and R740xd servers and the four-socket R840 server.

- Support for half-socket and full-socket SAP HANA virtual machines (VMs). Half-socket deployments provide better use of CPU resources for smaller SAP HANA

systems and enable two SAP HANA VMs to run on a single CPU socket, with each VM using 50 percent of the available physical cores. Full-socket SAP HANA VMs use a single socket (or multiple sockets) and all physical cores exclusively.

The SAP HANA HCI certification for PowerFlex systems does not include support for the two-layer deployment model, bare-metal deployments, or any hypervisors other than vSphere.

SAP HANA sizing considerations

Before you deploy SAP HANA on a PowerFlex system, determine the number of production and nonproduction SAP HANA systems that you require and the CPU, memory, and disk capacity requirements for these systems.

With SAP HANA TDI Phase 5, SAP introduced customer-workload-driven SAP HANA system sizing. SAP application performance standard (SAPS) requirements for specific customer workloads are used to determine the type and number of processors that are required to run SAP HANA. Customers use the SAP HANA Quick Sizer tool and sizing reports to determine the optimal number and types of PowerFlex nodes, CPU types, and memory sizes and disks for their SAP HANA environment, and then communicate the results to Dell.

For more information about SAPS, see SAP Standard Application Benchmarks.

For more information about the Quick Sizer tool, see Quick Sizer.

CPU considerations

In addition to the CPU and memory requirements for the SAP HANA VMs, consider the CPU and memory requirements for the vSphere hypervisor and PowerFlex system. For the hypervisor, add 10 percent of the CPU resources to the SAP HANA CPU requirements to account for the vSphere overhead. The PowerFlex platform requires one CPU socket on each physical node for the PowerFlex storage VM (SVM). SAP does not allow sharing of this socket with production SAP HANA VMs.

Customers can order PowerFlex systems with four-socket nodes (the R840 server) and two-socket nodes (R640 and R740xd servers).

- On a four-socket node, three sockets are available for SAP HANA (75 percent).
- On a two-socket node, one socket is available for SAP HANA (50 percent).

These restrictions apply to production SAP HANA VMs only. Nonproduction SAP HANA VMs or non HANA VMs can share a CPU socket with the PowerFlex SVM.

Note: SAP might review these restrictions in the future and allow the running of production SAP HANA VMs on the same socket as the PowerFlex SVM. Refer to the latest SAP Notes when planning your SAP HANA on PowerFlex solution (access requires login credentials).

PowerFlex nodes for SAP HANA require Intel Silver, Gold, or Platinum CPU models offering different clock speeds and core counts. PowerFlex allows the installation of half-socket SAP HANA VMs, while the smallest configurable SAP HANA VM requires at least eight physical cores. Consider using CPUs with at least 16 cores in the PowerFlex nodes for half-socket SAP HANA VM deployments.

Memory requirements

SAP has restrictive requirements for the physical memory configuration of an SAP HANA server. These restrictions also apply to virtualized servers such PowerFlex nodes. A symmetric, homogeneous assembly of Dual In-Line Memory Modules (DIMMs) (the memory chips in the server) with maximum utilization of all memory channels per processor is required. Figure 2 and Figure 3 show the permitted configurations of the PowerFlex nodes for SAP HANA.

The following figure shows the supported DIMM/memory configurations for R640 and R740xd servers:

DIMM GB / RAM GB	128	192	384	576	768	1536	3072
8	16	24	48		96	192	384
16	8	12	24	12+12	48	96	192
32	4	6	12		24	48	96
64	2	3	6		12	24	48
128	1	1,5	3		6	12	24

Figure 2.   DIMM/memory configurations for R640 and R740xd servers

The following figure shows the supported DIMM/memory configurations for R840 servers:

DIMM GB / RAM GB	128	192	384	768	1152	1536	3072	6144
8	16	24	48	96		192	384	768
16	8	12	24	48	24+24	96	192	384
32	4	6	12	24		48	96	192
64	2	3	6	12		24	48	96
128	1	1,5	3	6		12	24	48

Recommended	
Neutral	
Not allowed	

Figure 3.   DIMM/memory configurations for R840 servers

Disk considerations

The PowerFlex system organizes the storage cluster into protection domains and storage pools.

Protection domains

A protection domain is a subset of storage data servers (SDS), the functionality of the PowerFlex SVM that runs on each host. At least four hosts are required in a protection domain. All the hosts in the domain must have the same configuration—for example, one protection domain for the four-socket R840 nodes and another protection domain for the two-socket R640 nodes.

Storage pools

A storage pool is a subset of the physical storage devices in a protection domain. The persistent devices (data and log) of the SAP HANA production nodes require the capacity to be allocated on SSD storage pools. Each PowerFlex node requires at least eight SSDs. The size and number of SSDs depend on the capacity requirements, as described in Capacity requirements.

Capacity requirements

If the data that is written to the PowerFlex cluster is configured within a protection domain, it is mirrored on two different physical nodes (SDSs). Therefore, it is necessary to double the estimated capacity required for the SAP HANA deployments and add the required

spare capacity to determine the final capacity requirements. If vSphere VM snapshots or PowerFlex storage snapshots are used, this capacity must also be considered.

Spare capacity

In a PowerFlex cluster, spare capacity is reserved for the PowerFlex system to use when recovery from a failure is required. The spare capacity stores rebuilt data after a component failure and cannot be used for storage allocations. The spare capacity must be at least the size of the largest node.

Network

To support the SAP HANA bandwidth requirements, Dell requires 25 GbE network connectivity of the PowerFlex nodes.

In addition to the system networks required for vSphere and PowerFlex, configure the following networks for SAP HANA when the function is used:

- App server
- Clients
- System replication
- Internode
- Backup

We recommend creating a new vSphere distributed switch (VDS) for the networks that you require for SAP HANA and assigning additional 2 x 25 GbE network cards as uplinks to this new distributed switch. To meet the performance requirements of the SAP HANA internode network communication, set an MTU size of 9,000 on this distributed switch and in the operating system network settings of the SAP HANA VM for the internode network. For information about configuring networks on a distributed switch, see _Dell EMC PowerFlex Networking Best Practices and Design Considerations_.

Configuring the PowerFlex system for SAP HANA

PowerFlex benefits

PowerFlex software uses the local disks and LAN of existing servers to create a virtual SAN that has all the benefits of external storage at a fraction of the cost and complexity. PowerFlex turns the existing local storage devices into shared block storage. For many workloads, PowerFlex storage is comparable to, or better than, external shared block storage.

PowerFlex software enables administrators to add or remove servers and capacity on-the-fly. The software immediately responds to the changes by rebalancing the storage distribution and achieving a layout that optimally suits the new configuration.

PowerFlex is designed and implemented with enterprise-grade resilience. In addition, the PowerFlex software features an efficient distributed self-healing process that overcomes media and server failures without requiring administrator involvement.

PowerFlex virtual SAN components

The PowerFlex virtual SAN consists of the following software components:

- **Meta Data Manager (MDM)**—Configures and monitors the PowerFlex system.

- **Storage Data Server (SDS)**—Manages the capacity of a single server and acts as a back-end for data access. The SDS is installed on all servers that are contributing storage devices to the PowerFlex system. These devices are accessed through the SDS.

- **Storage Data Client (SDC)**—A lightweight device driver that resides in the ESXi kernel that exposes PowerFlex volumes as block devices to the vSphere hypervisor.

PowerFlex software in an ESXi-based system

In the VMware environment, the MDM and SDS components are installed on a dedicated SVM, while the SDC is installed directly on the ESXi host. The following figure shows the architecture:
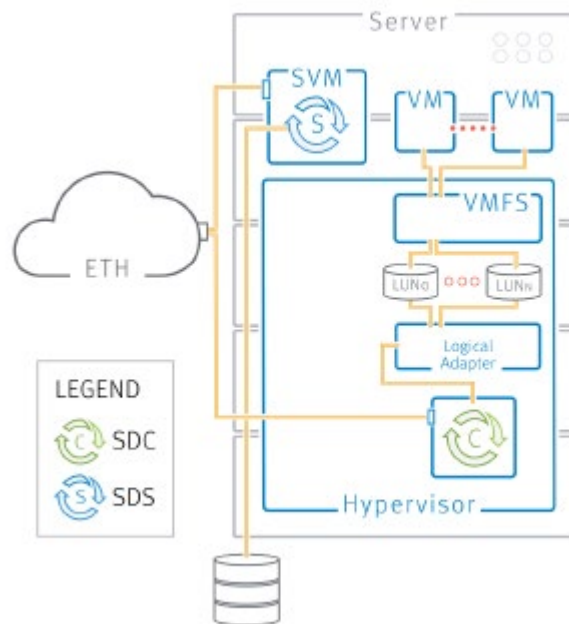
Figure 4.   PowerFlex on VMware ESXi

PowerFlex checksum feature

The PowerFlex system uses checksums to protect data that transits through the system. You can enable the checksum feature at the protection domain level and define it at the storage pool level. Depending on the type of workload, the checksum feature can have a significant impact on the performance.

SAP does not require a checksum feature for SAP HANA on the storage subsystem—the PowerFlex system, in this case. SAP HANA uses its own checksums to protect data that is stored on the persistence layer.

For the highest data protection requirements, the PowerFlex checksum feature is supported and can be enabled on the storage pool that is used for the SAP HANA persistence.

Note: The SAP HANA performance criteria for the persistence are met even with the checksum enabled.

PowerFlex management tools

You can provision, maintain, and monitor the PowerFlex system by using the following management clients:

- **Command Line Interface (CLI)**—The CLI enables you to perform the entire set of configuration, maintenance, and monitoring activities in the PowerFlex system. The CLI is available on the MDM node.

- **User Interface (UI)**—The PowerFlex UI is a stand-alone workstation application that enables you to perform standard configuration and maintenance activities and to monitor the health and performance of the storage system. You can use the UI to view the entire system and examine different elements as well as provision and modify many of the PowerFlex front-end and back-end components and features.

- **VMware plug-in**—The plug-in enables you to perform basic provisioning and maintenance activities in the VMware environment. The plug-in also provides a wizard to help you deploy the PowerFlex system in the VMware environment.

- **PowerFlex Installation Manager**—The PowerFlex Installation Manager is used for installing the PowerFlex system and upgrading and installing components as well as for running the `get-info` operation. The PowerFlex Installation Manager is installed as part of the PowerFlex gateway.

Note: You can perform many PowerFlex activities using more than one management tool.

SVM with SAP HANA VMs

The PowerFlex SVM is a Linux-based VM dedicated to the PowerFlex system. The PowerFlex SVM must run on a single CPU socket on each physical node within the PowerFlex cluster. The SVM requires eight vCPUs (cores). Although the CPU socket might have additional cores available, sharing this socket with production SAP HANA VMs is not permitted. The remaining cores on the socket can be used for nonproduction SAP HANA VMs and non SAP HANA VMs.

Configuring PowerFlex SVM on the first CPU socket

For performance reasons, you must associate the PowerFlex SVM to the first CPU socket (CPU 0) of each physical node.

Note: Before you can power off the SVM, you must put the PowerFlex node (SDS) into maintenance mode.

To associate the SVM with CPU 0:

1. Log in to the PowerFlex UI. Connect to the MDM management IP address or hostname using the username (the default is *admin*) and password you specified during installation.

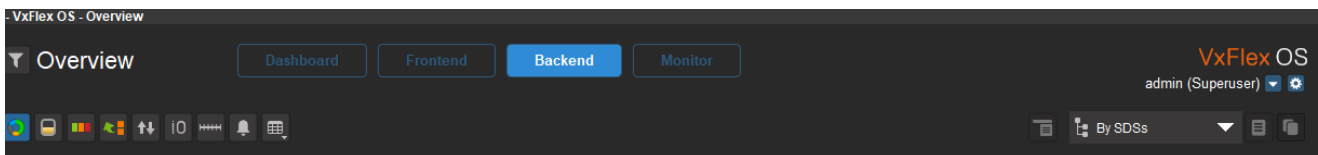2. Select the **Backend** view, as shown in the following figure:

Figure 5.   PowerFlex UI: Backend view

3. Select **By SDSs table filter**.

4. Right-click the SDS node with the SVM you are powering off, and then select **Enter Maintenance Mode**. Confirm this step.

5. Log in to the vCenter through the vSphere Web Client and locate the vSphere server.

6. Select the PowerFlex SVM, and then select **Actions** > **Power** > **Shut Down Guest OS**.

7. Select the SVM, and then select **Actions** > **Edit Settings**.

8. On the **VM Options** tab, expand **Advanced** and select **Configuration Parameters** > **Edit Configuration**.

9. Click **Add Row** to add a new option.

   a. In the **Name** column, enter **numa.nodeAffinity**.

   b. In the **Value** column, enter **0** to associate the VM with the first NUMA node (CPU socket 0).

10. Click **OK** twice.

11. Select the SVM and power it on manually.

12. In the **Backend** view, under **By SDSs table view**, right-click the SDS and select **Exit Maintenance Mode**. Confirm this step.

13. Repeat the preceding steps 2 to 12 for every SDS in the cluster to associate all SVMs to the first NUMA node.

PowerFlex storage provisioning

This section explains how to:

- Add a storage pool to a protection domain

- Create and map volumes

- Create a VMFS datastore

Adding a storage pool

A storage pool is always added to a protection domain. Each time that you add devices to the system, you must map the devices to storage pools. You can add storage pools both from the CLI and from the UI.

Storage pools in protection domains are preconfigured at the factory. If your environment requires that you create new storage pools, as an administrator, create a storage pool as follows:

1. In the PowerFlex UI, select **Backend**.

2. Right-click the **Protection Domain** that you want and select **Add Storage Pool**.

3. Enter a name for the new storage pool.

4. Select the storage pool configuration options that you want.

5. Using the vSphere plug-in for PowerFlex, in the vSphere Web Client select **Home** > **Inventories**, as shown in the following figure:
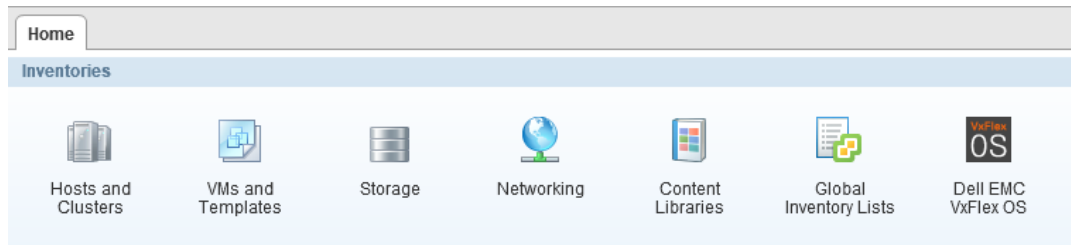
Figure 6.   Creating a storage pool

6.   Select the **Dell EMC PowerFlex** icon, and then double-click **Protection Domains** in the left pane.

7.   Select the protection domain that you want, and then select **Actions** > **Create Storage Pool**.

8.   Enter a name for the new storage pool.

9.   If you are prompted to enter a password, enter the PowerFlex authentication password.

Creating and mapping volumes

Volumes are created from storage pools and can be exposed to the hypervisor as a local storage device using the storage data client (SDC). When a volume is configured from a storage pool, it is distributed over all devices residing in that pool. Each volume block has two copies on two different SDSs, enabling the system to maintain data availability following a single point of failure.

Volumes can be:

- **Thick**—Capacity is allocated immediately, even if it is not used.

- **Thin**—Capacity is on reserve but not allocated until it is used.

Dell Technologies recommends using thick provisioning for creating PowerFlex volumes that will be used as datastores for the SAP HANA persistence.

As an administrator, add volumes as follows:

1.   In the PowerFlex UI, after selecting the **Volume** submenu from the **Frontend** tab, right-click the desired storage pool to create a volume.

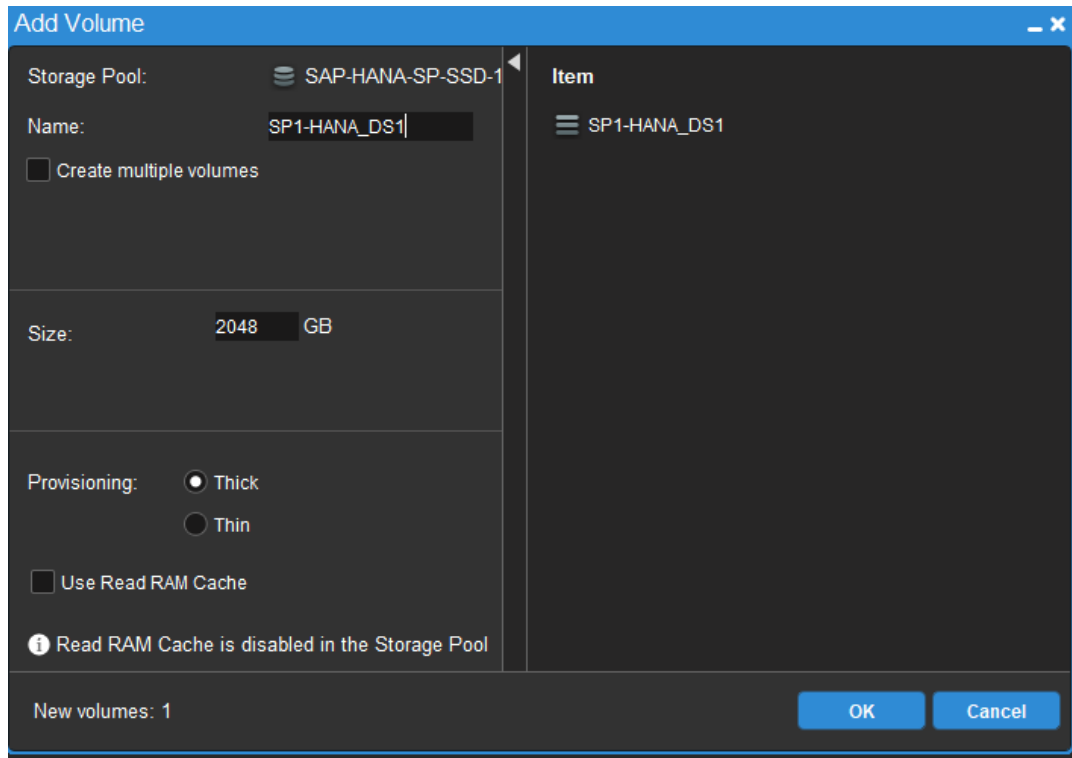The Add Volume page opens, as shown in the following figure:

Figure 7.   Adding a volume using the PowerFlex UI

2. To create more than one volume, enter the number of volumes that you want in the **Create multiple volumes** box.

    Multiple volumes are created with the same name. A number is appended to each volume.

Mapping volumes

Map PowerFlex volumes to all vSphere servers in the cluster before you create a datastore on the volumes. The PowerFlex SDC on the vSphere server presents the volumes to the ESXi operating system.

To map volumes using the PowerFlex UI:

1. Select the volumes in the **Frontend** > **Volumes** view.

2. From the context-sensitive menu, select **Map Volumes**.

    The Map Volumes window shows a list of the volumes that are to be mapped.

3. In the **Select Nodes** panel, select one or more SDCs to which you want to map the volumes and click **Map Volumes**.
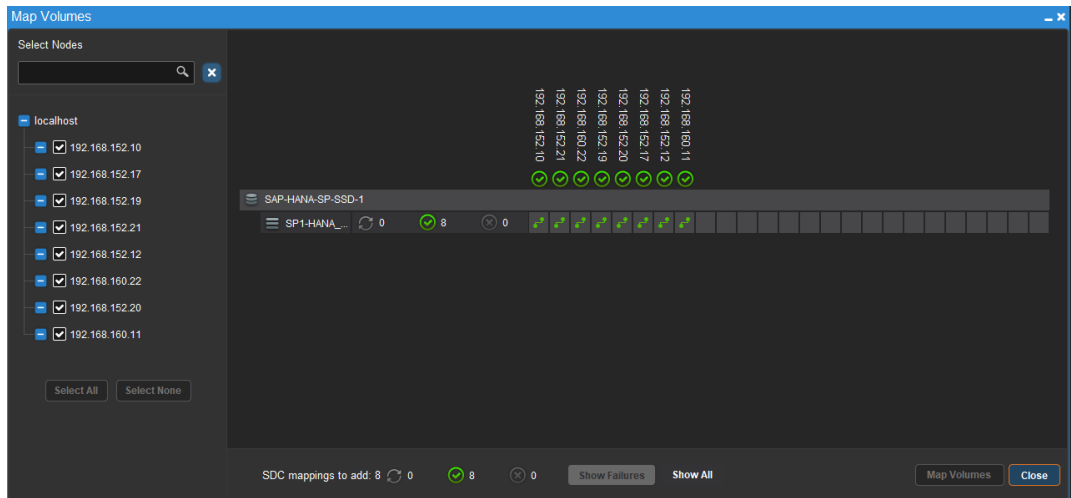
    The Map Volumes page opens.

Figure 8.   Mapping volumes to servers using the PowerFlex UI

4.   Keep the page open until the operation is complete and you can see the result.

5.   After you create a volume and map it to the desired SDCs, use the volume to provide storage capacity for the vSphere hypervisor. You can create a VMFS datastore on the hypervisor or expand an existing datastore.

Creating a VMFS datastore

To create a VMFS datastore on a PowerFlex volume:

1.   Select the volume in the PowerFlex UI and locate its ID in the properties sheet.

Note: You can also identify the unique ID in vSphere Web Client, which shows each PowerFlex device as **EMC Fibre Channel Disk** followed by an ID number that has **eui** as a prefix.

2.   In the vSphere Web Client, select the cluster in the left navigation pane, and then select the **New Datastore** icon on the **Datastores** tab.

The following figure shows the New Datastore page prior to completion of the create operation:
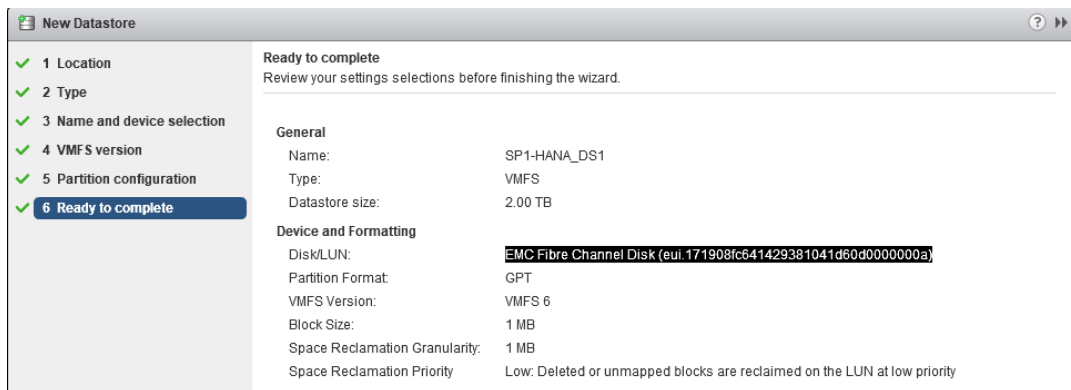

Figure 9.   Creating a vSphere datastore on a PowerFlex volume

3.   Click **Finish**.

The datastore that you created is connected to all the vSphere servers in the cluster and available to all VMs.

Configuration recommendations for SAP HANA virtual machines

Overview   The following section describes the configuration considerations and recommendations that apply to PowerFlex systems.

VMware best practices for SAP HANA deployments   VMware provides comprehensive information about best practices, requirements, and limitations for deploying SAP HANA VMs on the vSphere hypervisor. Refer to *Architecture Guidelines and Best Practices for Deployments of SAP HANA on VMware vSphere* when deploying SAP HANA VMs on PowerFlex systems with HCI nodes that are running vSphere.

SAP HANA deployment options   Scale-up (single-node) deployments are supported on PowerFlex systems that are running the vSphere hypervisor. Scale-up deployments are supported on both two-socket (R640/R740xd) and four-socket (R840) systems.

PowerFlex support for scale-out SAP HANA deployments is planned for the near future.

CPU and memory   Use SAP and VMware sizing guidelines to size CPU and memory requirements for the SAP HANA VMs. For more information, see SAP HANA sizing considerations. Overcommitment of CPU and memory resources is not allowed.

The SAP HANA HCI certification for PowerFlex systems includes support for half-socket and full-socket SAP HANA VMs. The differences are as follows:

- Half-socket deployments allow two SAP HANA VMs to run on a single CPU socket, with each VM using 50 percent of the available physical cores.

- Full-socket SAP HANA VMs use a single socket (or multiple sockets) and all physical cores exclusively.

Because CPU socket 0 (the first NUMA node) is reserved for the PowerFlex SVM, the remaining CPU sockets can be used for SAP HANA VMs. To prevent the SAP VMs from running on CPU 0, add the `numa.nodeAffinity` parameter to the configuration of the SAP HANA production VMs.

To add this parameter for each SAP HANA production VM:

1. Log in to the vCenter through the vSphere Web Client, and then select the SAP HANA VM.

2. From the **Power** menu, select **Actions** > **Shut Down Guest OS**.

3. After the guest operating system has shut down, select the SAP HANA VM again.

4. From the **Actions** menu, select **Edit Settings**.

5. On the **VM Options** tab, select **Advanced**.

6. Under **Configuration Parameters**, click the **Edit Configuration** button.

7. Click **Add Row** to add a new option.

    a. In the **Name** column, enter `numa.nodeAffinity`.

    b. In the **Value** column, enter `1` for a HANA VM running on a two-socket host or `1,2,3` for a HANA VM running on a four-socket host. The numbers **1,2,3**

enable the vSphere CPU scheduler to allocate CPU resources for a given SAP HANA virtual machine.

8.   Click **OK** twice to close the Edit VM dialog box.

9.   Power on the SAP HANA VM.

Note: When assigning memory to a production HANA VM, select the **Reserve all guest memory (All locked)** box.

CPU and memory sizes for SAP HANA VMs

With the NUMA architecture, each CPU in a server (NUMA node) has local memory assigned but can also access the nonlocal memory assigned to the other CPUs on the server. Because access to local memory is faster, we recommend configuring SAP HANA production VMs with the NUMA architecture in mind to avoid access to nonlocal memory.

The following table provides guidelines for configuring vCPU and memory for an SAP HANA production VM:

Table 2.   Configuration guidelines for vCPU and memory for an SAP production VM

SAP HANA VM size	vCPU (cores)	Memory	Maximum number of SAP HANA VMs per host
Small (half-socket)	50 percent of the cores of a NUMA node	Two-socket: 25 percent of the total RAM of the server Four-socket:12.5 percent of the total RAM of the server	Two-socket host: 2 Four-socket: 6
Medium (full-socket)	100 percent of the cores of a NUMA node	Two-socket: 50 percent of the total RAM of the server Four-socket: 25 percent of the total RAM of the server	Two-socket: 1 Four-socket: 3
Large (four-socket only)	All cores of two NUMA nodes	50 percent of the total RAM of the server	Four-socket: 1
X-Large (four-socket only)	All cores of three NUMA nodes	75 percent of the total RAM of the server	Four-socket: 1

Example:

A physical four-socket PowerEdge R840 server has a fully populated memory configuration with:

- 48 x 128 GB DIMMs (6 TB RAM), as shown in Figure 3.
- Intel Xeon Platinum 8180M CPUs with 28 cores per socket

The following table described the production SAP HANA VM sizes that are recommended:

Table 3.   Recommended SAP HANA VM sizes

VM	Configuration
Small	14 vCPUs, 1 socket, 768 GB RAM
Medium	28 vCPUs, 1 socket, 1,536 GB RAM
Large	56 vCPUs, 2 sockets (28 cores per socket), 3,072 GB RAM
X-Large	84 vCPUs, 3 sockets, (28 cores per socket), 4,608 GB RAM

The following figure shows a four-node PowerFlex cluster with the supported SAP HANA VM deployment options.

- Server 1 hosts a single X-Large SAP HANA VM using three sockets (NUMA nodes) and 75 percent of the available memory.

- Server 2 hosts one large and one medium VM.

- Server 3 hosts small half-socket VMs.

Note: The the SAP HANA VMs must use either full-sockets or half-sockets. Configurations with 1.5 or 2.5 sockets are not allowed.
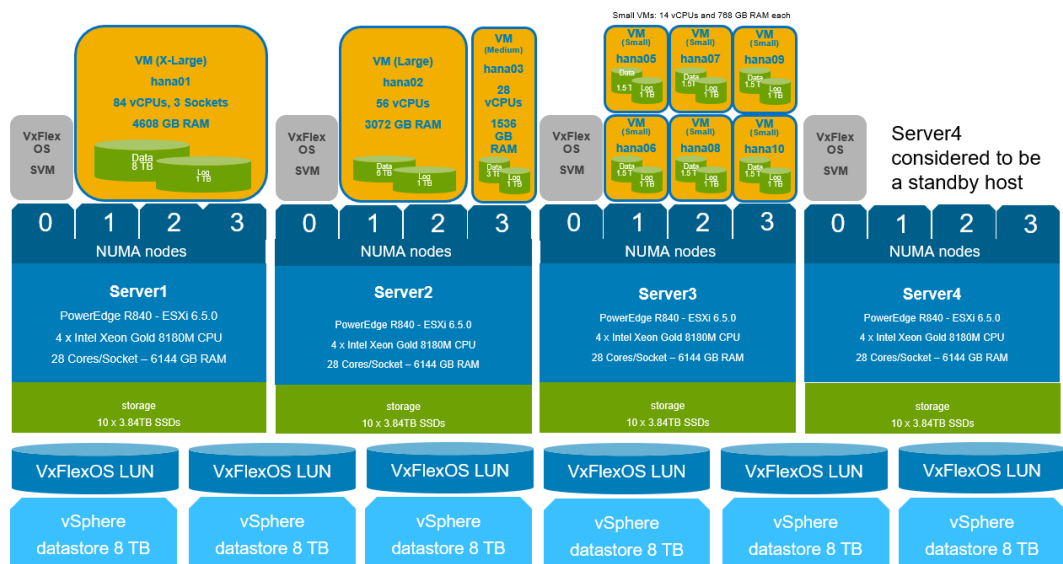

Figure 10.   Four-node PowerFlex cluster with SAP HANA VM deployment options

vCPUs and Hyperthreading

When hyperthreading is enabled on a physical PowerFlex node, the number of available vCPUs for a VM is twice the number of available physical cores. This technology allows a single processor core to execute two independent threads simultaneously. Hyperthreading does not double the performance of a system; instead, it increases the system's performance by better utilizing idle resources.

When you assign vCPUs to a production HANA VM, always use the physical core count even though the vSphere Web Client may allow twice that number when hyperthreading is

enabled. The VMware VMkernel CPU scheduler is designed to use scheduled threads against physical cores only.

Prepare the SAP HANA VM

The following section describes how to prepare:

- Operating system virtual disk
- SAP HANA persistence (data and log)
- Shared file system
- Formatting and mounting of the SAP HANA persistence

Operating system virtual disk

The virtual disk on which the operating system is installed can reside on any datastore that meets the functional requirements, specifically, support for vSphere vMotion.

SAP HANA persistence

Create the persistence of production SAP HANA VMs on a datastore that resides on PowerFlex volumes on an SSD storage pool.

Use the SAP Quick Sizer tool to determine the required capacity for the SAP HANA persistence. If the sizing details are not available, we recommend configuring the data and log device sizes based on the memory size of the SAP HANA VM:

- For data, configure 2 x RAM.
- For log, configure 0.5 x RAM (the maximum size is 1,024 GB).

For more information, see the Quick Sizer web page.

Note: Each SAP HANA VM uses its own data and log device and has access only to its pair of devices.

SAP HANA shared file system

SAP HANA installations require a file system to store the SAP HANA binaries, trace, and configuration files. This file system is mounted under the `/hana/shared` mount point.

Scale-up deployments

In SAP HANA scale-up deployments, this mount point can reside on the local system device. Alternatively, you can create a new disk on a datastore, format the disk using the XFS file system, and then mount it under the `/hana/shared` mount point. SAP HANA requires approximately 1 x the RAM memory capacity for the `/hana/shared` file system.

Formatting and mounting the SAP HANA persistence in the VM

After the SAP HANA VM starts, you must format and mount the persistence before you install the SAP HANA software.

Format the data and log devices as follows:

1. Run the following command, where `/dev/sdb` and `/dev/sdc` are the Linux devices for the persistence:

   ```
   mkfs.xfs /dev/sdb;mkfs.xfs /dev/sdc
   ```

2. Create the mount points by running the following commands:

```
mkdir -p /hana/data /hana/log

mount /dev/sdb /hana/data;mount /dev/sdc /hana/log
```

3. Add the SAP HANA mount points to the `/etc/fstab` file to ensure that the SAP HANA data and log volumes are mounted after a VM reboot.

High availability   The following VMware vSphere features support high availability (HA) in the SAP HANA system.

Automatic restart of SAP HANA VMs

Enable the vSphere HA feature in a vSphere cluster with SAP HANA. If one of the hosts in the cluster fails, this feature restarts a VM on another host that has enough free resources.

You can enable the SAP HANA service autostart feature in the SAP HANA software either at the time of the SAP HANA installation or by setting the autostart option to **1** in the `/hana/shared/<SID>/profile/<SID>_HDB<InstNo>_<hostname>` file.

Ensure that the minimum vSphere HA settings are in place, as shown in the following figure:

vSphere Availability   Edit...

vSphere HA is Turned ON

Runtime information for vSphere HA is reported under   vSphere HA Monitoring

Proactive HA is Turned OFF

Availability failure conditions and responses

Failure	Response	Details
Host failure	Restart VMs	Restart VMs using VM restart priority ordering.
Proactive HA	Disabled	Proactive HA is not enabled.
Host Isolation	Power off and restart VMs	VMs on isolated hosts will be powered off and restarted on available hosts.
Datastore with Permanent Device Loss	Disabled	Datastore protection for All Paths Down and Permanent Device Loss is disabled.
Datastore with All Paths Down	Disabled	Datastore protection for All Paths Down and Permanent Device Loss is disabled.
Guest not heartbeating	Reset VMs	VM monitoring enabled. VMs will be reset.

Figure 11.   vSphere HA settings

High availability for SAP HANA VMs

The SAP HANA auto-restart watchdog function automatically detects a failure and restarts the corresponding SAP HANA process: nameserver, indexserver, and so on. This service monitors the SAP HANA application and the associated services within a VM.

The VMware HA VM monitoring feature "Guest not heartbeating" restarts the VM guest operating system as well as SAP HANA on the same host. This feature manages operating system crashes if the SAP HANA autostart options are enabled.

Enable the heartbeat monitoring feature when you enable vSphere HA. We recommend setting **Heartbeat monitoring sensitivity** to **High**, as shown in the following figure:
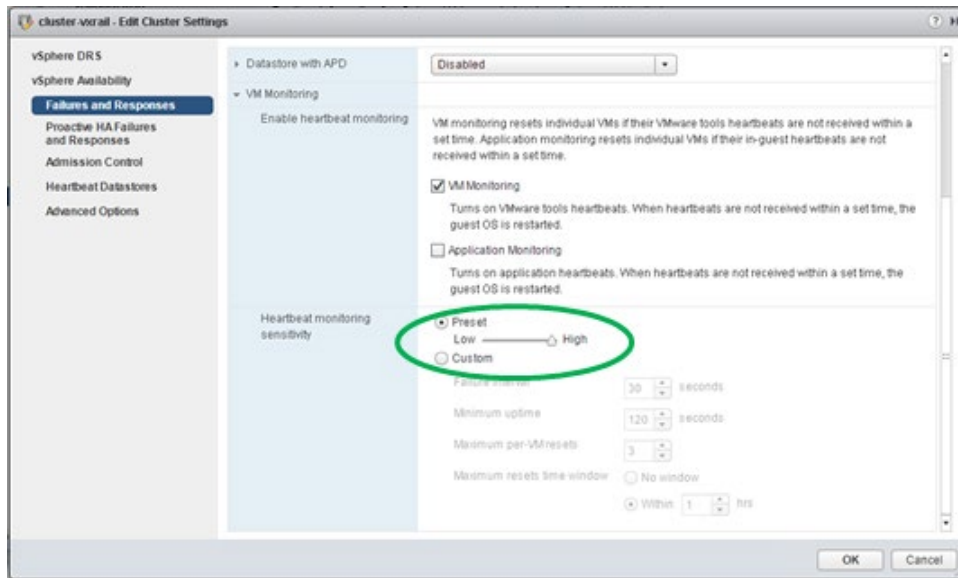
Figure 12.   vSphere heartbeat monitoring

VMware Tools must be installed and running in the VM. You can install VMware Tools as part of the operating system (`open-vm-tools`) or by using the vSphere web client and selecting **Guest OS** > **Install VMware Tools** in the VM context menu.

Optimizing file I/Os after the SAP HANA installation

Configure the SAP HANA file I/O layer with certain parameters to optimize file I/Os for a given storage solution and file system. The Linux XFS file system is used on all storage volumes for the SAP HANA persistence.

Setting file I/O layer parameters in SAP HANA 1.0

For SAP HANA 1.0 installations, run the SAP HANA `hdbparam` command as *<sid>*adm in the Linux shell:

```
su - <sid>adm
hdbparam –p          # lists current parameter setting
hdbparam --paramset fileio [DATA].max_parallel_io_requests=128
hdbparam --paramset fileio [LOG].max_parallel_io_requests=128
```

Setting file I/O layer parameters in SAP HANA 2.0

For SAP HANA 2.0 installations, use either `hdbsql` or the SQL function in SAP HANA Studio or cockpit and run the following SQL commands:

```
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio', 'max_parallel_io_requests[Data]') = '128' WITH
RECONFIGURE;
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio', 'max_parallel_io_requests[Log]') = '128' WITH
RECONFIGURE;
```

For more information, see SAP Note 2399079: Elimination of hdbparam in HANA 2 (access requires an SAP username and password).

Support requirements and considerations

SAP support calls

This section describes the checks you can perform before you report an SAP HANA issue with your PowerFlex HCI SAP HANA solution to SAP Support.

Identify a certified HCI solution

As part of the SAP HANA HCI certification requirements, SAP requires a method to detect if an SAP HANA VM is running in a certified hyperconverged environment on the vSphere hypervisor and the PowerFlex platform. SAP HANA Studio and the SAP HANA cockpit cannot report the hardware vendor and type when SAP HANA is running under VMware. It is important to report to SAP that SAP HANA is running on a PowerFlex HCI system.

A proven method to identify a certified PowerFlex system environment is to identify the ESXi host model. All PowerFlex hosts have a model name that starts with the "Dell Inc. PowerFlex-" prefix. The following hostnames are valid:

- Dell Inc. PowerFlex-R640

- Dell Inc. PowerFlex-R740xd

- Dell Inc. PowerFlex-R840

Look up the hostname of a PowerFlex system in the **Summary** tab of the vSphere UI. The following figure shows an example:
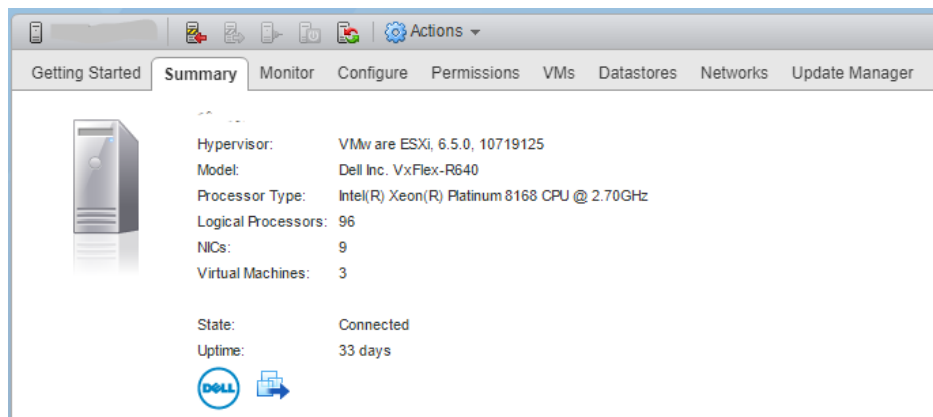

Figure 13.   vSphere UI Summary page: System models

Ensure that you report the physical host model name (BIOS name) from within the SAP HANA VMs.

The following VM option must be set for all SAP HANA VMs:

`SMBIOS.reflectHost=TRUE`

Set this option on the VM as follows:

1. Select **VM Hardware** > **Editing Settings** > **VM Options** > **Advanced** > **Edit Configuration**.

2. Enter `SMBIOS.reflectHost` in the **name** field and `TRUE` in the **value** field.

3. Click **Add**.

Health check with PowerFlex Manager

Before reporting a problem to SAP, check the physical health of the PowerFlex hardware. PowerFlex Manager provides management of IT operations for the PowerFlex system, increasing the system's efficiency by reducing time-consuming manual operations.

PowerFlex Manager includes the following features:

- A dashboard that provides system configuration details and communicates a health status for PowerFlex infrastructure elements and services

- Release Certification Matrix (RCM) compliance monitoring and remediation

- Hardware monitoring and alerting through Secure Remote Services to Dell Technologies Support

Using PowerFlex Manager, you can define and capture infrastructure requirements to support a specific business outcome. After the process is defined, it can be easily repeated through automation.

Health check with iDRAC

You can use the Integrated Dell Remote Access Controller (iDRAC) dashboard to verify the health of the system. The following figure shows an example of an iDRAC health check report:
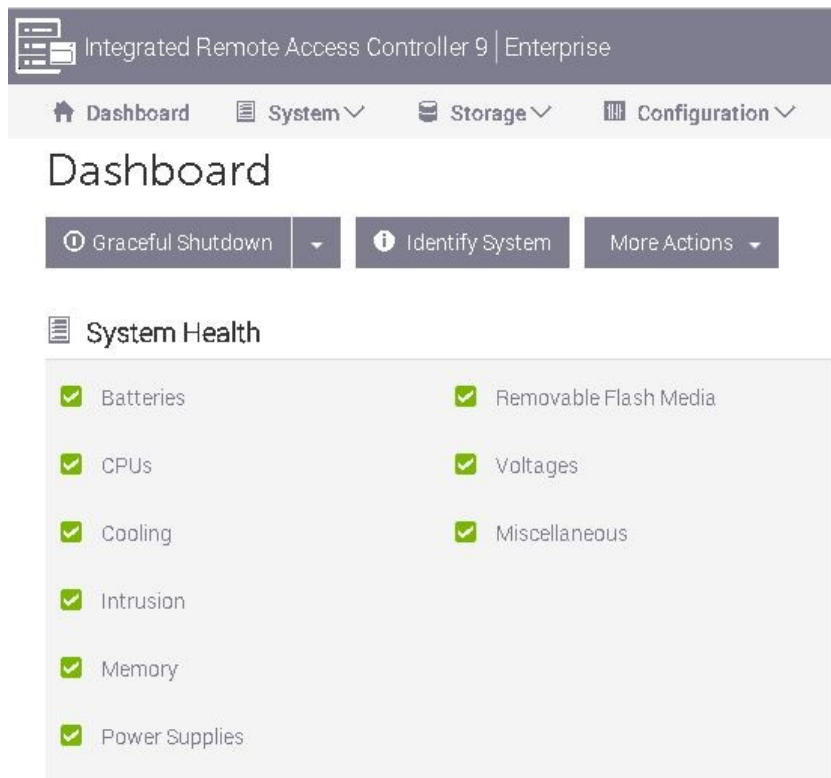

Figure 14.   iDRAC health check report

Ensure that all hardware-related issues are fixed before reporting an SAP HANA problem to SAP.

Dell Technologies support calls

Dell Technologies Online Support provides a single point of support for all PowerFlex hardware issues, PowerFlex related vSphere hypervisor issues, and PowerFlex software issues.

Generating a troubleshooting bundle

A troubleshooting bundle is a compressed file that contains logging information for components that are managed by PowerFlex Manager. The bundle includes PowerFlex Manager application logs, alert connector logs, PowerFlex gateway logs, and iDRAC lifecycle logs. Send the troubleshooting bundle to Dell Technologies Support for issue debugging.

Generate a troubleshooting bundle as follows:

1. In PowerFlex Manager, select **Settings** > **Virtual Appliance Management**.

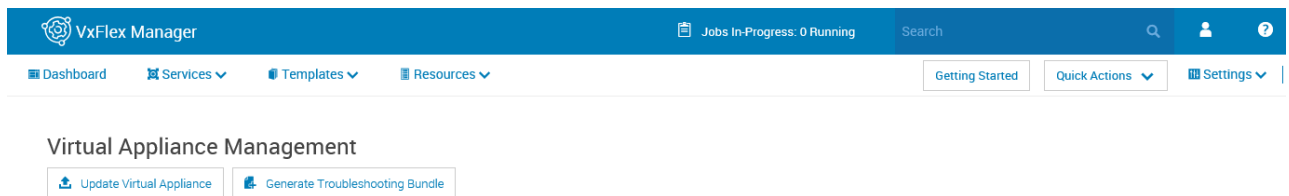   The Virtual Appliance Management page opens, as shown in the following figure:

Figure 15. Generating a troubleshooting bundle in PowerFlex Manager

2. Click **Generate Troubleshooting Bundle**.

    - If you are using Secure Remote Services, keep the default **Send to Configured Secure Remote Services** setting. You can configure Secure Remote Services in the Alert Connector section of the page.

    - If you are not using Secure Remote Services, select **Download Locally** to download the troubleshooting bundle to a local file.

References

Dell Technologies documentation

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

- Dell EMC PowerFlex Family Overview
- Dell EMC PowerFlex Architecture Overview
- Dell EMC PowerFlex Specification Sheet
- Dell EMC PowerFlex Networking Best Practices and Design Considerations
- Dell EMC PowerFlex product site
- Dell EMC PowerFlex support

Note: Deploying PowerFlex requires Dell Technologies implementation services. Detailed installation documentation is available only to Dell Technologies service personnel.

VMware documentation

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

- SAP Community Wiki: SAP HANA on VMware vSphere
- Virtualize Applications: SAP HANA on vSphere
- SAP and VMware: Key Offerings
- *Architecture Guidelines and Best Practices for Deployments of SAP HANA on VMware vSphere*
- *SAP Solutions on VMware Best Practices Guide*

SAP documentation

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

- Quick Sizer
- SAP HANA on VMware vSphere
- SAP Help Portal

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

- SAP Note 2399079—Elimination of hdbparam in HANA 2