

STORAGE CONFIGURATION BEST PRACTICES FOR SAP HANA TAILORED DATA CENTER INTEGRATION ON DELL EMC COMPELLENT STORAGE SYSTEMS

November 2017

Abstract

This solution guide describes best practices for SAP HANA in Tailored Data Center Integration (TDI) deployments on Dell EMC Compellent storage systems for use in a fully supported environment with existing data center infrastructures.

H16207.2

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 11/17 Solution Guide H16207.2.

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



Contents

Chapter 1	Executive Summary	5
	Business case.....	6
	Solution overview.....	6
	Document purpose.....	7
	Audience.....	7
	We value your feedback.....	7
Chapter 2	Compellent Fluid Data Architecture	8
	Overview.....	9
	Intelligence feature: Data Progression.....	9
	Flash-aware sub-LUN tiering.....	9
	Policy-based data management: Storage profiles.....	10
Chapter 3	SAP HANA TDI Solution	11
	Overview.....	12
	Implementing best practices.....	12
	Networking requirements and best practices.....	20
	Using /hana/shared volumes.....	21
Chapter 4	Compellent Storage Configuration	44
	Overview.....	45
	Configuring switch zoning.....	45
	Configuring QLogic HBA driver settings.....	45
	Setting up the storage array controller.....	45
Chapter 5	Accessing the Compellent SAN from SAP HANA Nodes	46
	Setting up multipath I/O.....	47
	Mounting SAP HANA data and log volumes.....	50
Chapter 6	Conclusion	52
	Summary.....	53
Chapter 7	References	54
	Dell EMC documentation.....	55
	SAP documentation.....	55
	Other documentation.....	55
Appendix A	Setting Up Cluster File Systems	56
	Configuring an SLES HA cluster for a shared OCFS2 volume.....	57

Configuring cman for GFS2 for shared volumes.....	58
Appendix B Setting Up an SC8000 Storage Array	62
Setting up the SC8000 controller.....	63
Logging in to Storage Center	64
Running the Startup Wizard	65
Configuring settings with storage array	73

Chapter 1 Executive Summary

This chapter presents the following topics:

Business case	6
Solution overview	6
Document purpose	7
Audience	7
We value your feedback	7

Business case

SAP HANA is an in-memory data processing platform with advanced capabilities. Initially, SAP HANA was provided in the form of fully integrated appliances with server, networking, and storage components that were delivered as one package. End-to-end testing and certification ensured that the SAP HANA appliances met the stringent criteria set by SAP, contributing significantly to the predictability and success of the platform. However, the appliance model has several limitations. In particular, the appliance model limits the customer's flexibility when choosing server, networking, and storage components, and does not enable the use of established data center operational best practices.

SAP HANA Tailored Data Center Integration (TDI) is an alternative to the existing SAP HANA appliance model. The SAP HANA TDI model allows customers to connect certified servers to shared storage resources. This model provides the following advantages:

- Enables reuse of existing hardware components
- Enables use of existing operational processes
- Leverages the existing ecosystem of server and storage vendors

As a result, customers can expect to realize reduced hardware and operational costs and increased flexibility in hardware vendor selection.

Solution overview

SAP HANA TDI deployments are customized solutions from which you can choose any of the certified SAP HANA server vendors along with any certified SAP enterprise storage.

Customers using SAP HANA TDI can:

- Choose their servers and storage
- Deploy well-understood data center operational best practices
- Tailor the storage capacity to business needs

This document describes configuration recommendations for deploying SAP HANA using Dell EMC™ PowerEdge servers and Dell EMC Compellent storage arrays. It describes the recommended storage design of Compellent SC5020, SC4020, SC7020, SC8000, and SC9000 storage arrays for use in an SAP HANA TDI deployment for up to 16 SAP HANA worker node solutions.

The configuration recommendations are based on extensive validation that was performed within Dell EMC labs. The configurations that are described in this document successfully meet the strict performance and availability requirements set by SAP.

Document purpose

This document provides information for deploying an SAP HANA solution using the TDI model. It provides an introduction to key technologies and to Dell EMC Fluid Data architecture, along with steps to integrate Compellent storage arrays with SAP HANA nodes.

Audience

This document is intended for system integrators, system or storage administrators, customers, partners, and anyone who must configure a Compellent storage array to be used 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 [Dell EMC Solution Feedback](#) with your comments.

Authors: Zafar Mahmood, Scott Saktanaset, Benjamin Simroth, Penelope Howe-Mailly

Chapter 2 Compellent Fluid Data Architecture

This chapter presents the following topics:

- Overview9**
- Intelligence feature: Data Progression9**
- Flash-aware sub-LUN tiering.....9**
- Policy-based data management: Storage profiles.....10**

Overview

Compellent is a radically simple, enterprise-grade storage platform. It is built on standards-based modular architecture that delivers next-generation storage virtualization. By virtualizing storage resources from their physical representation, the Compellent fluid data architecture can place the right data in the right place, at the right time, for the right cost.

Intelligence feature: Data Progression

Compellent Data Progression is a fully integrated feature of the virtualized storage platform. It intelligently classifies and migrates data to the best tier of storage by using real-time intelligence about how data is used throughout its lifecycle. This method allows businesses to retain frequently accessed data on high-performance drives and to move infrequently accessed data to lower-cost drives. By maintaining all data online and optimizing placement across multiple storage tiers, companies can purchase less storage overall and increase performance. Data progression further enhances value by extending the benefits of tiered storage to all-flash configurations. By using expensive Write Intensive drives for write operations and inexpensive Read Intensive drives for read operations, Compellent delivers the industry's first all-flash solution at the price of disk storage.

Flash-aware sub-LUN tiering

The storage virtualization and intelligence features work seamlessly across a wide range of applications and types of storage media. The foundation on which the fluid data architecture is built enables quick addition of new technologies. Whereas the traditional data progression algorithms run only once each day to optimize RAID levels and data tiering, flash-optimized data progression can move data across tiers throughout the day. For example, data that has recently been frozen in a replay (such as a snapshot) is now read-only and a good candidate to be moved to a read-intensive tier of storage. This process is handled in the background with little to no impact to the host. By moving the data to a RAID level and to storage media that has exceptional read I/O performance characteristics, the cost of storing the read-only data is significantly reduced. For example, data can be moved to space-efficient RAID 5 and inexpensive multi-level cell (MLC) solid state drives (SSDs).

Figure 1 shows the progression of data between tiers.

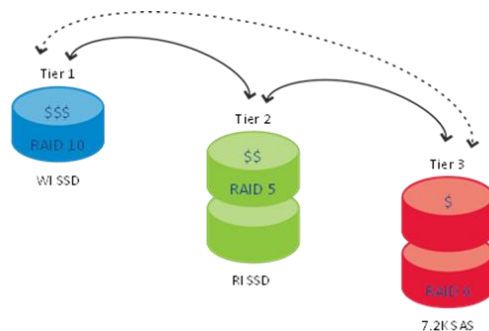


Figure 1. Compellent Data Progression

Policy-based data management: Storage profiles

Storage profiles are a tiering policy feature of the Compellent storage platform. When a storage profile is assigned to a volume, the volume inherits the tiering attributes that are defined in the storage profile. For example, a flash-optimized storage profile directs all write operations to high-performance Tier 1 storage (single-level cell (SLC) SSDs in RAID 10). On-demand data progression automatically converts replay and read-intensive data to RAID 5 and places it on Tier 2 read-intensive flash storage that has exceptional read performance characteristics. A cost-optimized storage profile enables the administrator to create a volume for applications such as backup and archive, as well as for low priority applications. The storage profiles feature radically simplifies tiering management by enabling storage administrators to assign flash storage to the most active data while retaining the flexibility of storing cold data in a more cost-effective tier.

Chapter 3 SAP HANA TDI Solution

This chapter presents the following topics:

Overview	12
Implementing best practices	12
Networking requirements and best practices	20
Using /hana/shared volumes	21

Overview

Table 1 provides a comparison between the appliance model and the SAP HANA TDI model.

Table 1. SAP HANA model comparison

Category	Appliance model	SAP HANA TDI model
Hardware selection	Little flexibility: <ul style="list-style-type: none"> The customer can choose between different appliance vendors. There is no option to replace certain hardware components already being used in the customer's data center. 	Preserves the IT budget and existing investments: <ul style="list-style-type: none"> Uses preferred storage Uses preferred network components
Implementation effort	Minimal customer effort due to preconfigured hardware with preinstalled software.	Only hardware is delivered: <ul style="list-style-type: none"> The customer installs the hardware, including the operating system (OS). Extensive documentation (guides, SAP notes, and so on) is available.
Solution validation	SAP and the hardware partner perform the solution validation together.	<ul style="list-style-type: none"> SAP AGS offers the SAP HANA Going-Live check. Self-managed infrastructure tests are available, such as SAP HANA Hardware Configuration Check Tool (HWCCT).
Support	SAP provides full support.	Individual support agreements with hardware partners are required.
OS service contract	The appliance vendor is the reseller of the OS service contract.	The customer obtains the OS service contract.

Note: Customers are responsible for deploying these TDI solutions with the help of SAP-certified consultants. Customers are also responsible for configuring system CPU, RAM, and disk so that the system meets the KPI requirements for implementing SAP HANA solutions without experiencing storage capacity issues.

Implementing best practices

Note: SAP certifies the minimum enterprise storage layout only, as described in this solution guide.

The Dell EMC SAP HANA TDI solution that is based on the Storage Center family can use the following options:

- A single Compellent SC5020/SC4020/SC7020/SC8000 storage array for up to eight active SAP HANA nodes with one standby, according to the disk sizing listed in [Table 4](#).
- A second Compellent SC5020/SC4020/SC7020/SC8000 storage array for up to 16 active SAP HANA Nodes with one or two standby nodes.

- A single Compellent SC9000 storage system can have up to 16 active SAP HANA nodes with one or two standby nodes.
- A storage system with a single disk folder, single storage profile for HANA, volumes, and disk storage.
- A building block approach to sizing Compellent storage arrays for SAP HANA nodes.

Note: A 16-node SAP HANA solution must have a minimum of one Compellent SC9000 storage array or two Compellent SC5020/SC4020/SC7020/SC8000 storage arrays.

Dell EMC PowerEdge Server BIOS settings

Use the following settings for the server BIOS:

- Boot settings—Select the UEFI Boot mode
- Memory operating mode—Optimizer mode

Recommended OS settings

See the following SAP notes for the OS settings:

- SUSE—[SAP NOTE 2205917: SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12](#)
- Red Hat Enterprise Linux (RHEL)—[SAP NOTE 2292690: SAP HANA DB: Recommended OS settings for RHEL 7](#)

Sizing for performance

Table 2 shows the sizing criteria for building SAP HANA solutions using Compellent storage arrays.

Table 2. Sizing criteria

Requirement	Description
CPU core to memory ratio	<ul style="list-style-type: none"> • 1 TB or 1.5 TB/88 CPU cores (8880/8890 E7 v4) with SAP HANA SPS 10 or earlier • 2 TB per 88/96 CPU cores (8880/8890 E7 v4) with SAP HANA SPS 11 or later • 3 TB per 112 CPU cores (with either 4 x Intel Xeon SKL Platinum 8180/M 2.5 G, 28C/56T, 10.4GT/s 3UPI, 38M cache or 4 x Intel Xeon SKL Platinum 8176/M 2.1G, 28C/56T, 10.4GT/s 3UPI, 38M cache)
Storage requirements (external shared storage)	<ul style="list-style-type: none"> • Log volume per SAP HANA node: 512 GB per 4 HANA worker nodes • Data volume per SAP HANA node: 3 x total memory • Shared volume: 1 x total memory per node or 1 x total memory per four nodes with SAP HANA SPS 11 or later

Requirement	Description
Network requirements	<ul style="list-style-type: none"> • For R930 systems: <ul style="list-style-type: none"> ▪ 1 x Intel X540 Dual Port Ethernet 10 Gb Base-T adapter for a backup network + I350 1 Gb Base-T Dual Port Network Daughter Card for SAP HANA Studio/Modeler for a management network ▪ 1 x Intel X540 Dual Port network adapter for a SAP HANA cluster network or a Network File System (NFS) network (bond2) ▪ 1 x Intel X540 Dual Port network adapter for a SAP HANA internal network (bond0) ▪ 1 x Intel X540 Dual Port network adapter for external sources ERP, BI Platform, SAP HANA, and data sources – PROD (bond1) ▪ Dedicated FC16 SAN between SAP HANA nodes and storage ▪ Redundant infrastructure for failover • For R940 systems: <ul style="list-style-type: none"> ▪ 1 x Intel Ethernet X710 DP 10 Gb DA/SFP+ for a backup network+ I350 1 Gb Base-T Dual Port Network Daughter Card for SAP HANA Studio/Modeler for management network ▪ 2 x Intel X710 Quad Port 10 Gb Direct Attach, SPF+, Converged Network Adapter for all bond networks: <ul style="list-style-type: none"> ○ bond0: Port-0—Internal HANA internal network ○ bond1: Port-1—external sources ERP, BI Platform, HANA and data sources ○ bond2: Port-2—HANA Shared NFS network ▪ Dedicated FC16 SAN between SAP HANA nodes and storage ▪ Redundant infrastructure for failover

Dell EMC SAP HANA scale-out solutions reference configurations

This section provides guidelines for building SAP HANA solutions using Compellent storage arrays.

Table 3 shows the specifications for 2 TB to 32 TB configurations.

Table 3. 2 TB to 48 TB configuration specifications

Component	Description
Server	Either: <ul style="list-style-type: none"> • 2–16 x PowerEdge R930 worker nodes + 1 or 2 x R930 nodes for standby • 2–16 x PowerEdge R940 worker nodes + 1 or 2 x R940 nodes for standby
CPU	Either: <ul style="list-style-type: none"> • On a R930 server—4 x Intel Xeon (BDW-EX) processors: E7-8880 v4 2.3GHz or E7-8890 v4 2.2GHz on R930 • On a R940 server—4 x Intel Xeon SKL Platinum 8180/M 2.5G, 28C/56T, 10.4GT/s 3UPI, 38M cache • On a R940 server—4 x Intel Xeon SKL Platinum 8176/M 2.1G, 28C/56T, 10.4GT/s 3UPI, 38M cache
Memory	<ul style="list-style-type: none"> • On a R930 server—1 TB, 1.5 TB, or 2 TB per node • On a R940 server—3 TB per node
OS	Either: <ul style="list-style-type: none"> • SUSE Linux Enterprise Server (SLES) 12 SP2 for SAP • Red Hat Enterprise Linux (RHEL) 7.3
Storage	1–2 x Compellent storage arrays with SSD and hard disk drives (HDDs) with a configuration option shown in Table 4 .

For a list of certified operating system configurations and storage, refer to [Find Certified Appliances](#) and [Find Certified Enterprise Storage](#).

Table 4 lists the valid array configurations for the SAP HANA TDI solution.

Table 4. Array configurations

Active Nodes	Compellent SC5020	Compellent SC4020	Compellent SC8000	Compellent SC7020/SC7020F	Compellent SC9000
2	6 x 960 GB RI SSD + 18 x 1.8 TB 10K RPM SAS	6 x 480 GB RI SSD + 12 x 1.8 TB 10K RPM SAS	7 x 480 GB RI SSD + 14 x 1.8 TB 10K RPM SAS	For SC7020: 6 x 480 GB RI SSD + 12 x 1.8 TB 10K RPM SAS For SC7020F: 12 x 1.9 GB RI SSD (for SC7020F)	6 x 480 GB RI SSD + 12 x 1.8 TB 10K RPM SAS
3	6 x 960 GB RI SSD + 30 x 1.8 TB 10K RPM SAS	6 x 480 GB RI SSD + 24 x 1.8 TB 10K RPM SAS	7 x 480 GB SSD + 14 x 1.8 TB 10K RPM SAS	For SC7020: 6 x 480 GB RI SSD + 18 x 1.8 TB 10K RPM SAS For SC7020F: 18 x 1.9 GB RI SSD	6 x 480 GB RI SSD + 24 x 1.8 TB 10K RPM SAS

Active Nodes	Compellent SC5020	Compellent SC4020	Compellent SC8000	Compellent SC7020/SC7020F	Compellent SC9000
4	6 x 960 GB RI SSD + 30 x 1.8 TB 10K RPM SAS	6 x 480 GB RI SSD + 24 x 1.8 TB 10K RPM SAS	12 x 480 GB RI SSD + 24 x 1.8 TB 10K RPM SAS	For SC7020: 6 x 480 GB RI SSD + 24 x 1.8 TB 10K RPM SAS For SC7020F: 24 x 1.9 GB RI	6 x 480 GB RI SSD + 24 x 1.8 TB 10K RPM SAS
5	12 x 960 GB RI SSD + 40 x 1.8 TB 10K RPM SAS	12 x 480 GB RI SSD + 30 x 1.8TB 10K RPM SAS	12 x 480 GB RI SSD + 24 x 1.8 TB 10K RPM SAS	For SC7020: 12 x 480 GB RI SSD + 30 x 1.8 TB 10K RPM SAS For SC7020F: 30 x 1.9 GB RI SSD	12 x 480 GB RI SSD + 30 x 1.8 TB 10K RPM SAS
6	12 x 960 GB RI SSD + 40 x 1.8 TB 10K RPM SAS	12 x 480 GB RI SSD + 30 x 1.8 TB 10K RPM SAS	18 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS	For SC7020: 12 x 480 GB RI SSD + 30 x 1.8 TB 10K RPM SAS For SC7020F: 36 x 1.9 GB RI SSD	12 x 480 GB RI SSD + 30 x 1.8 TB 10K RPM SAS
7	12 x 960 GB RI SSD + 54 x 1.8 TB 10K RPM SAS	12 x 480GB RI SSD + 30 x 1.8 TB 10K RPM SAS	18 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS	For SC7020: 12 x 480 GB RI SSD + 30 x 1.8 TB 10K RPM SAS For SC7020F: 36 x 1.9 GB RI SSD	12 x 480 GB RI SSD + 30 x 1.8 TB 10K RPM SAS
8	12 x 960 GB RI SSD + 54 x 1.8 TB 10K RPM SAS	12 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS	24 x 480 GB RI SSD + 48x 1.8 TB 10K RPM SAS	For SC7020: 12 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS For SC7020F: 42 x 1.9 GB RI SSD	12 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS
9	First array: 12 x 960 GB RI SSD + 54 x 1.8 TB 10K RPM SAS Second array: 6 x 960 GB RI SSD + 18 x 1.8 TB 10K RPM SAS	First array: 12 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS Second array: 6 x 480 GB RI SSD + 12 x 1.8 TB 10K RPM SAS	First array: 24 x 480 GB RI SSD + 48 x 1.8 TB 10K RPM SAS Second array: 7 x 480 GB RI SSD + 14 x 1.8TB 10K RPM SAS	First SC7020 array: 12 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS Second SC7020 array: 6 x 480 GB RI SSD + 12 x 1.8 TB 10K RPM SAS First SC7020F array: 42 x 1.9 GB RI SSD (for SC7020F) Second SC7020F array: 12 x 1.9 GB RI SSD	18 x 480 GB RI SSD + 48 x 1.8 TB 10K RPM SAS

Active Nodes	Compellent SC5020	Compellent SC4020	Compellent SC8000	Compellent SC7020/SC7020F	Compellent SC9000
10	<p>First array: 12 x 960 GB RI SSD + 54 x 1.8 TB 10K RPM SAS</p> <p>Second array: 6 x 960 GB RI SSD + 18 x 1.8 TB 10K RPM SAS</p>	<p>First array: 12 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS</p> <p>Second array: 6 x 480 GB RI SSD + 12 x 1.8 TB 10K RPM SAS</p>	<p>First array: 24 x 480 GB RI SSD + 48 x 1.8 TB 10K RPM SAS</p> <p>Second array: 7 x 480 GB RI SSD + 14 x 1.8 TB 10K RPM SAS</p>	<p>First SC7020 array: 12 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS</p> <p>Second SC7020 array: 6 x 480 GB RI SSD + 12 x 1.8 TB 10K RPM SAS</p> <p>First SC7020F array: 42 x 1.9 GB RI SSD (for SC7020F)</p> <p>Second SC7020F array: 12 x 1.9 GB RI SSD (for SC7020F)</p>	18 x 480 GB RI SSD + 48 x 1.8 TB 10K RPM SAS
11	<p>First array: 12 x 960 GB RI SSD + 54 x 1.8 TB 10K RPM SAS</p> <p>Second array: 6 x 960 GB RI SSD + 30 x 1.8 TB 10K RPM SAS</p>	<p>First array: 12 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS</p> <p>Second array: 6 x 480 GB RI SSD + 12 x 1.8 TB 10K RPM SAS</p>	<p>First array: 24 x 480 GB RI SSD + 48 x 1.8 TB 10K RPM SAS</p> <p>Second array: 12 x 480 GB RI SSD + 24 x 1.8 TB 10K RPM SAS</p>	<p>First SC7020 array: 12 x 480 GB RI SSD + 36 x 1.8TB 10K RPM SAS</p> <p>Second SC7020 array: 6 x 480 GB RI SSD + 12 x 1.8TB 10K RPM SAS</p> <p>First SC7020F array: 42 x 1.9 GB RI SSD (for SC7020F)</p> <p>Second SC7020F array: 18 x 1.9 GB RI SSD (for SC7020F)</p>	18 x 480 GB RI SSD + 48 x 1.8 TB 10K RPM SAS
12	<p>First array: 12 x 960 GB RI SSD + 54 x 1.8 TB 10K RPM SAS</p> <p>Second array: 6 x 960 GB RI SSD + 30 x 1.8 TB 10K RPM SAS</p>	<p>First array: 12 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS</p> <p>Second array: 6 x 480 GB RI SSD + 24 x 1.8 TB 10K RPM SAS</p>	<p>First array: 24 x 480 GB RI SSD + 48 x 1.8TB 10K RPM SAS</p> <p>Second array: 12 x 480 GB RI SSD + 24 x 1.8 TB 10K RPM SAS</p>	<p>First SC7020 array: 12 x 480 GB RI SSD + 36 x 1.8TB 10K RPM SAS</p> <p>Second SC7020 array: 6 x 480 GB RI SSD + 24 x 1.8TB 10K RPM SAS</p> <p>First SC7020F array: 42 x 1.9 GB RI SSD (for SC7020F)</p> <p>Second SC7020F array: 24 x 1.9 GB RI SSD (for SC7020F)</p>	18 x 480 GB RI SSD + 60 x 1.8 TB 10K RPM SAS

Active Nodes	Compellent SC5020	Compellent SC4020	Compellent SC8000	Compellent SC7020/SC7020F	Compellent SC9000
13	<p>First array: 12 x 960 GB RI SSD + 54 x 1.8 TB 10K RPM SAS</p> <p>Second array: 12 x 960 GB RI SSD + 40 x 1.8 TB 10K RPM SAS</p>	<p>First array: 12 x 480 GB SSD + 36 x 1.8 TB 10K RPM SAS</p> <p>Second array: 12 x 480 GB RI SSD + 30 x 1.8 TB 10K RPM SAS</p>	<p>First array: 24 x 480 GB RI SSD + 48 x 1.8 TB 10K RPM SAS</p> <p>Second array: 18 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS</p>	<p>First SC7020 array: 12 x 480 GB RI SSD + 36 x 1.8TB 10K RPM SAS</p> <p>Second SC7020 array: 12 x 480 GB RI SSD + 30 x 1.8TB 10K RPM SAS</p> <p>First SC7020F array: 42 x 1.9 GB RI SSD (for SC7020F)</p> <p>Second SC7020F array: 30 x 1.9 GB RI SSD (for SC7020F)</p>	24 x 480 GB RI SSD + 66 x 1.8 TB 10K RPM SAS
14	<p>First array: 12 x 960 GB RI SSD + 54 x 1.8 TB 10K RPM SAS</p> <p>Second array: 12 x 960 GB RI SSD + 40 x 1.8 TB 10K RPM SAS</p>	<p>First array: 12 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS</p> <p>Second array: 12 x 480GB RI SSD + 30 x 1.8 TB 10K RPM SAS</p>	<p>First array: 24 x 480 GB RI SSD + 48 x 1.8 TB 10K RPM SAS</p> <p>Second array: 18 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS</p>	<p>First SC7020 array: 12 x 480 GB RI SSD + 36 x 1.8TB 10K RPM SAS</p> <p>Second SC7020 array: 12 x 480 GB RI SSD + 30 x 1.8TB 10K RPM SAS</p> <p>First SC7020F array: 42 x 1.9 GB RI SSD (for SC7020F)</p> <p>Second SC7020F array: 36 x 1.9 GB RI SSD (for SC7020F)</p>	24 x 480 GB RI SSD + 66 x 1.8 TB 10K RPM SAS
15	<p>First array: 12 x 960 GB RI SSD + 54 x 1.8 TB 10K RPM SAS</p> <p>Second array: 12 x 960 GB RI SSD + 54 x 1.8 TB 10K RPM SAS</p>	<p>First array: 12 x 480 GB RI SSD + 36 x 1.8TB 10K RPM SAS</p> <p>Second array: 12 x 480 GB RI SSD + 30 x 1.8 TB 10K RPM SAS</p>	<p>First array: 24 x 480 GB RI SSD + 48 x 1.8 TB 10K RPM SAS</p> <p>Second array: 24 x 480 GB RI SSD + 48 x 1.8TB 10K RPM SAS</p>	<p>First SC7020 array: 12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS</p> <p>Second SC7020 array: 12 x 480 GB RI SSD + 30 x 1.8TB 10K RPM SAS</p> <p>First SC7020F array: 42 x 1.9 GB RI SSD (for SC7020F)</p> <p>Second SC7020F array: 36 x 1.9 GB RI SSD (for SC7020F)</p>	24 x 480GB RI SSD + 66 x 1.8 TB 10K RPM SAS

Active Nodes	Compellent SC5020	Compellent SC4020	Compellent SC8000	Compellent SC7020/SC7020F	Compellent SC9000
16	First array: 12 x 960 GB RI SSD + 54 x 1.8 TB 10K RPM SAS Second array: 12 x 960 GB RI SSD + 54 x 1.8 TB 10K RPM SAS	First array: 12 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS Second array: 12 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS	First array: 24 x 480 GB RI SSD + 48 x 1.8 TB 10K RPM SAS Second array: 24 x 480 GB RI SSD + 48 x 1.8 TB 10K RPM SAS	First SC7020 array: 12 x 480 GB RI SSD + 36 x 1.8 TB 10K RPM SAS Second SC7020 array: 12 x 480 GB RI SSD + 36 x 1.8TB 10K RPM SAS First SC7020F array: 42 x 1.9 GB RI SSD (for SC7020F) Second SC7020F array: 42 x 1.9 GB RI SSD (for SC7020F)	24 x 480 GB RI SSD + 72 x 1.8 TB 10K RPM SAS

Table 5 lists the recommended number of storage arrays for dual controller configurations.

Table 5. Recommended number of storage arrays

# Nodes	SC5020	SC4020	SC8000	SC7020	SC9000
2–8	1	1	1	1	1
9–16	2	2	2	2	1

Table 6 lists the recommended number of expansion enclosure disk shelves. One SC220 or SC420 disk shelf can hold up to 24 x 2.5" drives.

Table 6. Recommended number of SC220/SC420 disk shelves

# Nodes	SC5020	SC4020	SC8000	SC7020	SC9000
2	0 x SC420	1 x SC220	3 x SC220	0x SC420	3 x SC420
3	1 x SC420	1 x SC220	3 x SC220	0x SC420	3 x SC420
4	1 x SC420	1 x SC220	3 x SC220	0 x SC420	3 x SC420
5	1 x SC420	1 x SC220	3 x SC220	1 x SC420	3 x SC420
6	1 x SC420	1 x SC220	3 x SC220	1 x SC420	3 x SC420
7-8	2 x SC420	1 x SC220	3 x SC220	1 x SC420	3 x SC420
9–10	2 x SC420	2 x SC220	6 x SC220	1 x SC420	6 x SC420
11-13	3 x SC420	2 x SC220	6 x SC220	1 x SC420	6 x SC420
14	3 x SC420	2 x SC220	6 x SC220	2 x SC420	6 x SC420
15-16	4 x SC420	2 x SC220	6 x SC220	2 x SC420	6 x SC420

Networking requirements and best practices

To provide continuous availability, Dell EMC recommends redundant data paths for an SAP HANA scale-out appliance for both Ethernet networks and the SAN fabric.

Ethernet Network

When setting up Ethernet networks, note the following guidelines:

- The management/administrative network is reserved for external administration tools. A 1 GbE network connection is sufficient.
- For internal zone networks, 10 Gb/s non-blocking switches are required as well as one fully redundant 10 GbE network.
- A 10 GbE network connection is required for the client zone (bond1) network.
- Because SAP HANA system replication transfers large amounts of I/O traffic across two distanced data centers, an additional 10 GbE network connection is also required for the SAP HANA system replication network. The same principles apply to the SAP HANA internode network as to the network that SAP HANA system replication uses.
- Dell EMC recommends that you use jumbo frame settings (MTU=9000) for all 10 GbE network connections in an SAP HANA scale-out appliance.

SAN fabric

When using Fibre Channel for front-end connectivity, establish zones to ensure that the servers detect the storage. Refer to WWN zoning guidelines to plan the front-end connectivity before you start to cable the storage system.

You can apply zoning to either the ports on switches or to the World Wide Names (WWNs) of the end devices.

Use WWN zoning to create zones by using a single initiator host port and multiple Storage Center ports.

WWN zoning guidelines

In a WWN zoning configuration, a device can reside on any port or change physical ports and still be visible because the switch is seeking a WWN.

Table 7 lists guidelines for configuring WWN zones.

Table 7. WWN zone guidelines

Connection type	Guidelines
Virtual port mode (Recommended)	<ul style="list-style-type: none"> • Include all Storage Center virtual WWNs in a single zone. • Include all Storage Center physical WWNs in a single zone. • For each host server host bus adapter (HBA) port, create a zone that includes the single HBA WWN and multiple Storage Center virtual WWNs on the same switch.
Legacy mode	<ul style="list-style-type: none"> • Include all Storage Center front-end WWNs or ports in a single zone. • For each host server HBA port, create a zone that includes the single HBA WWN and multiple Storage Center virtual WWNs on the same switch.

For Fibre Channel replication, follow these guidelines:

- Include all Storage Center physical WWNs from Storage Center system A and Storage Center system B in a single zone.
- Include all Storage Center physical WWNs of Storage Center system A and the virtual WWNs of Storage Center system B on the particular fabric.
- Include all Storage Center physical WWNs of Storage Center system B and the virtual WWNs of Storage Center system A on the particular fabric.

Note: Some ports may not be used or dedicated for replication; however, ports that are used must be in these zones.

Using /hana/shared volumes

SAP HANA nodes require a `/hana/shared` volume across nodes, by using either of the following:

- An NFS share that is mounted on `/hana/shared` across all SAP HANA nodes
- A cluster file system such as the Oracle Cluster File System 2 (OCFS2) on SLES or Global File System 2 (GFS2) on RHEL

Using NFS for a /hana/shared volume

Using NFS on a `/hana/shared` volume requires the Dell FS8610 network attached storage (NAS) appliance to export the NFS share from a volume that is created on a Compellent storage array.

Building blocks for NFS option for /hana/shared

Dell EMC SAP HANA solutions are based on scalable building blocks. Table 8 lists the building blocks for NFS-based `/hana/shared`.

Table 8. Building blocks for scale-out with NFS-based /hana/shared

Module		
1	Infrastructure rack	<ul style="list-style-type: none"> • Network switches: <ul style="list-style-type: none"> ▪ 2 x Force10 S4048T switches 48 x 10 GbE ports for an internal SAP HANA network + an external SAP HANA network, SAP HANA-shared NFS network using VLANs ▪ 1 x Force10 S3048 switch, 48 x 1 GbE ports for a management network • SAN Fabric: <ul style="list-style-type: none"> ▪ 2 x Brocade 6510 24-36-48 Port FC8/16 switch • Management: <ul style="list-style-type: none"> ▪ 1 x PowerEdge R630 or R640 server as a management node + 1 x KMM console

Module		
2	Servers	<ul style="list-style-type: none"> • For R930 servers: <ul style="list-style-type: none"> ▪ 72 cores/1.5 TB or 2 TB with SAP HANA SPS 11 or later ▪ Intel X540 Dual Port Ethernet 10 Gb Base-T + I350 1Gb Base-T Dual Port Network Daughter Card ▪ 3 x Intel 10 GbE dual port RJ45 (X540-T2) connector ▪ 2 x Qlogic 2660 Single Port 16 Gb Optical Fibre Channel HBA • For R940 servers: <ul style="list-style-type: none"> ▪ 112 CPU cores/3 TB (either 4 x Intel Xeon SKL Platinum 8180/M 2.5G, 28C/56T,10.4GT/s 3UPI, 38M cache or 4x Intel Xeon SKL Platinum 8176/M 2.1G, 28C/56T,10.4GT/s 3UPI, 38M cache) <ul style="list-style-type: none"> ○ 1 x Intel Ethernet X710 DP 10 Gb DA/SFP+ for a backup network+ I350 1 Gb BT DP Network Daughter Card for SAP HANA Studio/Modeler for management network ○ 2 x Intel X710 Quad Port 10 Gb Direct Attach, SPF+, Converged Network Adapter for all bond networks ▪ 2 x Qlogic 2690 Single Port 16 Gb Optical Fibre Channel HBA
3	Compellent array	Any storage option from Table 4 .

Building block details

Table 9 lists building block details for NFS-based /hana/shared for R930 servers and SC4020 storage arrays.

Table 9. R930-based 3 TB or 4 TB scale-out solution using (SAP HANA SPS 11 or later) NFS /hana/shared

Component	Configuration
SAP HANA Nodes	2 x PowerEdge R930 active SAP HANA nodes + 1 or 2 x R930 standby SAP HANA nodes, each with: <ul style="list-style-type: none"> • SLES 12 SP2 for SAP or RHEL 7.3 for SAP • 1.5 TB RAM or 2 TB with SAP HANA SPS 11 or later • 4 x Intel Xeon 18 core E7-8880 v2 2.5GHz processors with 37.5 M L2 cache or E7-8890 v2 2.8GHz with 37.5M L2 cache
Management	1 x PowerEdge R630 server with: <ul style="list-style-type: none"> • SLES 12 SP2 for SAP or RHEL 7.3 for SAP • 1 x Force10 S3048 switch for an infrastructure management network • SAP HANA Studio
Storage	<ul style="list-style-type: none"> • FS8610 NAS appliance for /hana/shared refer to Dell EMC FS8610 NAS appliance for shared SAP HANA volumes) • A valid array configuration as listed in Table 4.
Storage Fabric	2 x Brocade 6510 24-36-48 Port FC16 switch

Component	Configuration
Network	<ul style="list-style-type: none"> 2 x Force 10 S4048T switches 48 x 10 GbE ports for an internal SAP HANA network + an external SAP HANA network (PROD) + NFS network using VLANs for up to 8 SAP HANA active nodes. An additional pair of 4048T switches is required for an additional SAP HANA active node. 1 x Force 10 S3048 switch, 48 x 1 GbE ports for a management network

Figure 2 shows a 3 TB or 4 TB configuration on an SC4020 array. Table 4 lists all valid storage options.




Figure 2. 3TB or 4TB scale-out solution using NFS for /hana/shared on an SC4020 array

Table 10 lists building block details for R930 servers and SC4020 storage arrays.

Table 10. R930-based 24 TB or 32 TB scale-out solution using (SAP HANA SPS 11 or later) NFS for /hana/shared

Component	Configuration
SAP HANA nodes	16 x PowerEdge R930 active SAP HANA nodes + 1 or 2 x R930 standby SAP HANA nodes, each with: <ul style="list-style-type: none"> SLES 12 SP2 for SAP or RHEL 7.3 for SAP 1.5 TB RAM or 2 TB with SAP HANA SPS 11 or later 4 x Intel Xeon 18 core E7-8880 v2 2.5GHz processors with 37.5M L2 cache or E7-8890 v2 2.8GHz with 37.5M L2 cache
Management	1 x PowerEdge R630 server with: <ul style="list-style-type: none"> SLES 12 SP2 for SAP or RHEL 7.3 for SAP 1 x Force10 S3048 switch for an infrastructure management network SAP HANA Studio

Component	Configuration
Storage	<ul style="list-style-type: none"> FS8610 NAS appliance for /hana/shared (refer to Dell EMC FS8610 NAS appliance for shared SAP HANA volumes) A valid array configuration as listed in Table 4.
Storage fabric	2 x Brocade 6510 24-36-48 Port FC16 switch
Network	<ul style="list-style-type: none"> 2 x Force 10 S4048T switches 48 x 10 GbE ports for an internal SAP HANA network + an external SAP HANA network (PROD) using VLANs 1 x Force 10 S3048 switch, 48 x 1GbE ports for a management network

Figure 3 shows a 24 TB or 32 TB configuration using NFS for /hana/shared on an SC4020 array. The callout numbers show the recommended rack placement and labeling for the 16 active and two standby nodes, as delivered from Dell. [Table 4](#) lists all valid storage options.




Figure 3. 24 TB or 32 TB scale-out solution using (SAP HANA SPS 11 or later) NFS for /hana/shared

Table 11 lists building block details for NFS-based /hana/shared for R940 servers and SC5020 storage arrays.

Table 11. R940-based 6 TB scale-out solution using (SAP HANA SPS 11 or later) NFS /hana/shared

Component	Configuration
SAP HANA nodes	2 x PowerEdge R940 active SAP HANA nodes + 1 or 2 x R940 standby SAP HANA nodes, each with: <ul style="list-style-type: none"> • SLES 12 SP2 for SAP • 3 TB RAM • Either 4 x Intel Xeon SKL Platinum 8180/M 2.5G, 28C/56T,10.4GT/s 3UPI, 38M cache or 4 x Intel Xeon SKL Platinum 8176/M 2.1G, 28C/56T,10.4GT/s 3UPI, 38M cache
Management	1 x PowerEdge R640 server with: <ul style="list-style-type: none"> • SLES 12 SP2 for SAP • 1 x Force10 S3048 switch for an infrastructure management network • SAP HANA Studio
Storage	<ul style="list-style-type: none"> • FS8610 NAS appliance for /hana/shared refer to (refer to Dell EMC FS8610 NAS appliance for shared SAP HANA volumes) • A valid array configuration as listed in Table 4.
Storage fabric	2 x Brocade 6510 24-36-48 Port FC16 switch
Network	<ul style="list-style-type: none"> • 2 x Force 10 S4048T switches 48 x 10 GbE ports for an internal SAP HANA network + an external SAP HANA network (PROD) + NFS network using VLANs for up to 8 SAP HANA active nodes. An additional pair of 4048T switches is required for an additional SAP HANA active node. • 1 x Force 10 S3048 switch, 48 x 1 GbE ports for a management network.

Figure 4 shows a 6 TB configuration on an SC5020 array. [Table 4](#) lists all valid storage options.




Figure 4. 6 TB scale-out solution using NFS for /hana/shared on an SC4020 array

Table 12 lists building block details for R940 servers and SC5020 storage arrays.

Table 12. R940-based 48TB scale-out solution using (SAP HANA SPS 11 or later) NFS for /hana/shared

Component	Configuration
SAP HANA nodes	16 x PowerEdge R940 active SAP HANA nodes + 1 or 2 x R940 standby SAP HANA nodes, each with: <ul style="list-style-type: none"> • SLES 12 SP2 for SAP • 3 TB RAM • Either 4 x Intel Xeon SKL Platinum 8180/M 2.5G, 28C/56T,10.4GT/s 3UPI,38M cache or with 4 x Intel Xeon SKL Platinum 8176/M 2.1G, 28C/56T,10.4GT/s 3UPI, 38M cache
Management	1 x PowerEdge R640 server with: <ul style="list-style-type: none"> • SLES 12 SP2 for SAP • 1 x Force10 S3048 switch for an infrastructure management network • SAP HANA Studio
Storage	<ul style="list-style-type: none"> • FS8610 NAS appliance for /hana/shared refer to (refer to Dell EMC FS8610 NAS appliance for shared SAP HANA volumes) • A valid array configuration as listed in Table 4.
Storage fabric	2 x Brocade 6510 24-36-48 Port FC16 switch
Network	<ul style="list-style-type: none"> • 2 x Force 10 S4048T switches 48 x 10 GbE ports for an internal SAP HANA network + an external SAP HANA network (PROD) using VLANs • 1 x Force 10 S3048 switch, 48 x 1GbE ports for a management network

Figure 5 shows a 48 TB configuration using NFS for /hana/shared on an SC5020 array. The callout numbers show the recommended rack placement and labeling for the 16 active and two standby nodes, as delivered from Dell EMC. [Table 4](#) lists all valid storage options.




Figure 5. 48 TB scale-out solution using (SAP HANA SPS 11 or later) NFS for /hana/shared

Dell EMC FS8610 NAS appliance for shared SAP HANA volumes

The FS8610 NAS appliance specification includes:

- 2U chassis NAS appliance with dual redundant controllers
- 10 Gb Ethernet client connectivity
- 8 Gb Fibre Channel back-end connectivity to the SC5020/SC4020 array

The following describes NFS-based /hana/shared for scale-out:

- The Storage Center provides the FS8610 scale-out NAS storage capacity, which is configured to be equal to the size of the scale-out appliance.
- The SAN network connects the FS8610 NAS appliance to the Storage Center and carries the block level traffic.
- You can add another FS8610 NAS appliance and an SC5020/SC4020 array to provide NFS shares for a backup/restore solution.
- The configuration can scale up to four FS8610 NAS appliances that are configured as a FluidFS cluster.
- The FS8610 NAS appliance includes a command-line interface (CLI) and optional EM graphical user interface (GUI) for management of NAS pools and exports.

Figure 6 provides an overview of the FS8610 scale-out NAS architecture.

The following figure shows an overview of the FS8600 scale-out NAS architecture.




Figure 6. FS8610 scale-out NAS architecture

The recommended mount options for `/hana/shared` in SAP HANA SPS 09, SAP HANA SPS 10, and SAP HANA SPS 11 are listed in [SAP Note 2099253: Creating a Network File System \(NFS\) for an SAP HANA Installation](#) (access requires an SPA username and password).

For information about how to set up FS8610 NAS for shared SAP HANA volumes, refer to the following guides:

- [Dell FluidFS Version 6.0 FS8600 Appliance Deployment Guide](#)
- [Dell FluidFS Version 6.0 FS8600 Appliance PreDeployment Requirements](#)
- [Dell FluidFS 6.0 FS8600 Appliance CLI Reference Guide](#)

Network configuration for scale-out solution using NFS for `/hana/shared`

An SAP HANA scale-out appliance can support up to 16 active and two standby SAP HANA nodes. The scale-out appliance includes the following networks:

- **Management network**—A 1 G Ethernet network using a Dell EMC Force 10 S3048 switch. A PowerEdge R640/R630 server, which is a dedicated management node, uses this network to manage all hardware components.
- An Integrated Dell EMC Remote Access Controlled (iDRAC) network is also part of the management network. It uses the onboard I350 1 Gb Base-T Dual Port Network Daughter Card.
- **SAP HANA-related networks**—One pair of 10G Force 10 S4048T switches in a Virtual Link Trunking (VLT) domain to support redundant data paths for different SAP HANA-related networks. Each S4048T switch has 48 x 10 G Base-T Ports + 4 x 40 Gb QSFP+ ports.

- **VLAN-10**—Dedicated SAP HANA shared NFS network (bond2 network). For a configuration larger than eight + two SAP HANA nodes, an additional pair of S4048T switches is required.
- **VLAN-20**—Dedicated SAP HANA Internal network (bond0 network)
- **VLAN-30**— SAP HANA prod/external network (bond1 network)

Other available ports can be used for other networks such as a backup or a required network to create a dedicated system replication network using the Intel X540 Dual Port 10 GbE network adapter on the daughter card.

- **Storage Area Network**— A dedicated and redundant FC16 SAN fabric to host all database-related volumes, using one pair of FC16 Brocade 6510 switches.

Figure 7 shows an NFS network configuration.

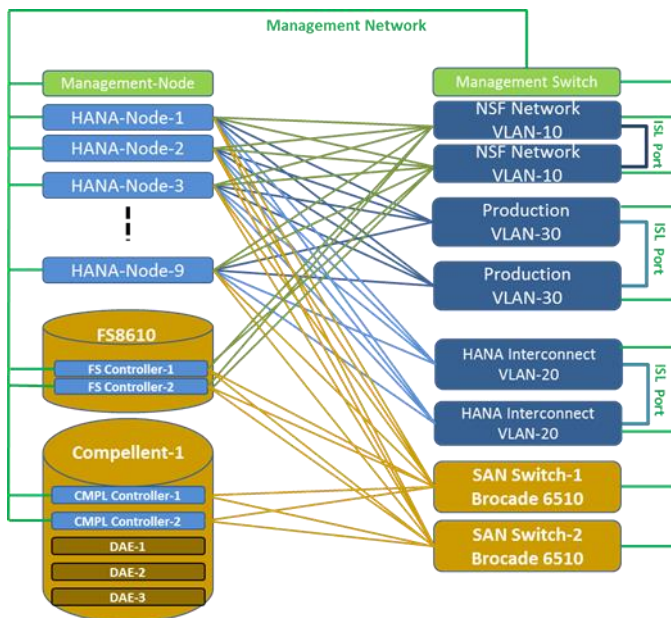


Figure 7. Network configuration using NFS for /hana/shared

Force10 Virtual Link Trunking

VLT is a dual active control plane implementation of a multisystem link aggregation group (LAG). It is a loop-free topology with active-active load sharing of links from access to core.

The Force 10 VLT feature provides Layer 2 multipathing between access network devices, such as switches or servers, and the core network, as shown in Figure 8.

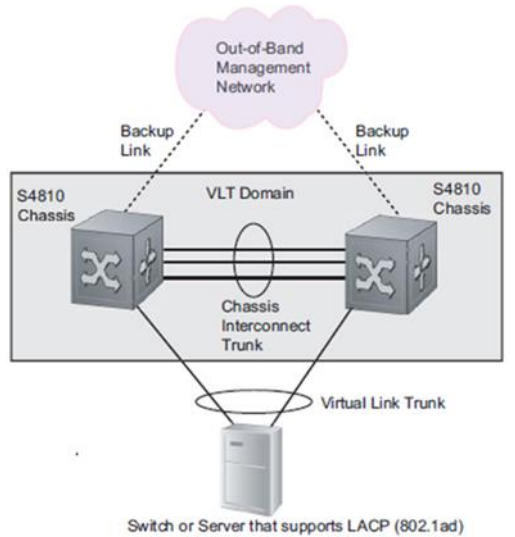


Figure 8. VLT configuration

The VLT domain configuration includes:

- **VLT interconnect (VLTi)**—Peer-link static port channel 128 (40 Gb port 56–port 60) for control and data exchange
- **VLT health check**—Out-of-band management port that provides a periodic keep-alive message

SAP HANA-related networks

One pair of Force 10 S4048T switches is configured in a VLT domain to host SAP HANA-related networks using different VLAN-IDs. The switch can support up to eight active and two standby SAP HANA nodes.

Supporting an SAP HANA scale-out appliance with more than eight active SAP HANA nodes requires an additional pair of Force 10 S4048T switches.

Figure 9 shows the network port assignment for an SAP HANA scale-out configuration with up to eight active and two standby SAP HANA nodes.

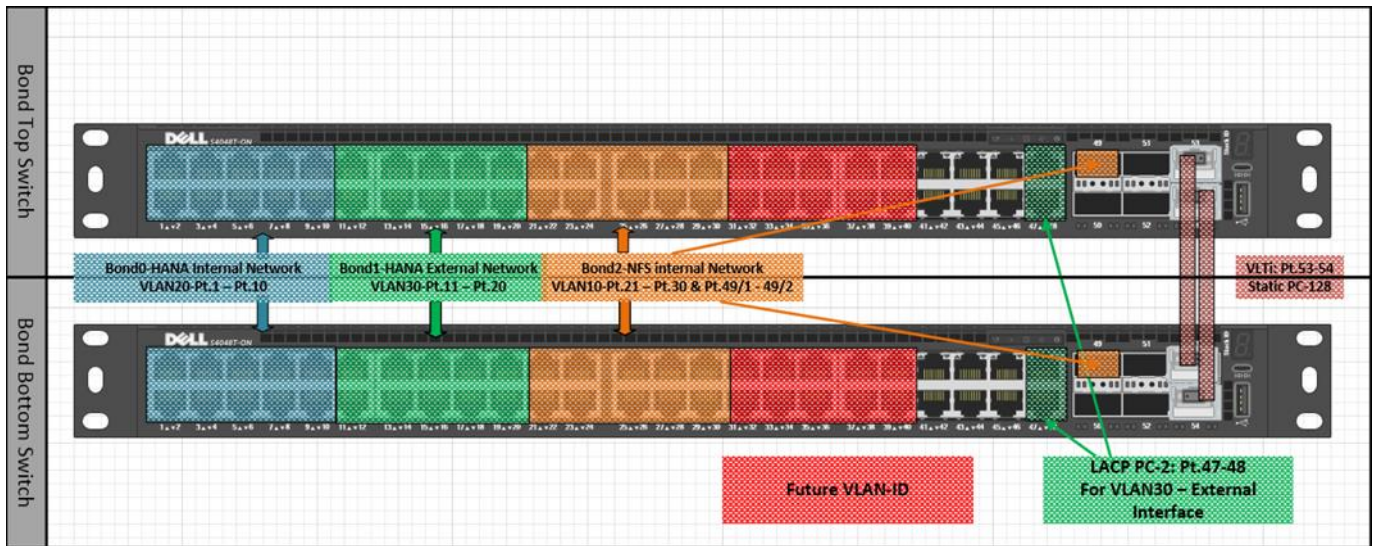


Figure 9. SAP HANA scale-out configuration with up to 10 active nodes

Figure 10 shows two pairs of Force 10 S4048T switches, in keeping with the requirement for a SAP HANA scale-out configuration with more than eight active SAP HANA nodes.

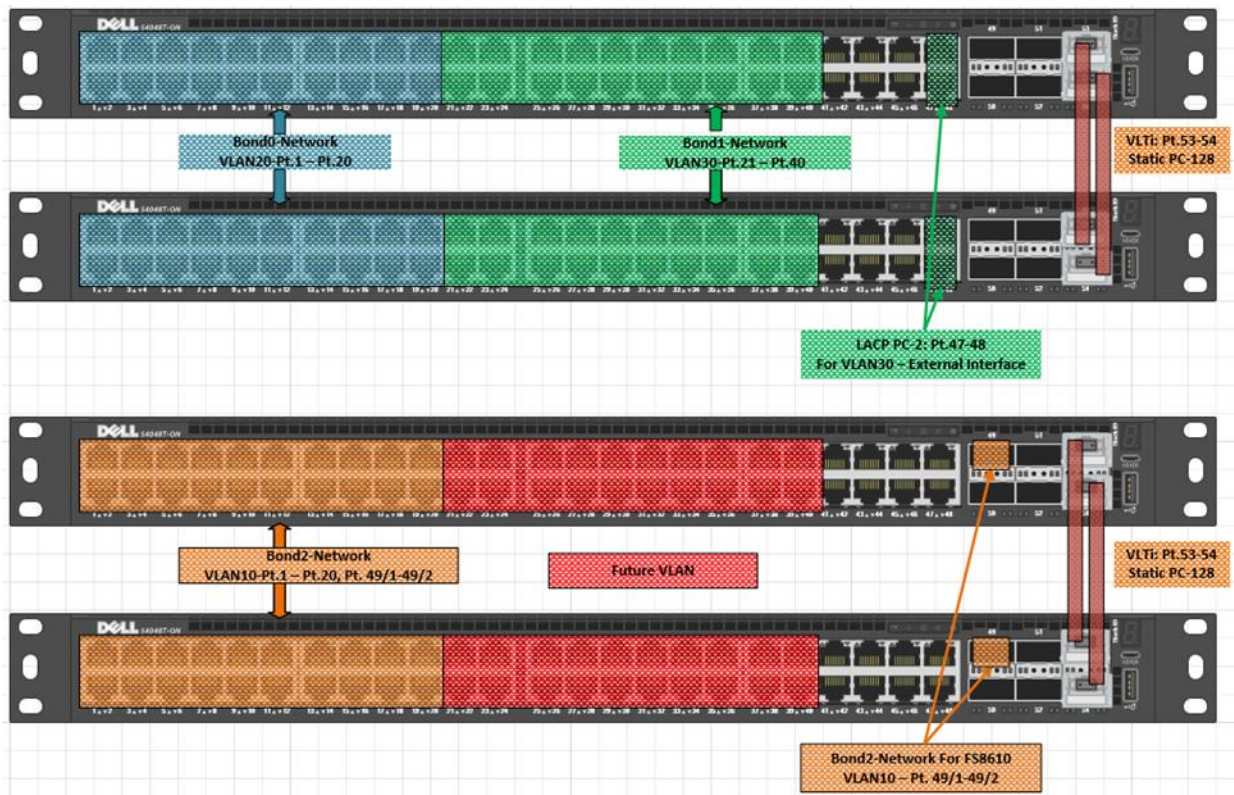
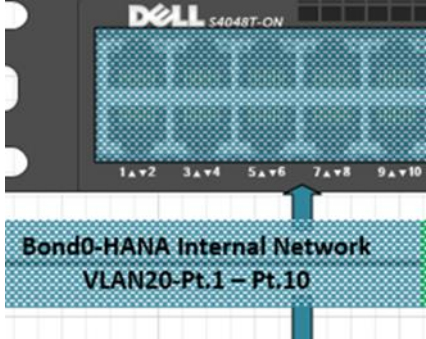
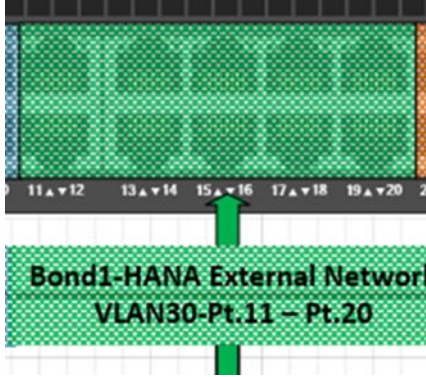
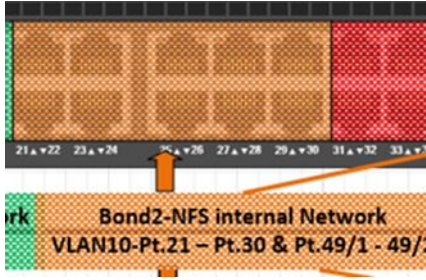


Figure 10. SAP HANA scale out configuration with more than 10 nodes

Table 13 shows detailed illustrations of network port assignment for a SAP HANA scale-out appliance with up to eight active and two standby SAP HANA nodes.

Table 13. Detailed network assignments

Network	Illustration
Dedicated SAP HANA internal network: <ul style="list-style-type: none"> • Bond0 • VLAN-20—port 1 to port 10 	
SAP HANA external/production network: <ul style="list-style-type: none"> • Bond1 • VLAN-30—port 11 to port 20 • Client and system replication networks 	
Dedicated NFS internal network: <ul style="list-style-type: none"> • Bond2 • VLAN-10—port 21 to port 30 and port 49/1 to port 49/2 	
Future network <ul style="list-style-type: none"> • For a backup or dedicated system replication network • Bondx • VLAN-<i>ID</i>—other available free ports 	

Network connection to the data center network infrastructure

The bond1 network connections include:

- External network:
 - VLAN-30—10 Gb port 11 to port 20
 - LACP port channel 2 for external interface—10 Gb port 47 to port 48

- vlt-peer-lag port-channel 2
- Tagged PC 2
- External switch:
 - The external switch's connections have the same LAG configuration:
 - PC 3 (any PC ID)
 - The same VLAN-ID:
 - VLAN-30
 - Same IP subnet: 10.68.204.x
- Tagged: tagged PC 3 (any port channel ID)

Verifying the hardware setup

Before installing SAP HANA, use the SAP HANA Hardware Configuration Check Tool (HWCCT) to verify the setup.

Use the HWCCT tool:

- Before production use.
- After each change in the landscape. If the system configuration is changed, run the HWCCT tool again to verify that the required KPI numbers are still within required thresholds.
- For regular health checks.

For information about the minimum KPI figures, and for instructions about how to run and configure HWCCT, refer to [SAP Note 1943937: Hardware Configuration Check Tool](#) (access requires an SPA username and password).

When validating the environment with the HWCCT tool 1.1 (fsperf), set the following parameters in the profile:

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

Setting HWCCT Parameters for SAP HANA 1.0

Use the hdbparam tool to configure the SAP HANA 1.0 database using the parameters that are set during HWCCT validation (refer to [Verifying the hardware setup](#)). The required steps are described in these SAP notes (access requires an SPA username and password):

- [SAP Note 2267798: Configuration of the SAP HANA Database during Installation Using hdbparam](#)
- [SAP Note 2186744: FAQ: SAP HANA Parameters](#)

Setting HWCCT Parameters for HANA 2.0

With SAP HANA 2.0, the hdbparam tool has been deprecated and is no longer installed as part of SAP HANA. Incorporate the HWCCT parameters that are set during HWCCT

validation (refer to [Verifying the hardware setup](#)) into the SAP HANA database environment using `global.ini` file as described in [SAP Note 2399079: Elimination of hdbparam in HANA 2](#) (access requires an SPA username and password).

Example: `/hana/shared/config/global.ini:`

```
.
.
.
[storage]
ha_provider = fcClient
ha_provider_path = /hana/shared/config/hdb_ha
partition_*_*_prtype = 5
partition_*_*_mountOptions = -t xfs
partition_1_data__alias = datavol1
partition_1_log__alias = logvol1
partition_2_data__alias = datavol2
partition_2_log__alias = logvol2
partition_3_data__alias = datavol3
partition_3_log__alias = logvol3
partition_4_data__alias = datavol4
partition_4_log__alias = logvol4
```

Dell EMC SAP HANA storage configuration

Meet these requirements when configuring Dell EMC SAP HANA storage:

- One Compellent SC5020/SC4020/SC8000/SC7020 array can serve up to eight SAP HANA nodes (16 TB with SAP HANA SPS 11 or later)
- One Compellent SC9000 Array can be configured to serve up to eight SAP HANA nodes (16 TB with SAP HANA SPS 11 or later)
 - 1 x SC5020/SC4020/SC8000/SC7020 dual controller array in a redundant configuration
 - 4 x SAS, 12 Gb, 4 Wide-Port, PCI-e backend (two per controller)
 - 8 Gb/16 Gb FC ports (four per controller)
 - For disk population and count, refer to [Table 4](#)
 - Compellent Fluid Data tiers
- Required Volume Space (8 x 3 TB per node):
 - 8 x 512 GB LOG volumes per array (approximately 4 TB)
 - 8 x 3 x 3 TB DATA volumes per array (approximately 74 TB)
 - 1 x `/hana/shared` volume/appliance per scale-out appliance on Array1 approximately 3 TB for 4 nodes = 2 x 3 TB = 6 TB for 8 x 2 TB nodes
 - Total required by volumes:
 - 4 TB (LOG) + 74 TB (DATA) + 4 TB (`/hana/shared`) approximately 84 TB for 8 x 2 TB node appliance
 - All writes are assigned Tier 1 SSDs (based on the storage profile)

SAP HANA storage connector API

Note: The SAP HANA Storage Connector API for Fibre Channel is the only supported API for use with this solution.

For the SAP HANA storage connector API, note the following:

- All data and log volumes use XFS.
- Only the owner node mounts its data and log volume.
- The standby node only mounts `/hana/shared` (NFS).
- The SAP HANA storage connector uses SCSI-3 Persistent Group Reservation (PGR) to guarantee file system integrity.
- During a node failure or a graceful shutdown, the reservation is released. The standby node then registers and acquires the reservation on the orphaned data and log volumes.

Using cluster file systems for a `/hana/shared` volume

A Dell EMC SAP HANA scale-out appliance using SC8000 storage arrays only supports two cluster file systems:

- The OCFS2 file system for SLES 11 SP3
- The GFS2 for RHEL 6.5

For more information about OCFS2 and GFS2 cluster setup procedures, refer to [Appendix A](#).

Building blocks for OCFS2

Dell EMC SAP Hana solutions are based on scalable building blocks. Table 14 lists the building blocks for an OCFS2 configuration.

Table 14. Building blocks for OCFS2

Module		
1	Infrastructure rack	<ul style="list-style-type: none"> • Network switches: <ul style="list-style-type: none"> ▪ 2 x Force 10 S4820T switches 48x10gE ports for internal SAP HANA network + external SAP HANA network ▪ (PROD) using VLANs + 1 x Force 10 S3048 switch, 48 x 1 GbE ports – management network + 2 x Force 10 S60 switches, 48 x 1 GbE ports – cluster network • SAN Fabric: <ul style="list-style-type: none"> ▪ 2 x Brocade 6510 24-36-48 Port FC8/16 Switch ▪ Management: <ul style="list-style-type: none"> ▪ 1 x PowerEdge R630 server as a management node + 1 x KMM console

Module		
2	Servers	<ul style="list-style-type: none"> • R930 72 cores/1.5 TB • Intel X540 Dual Port Ethernet 10 Gb BT + I350 1Gb BT Dual Port Network Daughter Card • 3 x Intel 10 GbE dual port RJ45 (X540-T2) • 2 x Qlogic 2660 Single Port 16 Gb Optical Fibre Channel HBA
3	Compellent array	One Compellent Array for 8 SAP HANA active nodes (8 TB) + 1 SAP HANA standby node

Building block details

Table 15 lists building block details for an OCFS2 configuration.

Table 15. 3 TB Configuration OCFS2/GFS2

Component	Configuration
SAP HANA Nodes	2 x PowerEdge R930 server active SAP HANA nodes + 1 or 2 x R930 server Standby SAP HANA node each with: <ul style="list-style-type: none"> • SLES 11 SP3 for SAP or RHEL 6.6 • 1.5 TB RAM • 4 x Intel Xeon 18 core E7-8880 v2 2.5GHz with 37.5M L2 cache or E7-8890 v2 2.8GHz with 37.5M L2 cache
Management	1 x PowerEdge R630 server with: <ul style="list-style-type: none"> • SLES 11 SP3 for SAP or RHEL 6.6 • 1 x Force10 S60 switch for Infrastructure management network • SAP HANA Studio
Storage	1 x Compellent SC8000, 6 Gb SAS Dual Controller Array with: <ul style="list-style-type: none"> • 2 x SAS, 6 Gb, 4 Wide-Port, PCI-e backend/controller (dual chain) • 2 x FC16, 2-Port QLE2662, PCI-e front end per controller 1 x 24 bay 2.5" 400G SAS SSD with 7 total disks per array • 1 x 24 bay 2.5" 400G SAS SSD with 7 total disks per array • 2 x 24 bay 2.5" 1.8TB 10K RPM SAS with 14 (7+7) spindles per array
Storage Fabric	2 x Brocade 6510 24-36-48 Port FC16 Switch
Network	2 x Force 10 S4820T switches 48x10gE ports for internal SAP HANA network + external SAP HANA network (PROD) using VLANs 2 x Force 10 S60 switches, 48 x 1 GbE ports – Cluster network 1 x Force 10 S60 switch, 48 x 1 GbE ports – management network

Figure 11 shows a 3 TB OCFS2 configuration on an SC4020 array.




Figure 11. 3 TB OCFS2 configuration

Compellent SC8000 storage array configuration for SAP HANA

Follow these guidelines when configuring Dell EMC SAP HANA storage using SC8000 storage arrays:

- One Compellent array can attach up to eight SAP HANA nodes (8 TB):
 - 2 x SC8000 controllers with 64 GB RAM each in a redundant configuration
 - 2 x SAS, 6 Gb, 4 Wide-Port, PCI-e back-end/controller
 - 2 x FC16, 2-Port QLE2662, PCI-e front-end/controller
 - 2 x 24 bay 2.5" 1.8 TB 10K RPM SAS with up to 48 spindles per array
 - 1 x 24 bay 2.5" 400 G SAS SSD with up to 24 total disks per array
- Compellent Fluid Data tiers for a fully populated array:
 - Tier 1: 24 x 400 GB SSD (1HS) R10 (approximately 4.8 TB) usable space per array
 - Tier 2: 0 x 15K RPM SAS: Not applicable
 - Tier 3: 48 x 1.8 TB 10K SAS R5 (approximately 72TB)
 - Single disk folder with approximately 76 TB (1.8 TB SAS) usable space per array
 - All writes are assigned to Tier 1 SSDs (based on the storage profile)

For information about tiering, refer to [Flash-aware sub-LUN tiering](#).

- Required volume space (for example, for 8 x 1.5 TB SAP HANA worker nodes):

- 8 x 512 GB log volumes per array (approximately 4 TB)
- 8 x 4.5 TB data volumes per array (approximately 36 TB)
- 1 x `/hana/shared` volume per scale-out appliance on Array1 for an approximate total appliance size of 12 TB for 1.5 TB nodes
- Total required by volumes
 - Approximately 4 TB + 36 TB +12 TB (`/hana/shared`) approximately 52 TB for an 8 x 1.5 TB node scale-out appliance
- All writes are assigned to Tier 1 SSDs (based on the storage profile)

Network configuration for a cluster file system

An SAP HANA 3.0 scale-out appliance can support up to 16 active and two standby SAP HANA nodes. The scale-out appliance includes the following networks.

- **Management network**—A 1 G Ethernet network using a Force 10 S60 switch. A PowerEdge R630 server, which is a dedicated management node, uses this network to manage all hardware components. A Dell iDRAC network is also part of the management network. It uses the onboard I350 1 Gb Base-T Dual Port Network Daughter Card.
- **Internal Cluster network**—A dedicated and redundant 1 G communication channel for the internal heartbeat network, using a pair of Force 10 S60 switches (bond2 network)
- **SAP HANA-related networks**—One pair of 10 G Force 10 S4820T switches in a VLT domain to supports redundant data paths for different SAP HANA-related networks. Each S4820T switch has 48 x 10 G Base-T Ports + 4 x 40 Gb QSFP+ ports.
 - **VLAN-20**—Dedicated SAP HANA Internal network (bond0 network)
 - **VLAN-30**— SAP HANA prod/external network (bond1 network)

Other available ports can be used for other networks such as a backup or a required network, to create a dedicated system replication network using the Intel X540 Dual Port 10 Gb Ethernet network adapter on the daughter card.

- **Storage Area Network**— dedicated and redundant FC16 SAN fabric to host all database related volumes, using one pair of FC16 Brocade 6510 switches

Figure 12 shows a clustered-file system network configuration.

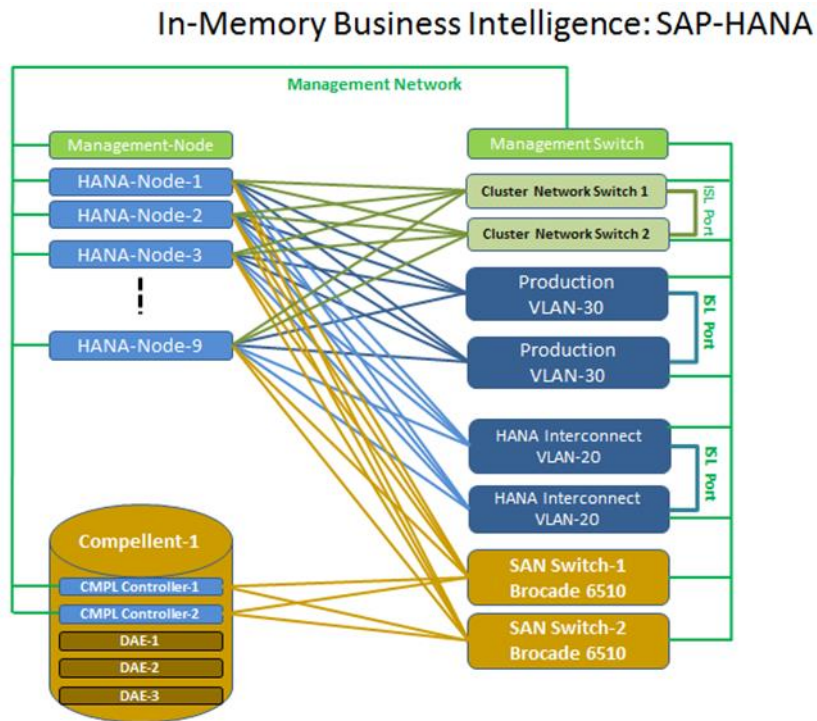


Figure 12. Network configuration using a cluster file system

For information about Force 10 VLT, refer to [Force10 Virtual Link Trunking](#).

SAP HANA-related networks

One pair of Force 10 S4820T switches is configured in a VLT domain to host SAP HANA-related networks using different VLAN-IDs. The switch can support up to 16 active and two standby SAP HANA nodes.

Up to 16 nodes using copper RJ-45 ports, SFP+ ports (port 50 to port 55) are required for deployments with more than 16 nodes.

Figure 13 shows the network port assignment for a SAP HANA scale-out configuration with up to 16 active and two standby SAP HANA nodes.

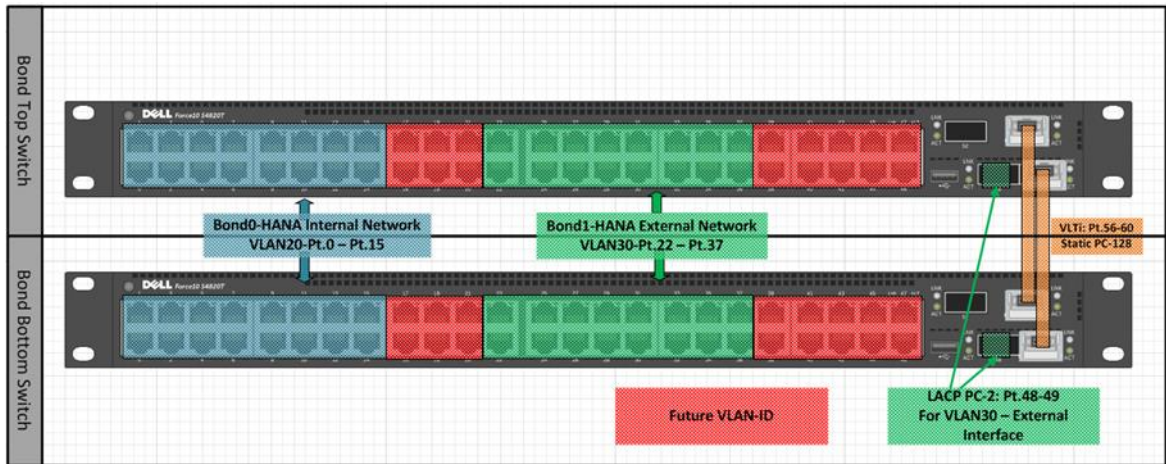


Figure 13. SAP HANA scale-out configuration with up to 16 active nodes

Table 16 shows detailed illustrations of network port assignment for a SAP HANA scale-out appliance with up to eight active and two standby SAP HANA nodes

Table 16. Detailed network assignments

Network	Port assignments
Dedicated SAP HANA internal network: <ul style="list-style-type: none"> Bond-0: VLAN-20 – port 0–port 15 	
SAP HANA External/Prod Network: <ul style="list-style-type: none"> Bond-1: VLAN-30 - port 22– port 37 Client and system replication networks 	
Dedicated network for backup or system replication: <ul style="list-style-type: none"> Bond-x: VLAN-ID – other available ports on the switch 	

Volume layout and mappings for HANA

Table 17 shows the volume layout and mappings to each SAP HANA node on an array.

Table 17. Volume layout and mappings

Array	Component	Configuration
SC8000 Array1	Data volumes	<ul style="list-style-type: none"> • /hana/data/HAN/mnt00001 ... mnt00008: • File system: XFS using SAP HANA storage connector
	Log volumes	<ul style="list-style-type: none"> • /hana/log/HAN/mnt00001 .. mnt00008: • File system: XFS using SAP HANA storage connector
	SAPMNT	<ul style="list-style-type: none"> • hana/shared (shared SAP HANA home binaries and executables available on all nodes) • File system: OCFS2/GFS2
SC8000 Array2	Data volumes	<ul style="list-style-type: none"> • /hana/data/HAN/mnt00009 ... mnt000016:
	Log volumes	<ul style="list-style-type: none"> • /hana/log/HAN/mnt00009 .. Mnt000016: • File system: XFS using SAP HANA storage connector

Each Compellent SC8000 storage array can host a maximum of 8 x 3 TB data volumes and 8 x 512 GB log volumes. Additionally, the first array also hosts the /hana/shared mount point for shared SAP HANA binaries. The array is sized equal to the total scale-out appliance size in TB and uses OCFS2 or GFS2 as the shared file system.

Note: The GFS2 file system is only certified for the Compellent SC8000 as a scale-out appliance.

SAP HANA shared file system on Compellent SC8000 arrays

SAP HANA database binaries are installed on a shared file system. In a cluster file system scenario, a separate NFS server is not required to provide a shared file system infrastructure for the /hana/shared volume, even though some complexity of cluster file system management is involved.

OCFS2 configuration on SLES 11

Dell EMC recommends that you use the OCFS2 file system to provide the shared mount point hana/shared for SAP HANA binaries.

OCFS2 uses the SLES High Availability (HA) Extension, which is based on open source clustering technologies:

- Corosync/OpenAIS as the messaging infrastructure (node membership)
- PCMK as the OCFS2 cluster stack
- STONITH Block Device (SBD) for fencing

Use the following SBD settings:

```
Timeout (watchdog): 30
Timeout (msgwait) : 60
stonith-timeout="120s
stonith-enabled="true"
```

The OCFS2 mount operation must wait for the distributed lock manager (DLM) to start up before it mounts the volume. [Table 17](#) shows the configuration for each SC8000 array.

GFS2 configuration on RHEL 6.5 (SC8000)

On an SC8000 array on RHEL 6.5, `/hana/shared` uses GFS2 to provide a shared mount point for SAP HANA binaries by using:

- The cluster manager (CMAN) in the High Availability Add-on for RHEL 6.5
- Clustered Logical Volume Manager (CLVMD)
 - CLVMD uses an Intelligent Platform Management Interface (IPMI) LAN for node fencing, as shown in [Figure 14](#):

```
agent="fence_ipmilan"
ipaddr="<iDRAC IP>"
lanplus="1"
power_wait="4"
```

```
[root@dellhana1 ~]# ccs_tool lsnode
Cluster name: HANAcluster, config_version: 20
Nodename          Votes Nodeid Fencetype
dellhana1         1      1  device_node1
dellhana2         1      2  device_node2
dellhana3         1      3  device_node3
dellhana4         1      4  device_node4
dellhana5         1      5  device_node5
dellhana6         1      6  device_node6
dellhana7         1      7  device_node7
dellhana8         1      8  device_node8
```

Figure 14. Node fences

Chapter 4 Compellent Storage Configuration

This chapter presents the following topics:

Overview	45
Configuring switch zoning	45
Configuring Qlogic HBA driver settings.....	45
Setting up the storage array controller	45

Overview

An SAP HANA workload requires special consideration for sizing and allocating volumes for data and log volumes, which provide the persistence layer for the SAP HANA in-memory database engine. SAP HANA requires access to all data and log volumes from all nodes as shared storage although each worker node owns and operates its own set of data and log volumes.

The workload generates random I/O on the data volume and sequential I/O on the log volume. Data volume access from SAP HANA can range from 4K, 16K, 64K, 1M, and 16M block sizes. Log volume access from SAP HANA can range from 4K, 16K, and 1M block sizes.

Configuring switch zoning

Configure switch zoning so that each host can detect all virtual ports on all storage arrays. For detailed zoning best practices, refer to the documentation for your storage array that is listed in [References](#).

Configuring Qlogic HBA driver settings

To avoid any impact to SAP HANA operation, configure the settings to process controller failover in dual controller configurations.

Add the following entry to the `/etc/modprobe.conf.local` file:

```
options qla2xxx qlport_down_retry=5
```

Note that this option is required.

Setting up the storage array controller

For details about setting up SC8000 storage for SAP HANA, refer to [Appendix A](#).

For information about how to initialize and set up the storage arrays, refer to the documentation for your storage array that is listed in [References](#).

Chapter 5 Accessing the Compellent SAN from SAP HANA Nodes

This chapter presents the following topics:

Setting up multipath I/O	47
Mounting SAP HANA data and log volumes	50

Setting up multipath I/O

To access the block devices from SAP HANA nodes, enable native Linux DM-Multipathing (DM-MPIO):

1. To set up DM-MPIO, create the `multipath.conf` configuration file.

The following example shows a sample `multipath.conf` configuration file for an eight-node (worker nodes) SAP HANA solution. Use the Compellent Storage Center GUI to find the World Wide Identifiers (WWIDs).

```
defaults {
    user_friendly_names yes
    rr_min_io          10
}
blacklist {
    wwid 36d4ae5206ecd5a0016b189780d280cc9
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
    devnode "^hd[a-z]"
}
multipaths {
    multipath {
        wwid 3600d31000ec990000000000000000016
        alias sapmnt
    }
    multipath {
        wwid 3600d31000ec990000000000000000012
        alias logvol1
    }
    multipath {
        wwid 3600d31000ec990000000000000000013
        alias logvol2
    }
    multipath {
        wwid 3600d31000ec990000000000000000014
        alias logvol3
    }
    multipath {
        wwid 3600d31000ec990000000000000000015
        alias logvol4
    }
    multipath {
        wwid 3600d31000ec990000000000000000017
        alias logvol5
    }
    multipath {
        wwid 3600d31000ec990000000000000000018
        alias logvol6
    }
    multipath {
        wwid 3600d31000ec990000000000000000019
        alias logvol7
    }
}
```

```

    }
multipath {
    wwid 36000d31000ec99000000000000000001a
    alias logvol8
}
multipath {
    wwid 36000d31000ec99000000000000000000e
    alias datavol1
}
multipath {
    wwid 36000d31000ec99000000000000000000f
    alias datavol2
}
multipath {
    wwid 36000d31000ec990000000000000000010
    alias datavol3
}
multipath {
    wwid 36000d31000ec990000000000000000011
    alias datavol4
}
multipath {
    wwid 36000d31000ec99000000000000000001b
    alias datavol5
}
multipath {
    wwid 36000d31000ec99000000000000000001c
    alias datavol6
}
multipath {
    wwid 36000d31000ec99000000000000000001d
    alias datavol7
}
multipath {
    wwid 36000d31000ec99000000000000000001e
    alias datavol8
}
}

```

2. To activate settings and rescan new devices, restart the daemon:

```
service multipathd restart
```

The volumes are now visible on each host.

3. To show that each device has four paths to the array, run the following command:

```

dellhana1:/ # multipath -ll
datavol1 (36000d31000ec990000000000000000011) dm-12
COMPELNT,Compellent Vol
size=1.6T features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active

```



```

|- 19:0:0:4  sdd  8:48  active ready running
|- 19:0:2:4  sdu  65:64  active ready running
|- 20:0:0:4  sdal 66:80  active ready running
`- 20:0:2:4  sdbc 67:96  active ready running
logvol1 (36000d31000ec99000000000000000015) dm-10
COMPELNT,Compellent Vol
size=512G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
|- 19:0:0:8  sdf  8:80  active ready running
|- 19:0:2:8  sdw  65:96  active ready running
|- 20:0:0:8  sdan 66:112 active ready running
`- 20:0:2:8  sdbe 67:128 active ready running
datavol2 (36000d31000ec99000000000000000010) dm-21
COMPELNT,Compellent Vol
size=1.6T features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
|- 19:0:1:3  sdl  8:176  active ready running
|- 19:0:3:3  sdac 65:192 active ready running
|- 20:0:1:3  sdat 66:208 active ready running
`- 20:0:3:3  sdbk 67:224 active ready running
logvol2 (36000d31000ec99000000000000000014) dm-17
COMPELNT,Compellent Vol
size=512G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
|- 19:0:1:7  sdn  8:208  active ready running
|- 19:0:3:7  sdae 65:224 active ready running
|- 20:0:1:7  sdav 66:240 active ready running
`- 20:0:3:7  sdbm 68:0  active ready running
datavol3 (36000d31000ec9900000000000000000f) dm-8
COMPELNT,Compellent Vol
size=1.6T features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
|- 19:0:0:2  sdc  8:32  active ready running
|- 19:0:2:2  sdt  65:48  active ready running
|- 20:0:0:2  sdak 66:64  active ready running
`- 20:0:2:2  sdbb 67:80  active ready running
logvol3 (36000d31000ec99000000000000000013) dm-11
COMPELNT,Compellent Vol
size=512G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
|- 19:0:0:6  sde  8:64  active ready running
|- 19:0:2:6  sdv  65:80  active ready running
|- 20:0:0:6  sdam 66:96  active ready running
`- 20:0:2:6  sdbd 67:112 active ready running
datavol4 (36000d31000ec9900000000000000000e) dm-16
COMPELNT,Compellent Vol
size=1.6T features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
|- 19:0:1:1  sdk  8:160  active ready running
|- 19:0:3:1  sdab 65:176 active ready running

```

```

|- 20:0:1:1 sdas 66:192 active ready running
`- 20:0:3:1 sdbj 67:208 active ready running
logvol14 (36000d31000ec990000000000000000012) dm-20
COMPELNT,Compellent Vol
size=512G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
|- 19:0:1:5 sdm 8:192 active ready running
|- 19:0:3:5 sdad 65:208 active ready running
|- 20:0:1:5 sdau 66:224 active ready running
`- 20:0:3:5 sdbl 67:240 active ready running
sapmnt (36000d31000ec990000000000000000016) dm-24
COMPELNT,Compellent Vol
size=8.0T features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
|- 19:0:1:9 sdo 8:224 active ready running
|- 19:0:3:9 sdaf 65:240 active ready running
|- 20:0:1:9 sdaw 67:0 active ready running
`- 20:0:3:9 sdbn 68:16 active ready running

```

Mounting SAP HANA data and log volumes

SAP strongly recommends that you use the XFS file system for data and log volumes. Because the SAP HANA Storage Connector API does not support partitioned devices, the volumes must be unpartitioned.

To mount the SAP HANA data and log volumes, use the XFS file system with the SAP HANA Storage Connector API:

1. On the master and worker nodes, specify the following options in the `/etc/suoders` file:

```

<SID>adm ALL=NOPASSWD: /sbin/multipath, /sbin/multipathd,
/etc/init.d/multipathd, /usr/bin/sg_persist, /bin/mount,
/bin/umount, /bin/kill, /usr/bin/lsof

```

where *SID* is the SAP HANA system ID.

2. Format all volumes to be used for persistence with the XFS file system:

```

mkfs.xfs /dev/mapper/datavol1
mkfs.xfs /dev/mapper/logvol1

```

Based on the configuration that is provided in the `global.ini` file, the SAP HANA Storage Connector API automatically mounts the volumes on worker nodes and on a standby node (in case a worker node fails):

```

[communication]
listeninterface = .global
[persistence]
basepath_datavolumes=/hana/data/ANA
basepath_logvolumes=/hana/log/ANA
use_mountpoint = yes
basepath_shared = no

```

```
[public_hostname_resolution]
use_default_route = no
map_dellhanapri1 = dellhanapri1-prod
map_dellhanapri2 = dellhanapri2-prod
map_dellhanapri3 = dellhanapri3-prod
[storage]
ha_provider = hdb_ha.fcClient
partition_*_*__prtype = 5
partition_*_*__mountOptions = -t xfs
partition_1_data__alias = datavol1
partition_1_log__alias = logvol1
partition_2_data__alias = datavol2
partition_2_log__alias = logvol2
partition_3_data__alias = datavol3
partition_3_log__alias = logvol3
partition_4_data__alias = datavol4
partition_4_log__alias = logvol4
```

For information about how to install and configure the SAP HANA environment, refer to the [SAP HANA Server Installation and Update Guide](#).

Chapter 6 Conclusion

This chapter presents the following topics:

Summary53

Summary

The SAP TDI model enables customers to achieve cost savings and flexibility by deploying SAP HANA. The TDI approach allows customers to reduce hardware costs, leverage existing hardware investments, apply operational best practices that are used elsewhere in the data center, and consolidate SAP and non-SAP applications on a common infrastructure. Shared enterprise storage is a key enabling technology to achieve these advantages.

The Compellent line of storage systems provides an excellent foundation for building robust and high performance SAP HANA infrastructures. The enterprise-grade software features from Compellent such as Data Progression, replays (snapshots), and replication capabilities further enhance the value of the solution.

Joint Dell EMC and SAP testing showed that the storage design that is described in this document meets the KPI requirements to implement SAP HANA solutions on Compellent storage arrays.

Chapter 7 References

This chapter presents the following topics:

Dell EMC documentation	55
SAP documentation	55
Other documentation	55

Dell EMC documentation

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

- [Dell Compellent SC4020 manuals and documents](#)
- [Dell Storage SC7020 manuals and documents](#)
- [Dell Storage SC8000](#)
- [Dell Storage SC9000](#)
- [New economies of storage with the Compellent Flash-optimized solutions](#)
- [Dell SAP/ SAP HANA Solutions](#)

SAP documentation

The following search results on the SAP website provide additional and relevant information:

- [Dell certified appliances Scale-Up \(for Single Node\)](#)
- [Dell certified appliances Scale-Up \(for Suite on SAP HANA\)](#)
- [Dell certified appliances Scale-Out](#)

The following documentation on the SAP website provides additional and relevant information:

- [SAP HANA Server Installation and Update Guide](#)
- [SAP HANA Tailored Data Center Integration](#)
- [Learn more about SAP HANA Tailored Data Center Integration \(TDI\)](#)

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

- [SAP Note 2009879: SAP HANA Guidelines for Red Hat Enterprise Linux \(RHEL\) Operating System](#)
- [SAP Note 1944799: SAP HANA Guidelines for SLES Operating System Installation](#)

Other documentation

The following Red Hat and SUSE documentation provides additional and relevant information:

- [Red Hat Enterprise Linux for SAP HANA](#)
- [SUSE Linux Enterprise Server for SAP Applications](#)

Appendix A Setting Up Cluster File Systems

This appendix presents the following topics:

Configuring an SLES HA cluster for a shared OCFS2 volume	57
Configuring cman for GFS2 for shared volumes	58

Configuring an SLES HA cluster for a shared OCFS2 volume

The following example setup shows how to use cluster file systems for `/hana/shared` volumes on SAP HANA nodes.

To configure the OCFS2 file system, configure the SLES HA cluster that provides the heartbeat, node membership, and fencing for SAP HANA nodes:

1. Create a 1 GB partition on a `/dev/mapper/sapmnt` volume to be used as a Shoot The Other Node In The Head (STONITH) block device.
2. Configure the `/etc/corosync/corosync.conf` file on each node. Use the default file and edit the parameters that are based on the configuration being used for the cluster network on your environment:

```
bindnetaddr: 192.168.0.0
mcastaddr: 225.18.7.1
mcastport: 5405
```

3. Configure `/etc/sysconfig/sbd` on each node, specifying the STONITH block device:

```
SBD_DEVICE="/dev/mapper/sapmnt_part1"
SBD_OPTS="-W"
```

4. To load the Cluster Information Base (CIB), type the following command (the `configuration_ha` file contains the CIB details):

```
crm configure load replace configuration_ha

primitive dlm ocf:pacemaker:controld \
    op start timeout="90s" \
    op stop timeout="100s" \
    op monitor interval="60" timeout="60"
primitive o2cb ocf:ocfs2:o2cb \
    op start timeout="90s" \
    op stop timeout="100s" \
    op monitor interval="60" timeout="60"
primitive sbd_stonith stonith:external/sbd \
    params sbd_device="/dev/mapper/sapmnt_part1"
group base-group dlm o2cb
clone base-clone base-group \
    meta interleave="true" target-role="started"
property $id="cib-bootstrap-options" \
    dc-version="1.1.7-
77eeb099a504ceda05d648ed161ef8b1582c7daf" \
    cluster-infrastructure="openais" \
    stonith-timeout="120s" \
    stonith-enabled="true" \
    expected-quorum-votes="1" \
    no-quorum-policy="freeze" \
    maintenance-mode="false" \
    last-lrm-refresh="1336399136"
```

```
op_defaults $id="op_defaults-options" \
    record-pending="false"
```

5. To start DLM on all nodes, start the OpenAIS service:

```
chkconfig openais on
service openais start
```

6. To configure the OCFS2 file system and to mount it as a `/hana/shared` volume on all nodes, create another partition on `/dev/mapper/sapmnt`, taking all the available space:

```
mkfs.ocfs2 -b 2k -C 1M -L "sapmnt" -N 17 /dev/mapper/sapmnt_part2
```

7. Create an `/etc/fstab` entry on all nodes for `/hana/shared` as follows:

```
/dev/mapper/sapmnt_part2    /hana/shared
ocfs2 _netdev,nointr,data=writeback,noatime
```

8. To mount the `/hana/shared` volume to use for the SAP HANA binaries, issue the following command on all nodes:

```
mount -a -t ocfs2
```

Configuring cman for GFS2 for shared volumes

The following example setup shows how to configure the Cluster Manager (cman) utility for configuring the GFS2 file system for RHEL6.5 on a Compellent SC800 storage controller.

1. Create a single partition on `/dev/mapper/sapmnt`, taking all the available space, to configure the GFS2 file system and to mount it as a `/hana/shared` volume on all nodes.
2. Enable the cluster volume manager on all nodes:

```
lvmconf --enable-cluster
```

3. Create a cluster configuration file, using values that are based on your environment:

```
vi /etc/cluster/cluster.conf
```

Note: This configuration file uses IPMI fencing for cluster issues.

```
<?xml version="1.0"?>
<cluster config_version="20" name="HANACluster">
  <fence_daemon clean_start="0" post_fail_delay="1"
post_join_delay="3"/>
  <clusternodes>
    <clusternode name="dellhana1" nodeid="1">
      <fence>
        <method name="1">
          <device domain="dellhana1" name="device_node1"/>

```

```

        </method>
    </fence>
</clusternode>
<clusternode name="dellhana2" nodeid="2">
    <fence>
        <method name="1">
            <device domain="dellhana2" name="device_node2"/>
        </method>
    </fence>
</clusternode>
<clusternode name="dellhana3" nodeid="3">
    <fence>
        <method name="1">
            <device domain="dellhana3" name="device_node3"/>
        </method>
    </fence>
</clusternode>
<clusternode name="dellhana4" nodeid="4">
    <fence>
        <method name="1">
            <device domain="dellhana4" name="device_node4"/>
        </method>
    </fence>
</clusternode>
<clusternode name="dellhana5" nodeid="5">
    <fence>
        <method name="1">
            <device domain="dellhana5" name="device_node5"/>
        </method>
    </fence>
</clusternode>
<clusternode name="dellhana6" nodeid="6">
    <fence>
        <method name="1">
            <device domain="dellhana6" name="device_node6"/>
        </method>
    </fence>
</clusternode>
<clusternode name="dellhana7" nodeid="7">
    <fence>
        <method name="1">
            <device domain="dellhana7" name="device_node7"/>
        </method>
    </fence>
</clusternode>
<clusternode name="dellhana8" nodeid="8">
    <fence>
        <method name="1">
            <device domain="dellhana8" name="device_node8"/>
        </method>
    </fence>
</clusternode>

```

```

        </fence>
    </clusternode>
</clusternodes>
<cman/>
<fencedevices>
    <fencedevice agent="fence_ipmilan" ipaddr="172.16.4.41"
lanplus="1" login="root" name="device_node1" passwd="calvin"
power_wait="4"/>
    <fencedevice agent="fence_ipmilan" ipaddr="172.16.4.42"
lanplus="1" login="root" name="device_node2" passwd="calvin"
power_wait="4"/>
    <fencedevice agent="fence_ipmilan" ipaddr="172.16.4.43"
lanplus="1" login="root" name="device_node3" passwd="calvin"
power_wait="4"/>
    <fencedevice agent="fence_ipmilan" ipaddr="172.16.4.44"
lanplus="1" login="root" name="device_node4" passwd="calvin"
power_wait="4"/>
    <fencedevice agent="fence_ipmilan" ipaddr="172.16.4.45"
lanplus="1" login="root" name="device_node5" passwd="calvin"
power_wait="4"/>
    <fencedevice agent="fence_ipmilan" ipaddr="172.16.4.46"
lanplus="1" login="root" name="device_node6" passwd="calvin"
power_wait="4"/>
    <fencedevice agent="fence_ipmilan" ipaddr="172.16.4.47"
lanplus="1" login="root" name="device_node7" passwd="calvin"
power_wait="4"/>
    <fencedevice agent="fence_ipmilan" ipaddr="172.16.4.48"
lanplus="1" login="root" name="device_node8" passwd="calvin"
power_wait="4"/>
</fencedevices>
<rm>
    <failoverdomains/>
    <resources/>
</rm>
</cluster>

```

4. Validate the cluster configuration file for any errors:

```
ccs_config_validate
```

5. Copy the `cluster.conf` file to all nodes under `/etc/cluster`.

6. Start the `cman` cluster service on all nodes:

```
chkconfig cman on
service cman start
```

7. Check the status of the service:

```
service cman status
```

8. Start the cluster volume manager service on all nodes:

```
chkconfig clvmd on
```

```
service clvmd start
```

9. **Configure the cluster logical volume on the /dev/mapper/sapmntp1 partition to be used for the /hana/shared volume:**

```
pvccreate /dev/mapper/sapmntp1
vgcreate -cy vg_hanashared /dev/mapper/sapmntp1
lvcreate -L <size>T -n lv_hanashared vg_hanashared
```

10. **Create a GFS2 file system:**

```
mkfs.gfs2 -t HANAcluster:shared -j 18
/dev/vg_hanashared/lv_hanashared
```

11. **Create a /hana/shared mount point on all nodes and add the entry to /etc/fstab:**

```
mkdir -p /hana/shared
/dev/mapper/vg_hanashared-lv_hanashared /hana/shared gfs2
noatime,nodiratime 0 0
```

12. **Activate the volume group on all nodes:**

```
vgchange -aly vg_hanashared
```

13. **Start the GFS2 service on all nodes:**

```
chkconfig gfs2 on
service gfs2 start
```

14. **Configure iDRAC for IPMI-based node fencing on all nodes:**

```
racadm config -g cfgIpmiLan -o cfgIpmiLanEnable 1
racadm config -g cfgIpmiLan -o cfgIpmiLanPrivilegeLimit 4
```

Where:

- 2 = User
- 3 = Operator
- 4 = Administrator

Appendix B Setting Up an SC8000 Storage Array

This appendix presents the following topics:

Setting up the SC8000 controller	63
Logging in to Storage Center	64
Running the Startup Wizard	65
Configuring settings with storage array.....	73

Setting up the SC8000 controller

Initialize controllers

Use a serial cable to initialize both controllers. Set the following parameters:

- **Emulation**—VT220
- **Column Mode**—132
- **Line Wrapping**—Off
- **Connection**—Serial Port
- **Connection Type**—Direct
- **Baud Rate**—115,200
- **Parity**—None

Assign IP addresses to controllers

Storage Center controllers use static IP addresses:

- The eth0 Ethernet interface supports system login and access to the software. It is used to send email, alerts, SNMP traps, and phone home data. For more information about the Phone Home feature, refer to [Configuring the Phone Home feature](#).
- The eth1 Ethernet interface is used for dedicated Inter-Process Communication (IPC) between controllers in a clustered-controller system. There is no default gateway for eth1 and it does not have to be set.

Note: You rarely have to change the eth1 address from its factory setting because it is only used for IPC between controllers. You might have to change the address during setup if both controllers are factory-configured with the same eth1 address; this scenario is unlikely. Installation team members do not change eth1 settings unless requested by the customer. If changed, the settings must have a different subnet than eth0.

To set the IP address, netmask, and gateway type the following commands. The options are separated by a single space.

Note: The second DNS server IP address is optional.

```
controller ipconfig eth0 [IP address] [netmask] [Gateway]
controller dnsserver [DNS server 1 IP address] [DNS server 2 IP
address]
controller domainname [domain name]
```

The following is an example of IP settings:

```
controller ipconfig eth0 172.16.5.21 255.255.0.0 172.16.0.1
controller dnsserver [Return key]
controller domainname [Return key]
```

Note: In a dual-controller system, set these options for Controller 2, also. For the second controller, use the following IP address: eth0 172.16.5.22 255.255.0.0 172.16.0.1.

Verify settings To verify the settings, type:

`controller show`

Figure 15 shows example results.

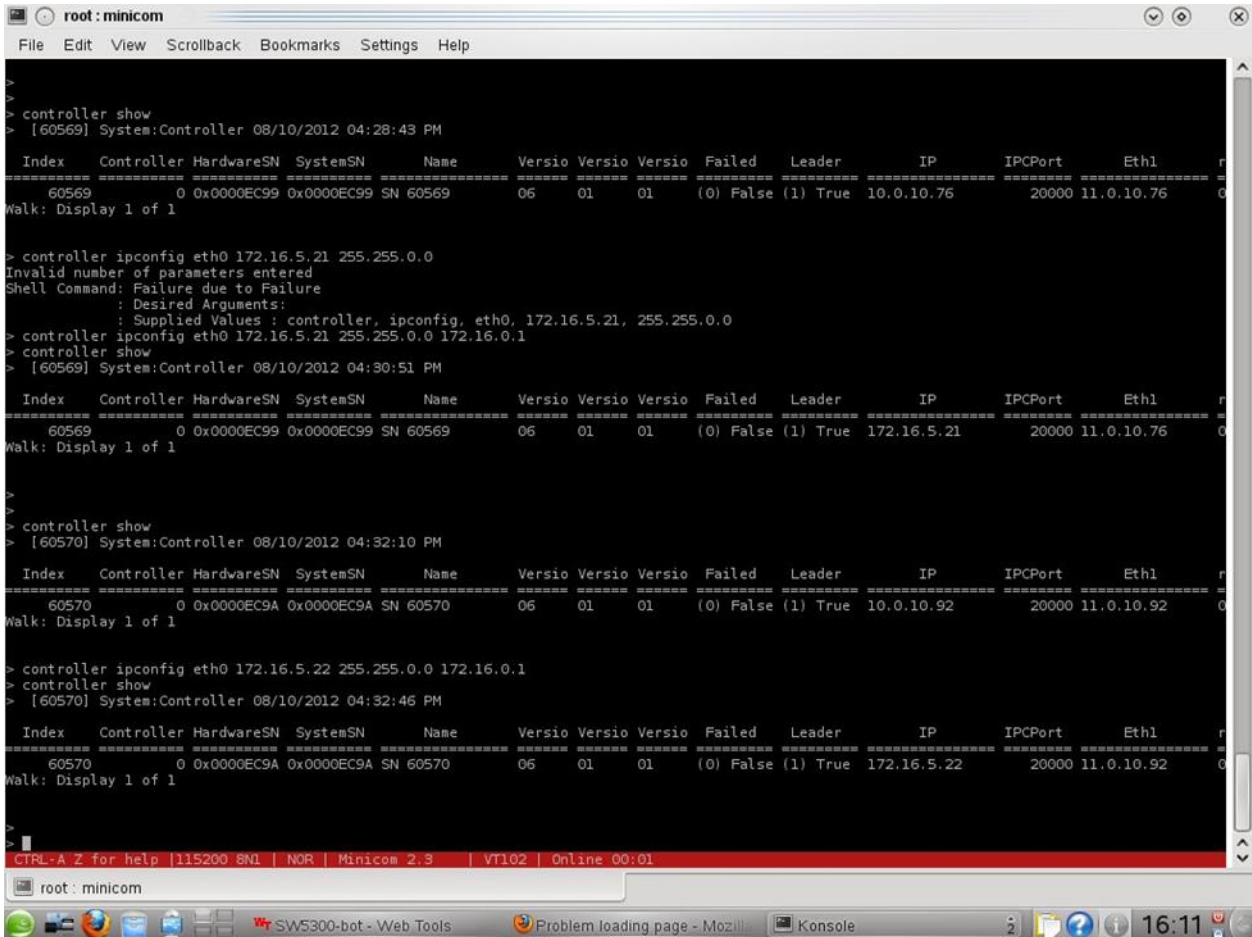


Figure 15. `controller show` command results

Logging in to Storage Center

Note: The Storage Center System Manager cannot load with the following unique combination of applications: Windows 2008 (64 bit), FireFox 3.0, and Java Runtime Environment 6_10.

To log in to Storage Center:

1. Open a browser on a workstation or PC with an Ethernet connection to the controller.

The following browsers are supported:

- Microsoft Windows Internet Explorer Versions 7, 8, and 9
- Mozilla Firefox Version 3 and Version 4 on Microsoft Windows

2. Type the eth0 IP address of Controller 1.

3. Click **Yes** to acknowledge all security alerts.

Note: Alerts differ, depending on your browser settings and Java version.

4. Type your username and password in the Storage Center Login window. The Startup Wizard appears.

Running the Startup Wizard

Loading the license

The license file is sent as an email attachment. Save the file in a location that is accessible from the PC or workstation that is used for the installation.

End-user license files have the following naming convention:

snxxx_35_date.lic

where:

- *snxxx* is the serial number of the Storage Center. This number must match the serial number of the controller from which you start the wizard.
- *35* indicates that the system runs post-3.5 code.
- *date* is the license generation date in YYMMDD format.
- *lic* is the file name extension.

On the **Load License** page:

1. Browse to and select the license file, and then click **Load License**, as shown in Figure 16.




Figure 16. Load the license

The Startup Wizard notifies you when the license is successfully loaded.

2. Click **Continue**.

Creating a disk folder

To create a disk folder:

1. From **Create Disk Folder** page, the under **Select the disk to be included in the Disk Folder**, select the disks, and then click **Continue**, as shown in Figure 17.

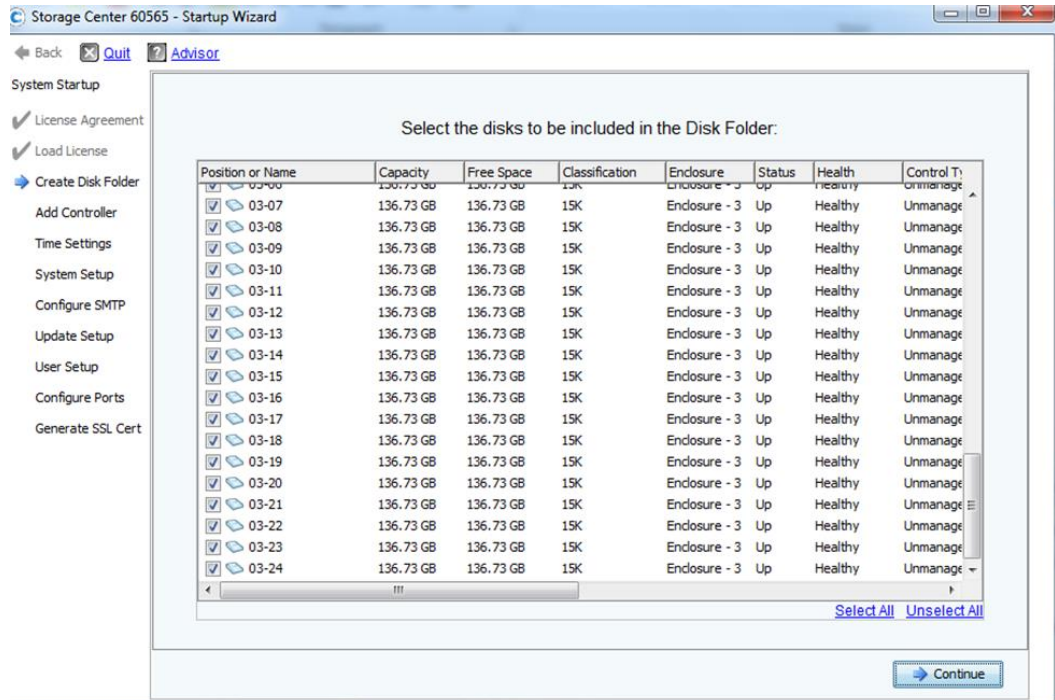


Figure 17. Create a disk folder

2. Under **Select the disks to be designated as Hot Spares**, select one spare disk per enclosure, and then click **Continue**, as shown in Figure 18.

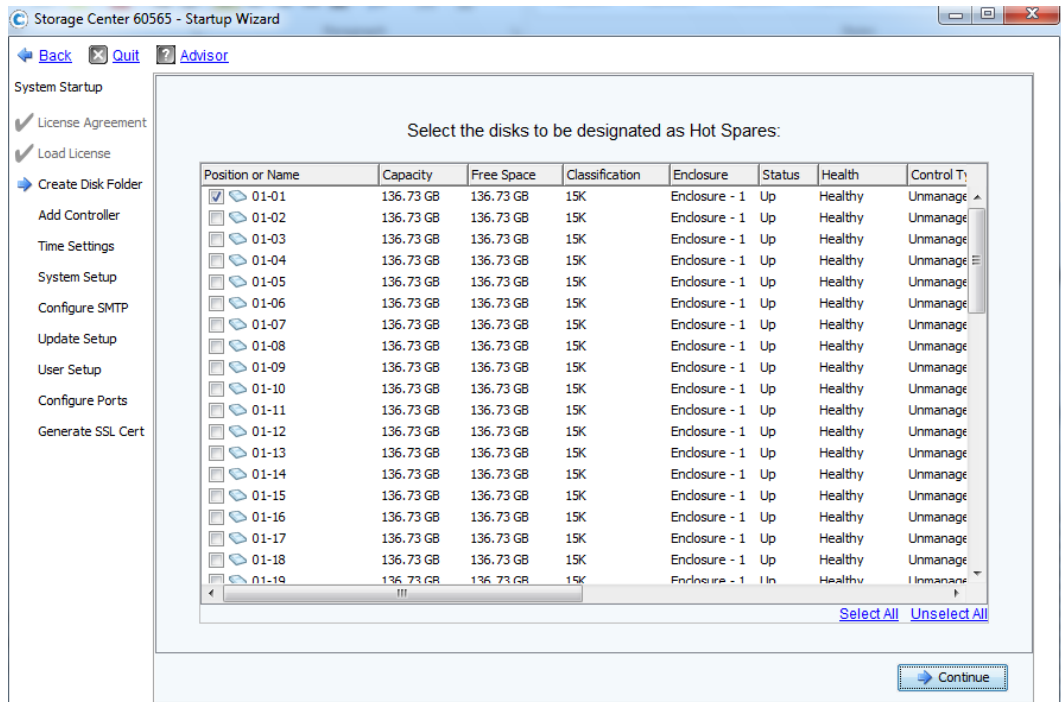


Figure 18. Select a spare disk

Figure 19 shows the attributes for the folder to be created. By default the disk folder automatically assigns disks to each tier in Compellent storage:

- Tier 1 includes SSDs
- Tier 2 includes 15 K RPM disks
- Tier 3 includes 10 K RPM disks

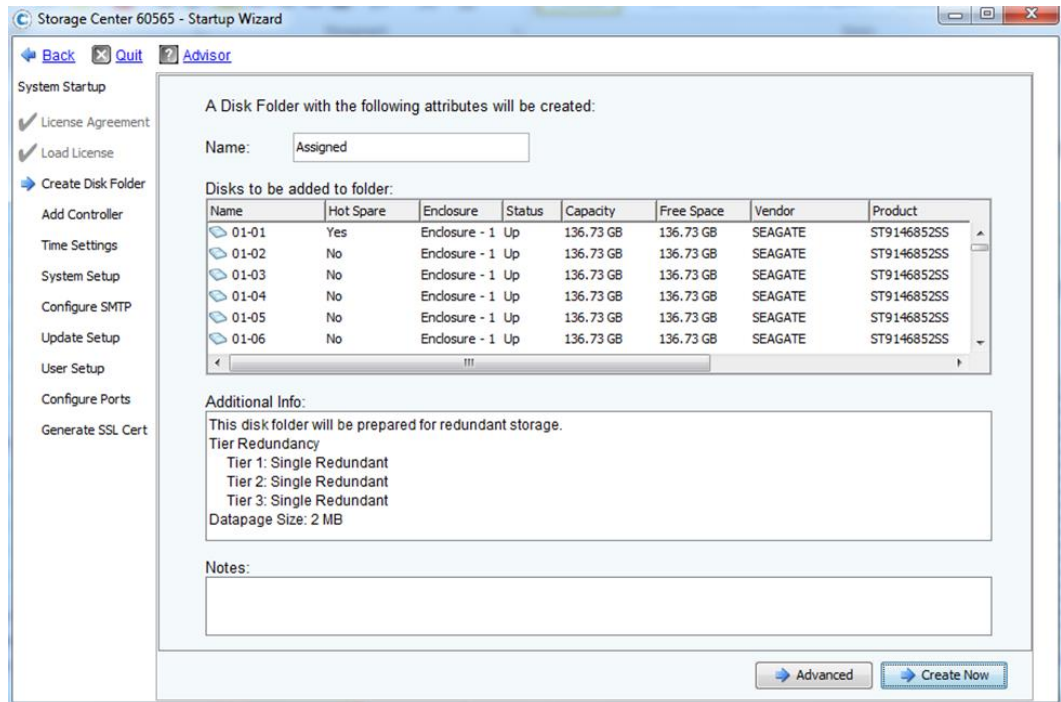


Figure 19. Disk folder attributes

3. Click **Create Now**.

Adding a controller

To add a controller:

1. From the **Add Controller** page, click **Add Controller**.
2. Type an IP address in the **Ether 1 (IPC) Interface** text box, and then click **Continue**, as shown in Figure 20. To find the IP address, run `controller show` on the second controller.

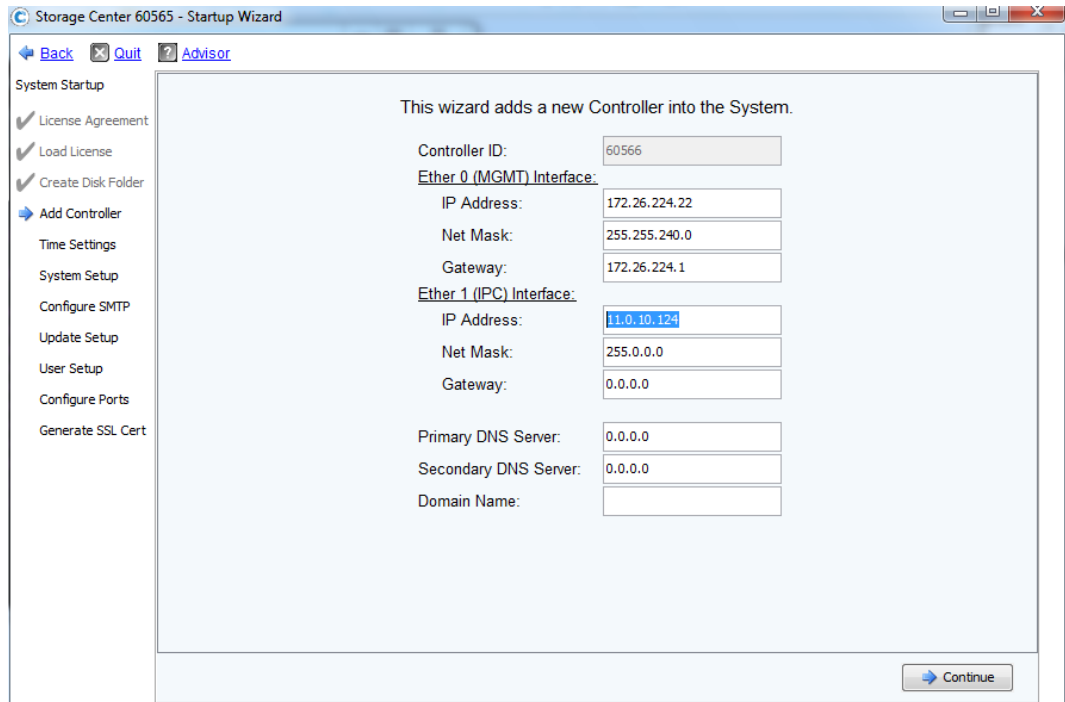


Figure 20. Provide the Ether 1 (IPC) Interface IP address

3. Verify the information and click **Join Now**.
4. Wait a few minutes while the newly added controller restarts.

Configuring time settings

On the **Time Settings** page, set the region, the time zone, and the time for the array's location.

System Setup

To set up the system:

1. From the **System Setup** page, type **172.16.5.20** in the **Management IP Address** text box, and then click **Continue**.
2. If there is a second storage array, type **172.16.5.24** in the **Management IP Address** text box, and then click **Continue**, shown in Figure 21. Otherwise, leave the **Management IP Address** text box empty and then click **Continue**.

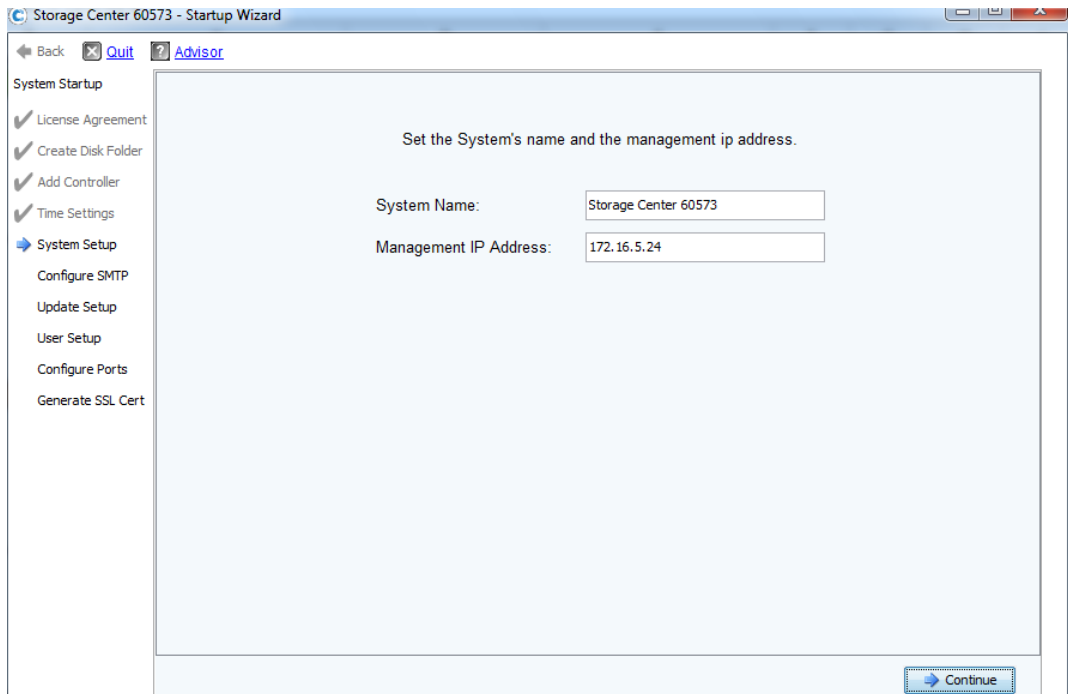


Figure 21. Set up the IP addresses

3. Disable read and write cache, and then click **Continue**.

Setting up email

In the **Configure SMTP** page:

1. Optionally, configure SMTP.
2. Click **Continue**.

Setting up updates

On the **Update Setup** page, select **Notify me of a software update but do not download automatically** from the list box, and then click **Continue**.

Setting timeout values and alert notifications

To set up timeout values and alert notifications: From the **User Setup** page:

1. From the **User Setup** page, from the list menu, select the session timeout value.
2. Type the email addresses for alerts, and then click **Continue**.

Configuring Ports

Before proceeding to this step in the wizard, ensure that all the FC ports are connected to FC switches.

To configure ports:

1. From the **Configure Ports** page, select **Convert Virtual Port Mode**, and then click **Continue**.
2. Click **Continue** again to proceed with the configuration.
3. Click **Configure Local Ports**.
4. Click **Edit Fault Domains**.

5. Click **Edit Fault Domain**, as shown in Figure 22.




Figure 22. Edit fault domain

6. To change the **Domain 0** name to **Top**, type **Top** and then click **OK**, as shown in Figure 23.

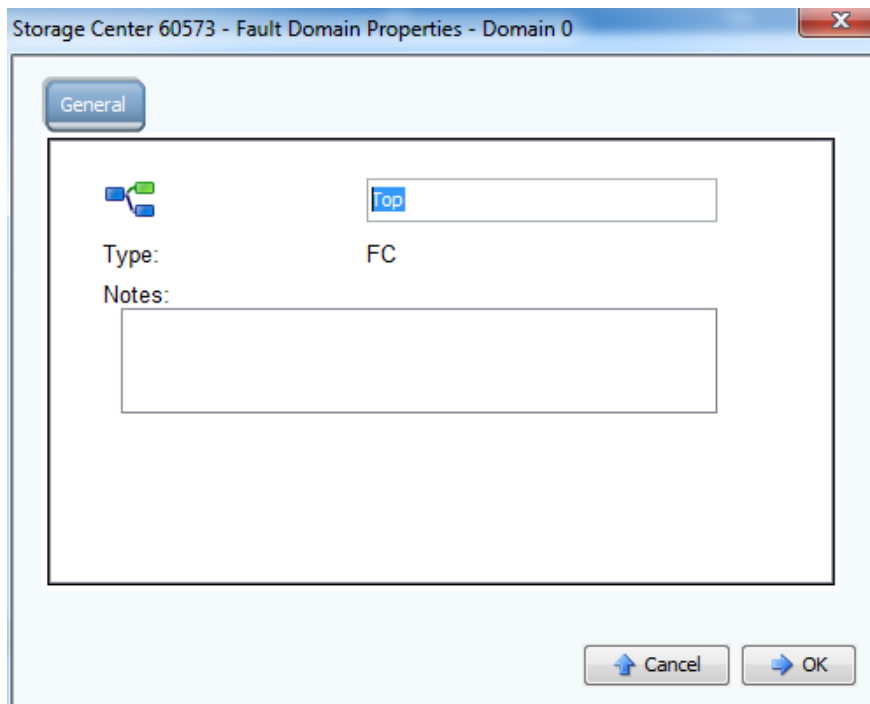


Figure 23. Change the fault domain name

7. Click **Create Fault Domain**, as shown in Figure 24.




Figure 24. Create fault domain

8. Create the **Bottom** fault domain, and then click **Continue**, as shown in Figure 25.

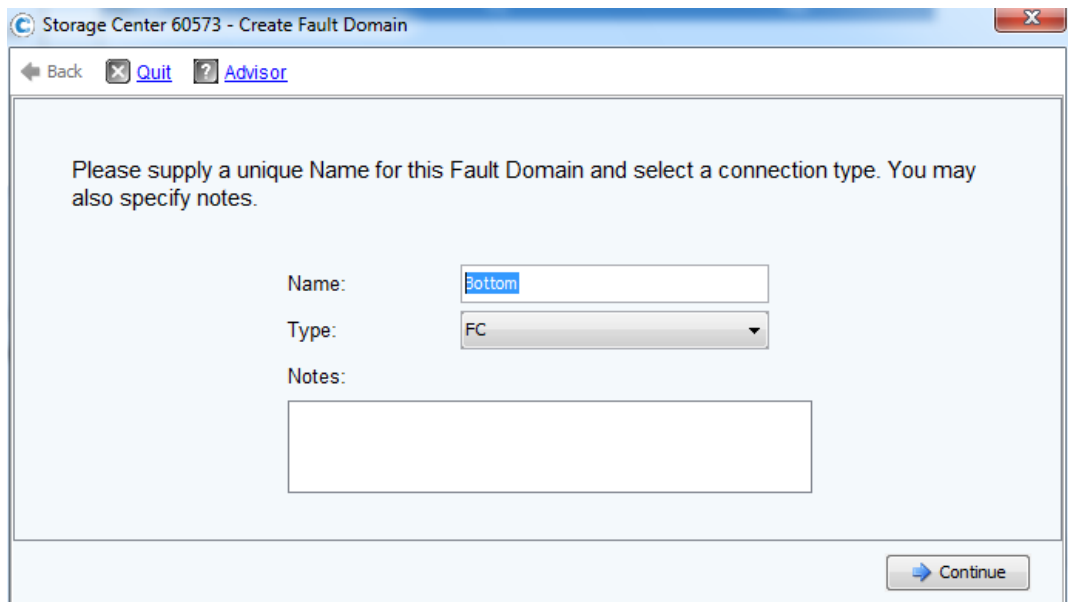


Figure 25. Create Bottom fault domain

9. Click **Create Now**.
10. Click **Return** to return to the **Configure Ports** page.
11. On the **Configure Ports** page (see Figure 26), do the following:
 - a. From the **Purpose** list box, select **Front End**.

- b. From the **Fault Domain** list box, set Port-1 to **Top**.
- c. From the **Fault Domain** list box, set Port-2 to **Bottom**.




Figure 26. Configure ports

- 12. Ensure settings look like Figure 27, and then click **Assign Now**.




Figure 27. Configured ports

- 13. Click **Continue**.

Generating or importing certificates

On the **Generate SSL Certificate** page, do either of the following:

- Generate a new Secure Socket Layer (SSL) certificate for the Storage Center
- Import an existing certificate for the Storage Center

The SSL certificate must match the Storage Center IP Address or DNS hostname.

Configuring settings with storage array

After you complete the Startup Wizard, use Storage Center to configure Compellent storage array settings such as the storage profile, the replay profile, and the data progression schedule. Also, set up volumes and hosts.

Configuring the Phone Home feature

For information about how to configure the Compellent Phone Home feature, refer to the documentation for your storage array that is listed in [References](#).

Enabling manual storage mode

To enable manual storage mode:

1. Log in to the Storage Center using the following credentials:
 - Username: **Admin**
 - Password: **“mmm”**
2. To enable manual storage mode, select **System > Setup > Enable Manual Storage Mode**.
3. Click **Continue**, as shown in Figure 28.

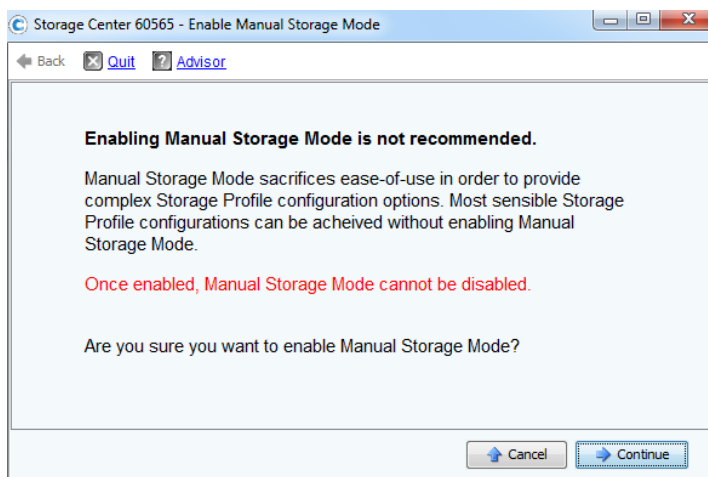


Figure 28. Enable Manual Storage Mode

Configuring the storage profile

Data and log volumes both have storage profiles.

To create a profile for a SAP HANA volume:

1. From Storage Center, click **Create Storage Profile**, as shown in Figure 29.

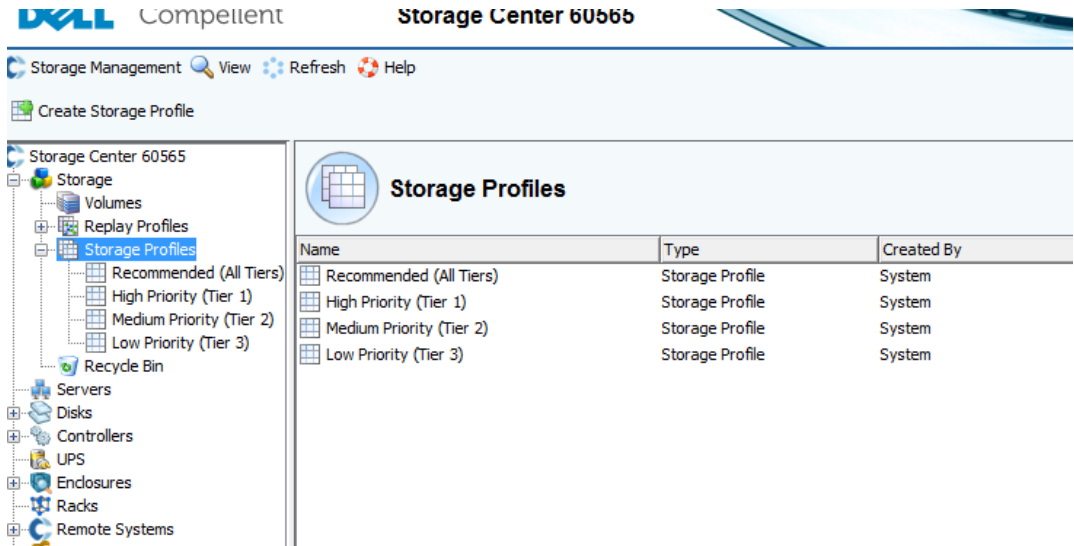


Figure 29. Create a storage profile

2. Enable and disable the **Writable Data** and **Replay Data** check boxes, as shown in Figure 30, and then click **Continue**.

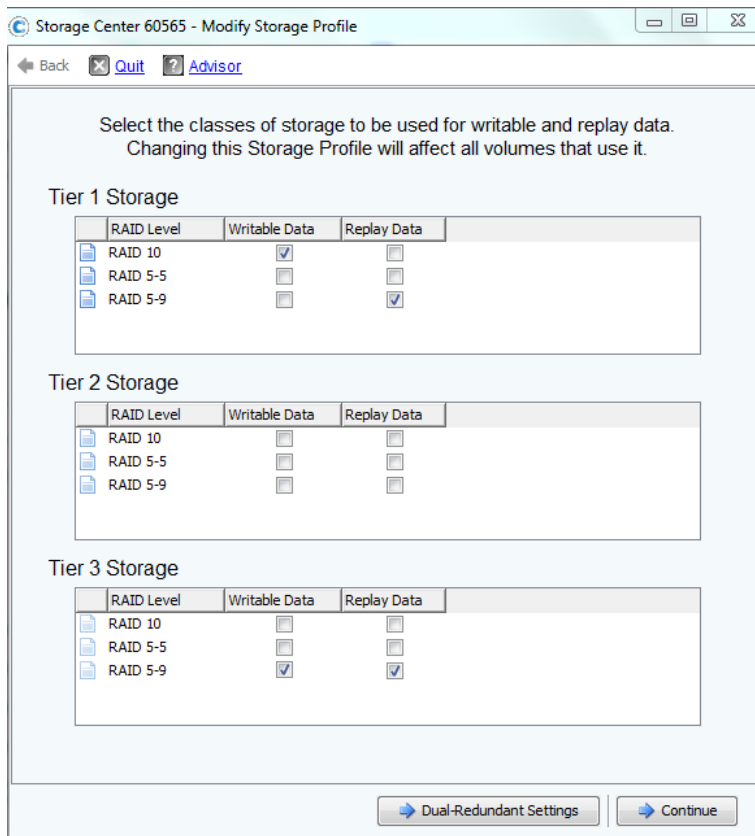


Figure 30. Set Writable Data and Replay Data settings

Note: Because both Tier-1 and Tier-3 storage are configured as Single Redundant, the check boxes shown are for the initial configuration (as Single Redundant). If modified later, the check boxes differ for Single Redundant and Dual Redundant settings. Ignore a Dual Redundant setting because it is not applicable.

3. Type a name for a storage profile, such as **Profile for HANA Vol.**
4. Click **Apply Changes.**

Configuring the Replay profile

To create a Replay profile for the storage array:

1. Click **Create Replay Profile.**
2. Type a name for the replay profile, such as Replay Profile for HANA Vol, and then click **Add Rule**, as shown in Figure 31.

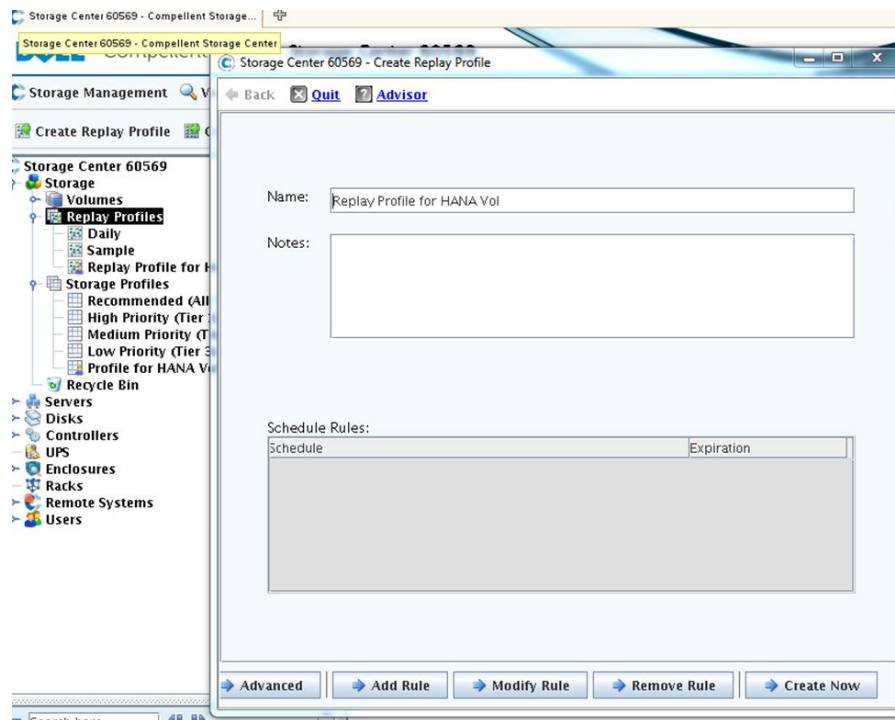


Figure 31. Create the replay profile

3. To set the schedule type, start time, and expiration, do the following:
 - a. From the **Schedule Type** list menu, select **Daily**.
 - b. From the **at** list menu, select **7:00 PM** for the time to start the replay.
 - c. From the **Expiration** list menu, type **1** and select **days**.
 - d. Click **Continue**.
4. Review the schedule rules and then click **Create Now**.
5. Click **Close**.
6. If there is a second array, repeat these steps.

Configuring data progression

To configure the data progression:

1. To open the System Properties page, right-click **Storage Center**.

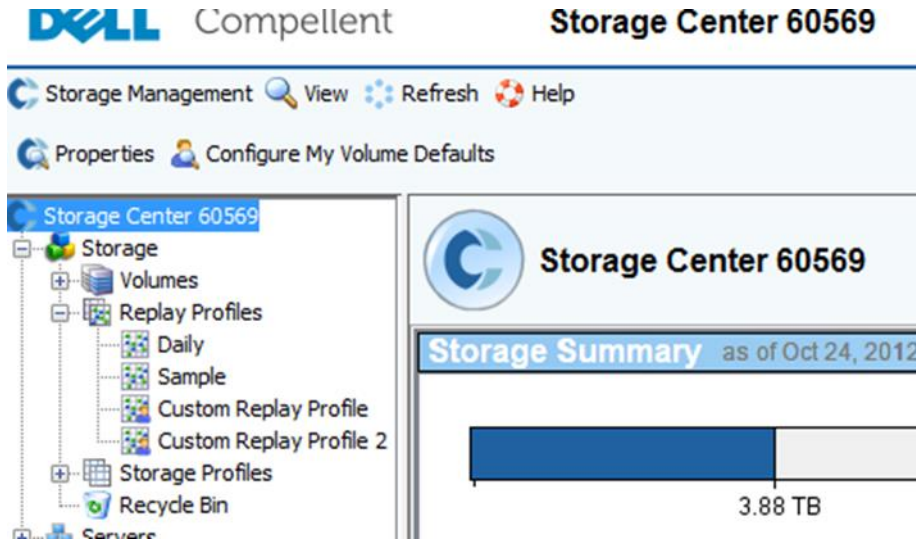


Figure 32.

2. Click **Data Progression**.

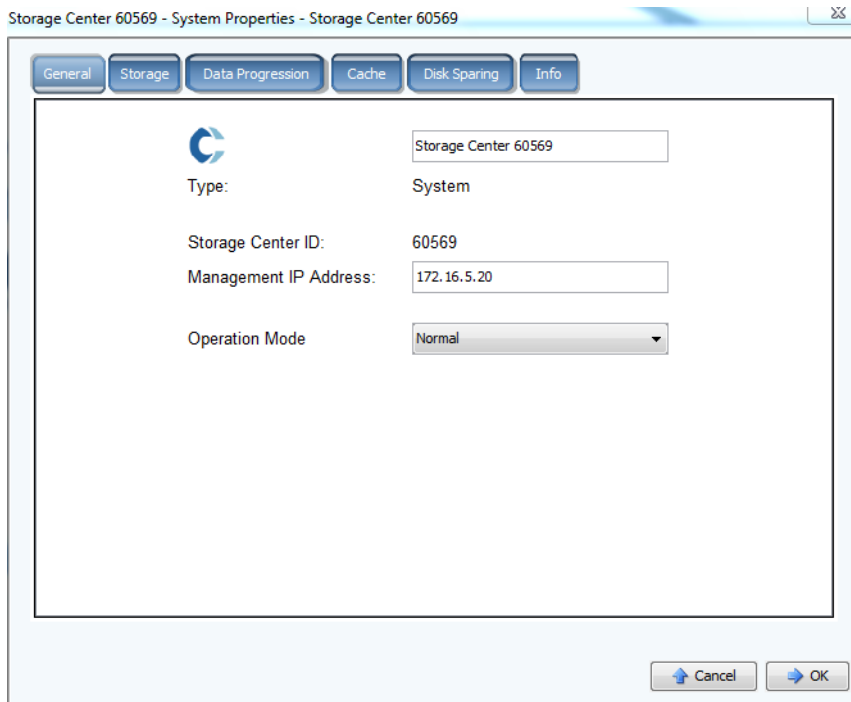


Figure 33. Set data progression

3. Click **OK**.

Setting up volumes

On each storage array, create up to eight data volumes (3 x RAM per node) and eight log volumes (512 GB each) for eight SAP HANA worker nodes. On the first array, create an additional **sapmnt** LUN sized at 512 GB for a set of four SAP HANA nodes, as shown in

Figure 34. For example, for eight nodes, create a **sapmnt** LUN sized at 1 TB. This additional LUN is used for `/hana/shared` as the NFS mount point.

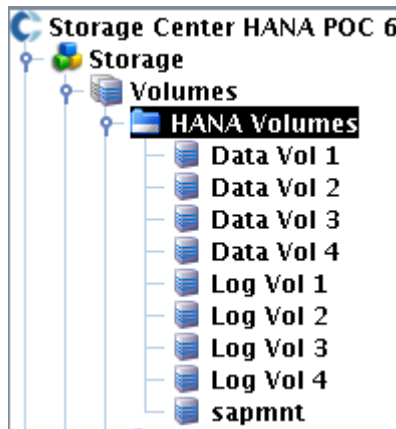


Figure 34. Volume setup

If the deployment does not use the FS8610 NAS appliance for NFS, the LUN is formatted with OCFS2 for an SLES deployment and GFS2 for a RHEL deployment.

Setting up hosts

Create a server cluster that consists of all SAP HANA nodes that access the volumes on the storage array, as shown in Figure 35.

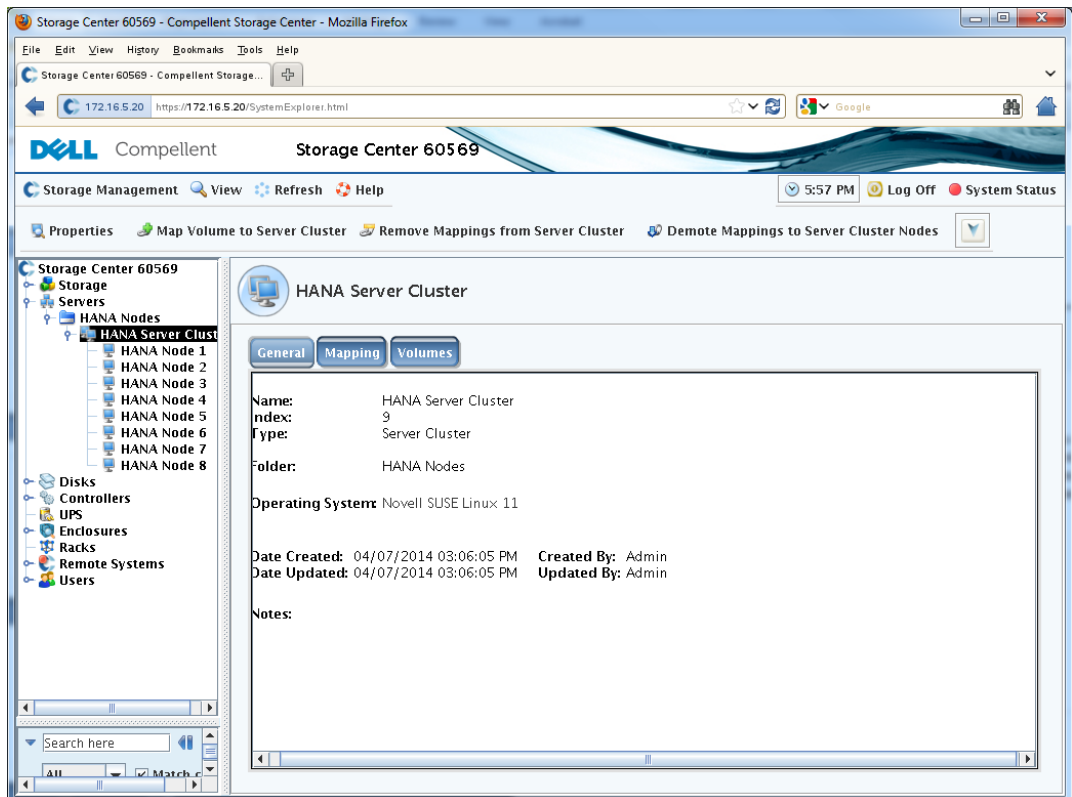


Figure 35. Create a server cluster

