



# SAP HANA Tailored Datacenter Integration with Dell Compellent Storage Series

Dell Global Solutions Engineering  
September 2016

A Dell Reference Architecture

## Revisions

Date	Description
April 2014	Initial release
February 2015	Second release
December 2015	Third release (Haswell E7 and SC9000)
September 2016	Fourth release (Broadwell E7 and SC7020)
February 2017	To leverage Mainstream Read Intensive SSDs Refresh HW Components

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2013 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell™, the Dell logo, Dell Boomi™, Dell Precision™, OptiPlex™, Latitude™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, Compellent™, KACE™, FlexAddress™, Force10™, Benchmark Factory for Databases® and Vostro™ are trademarks of Dell Inc. Other Dell trademarks may be used in this document. Cisco Nexus®, Cisco MDS®, Cisco NX-OS®, and other Cisco Catalyst® are registered trademarks of Cisco System Inc. EMC VNX®, and EMC Unisphere® are registered trademarks of EMC Corporation. Intel®, Pentium®, Xeon®, Core® and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™ and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat® and Red Hat® Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE® are registered trademarks of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Citrix®, Xen®, XenServer® and XenMotion® are either registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries. VMware®, Virtual SMP®, vMotion®, vCenter® and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries. IBM® is a registered trademark of International Business Machines Corporation. Broadcom® and NetXtreme® are registered trademarks of Broadcom Corporation. Qlogic is a registered trademark of QLogic Corporation.



# Table of contents

Revisions .....	2
Acknowledgements .....	5
Executive summary .....	6
1 Introduction .....	7
1.1 Scope .....	7
1.2 Audience .....	7
2 SAP HANA TDI Solution Overview .....	8
2.1 SAP HANA TDI Solution Overview and Best Practices .....	8
2.1.1 Sizing for Performance .....	9
2.1.2 Dell SAP HANA Scale-out Solutions reference configurations .....	10
2.1.3 Sizing Overview (1TB, 1.5TB or 2TB nodes) .....	10
2.1.4 Networking Requirement and Best Practices .....	13
2.2 NFS Approach for /hana/shared volume only .....	14
2.2.1 Modular/Building Block Approach for NFS .....	14
2.2.2 Building Block Details .....	15
2.2.3 Dell FS8610 NAS for shared HANA volume .....	17
2.2.4 Network Configuration .....	18
2.2.5 Network Connection to the data center network infrastructure .....	22
2.2.6 Guidelines for deploying on RedHat Enterprise Linux .....	23
2.2.7 Guidelines for deploying on SUSE Enterprise Linux .....	23
2.2.8 Dell SAP HANA Appliance Hardware Certification Parameters .....	24
2.2.9 Dell SAP HANA Storage Configuration .....	25
2.2.10 SAP HANA Storage Connector API .....	25
2.3 Cluster File System Approach for /hana/shared .....	26
2.3.1 Modular/Building Block Approach for OCFS2 .....	26
2.3.2 Building Block Details .....	27
2.3.3 Dell Compellent SC8000 storage Array Configuration for SAP HANA .....	28
2.3.4 Network Configuration .....	29
2.3.5 Guidelines for deploying on RedHat Enterprise Linux .....	32
2.3.6 Guidelines for deploying on SUSE Enterprise Linux .....	32
2.3.7 Volume layout and mappings for HANA .....	33



3	Dell Compellent Fluid Data Architecture Overview .....	36
3.1	Intelligence .....	36
3.2	Flash Aware Sub-Lun Tiering .....	36
3.3	Policy based data management .....	37
4	Dell Compellent storage configuration overview .....	38
4.1	Switch Zoning Configuration .....	38
4.2	Qlogic Host Bus Adapter driver settings .....	38
4.3	Storage Array Controller Setup .....	38
5	Accessing Compellent SAN from SAP HANA nodes .....	40
5.1	Enable native Linux Multipathing (DM-MPIO) .....	40
5.2	Using XFS file system with SAP Storage Connector API to mount HANA DATA and LOG volumes .....	43
6	Conclusion .....	45
	Additional resources .....	46
A	APPENDIX .....	47
	Configure SLES HA cluster for shared OCFS2 volume .....	47
	Configure CMAN for GFS2 for shared volume for RHEL 6.5 (SC8000) .....	48
B	APPENDIX .....	52
	SC8000 Controller Setup .....	52
	Startup Wizard .....	53
	Load License .....	55
	Create Disk Folder .....	56
	Add Controller .....	59
	Time Setting .....	62
	System Setting .....	63
	Configure SMTP .....	65
	Update Setup .....	66
	User Setup .....	67
	Configure Ports .....	68
	Post Settings for Compellent storage array .....	79
	Configuring Storage Profile .....	80
	Configuring Replay Profile .....	83
	Setting up Volumes .....	90



Setting up Hosts .....	91
------------------------	----

## Acknowledgements

Zafar Mahmood, Dell Global Solutions Engineering

Scott Saktanaset, Dell Storage Solutions Manager

Benjamin Simroth, Dell Global Solutions Engineering



## Executive summary

SAP HANA is the industry's leading in memory data processing platform with advanced capabilities. The initial roll out of SAP HANA was in the form of fully integrated appliances where server, networking and storage components were delivered as one package. This end to end testing and certification ensured that the SAP HANA appliances met the stringent criteria set by SAP contributing in a major way towards the predictability and success of the platform. Despite this, the appliance model is characterized with several drawbacks. The major drawback being that the appliance model limits customer flexibility in choosing server, networking and storage components. Additionally, the appliance model does not enable the use of established data center operational best practices.

The SAP HANA Tailored Datacenter Integration (TDI) program is an additional option to existing SAP HANA appliance model. This model allows connecting certified servers with shared storage resources. The SAP HANA TDI program delivers the following advantages

- Enables the reuse of existing hardware components
- Enables the use of existing operational processes
- Leverage 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.

This paper describes a reference architecture for deploying SAP HANA using Dell PowerEdge servers and Dell Compellent storage. The reference architecture is developed based on extensive validation performed within Dell labs. The configurations described in this whitepaper successfully meets the strict performance and availability requirements set by SAP.



# 1 Introduction

The purpose of this whitepaper is to provide a reference architecture for deploying SAP HANA solution using the Tailored Datacenter Integration option. This option enables customers the following advantages

- Greater degree of flexibility in choosing servers and storage
- Ability to deploy well understood data center operational best practices
- Ability to tailor the storage capacity based on need

## 1.1 Scope

The scope of this paper includes:

- Introduction to key technologies
- Introduction to Dell Fluid Data Architecture
- Steps to integrate Dell Compellent Storage array with SAP HANA nodes

## 1.2 Audience

This document is intended for system/storage admins, customers, partners, system integrators or anyone seeking to configure Dell Compellent storage array in a SAP HANA solution with TDI option.



## 2 SAP HANA TDI Solution Overview

Customers are responsible to deploy these TDI solutions with the help of SAP certified consultants. They are also responsible to configure the systems in terms of CPU, RAM and disk such that the system meets the KPI requirements to implement SAP HANA solutions without running into storage capacity issues.

Table 1 TDI vs Appliance Model

Category	Appliance Delivery Model	SAP HANA TDI Approach
Hardware Selection	Little flexibility only <ul style="list-style-type: none"><li>• Customer can choose between different appliance vendors</li><li>• No option to replace certain HW components already being used in the customer's data center</li></ul>	Save IT budget and existing investments <ul style="list-style-type: none"><li>• Use preferred storage</li><li>• Use preferred network components</li></ul>
Implementation Effort	Low for customer <ul style="list-style-type: none"><li>• Preconfigured HW plus preinstalled SW</li></ul>	Only HW is delivered <ul style="list-style-type: none"><li>• Installation to be done by customer (including the OS)</li><li>• Extensive documentation (guides, SAP notes) available</li></ul>
Solution Validation	Done together by SAP and HW Partner	<ul style="list-style-type: none"><li>• SAP HANA Going-Live check offered by SAP AGS</li><li>• Self-managed infrastructure tests possible<ul style="list-style-type: none"><li>◦ SAP HANA HW Configuration Check Tool</li></ul></li></ul>
Support	Fully provided by SAP	Individual support agreement(s) with HW partner(s) required
OS Service Contract	Appliance vendor is the reseller of OS service contract	Customer to take care for OS service contract

### 2.1 SAP HANA TDI Solution Overview and Best Practices

SAP HANA Tailored Datacenter Integration deployments are customized solutions where you can choose any of the certified SAP HANA server vendors along with any certified SAP enterprise storage.

This reference architecture guide describes the recommended storage design of Dell Compellent SC4020, SC7020, SC8000 and SC9000 storage array for use in SAP HANA Tailored Datacenter Integration for up to 16 HANA worker node solutions.

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





When implementing SAP HANA solutions using Dell CML storage arrays, following considerations should be observed:

- SAP certifies the minimum enterprise storage layout only, as described in this document.
- The SAP HANA Storage Connector API for Fiber Channel is the only supported API for use with this solution.
- Each Dell Compellent SC4020 storage array may have a maximum of 8 active SAP HANA nodes.
- Each Dell Compellent SC8000 storage array may have a maximum of 8 active SAP HANA nodes.
- Each Dell Compellent SC7020 storage array may have a maximum of 8 active SAP HANA nodes.
- Each Dell Compellent SC9000 storage array may have a maximum of 16 active SAP HANA nodes.
- A 16 node SAP HANA solution will have a minimum of 1 Dell Compellent SC9000 or 2 Dell Compellent SC4020/SC7020/SC8000 storage arrays

The Dell SAP HANA TDI solution based on SC storage family may use the following options:

- A single Dell Compellent SC4020, SC7020, OR SC8000 storage array for up to 8 active SAP HANA nodes with one standby according to the disk sizing as listed in table [2.1.3](#). The SC9000 can handle up to 16 active SAP HANA Nodes with one or two standby nodes.
- A second Dell Compellent SC4020, SC7020 OR SC8000 storage array for up to 16 active SAP HANA Nodes with one or two standby nodes.
- Build the storage system with single disk folder, single storage profile for HANA, volumes, disk storage.
- Use following building block approach to size Dell Compellent storage arrays for HANA nodes

### 2.1.1 Sizing for Performance

The following sizing criteria, must be observed while building SAP HANA solutions using Dell Compellent storage arrays:

- **CPU core to Memory ratio**
  - 1TB OR 1.5TB / 88 CPU cores (8880/8890 E7 v4) with HANA SPS10 or lower
  - 2TB / 88/96 CPU cores (8880/8890 E7 v4) with HANA SPS11 or higher
- **Storage requirements (external shared storage)**
  - Log volume/HANA node : 512GB (SSD RAID)
  - Data volume/HANA node: 3 x total memory (SAS drives / FC drives)
  - Shared volume: 1 x total memory per node or since HANA SPS11 1 x total memory per four nodes
- **Network requirements**
  - 1x Intel Ethernet X540 DP 10Gb BT for Backup network+ I350 1Gb BT DP Network Daughter Card for HANA studio/modeler for management network
  - 1x Dual Port Intel X540 for HANA cluster network or NFS network (bond2)



- 1x Dual Port Intel X540 for Internal HANA internal network (bond0)
- 1x Dual Port Intel X540 for external sources ERP, BI Platform, HANA & data sources – PROD (Bond1)
- Dedicated FC16 SAN between HANA nodes & storage
- Redundant infrastructure for Failover

## 2.1.2 Dell SAP HANA Scale-out Solutions reference configurations

Following information is provided here as a guideline to build SAP HANA solutions using Dell Compellent storage arrays:

### (2-32TB) Configuration

- Server: 2-16 x R930 workers + 1 or 2 x R930 for standby
- CPU: 4x Intel Xeon (BDW-EX): E7-8880 v4 2.3GHz OR E7-8890 v4 2.2GHz
- Memory: 1TB, 1.5TB or 2TB / node
- OS: SLES 11 SP4 for SAP or RHEL 6.7
- 1-2x Compellent Storage Arrays with SSD and HDDs with the configuration options below.

A list of all certified OS configurations could be found at <http://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/appliances.html#categories=Dell,Scale-out&viewcount=100> and all possible storages are listed at <http://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/enterprise-storage.html#categories=Dell>

## 2.1.3 Sizing Overview (1TB, 1.5TB or 2TB nodes)

Active Nodes	Compellent SC4020	Compellent SC8000	Compellent SC7020/SC7020F	Compellent SC9000
2	6 x 480GB RI SSD + 12 x 1.8TB 10K RPM SAS	7 x 480GB RI SSD + 14 x 1.8TB 10K RPM SAS	6 x 480GB RI SSD + 12 x 1.8TB 10K RPM SAS 12 x 1.9GB RI SSD (for SC7020F)	6 x 480GB RI SSD + 12 x 1.8TB 10K RPM SAS
3	6 x 480GB RI SSD + 24 x 1.8TB 10K RPM SAS	7 x 480GB SSD + 14 x 1.8TB 10K RPM SAS	6 x 480GB RI SSD + 18 x 1.8TB 10K RPM SAS 18 x 1.9GB RI SSD (for SC7020F)	6 x 480GB RI SSD + 12 x 1.8TB 10K RPM SAS
4	6 x 480GB RI SSD + 24 x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 24 x 1.8TB 10K RPM SAS	6 x 480GB RI SSD + 24 x 1.8TB 10K RPM SAS 24 x 1.9GB RI SSD (for SC7020F)	6 x 480GB RI SSD + 24 x 1.8TB 10K RPM SAS
5	12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 24 x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS 30 x 1.9GB RI SSD (for SC7020F)	12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS



<b>6</b>	12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS	18 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS <b>36 x 1.9GB RI SSD (for SC7020F)</b>	12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS
<b>7</b>	12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS	18 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS <b>36 x 1.9GB RI SSD (for SC7020F)</b>	12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS
<b>8</b>	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS	24 x 480GB RI SSD + 48x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS <b>42 x 1.9GB RI SSD (for SC7020F)</b>	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS

<b>Active Nodes</b>	<b>Compellent SC4020</b>	<b>Compellent SC8000</b>	<b>Compellent SC7020</b>	<b>Compellent SC9000</b>
<b>9</b>	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS  6 x 480GB RI SSD + 12 x 1.8TB 10K RPM SAS	24 x 480GB RI SSD + 48 x 1.8TB 10K RPM SAS  7 x 480GB RI SSD + 14 x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS <b>42 x 1.9GB RI SSD (for SC7020F)</b>  6 x 480GB RI SSD + 12 x 1.8TB 10K RPM SAS <b>12 x 1.9GB RI SSD (for SC7020F)</b>	18 x 480GB RI SSD + 48 x 1.8TB 10K RPM SAS
<b>10</b>	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS  6 x 480GB RI SSD + 12 x 1.8TB 10K RPM SAS	24 x 480GB RI SSD + 48 x 1.8TB 10K RPM SAS  7 x 480GB RI SSD + 14 x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS <b>42 x 1.9GB RI SSD (for SC7020F)</b>  6 x 480GB RI SSD + 12 x 1.8TB 10K RPM SAS <b>12 x 1.9GB RI SSD (for SC7020F)</b>	18 x 480GB RI SSD + 48 x 1.8TB 10K RPM SAS
<b>11</b>	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS  6 x 480GB RI SSD + 12 x 1.8TB 10K RPM SAS	24 x 480GB RI SSD + 48 x 1.8TB 10K RPM SAS  12 x 480GB RI SSD + 24 x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS <b>42 x 1.9GB RI SSD (for SC7020F)</b>  6 x 480GB RI SSD + 12 x 1.8TB 10K RPM SAS <b>18 x 1.9GB RI SSD (for SC7020F)</b>	18 x 480GB RI SSD + 48 x 1.8TB 10K RPM SAS
<b>12</b>	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS  6 x 480GB RI SSD + 24 x 1.8TB 10K RPM SAS	24 x 480GB RI SSD + 48 x 1.8TB 10K RPM SAS  12 x 480GB RI SSD + 24 x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS <b>42 x 1.9GB RI SSD (for SC7020F)</b>  6 x 480GB RI SSD + 24 x 1.8TB 10K RPM SAS <b>24 x 1.9GB RI SSD (for SC7020F)</b>	18 x 480GB RI SSD + 60 x 1.8TB 10K RPM SAS
<b>13</b>	12 x 480GB SSD + 36 x 1.8TB 10K RPM SAS  12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS	24 x 480GB RI SSD + 48 x 1.8TB 10K RPM SAS  18 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS <b>42 x 1.9GB RI SSD (for SC7020F)</b>  12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS <b>30 x 1.9GB RI SSD (for SC7020F)</b>	24 x 480GB RI SSD + 66 x 1.8TB 10K RPM SAS



<b>14</b>	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS  12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS	24 x 480GB RI SSD + 48 x 1.8TB 10K RPM SAS  18 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS 42 x 1.9GB RI SSD (for SC7020F)  12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS 36 x 1.9GB RI SSD (for SC7020F)	24 x 480GB RI SSD + 66 x 1.8TB 10K RPM SAS
<b>15</b>	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS  12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS	24 x 480GB RI SSD + 48 x 1.8TB 10K RPM SAS  24 x 480GB RI SSD + 48 x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS 42 x 1.9GB RI SSD (for SC7020F)  12 x 480GB RI SSD + 30 x 1.8TB 10K RPM SAS 36 x 1.9GB RI SSD (for SC7020F)	24 x 480GB RI SSD + 66 x 1.8TB 10K RPM SAS
<b>16</b>	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS  12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS	24 x 480GB RI SSD + 48 x 1.8TB 10K RPM SAS  24 x 480GB RI SSD + 48 x 1.8TB 10K RPM SAS	12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS 42 x 1.9GB RI SSD (for SC7020F)  12 x 480GB RI SSD + 36 x 1.8TB 10K RPM SAS 42 x 1.9GB RI SSD (for SC7020F)	24 x 480GB RI SSD + 72 x 1.8TB 10K RPM SAS

#### First Compellent Array Pair

#### Second Compellent Array Pair

#### Recommended number of Storage Arrays (dual controller configuration)

# Nodes	SC4020	SC8000	SC7020	SC9000
2-8	1	1	1	1
9-16	2	2	2	1

#### Recommended number of SC220/SC420 disk shelf's

# Nodes	SC4020	SC8000	SC7020	SC9000
2-3	1xSC220	3xSC220	0xSC420	3xSC420
4	1xSC220	3xSC220	0 xSC420	3xSC420
5	1xSC220	3xSC220	1 xSC420	3xSC420
6-8	1xSC220	3xSC220	1 xSC420	3xSC420
9-11	2xSC220	6xSC220	1 xSC420	6xSC420
12	2xSC220	6xSC220	1 xSC420	6xSC420
13	2xSC220	6xSC220	1 xSC420	6xSC420
14-16	2xSC220	6xSC220	2 xSC420	6xSC420

One SC220 or SC420 disk shelf can hold up to 24 x 2.5" drives



## 2.1.4 Networking Requirement and Best Practices

For reasons of availability, redundant data paths are recommended for SAP HANA scale-out appliance for both Ethernet networks and SAN fabric.

### 2.1.4.1 Ethernet Network

- Management/Administrative network is reserved for external administration tool. 1GbE network connection would be sufficient.
- For Internal Zone Networks, 10Gbit/s non-blocking switches are required as well as one fully redundant 10GbE network.
- 10GbE network connection is required for Client Zone (bond1) Network.
- Addition 10GbE network connection is also required for HANA system replication network since SAP HANA system replication is a solution that involves the transfer of large amount of IO traffic across two distanced data centers. The same principles apply to the HANA internode network as to the network used by SAP HANA system replication.
- We recommend to use jumbo frame settings (MTU=9000) for all involved 10GbE network connections in SAP HANA scale-out appliance

### 2.1.4.2 SAN Fabric

When using Fibre Channel for front-end connectivity, zones must be established to ensure that storage is visible to the servers. Use the zoning concepts discussed in this section to plan the front-end connectivity before starting to cable the storage system.

Zoning can be applied to either the ports on switches or to the World Wide Names (WWNs) of the end devices. It is recommended to use WWN zoning and to create zones using a single initiator host port and multiple Storage Center ports

#### WWN Zoning Guidelines

When WWN zoning is configured, a device may reside on any port or change physical ports and still be visible because the switch is seeking a WWN.

Connection Type	Guidelines
Virtual port mode (Recommended)	<ul style="list-style-type: none"><li>• Include all Storage Center virtual WWNs in a single zone.</li><li>• Include all Storage Center physical WWNs in a single zone.</li><li>• 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</li></ul>
Legacy mode	<ul style="list-style-type: none"><li>• Include all Storage Center front-end WWNs or ports in a single zone.</li><li>• 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</li></ul>



### For Fibre Channel replication:

- 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.

## 2.2 NFS Approach for /hana/shared volume only

HANA nodes require /hana/shared volume across nodes. This can either be achieved by using a cluster file system such as OCFS2 on SLES or GHF2 on RHEL. Or this can be achieved by providing an NFS share mounted on /hana/share across all HANA nodes. This section lists requirements to use NFS option. This option requires Dell FS8610 NAS appliance which is used to export this NFS share from a volume created on CML storage array.

### 2.2.1 Modular/Building Block Approach for NFS

- **Module1:Infrastructure Rack**

#### **Network switches:**

- 2 x Force 10 S4048T 48x10gE ports for internal HANA network + external HANA network, HANA shared NFS network using VLAN's
- 1 x Force 10 S3048, 48 x 1GbE ports – management network

#### **SAN Fabric:**

2 x Brocade 6510 24-36-48 Port FC8/16 Switch

#### **Management:**

1xPER630 as management node + 1 x KMM console

#### **Module2: Servers**

- R930 72 cores/1.5 TB or since HANA SPS11 2 TB
- Intel Ethernet X540 DP 10Gb BT + I350 1Gb BT DP Network Daughter Card
- 3 x Intel10GbE dual port RJ45 (X540-T2)



- 2 x Qlogic 2660 Single Port 16Gb Optical Fiber Channel HBA

### Module 3: Compellent Array


- The sizing is described in table [2.1.3](#)

## 2.2.2 Building Block Details

### 2.2.2.1 3 TB or 4 TB (HANA SPS11 or later) Configuration NFS

Component	Configuration
HANA Nodes	<ul style="list-style-type: none"> <li>• 2 x Dell PowerEdge R930 active HANA nodes + 1 or 2 x R930 Standby HANA node each with:</li> <li>• SLES 11 SP3 for SAP or RHEL 6.6</li> <li>• 1.5TB RAM or 2 TB with HANA SPS11 or later</li> <li>• 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</li> </ul>
Management	1 x Dell PowerEdge R630 with: <ul style="list-style-type: none"> <li>• SLES 11 SP3 for SAP or RHEL 6.6</li> <li>• 1xForce10 S3048 for Infrastructure management network</li> <li>• SAP HANA Studio</li> </ul>
Storage	<ul style="list-style-type: none"> <li>• FS8610 NFS appliance for /hana/shared</li> <li>• Table <a href="#">2.1.3</a></li> </ul>
Storage Fabric	2 x Brocade 6510 24-36-48 Port FC16 Switch
Network	<ul style="list-style-type: none"> <li>• 2 x Force 10 S4048T 48x10gE ports for internal HANA network + external HANA network (PROD) + NFS network using VLAN's for up to 8 HANA active nodes. Addition pair of 4048T switches is required for addition HANA active node.</li> <li>•</li> <li>• 1 x Force 10 S3048, 48 x 1GbE ports – management network</li> </ul>






NOTE: The diagram shows SC4020. Any storage option from [Table 2.1.3](#) may be selected

### 2.2.2.2 24 TB or 32 TB (HANA SPS11 or later) Configuration NFS

Component	Configuration
HANA Nodes	<ul style="list-style-type: none"> <li>2 x Dell PowerEdge R930 active HANA nodes + 1 or 2 x R930 Standby HANA node each with:</li> <li>SLES 11 SP3 for SAP or RHEL 6.6</li> <li>1.5TB RAM or 2 TB with HANA SPS11 or later</li> <li>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</li> </ul>
Management	1 x Dell PowerEdge R630 with: <ul style="list-style-type: none"> <li>SLES 11 SP3 for SAP or RHEL 6.6</li> <li>1xForce10 S3048 for Infrastructure management network</li> <li>SAP HANA Studio</li> </ul>
Storage	<ul style="list-style-type: none"> <li>FS8610 NFS appliance for /hana/shared</li> <li>Table <a href="#">2.1.3</a></li> </ul>
Storage Fabric	2 x Brocade 6510 24-36-48 Port FC16 Switch
Network	<ul style="list-style-type: none"> <li>2 x Force 10 S4048T 48x10gE ports for internal HANA network + external HANA network (PROD) using VLAN's</li> <li>1 x Force 10 S3048, 48 x 1GbE ports – management network</li> </ul>






NOTE: The diagram shows SC4020. Any storage option from [Table 2.1.3](#) may be selected

### 2.2.3 Dell FS8610 NAS for shared HANA volume

- **FS8610 Specification:**
  - 2U chassis NAS Appliance with dual redundant controllers
  - 10Gb Ethernet client connectivity
  - 8Gb Fibre Channel back-end connectivity to the SC4020
- **NFS based /hana/share for Scale-out**
  - The Storage Center provides the FS8600 scale-out NAS storage capacity which is configured to be equal to the size of the appliance
  - The SAN network connects the FS8600 to the Storage Center and carries the block level traffic
  - Another FS8610 and an SC4020 may be added to provide NFS shares for backup/restore solution
  - Can be scale-up to four FS8600 appliances configured as a FluidFS cluster
  - CLI and optional EM GUI for management of NAS pools and exports.

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



The recommended mount options for /hana/shared in HANA SPS09, 10 and 11 are listed in [SAP note 2099253](#)

Please refer to the following guides for FS8610 setup for HANA nodes /hana/shared volume NFS setup:

[ftp://ftp.dell.com/Manuals/all-products/esuprt\\_ser\\_stor\\_net/esuprt\\_compellent/dell-compellent-fs8600\\_Deployment%20Guide\\_en-us.pdf](ftp://ftp.dell.com/Manuals/all-products/esuprt_ser_stor_net/esuprt_compellent/dell-compellent-fs8600_Deployment%20Guide_en-us.pdf)

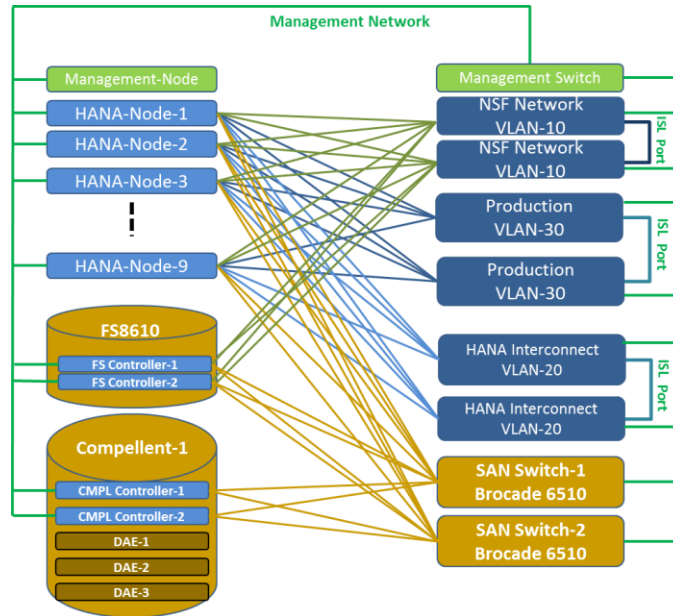
[ftp://ftp.dell.com/Manuals/all-products/esuprt\\_ser\\_stor\\_net/esuprt\\_compellent/dell-compellent-fs8600\\_Reference%20Guide\\_en-us.pdf](ftp://ftp.dell.com/Manuals/all-products/esuprt_ser_stor_net/esuprt_compellent/dell-compellent-fs8600_Reference%20Guide_en-us.pdf)

## 2.2.4 Network Configuration

SAP HANA 3.0 Scale-out appliance can support up to 16 active and 2 stand-by HANA nodes. It includes the following networks.

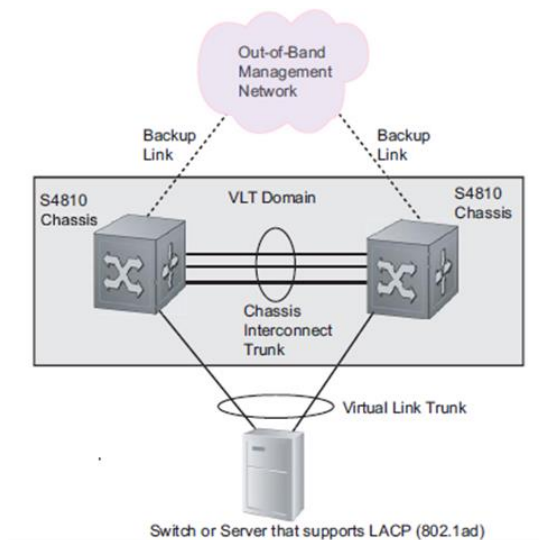
1. **Management network** F10 S3048 switch (1G). All hardware components can be managed via this network by a dedicated management node, a R630. Dell iDRAC network is also part of the management network. It uses the onboard I350 1Gb BT DP Network Daughter Card
2. **HANA related networks**: one pair of 10G Dell F10 S4048T in a Virtual Link Trunking (VLT) domain to support redundant data paths for different HANA related networks. Each S4048T has 48 x 10GBase T Ports + 4 x 40Gb QSFP+ ports
  - VLAN-10: dedicated HANA shared NFS network – bond2 network: For configuration larger than 8+2 HANA nodes, additional pair of S4048T switch required

- VLAN-20: dedicated HANA Internal network – bond0 network
  - VLAN-30: HANA prod/external network – bond1 network
  - Other available ports can be used for other network such as backup or required, to create a dedicated system replication network using the Intel Ethernet X540 DP 10Gb on the daughter card
4. **Storage Area Network:** dedicated and redundant FC16 SAN fabric to host all database related volumes, using one pair of FC16 Brocade 6510 switches



#### 2.2.4.1 Force10 Virtual Link Trunking (VLT):

- Virtual Link Trunking feature is a mechanism to provide **layer 2 multi-pathing** between access network devices (switches or servers) and the core network.
- VLT is a Dual Active Control Plane implementation of Multi System LAG
- VLT provides a **loop-free topology with active-active load-sharing** of links from access to core



#### VLT domain configuration:

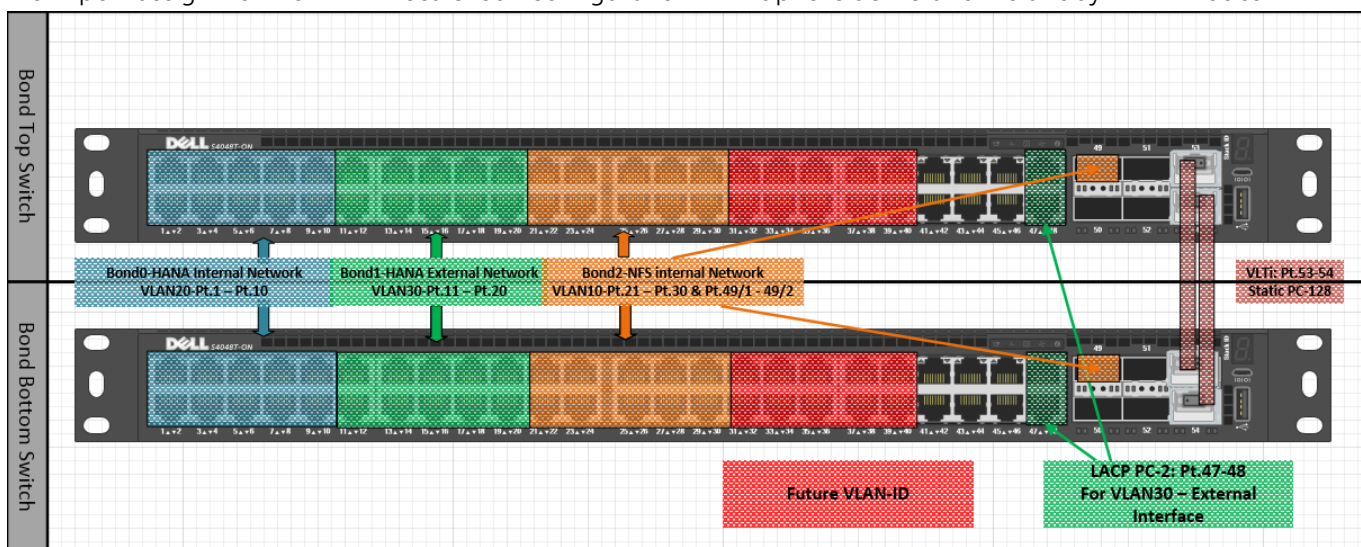
- VLTi (Peer-link): Static port channel 128 (40Gb ports 56 – 60) for control and data exchange
- VLT Health Check: management port (OOB Port). Periodic keep alive message

#### 2.2.4.2 HANA related networks:

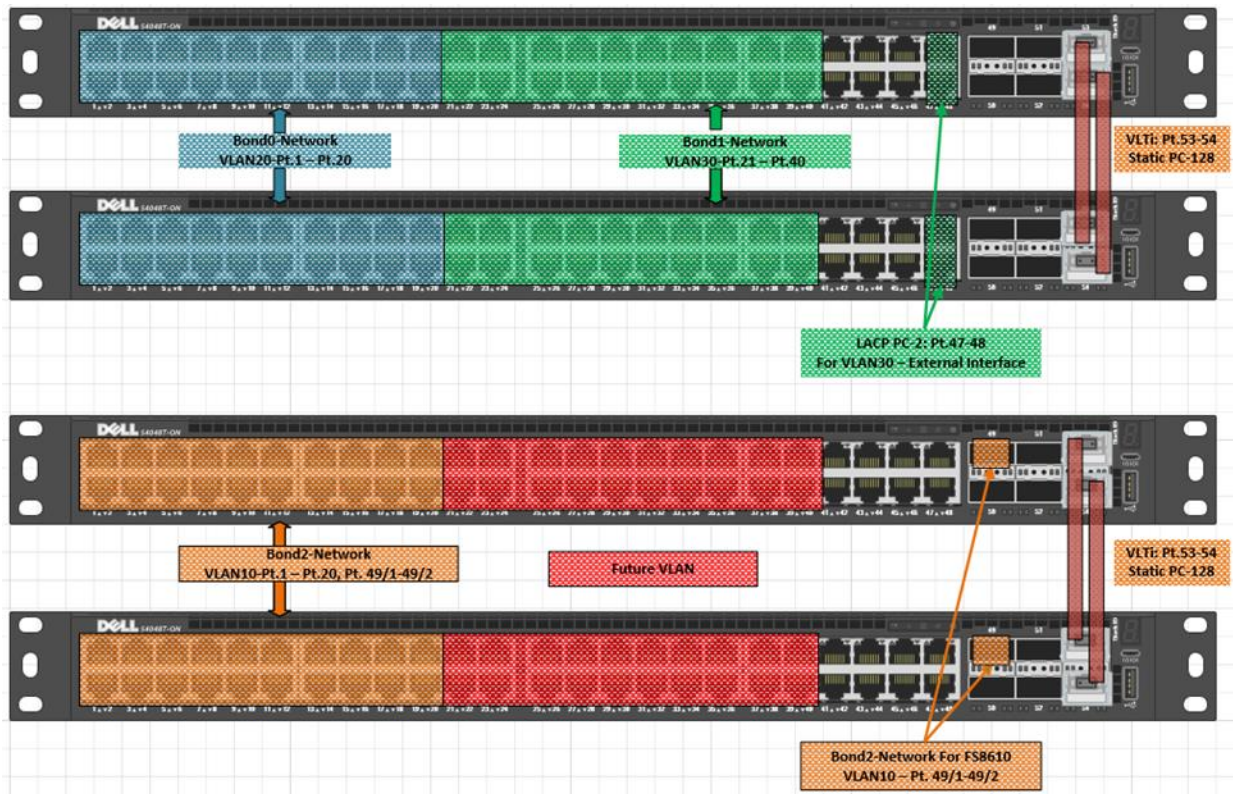
One pair of F10 S4048T is configured in a VLT domain to host HANA related networks via different VLAN-IDs. It can support up to 8 active and 2 stand-by HANA nodes.

- Additional pair of F10 S4048T is required to support a HANA Scale-out appliance with more than 8 active HANA nodes.

Below is network port assignment for HANA scale-out configuration with up to 8 active and 2 standby HANA nodes




Two pairs of F10 S4048T switches are required for HANA scale-out configuration with more than 8 active HANA nodes as shown below.




Below are detailed illustrations of network port assignment for HANA scale-out appliance with up to 8 active and 2 standby HANA nodes.

- **Dedicated HANA Internal Network:**
  - Bond-0: VLAN-20 – port 1- 10






- HANA External/Prod Network:
  - Bond-1: VLAN-30 - port 11-20
  - Client and System Replication networks



- Dedicated NFS Internal Network:
  - Bond-2: VLAN-10 - port 21- 30 & 49/1 - 49/2



- Future Network such as backup or dedicated system replication Network:
  - Bond-X: VLAN-ID - other available ports


## 2.2.5 Network Connection to the data center network infrastructure

### Bond1 Network: External Network:

- VLAN-30: 10Gb port 11 - 20
- LACP port channel 2 for External Interface: 10Gb port 47-48
- vlt-peer-lag port-channel 2
- Tagged PC 2

### External switch:

- 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)



## 2.2.6 Guidelines for deploying on RedHat Enterprise Linux

Please refer to the following guidelines from SAP and Red Hat while deploying HANA:

Introduction to Red Hat Enterprise Linux for SAP HANA

<http://www.redhat.com/en/resources/red-hat-enterprise-linux-sap-hana>

SAP HANA DB: Recommended OS settings for RHEL 6.5

<http://help.sap.com/disclaimer?site=https://service.sap.com/sap/support/notes/2009879>

SAP HANA Server Installation and Update Guide

[http://help.sap.com/hana/SAP\\_HANA\\_Server\\_Installation\\_Guide\\_en.pdf](http://help.sap.com/hana/SAP_HANA_Server_Installation_Guide_en.pdf)

## 2.2.7 Guidelines for deploying on SUSE Enterprise Linux

Introduction to SUSE Enterprise Linux for SAP HANA

<https://www.suse.com/products/sles-for-sap/>

SAP HANA Guidelines for SLES Operating System

<http://help.sap.com/disclaimer?site=https://service.sap.com/sap/support/notes/1944799>

SAP HANA Server Installation and Update Guide

[http://help.sap.com/hana/SAP\\_HANA\\_Server\\_Installation\\_Guide\\_en.pdf](http://help.sap.com/hana/SAP_HANA_Server_Installation_Guide_en.pdf)

## 2.2.8 Dell SAP HANA Appliance Hardware Certification Parameters

Before installing HANA please make sure that the setup is verified with the SAP HANA Hardware Configuration Check Tool (HWCCT).

The HWCCT tool is used in the following scenarios:

- The tool must be run before productive use
- In case the system configurations is changed for any reason, it should be run again to verify that that the required KPI numbers are still within required thresholds SHOULD be used after each change in the landscape
- The tool can also be used for regular health checks

The minimum KPI figures and instructions about how to run and configure HWCCT can be found in SAP Note 1943937.

When validating the environment with the HWCCT tool 1.1 (fsperf) use these parameters in the profile.

`async_write_submit_active=on`

`async_write_submit_blocks=all`

`async_read_submit=on`

`max_parallel_io_requests=256`

### **Configuring HANA 1.0 with HWCCT Parameters:**

Please configure the HANA 1.0 database with hdbparm tool with the parameters used during HWCCT validation provided above. The needed steps are described in:

<http://service.sap.com/sap/support/notes/2267798> and

<http://service.sap.com/sap/support/notes/2186744>.

### **Configuring HANA 1.0 with HWCCT Parameters:**





With HANA2, the hdbparam tool has been deprecated. It is no longer installed as part of HANA. The above mentioned HWCCT parameters used for validation should be incorporated into the HANA database environment using global.ini file as described in SAP NOTE:

<https://launchpad.support.sap.com/#/notes/2399079>

## 2.2.9 Dell SAP HANA Storage Configuration

- **One Compellent SC4020/SC8000/SC7020 Array** can serve up to **8 HANA nodes (16TB) since HANA SPS11)**
- One **Compellent SC9000 Array can be configured to** serve up to **8 HANA nodes (16TB) since HANA SPS11)**
  - 1 x SC4020/SC8000/SC7020 dual controller array in redundant configuration
  - 4 x SAS, 12Gb, 4 Wide-Port, PCI-e backend (2 per controller)
  - 8Gb/16Gb FC ports (4 per controller)
  - For disk population and count, please refer to [Table 2.1.3](#)
  - Dell Compellent Fluid Data Tiers
- Single disk folder  $\approx$  62TB useable space/per array
- **Required Volume Space (8 x 2TB/node):**
  - 8x512GB LOG volumes/per array  $\approx$  4 TB
  - 8x6TB DATA volumes/per array  $\approx$  48 TB
  - 1x /hana/shared volume/appliance on array1  $\approx$  2TB for 4 nodes = 2x2TB = 4TB for 8x2TB nodes
  - Total required by volumes
    - 4TB (LOG) + 48TB (DATA) +4TB (/hana/shared)  $\approx$ 56TB for 8x2TB node appliance
  - All writes  $\rightarrow$  tier1 SSD's (based on storage profile)

## 2.2.10 SAP HANA Storage Connector API

- All DATA and LOG volumes use XFS
- Only the owner node mounts its DATA and LOG volume
- Standby node only mounts /hana/shared (NFS)
- SAP HANA storage connector uses SCSI3 PGR mechanism to guarantee FS integrity
- Reservation released upon node failure or graceful shutdown for the SB node to register  $\rightarrow$  acquire reservation on the orphaned DATA and LOG volumes



### 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
```

## 2.3 Cluster File System Approach for /hana/shared

There are two cluster file system supported with cluster FS based Dell|SAP HANA scale-out appliance using SC8000 storage arrays only. OCFS2 is for SLES 11 SP3 and GFS2 is for RHEL 6.5. Appendix A describes further details about OCFS2 and GFS2 cluster setup procedures.

### 2.3.1 Modular/Building Block Approach for OCFS2

- **Module1:Infrastructure Rack**

**Network switches:**

- 2 x Force 10 S4820T 48x10gE ports for internal HANA network + external HANA network
- (PROD) using VLAN's + 1 x Force 10 S3048, 48 x 1GbE ports – management network + 2 x Force 10 S60, 48 x 1GbE ports – cluster network

**SAN Fabric:**

2 x Brocade 6510 24-36-48 Port FC8/16 Switch

**Management:**

1xPER630 as management node + 1 x KMM console

**Module2: Servers**



- R930 72 cores/1.5TB
- Intel Ethernet X540 DP 10Gb BT + I350 1Gb BT DP Network Daughter Card
- 3 x Intel10GbE dual port RJ45 (X540-T2)
- 2 x Qlogic 2660 Single Port 16Gb Optical Fiber Channel HBA

#### Module 3: Compellent Array


- One Compellent Array for 8 HANA active nodes (8TB) + 1 HANA standby node

## 2.3.2 Building Block Details

### 2.3.2.1 3 TB Configuration OCFS2/GFS2

Component	Configuration
HANA Nodes	<ul style="list-style-type: none"> <li>• 2 x Dell PowerEdge R930 active HANA nodes + 1 or 2 x R930 Standby HANA node each with:</li> <li>• SLES 11 SP3 for SAP OR RHEL 6.6</li> <li>• 1.5TB RAM</li> <li>• 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</li> </ul>
Management	1 x Dell PowerEdge R630 with: <ul style="list-style-type: none"> <li>• SLES 11 SP3 for SAP or RHEL 6.6</li> <li>• 1xForce10 S60 for Infrastructure management network</li> <li>• SAP HANA Studio</li> </ul>
Storage	1 x Dell Compellent SC8000, 6Gb SAS Dual Controller Array with: <ul style="list-style-type: none"> <li>• 2 x SAS, 6Gb, 4 Wide-Port, PCI-e backend/per controller (dual chain)</li> <li>• 2 x FC16, 2-Port QLE2662, PCI-e front end/per controller</li> </ul>
	<ul style="list-style-type: none"> <li>• 1 x 24 bay 2.5" 400G SAS SSD with 7 total disks/per array</li> <li>• 2 x 24 bay 2.5" 1.8TB 10K RPM SAS with 14 (7+7) spindles/per array</li> </ul>
Storage Fabric	2 x Brocade 6510 24-36-48 Port FC16 Switch
Network	<ul style="list-style-type: none"> <li>• 2 x Force 10 S4820T 48x10gE ports for internal HANA network + external HANA network (PROD) using VLAN's</li> <li>• 2 x Force 10 S60, 48 x 1GbE ports – Cluster network</li> <li>• 1 x Force 10 S60, 48 x 1GbE ports – management network</li> </ul>






### 2.3.3 Dell Compellent SC8000 storage Array Configuration for SAP HANA

- **One Compellent Array can attach up to 8 HANA nodes (8TB)**
  - 2 x SC8000 controllers with 64GB RAM each in redundant configuration
  - 2 x SAS, 6Gb, 4 Wide-Port, PCI-e backend/per controller
  - 2 x FC16, 2-Port QLE2662, PCI-e front end/per controller
  - 2 x 24 bay 2.5" 1.8TB 10K RPM SAS with up to 48 spindles/per array
  - 1 x 24 bay 2.5" 400G SAS SSD with up to 24 total disks/per array
- **Dell Compellent Fluid Data Tiers for fully populated array**
  - Tier1: 24x 400GB SSD (1HS) R10 ≈ 4.8TB useable space/per array
  - Tier2: 0x15K RPM SAS: N/A
  - Tier3: 48x1.8TB 10K SAS R5 ≈ 72TB
  - Single disk folder ≈ 76TB (1.8TB SAS) useable space/per array
  - All writes → tier1 SSD's (based on storage profile)
- **Required Volume Space (example for 8x1.5TB HANA worker nodes):**
  - 8x512GB LOG volumes/per array ≈ 4 TB
  - 8x4.5TB DATA volumes/per array ≈ 36 TB

- 1x /hana/shared volume/appliance on array1 ≈ total appliance size 12TB for 1.5TB nodes
- Total required by volumes
- ≈ 4TB+ 36TB+12TB(/hana/shared) ≈ 52TB for 8x1.5TB node appliance
- All writes go to tier1 SSD's (based on storage profile)

## DATA PROGRESSION

- Data Progression's data mover runs once per day
  - Default start time is 7pm
  - Data is moved per page
  - Historical Replay pages eligible to move to lowest tier immediately
  - Default setting – 12 days down, 3 up
  - Data Progression will run against multiple volumes at the same time
- RAID Restripe occurs:
  - When adding additional drives
  - When RAID extents score low

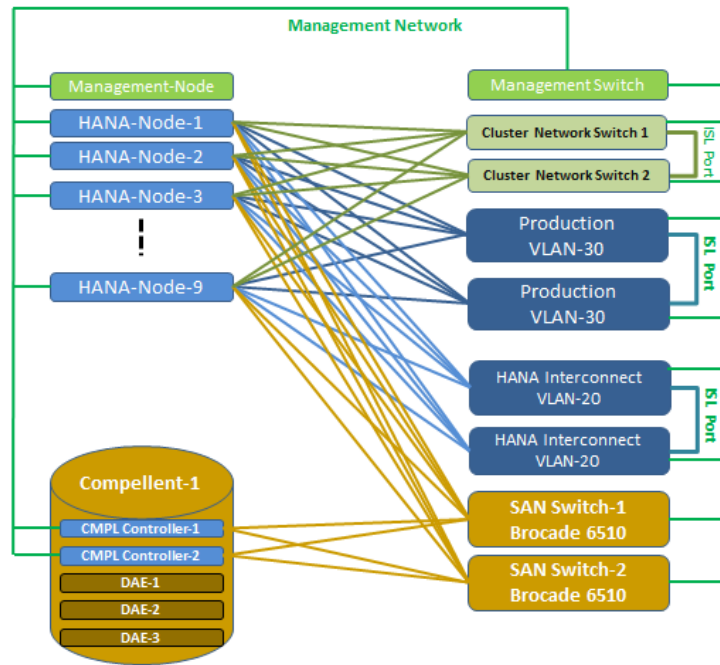


## 2.3.4 Network Configuration

SAP HANA 3.0 Scale-out appliance can support up to 16 active and 2 stand-by HANA nodes. It includes the following networks.


3. **Management network** is 1G Ethernet network using Dell Force10 (F10) S60. All hardware components can be managed via this network by a dedicated management node, a R630. Dell iDRAC network is also part the management network. It uses the onboard I350 1Gb BT DP Network Daughter Card
4. **Internal Cluster network** is a dedicated and redundant 1G communication channel for internal heart beat network, using a pair of F10 S60 – bond2 network
5. **HANA related networks**: one pair of 10G Dell F10 S4820T in a Virtual Link Trunking (VLT) domain to supports redundant data paths for different HANA related networks. Each S4820T has 48 x 10GBase T Ports + 4 x 40Gb QSFP+ ports
  - VLAN-20: dedicated HANA Internal network – bond0 network
  - VLAN-30: HANA prod/external network – bond1 network
  - Other available ports can be used for other network such as backup or required, to create a dedicated system replication network using the Intel Ethernet X540 DP 10Gb on the daughter card
5. **Storage Area Network**: dedicated and redundant FC16 SAN fabric to host all database related volumes, using one pair of FC16 Brocade 6510 switches

## In-Memory Business Intelligence: SAP-HANA



### 2.3.4.1 Force10 Virtual Link Trunking (VLT):

- Virtual Link Trunking feature is a mechanism to provide **layer 2 multi-pathing** between access network devices (switches or servers) and the core network.
- VLT is a Dual Active Control Plane implementation of Multi System LAG
- VLT provides a **loop-free topology with active-active load-sharing** of links from access to core



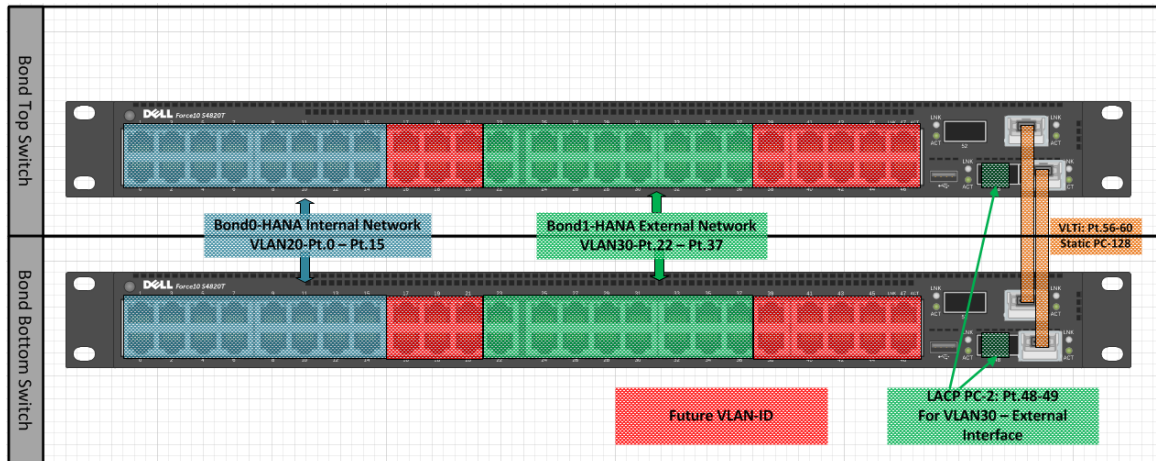
### VLT domain configuration:

- VLTi (Peer-link): Static port channel 128 (40Gb ports 56 – 60) for control and data exchange
- VLT Health Check: management port (OOB Port). Periodic keep alive message

### 2.3.4.2 HANA related networks:


One pair of F10 S4820T are configured in a VLT domain to host HANA related networks via different VLAN-IDs. It can support up to 16 active and 2 stand-by HANA nodes.

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



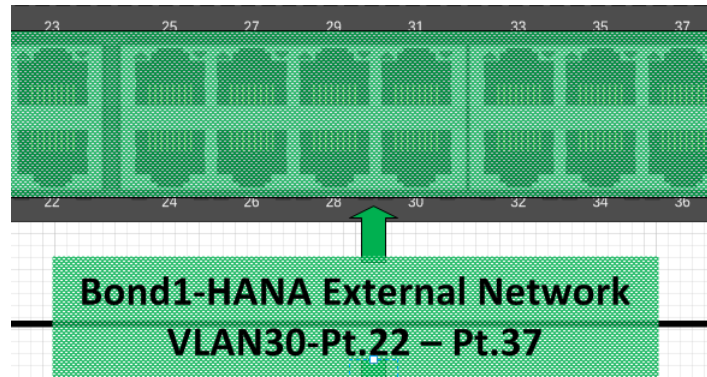
- **Dedicated HANA Internal Network:**

- Bond-0: VLAN-20 – port 0-15



- **HANA External/Prod Network:**

- Bond-1: VLAN-30 – port 22-37
- Client and System Replication networks



- **Dedicated Network for backup or system replication:**
  - Bond-X: VLAN-ID – other available ports on the switch

### 2.3.5 Guidelines for deploying on RedHat Enterprise Linux

Please refer to the following guidelines from SAP and Red Hat while deploying HANA:

Introduction to Red Hat Enterprise Linux for SAP HANA

<http://www.redhat.com/en/resources/red-hat-enterprise-linux-sap-hana>

SAP HANA DB: Recommended OS settings for RHEL 6.5

<http://help.sap.com/disclaimer?site=https://service.sap.com/sap/support/notes/2009879>

SAP HANA Server Installation and Update Guide

[http://help.sap.com/hana/SAP\\_HANA\\_Server\\_Installation\\_Guide\\_en.pdf](http://help.sap.com/hana/SAP_HANA_Server_Installation_Guide_en.pdf)

### 2.3.6 Guidelines for deploying on SUSE Enterprise Linux

Introduction to SUSE Enterprise Linux for SAP HANA

<https://www.suse.com/products/sles-for-sap/>

SAP HANA Guidelines for SLES Operating System

<http://help.sap.com/disclaimer?site=https://service.sap.com/sap/support/notes/1944799>

SAP HANA Server Installation and Update Guide

[http://help.sap.com/hana/SAP\\_HANA\\_Server\\_Installation\\_Guide\\_en.pdf](http://help.sap.com/hana/SAP_HANA_Server_Installation_Guide_en.pdf)



2.3.7 Volume layout and mappings for HANA

The following tables describe volume layout and mappings to each HANA node on an array. Each Dell Compellent SC8000 storage array may host a maximum of 8x3TB volume for DATA and 8x512GB volume for LOG. Additionally, the first array also hosts /hana/shared mount point for shared hana binaries and is sized equal to the total appliance size in TB and uses OCFS2 or GFS as the shared file system.

2.3.7.1 SAP HANA shared file system on Dell 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 /hana/shared volume at the expense of some complexity of cluster file system management.

**Please note that the file system GFS2 is at that point in time only certified for the Dell Compellent SC8000 as an appliance.**

2.3.7.2 OCFS2 Configuration on SLES 11

The following settings are recommended to be used to configure OCFS2 file system for /hana/shared:

- Use OCFS2 to provide shared mount point hana/shared for hana binaries
- OCFS2 uses SLES HA Extension: Based on open source clustering technologies
  - It uses Corosync/OpenAIS as messaging infrastructure (node membership)
  - PCMK as the OCFS2 cluster stack
  - Configuration: STONITH Block Device (SBD) is used for fencing with following settings:

Timeout (watchdog): 30

Timeout (msgwait) : 60

stonith-timeout="120s

stonith-enabled="true"

OCFS2 mount operation should wait for DLM to startup before mounting the volume

SC8000 Array1	<p><b>Data Volumes:</b> /hana/data/HAN/mnt00001 ... mnt00008: File System: XFS using HANA storage connector</p> <p><b>Log Volumes:</b></p>
---------------	--



	/hana/log/HAN/mnt00001 .. mnt00008: File System: XFS using HANA storage connector  <b>SAPMNT:</b> /hana/shared (shared HANA home binaries and executables available on all nodes) File System: OCFS2/GFS2
SC8000 Array2	<b>Data Volumes:</b>  /hana/data/HAN/mnt00009 ... mnt000016:  <b>Log Volumes:</b>  /hana/log/HAN/mnt00009 .. Mnt000016: File System: XFS using HANA storage connector

### 2.3.7.3 GFS2 Configuration on RHEL 6.5 (SC8000)

- /hana/shared uses GFS2 to provide shared mount point for hana binaries
- Using RHEL HA CMAN
- CLVMD for cluster logical volume
  - IPMI LAN for node fencing
    - 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
```



## 3 Dell Compellent Fluid Data Architecture Overview

Dell 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 physical storage resources from their representation, the fluid data architecture can place the right data in the right place, at the right time for the right cost. Dell Compellent is a true enterprise grade storage platform that delivers


### 3.1 Intelligence

Award-winning and patented, Dell Compellent Data Progression intelligently classifies and migrates data to the optimum tier of storage using real-time intelligence about how data is used throughout its lifecycle. This allows businesses to retain frequently accessed data on high-performance drives and 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 the value by extending the benefits of tiered storage to all flash configurations. By using expensive Write Intensive drives for writes and inexpensive Read Intensive drives for reads, Compellent delivers the industry's first all flash solution at the price of disk.

### 3.2 Flash Aware Sub-Lun Tiering

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





**Figure 1** Compellent data progression

### 3.3 Policy based data management

Storage profiles are a feature of the Dell Compellent storage platform that radically simplifies tiering policy management. By assigning a storage profile to a volume, the volume inherits the tiering attributes defined in the storage profile. For example a Flash-optimized storage profile directs all writes to high-performance Tier 1 storage (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 allows the administrator to create a volume for cost-optimized applications such as backup, archive and low priority applications. The storage profiles feature simplify tiering management which enables 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.

## 4 Dell Compellent storage configuration overview

SAP HANA workload requires special consideration in terms of sizing and allocation of volumes for DATA and LOG volumes which provide persistence layer for HANA in-memory database engine. HANA requires access to all DATA and LOG volumes from all nodes as shared storage even though each worker node owns and operates its own set of DATA and a LOG volume.

The workload exhibits random IO on DATA volume and sequential IO on the LOG volume. DATA volume access from HANA can range from 4K, 16K, 64K, 1M and 16M block sizes while LOG volume access from HANA can range from 4K, 16K and 1M block sizes.

### 4.1 Switch Zoning Configuration

Switch zoning should be configured such that each host should be able to see all virtual ports on all storage arrays. For detailed zoning best practices, please refer to Dell Compellent storage administration guide.

### 4.2 Qlogic Host Bus Adapter driver settings

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

```
options qla2xxx qlport_down_retry=5
```

This option is required to process controller failover in dual controller configurations to avoid any impact to SAP HANA operation.

### 4.3 Storage Array Controller Setup

Appendix A provides details about setting up an SC8000 storage for HANA.

Please refer to the following storage array documentation to initialize and setup the storage arrays:

SC4020:

[ftp://ftp.dell.com/Manuals/all-products/esuprt\\_ser\\_stor\\_net/esuprt\\_compellent/dell-compellent-sc4020\\_Administrator%20Guide5\\_en-us.pdf](ftp://ftp.dell.com/Manuals/all-products/esuprt_ser_stor_net/esuprt_compellent/dell-compellent-sc4020_Administrator%20Guide5_en-us.pdf)

SC7020:

[ftp://ftp.dell.com/Manuals/all-products/esuprt\\_ser\\_stor\\_net/esuprt\\_dell\\_storage/dell-storage-sc7020\\_Administrator%20Guide\\_en-us.pdf](ftp://ftp.dell.com/Manuals/all-products/esuprt_ser_stor_net/esuprt_dell_storage/dell-storage-sc7020_Administrator%20Guide_en-us.pdf)



SC8000:

[ftp://ftp.dell.com/Manuals/all-products/esuprt\\_ser\\_stor\\_net/esuprt\\_compellent/dell-compellent-sc8000\\_Administrator%20Guide4\\_en-us.pdf](ftp://ftp.dell.com/Manuals/all-products/esuprt_ser_stor_net/esuprt_compellent/dell-compellent-sc8000_Administrator%20Guide4_en-us.pdf)

[ftp://ftp.dell.com/Manuals/all-products/esuprt\\_ser\\_stor\\_net/esuprt\\_compellent/dell-compellent-sc8000\\_Administrator%20Guide\\_en-us.pdf](ftp://ftp.dell.com/Manuals/all-products/esuprt_ser_stor_net/esuprt_compellent/dell-compellent-sc8000_Administrator%20Guide_en-us.pdf)

SC9000:

[ftp://ftp.dell.com/Manuals/all-products/esuprt\\_ser\\_stor\\_net/esuprtl\\_storage/storage-sc9000\\_Install%20Guide\\_en-us.pdf](ftp://ftp.dell.com/Manuals/all-products/esuprt_ser_stor_net/esuprtl_storage/storage-sc9000_Install%20Guide_en-us.pdf)

[ftp://ftp.dell.com/Manuals/all-products/esuprt\\_ser\\_stor\\_net/esuprtl\\_storage/storage-sc9000\\_Administrator%20Guide5\\_en-us.pdf](ftp://ftp.dell.com/Manuals/all-products/esuprt_ser_stor_net/esuprtl_storage/storage-sc9000_Administrator%20Guide5_en-us.pdf)

[ftp://ftp.dell.com/Manuals/all-products/esuprt\\_ser\\_stor\\_net/esuprt\\_dell\\_storage/storage-sc9000\\_Administrator%20Guide4\\_en-us.pdf](ftp://ftp.dell.com/Manuals/all-products/esuprt_ser_stor_net/esuprt_dell_storage/storage-sc9000_Administrator%20Guide4_en-us.pdf)



## 5 Accessing Compellent SAN from SAP HANA nodes

### 5.1 Enable native Linux Multipathing (DM-MPIO)

To access the block devices from HANA nodes, first we need to setup multipath IO. Following is an example multipath.conf configuration file for an 8 node HANA solution which can be used to setup DM-MPIO. The WWID's can be found using CML storage center GUI.

```
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 36000d31000ec99000000000000000016
        alias sapmnt
    }
    multipath {
        wwid 36000d31000ec99000000000000000012
        alias logvol1
    }
    multipath {
        wwid 36000d31000ec99000000000000000013
        alias logvol2
    }
    multipath {
        wwid 36000d31000ec99000000000000000014
        alias logvol3
    }
    multipath {
        wwid 36000d31000ec99000000000000000015
        alias logvol4
    }
    multipath {
        wwid 36000d31000ec99000000000000000017
        alias logvol5
    }
}
```





```

multipath {
    wwid 36000d31000ec99000000000000000018
    alias logvol6
}
multipath {
    wwid 36000d31000ec99000000000000000019
    alias logvol7
}
multipath {
    wwid 36000d31000ec9900000000000000001a
    alias logvol8
}
multipath {
    wwid 36000d31000ec9900000000000000000e
    alias datavol1
}
multipath {
    wwid 36000d31000ec9900000000000000000f
    alias datavol2
}
multipath {
    wwid 36000d31000ec99000000000000000010
    alias datavol3
}
multipath {
    wwid 36000d31000ec99000000000000000011
    alias datavol4
}
multipath {
    wwid 36000d31000ec9900000000000000001b
    alias datavol5
}
multipath {
    wwid 36000d31000ec9900000000000000001c
    alias datavol6
}
multipath {
    wwid 36000d31000ec9900000000000000001d
    alias datavol7
}
multipath {
    wwid 36000d31000ec9900000000000000001e
    alias datavol8
}

```



```
}
```

After creating the `multipath.conf` file, we need to restart daemon to activate settings and rescan new devices:

### **service multipathd restart**

The volumes should now be visible on each host. Each device should have 4 paths to the array as shown by running the following command:

#### **dellhana1:/ # multipath -ll**

```
datavol1 (36000d31000ec99000000000000000011) 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 (36000d31000ec99000000000000000000e) 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
logvol4 (36000d31000ec9900000000000000000012) 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 (36000d31000ec99000000000000000000016) 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

```

## 5.2 Using XFS file system with SAP Storage Connector API to mount HANA DATA and LOG volumes

### 1. /etc/sudoers File Configuration

On the master and worker nodes, the /etc/suoders file must have the following options specified:

```
<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
```

Here SID should be replaced by the HANA system ID.

2. SAP strongly recommends using XFS file system for DATA and LOG volumes. The volumes should be un-partitioned as SAP storage connector API does not support partitioned devices at this time.

To format devices with XFS file system:

**mkfs.xfs /dev/mapper/datavol1**

**mkfs.xfs /dev/mapper/logvol1**

For all volumes to be used for persistence.



3. HANA storage connector API automatically mounts the volumes on worker nodes and on standby node in case a worker node fails based on the configuration provided in the global.ini file shown below:

```
[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
```

4. Please refer to HANA Installation instructions to install and configure SAP HANA environment:

SAP HANA Server Installation and Update Guide

[http://help.sap.com/hana/SAP\\_HANA\\_Server\\_Installation\\_Guide\\_en.pdf](http://help.sap.com/hana/SAP_HANA_Server_Installation_Guide_en.pdf)



## 6 Conclusion

SAP Tailored Data Integration (TDI) program provides a way for customers to achieve cost savings and flexibility in deploying SAP HANA. The TDI option allows customers to reduce hardware costs, apply operational best practices used elsewhere in the data center, leverage existing hardware investments and consolidate SAP and non-SAP applications on common infrastructure. A key enabling technology for this is shared enterprise storage.

The Compellent line of storage systems from Dell 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.



## Additional resources

Support.dell.com is focused on meeting your needs with proven services and support.

DellTechCenter.com is an IT Community where you can connect with Dell Customers and Dell employees for the purpose of sharing knowledge, best practices, and information about Dell products and installations.

Referenced or recommended Dell/SAP publications:

- [New economies of storage with the Compellent Flash-optimized solutions](#)
- [SAP HANA Server Installation Guide:](#)
- [SAP HANA Tailored Data Center Integration](#)
- <http://dell.com/hana>
- [Learn more about our SAP HANA offerings](#)
- [Dell certified appliances Scale-Up \(for Single Node\)](#)
- [Dell certified appliances Scale-Up \(for Suite on SAP HANA\)](#)
- [Dell certified appliances Scale-Out](#)
- [Learn more about SAP HANA Tailored Datacenter Integration \(TDI\)](#)



## A APPENDIX

The example setup provided below is for using cluster file systems for /hana/shared volume on HANA nodes.

### Configure SLES HA cluster for shared OCFS2 volume

To configure OCFS2 file system, we need to configure the SLES HA cluster which provides the heartbeat, node membership and fencing for HANA nodes:

1. create a 1GB partition on /dev/mapper/sapmnt volume to be used as STONITH block device
2. Configure /etc/corosync/corosync.conf file on each node. One can use the default file by editing the following parameters based on the network configuration being used for the cluster network:

# The following values need to be set based 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. Load the CIB (Cluster Information Base) CRM:  
crm configure load replace configuration\_ha (where configuration\_ha file contains the CIB details)

```
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. Start openais service to start DLM on all nodes  
**chkconfig openais on**  
**service openais start**
6. Create another partition on /dev/mapper/sapmnt taking all rest of the available space to configure OCFS2 file system and to mount it as /hana/shared volume on all nodes:

```
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. Issue "**mount -a -t ocfs2**" on all nodes to mount the /hana/shared volume which would be used for HANA binaries

## Configure CMAN for GFS2 for shared volume for RHEL 6.5 (SC8000)

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

```
lvnconf --enable-cluster
```

3. Create a cluster configuration file for example: Please adopt according to your environment:  
vi /etc/cluster/cluster.conf

**NOTE:** The configuration file provided below 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>
</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 cman cluster service on all nodes:  
chkconfig cman on  
service cman start  
  
to check the status of the service:  
'service cman status
7. On all nodes, start the cluster volume manager service:  
chkconfig clvmd on  
service clvmd start
8. Configure Cluster Logical Volume on /dev/mapper/sapmntp1 partition to be used for /hana/shared volume:  
pvcreate /dev/mapper/sapmntp1  
vgcreate -cy vg\_hanashared /dev/mapper/sapmntp1  
lvcreate -L <size>T -n lv\_hanashared vg\_hanashared
9. Create GFS2 file system:  
mkfs.gfs2 -t HANACluster:shared -j 18 /dev/vg\_hanashared/lv\_hanashared
10. Create mount point /hana/shared on all nodes and add the entry to the /etc/fstab:  
mkdir -p /hana/shared  
/dev/mapper/vg\_hanashared-lv\_hanashared /hana/shared gfs2 noatime,nodiratime 0 0
11. Activate the volume group on all nodes:  
vgchange -aly vg\_hanashared
12. Start the GFS2 service on all nodes:  
chkconfig gfs2 on  
service gfs2 start
13. 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)



## B APPENDIX

Example setup procedure for an SC8000 storage array.

### SC8000 Controller Setup

1. Initialize both controllers using serial cable:

**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.

- **eth0** supports system login and access for the software. It is used to send email, alerts, SNMP traps, and phone home data.
- **eth1** 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 need to be set.

**Note** There is rarely a need to change the eth1 address from its factory setting as this address is only used for IPC between controllers. One reason you might need to change the address during setup is in the unlikely scenario that both controllers were factory configured with the same eth1 address.

Installation team members do not change eth1 settings unless requested by the customer. If changed, the settings can NEVER reside in the same subnet as eth0.

Enter the following commands to set the IP address, netmask, and gateway. The options are separated by a single space. (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]
```

Verify settings by entering:



## controller show

**Note** In a dual-controller system, repeat steps 1 through 6 for Controller 2.

For the second controller, please use this IP address as below

**eth0 172.16.5.22 255.255.0.0 172.16.0.1**

The below picture shows a screenshot after a "Setup IP Address" task completed.



```
root: minicom
File Edit View Scrollback Bookmarks Settings Help

> controller show
[60569] System:Controller 08/10/2012 04:28:43 PM
=====
Index Controller HardwareSN SystemSN Name Versio Versio Versio Failed Leader IP IPCPort Eth1
=====
60569 0 0x0000EC99 0x0000EC99 SN 60569 06 01 01 (0) False (1) True 10.0.10.76 20000 11.0.10.76
Walk: Display 1 of 1

> controller ipconfig eth0 172.16.5.21 255.255.0.0
Invalid number of parameters entered
Shell Command: Failure due to Failure
: Desired Arguments:
: Supplied Values : controller, ipconfig, eth0, 172.16.5.21, 255.255.0.0
> controller ipconfig eth0 172.16.5.21 255.255.0.0 172.16.0.1
> controller show
[60569] System:Controller 08/10/2012 04:30:51 PM
=====
Index Controller HardwareSN SystemSN Name Versio Versio Versio Failed Leader IP IPCPort Eth1
=====
60569 0 0x0000EC99 0x0000EC99 SN 60569 06 01 01 (0) False (1) True 172.16.5.21 20000 11.0.10.76
Walk: Display 1 of 1

> controller show
[60570] System:Controller 08/10/2012 04:32:10 PM
=====
Index Controller HardwareSN SystemSN Name Versio Versio Versio Failed Leader IP IPCPort Eth1
=====
60570 0 0x0000EC9A 0x0000EC9A SN 60570 06 01 01 (0) False (1) True 10.0.10.92 20000 11.0.10.92
Walk: Display 1 of 1

> controller ipconfig eth0 172.16.5.22 255.255.0.0 172.16.0.1
> controller show
[60570] System:Controller 08/10/2012 04:32:46 PM
=====
Index Controller HardwareSN SystemSN Name Versio Versio Versio Failed Leader IP IPCPort Eth1
=====
60570 0 0x0000EC9A 0x0000EC9A SN 60570 06 01 01 (0) False (1) True 172.16.5.22 20000 11.0.10.92
Walk: Display 1 of 1

CTRL-A Z for help |115200 8N1 | NOR | Minicom 2.3 | VT102 | Online 00:01
root: minicom
```

## Startup Wizard

**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.

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. Enter the eth0 IP address of Controller 1.


Click **Yes** to acknowledge all security alerts

**Note** Alerts that appear differ depending on your browser settings and Java version.

The **Storage Center Login** window appears.

The screenshot shows the 'Storage Center Login' interface. The top header is black with the Dell Compellent logo on the left and the title 'STORAGE CENTER LOGIN' on the right. The main content area is white and features a login form with a 'User' field containing 'Admin' and an empty 'Password' field, with a blue 'LOGIN' button to the right. To the right of the login form are three stacked images: a server rack, a water splash, and a server component. At the bottom left is the 'goahead WEB SERVER' logo, and at the bottom right is the copyright notice '© Compellent Technologies, Inc. 2003-2011'.

## Load License



**Note** The license file is sent as an email attachment. Upon receipt, you should save the file in a location accessible from the PC used for installation.

The **Load License** step asks you to browse to the location of the Storage Center license file. 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 should match the serial number of the controller from which you start the wizard.
- 35 indicates the system runs post-3.5 code.
- License generation date is shown in YYMMDD format.
- .lic is the file extension.

To load the license:

1. Browse to and select the license file, and click **Load License**. The Startup Wizard informs you when the license is successfully loaded.
2. Click **Continue**. The **Create Disk Folder** step appears.

## Create Disk Folder

Storage Center 60565 - Startup Wizard

Back Quit Advisor

System Startup

License Agreement

Load License

Create Disk Folder

Add Controller

Time Settings

System Setup

Configure SMTP

Update Setup

User Setup

Configure Ports

Generate SSL Cert

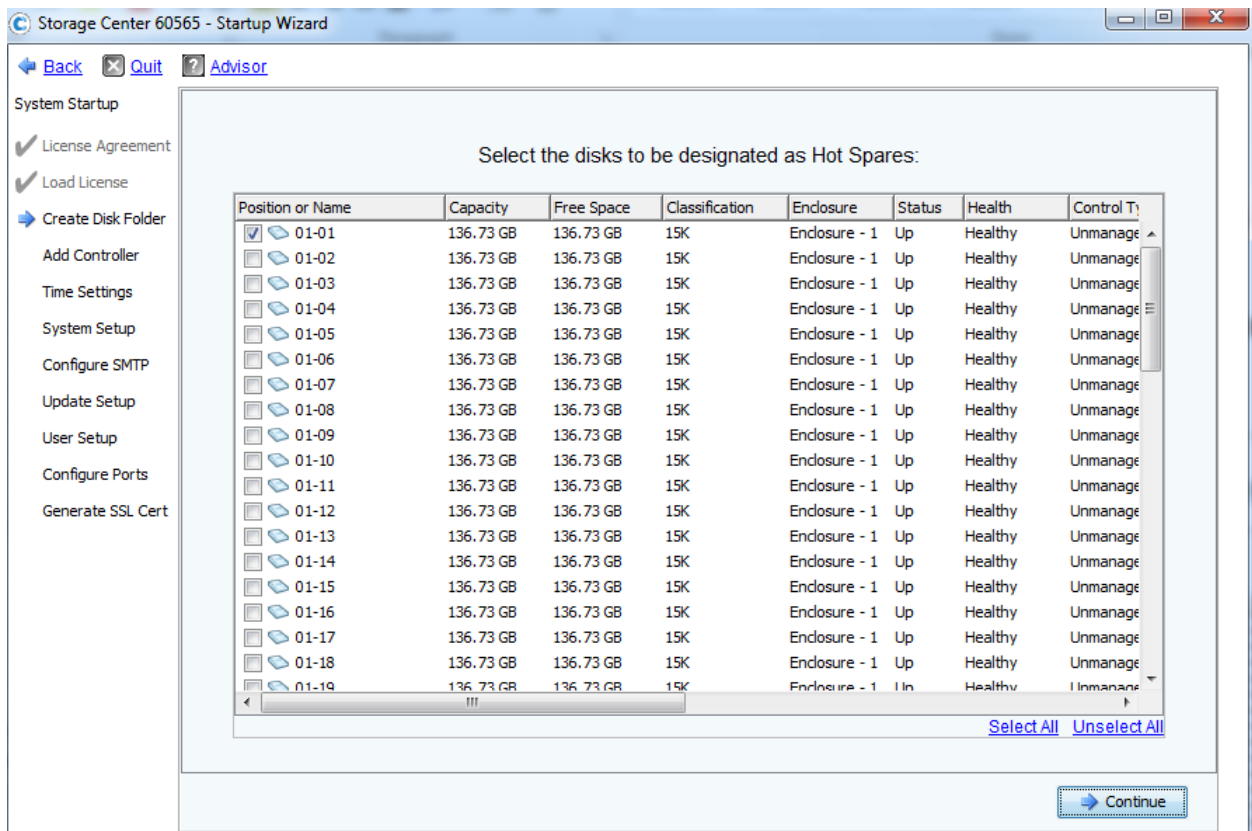
Select the disks to be included in the Disk Folder:

Position or Name	Capacity	Free Space	Classification	Enclosure	Status	Health	Control T
03-00	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-07	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-08	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-09	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-10	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-11	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-12	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-13	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-14	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-15	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-16	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-17	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-18	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-19	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-20	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-21	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-22	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-23	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage
03-24	136.73 GB	136.73 GB	15K	Enclosure - 3	Up	Healthy	Unmanage

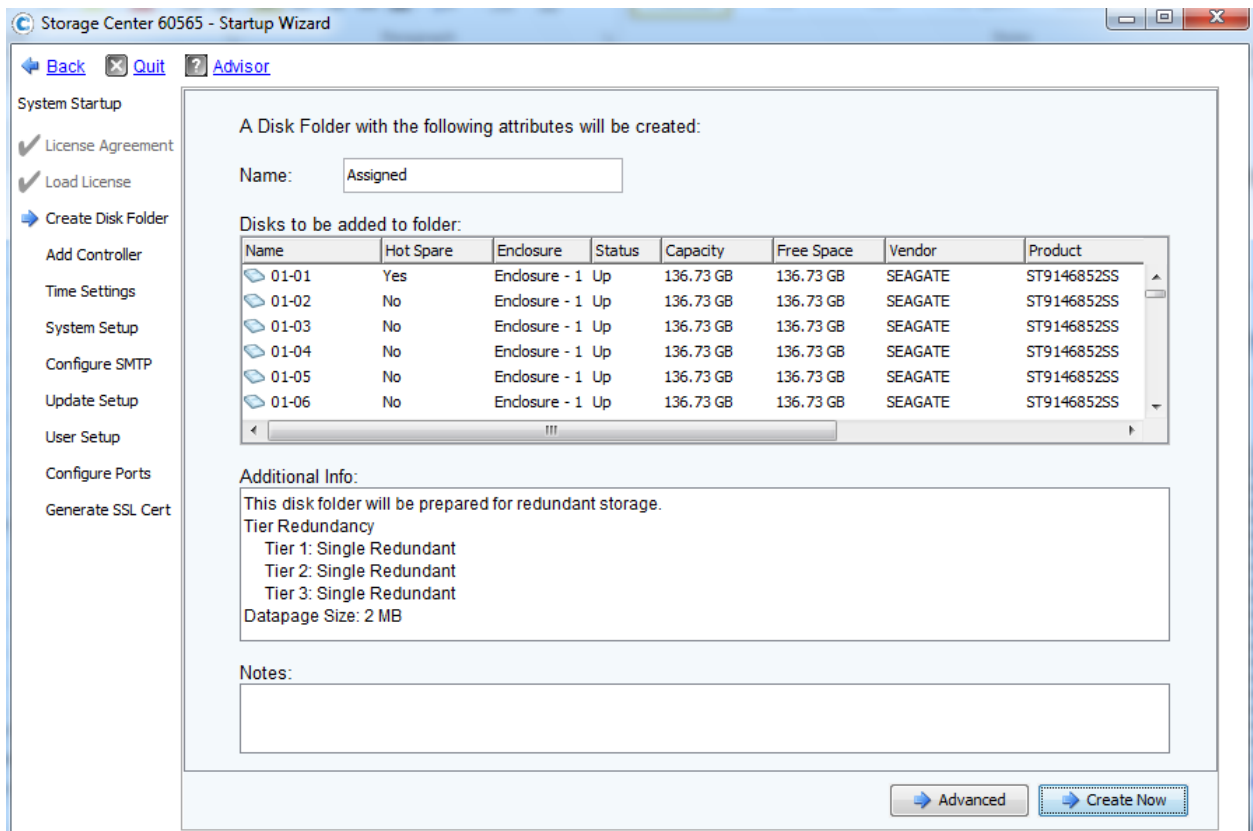
Select All Unselect All

Continue





Select 1 spare disk per enclosure




Tier1 = Single Redundant

Tier3 = Single Redundant

Then Click "Create Now" to proceed.

## Add Controller



Click "Add Controller"

Storage Center 60565 - Startup Wizard

[Back](#) [Quit](#) [Advisor](#)

System Startup

- ✓ License Agreement
- ✓ Load License
- ✓ Create Disk Folder
- ➔ Add Controller
- Time Settings
- System Setup
- Configure SMTP
- Update Setup
- User Setup
- Configure Ports
- Generate SSL Cert

This wizard adds a new Controller into the System.

Controller ID: 60566

Ether 0 (MGMT) Interface:

IP Address: 172.26.224.22

Net Mask: 255.255.240.0

Gateway: 172.26.224.1

Ether 1 (IPC) Interface:

IP Address: 11.0.10.124

Net Mask: 255.0.0.0

Gateway: 0.0.0.0

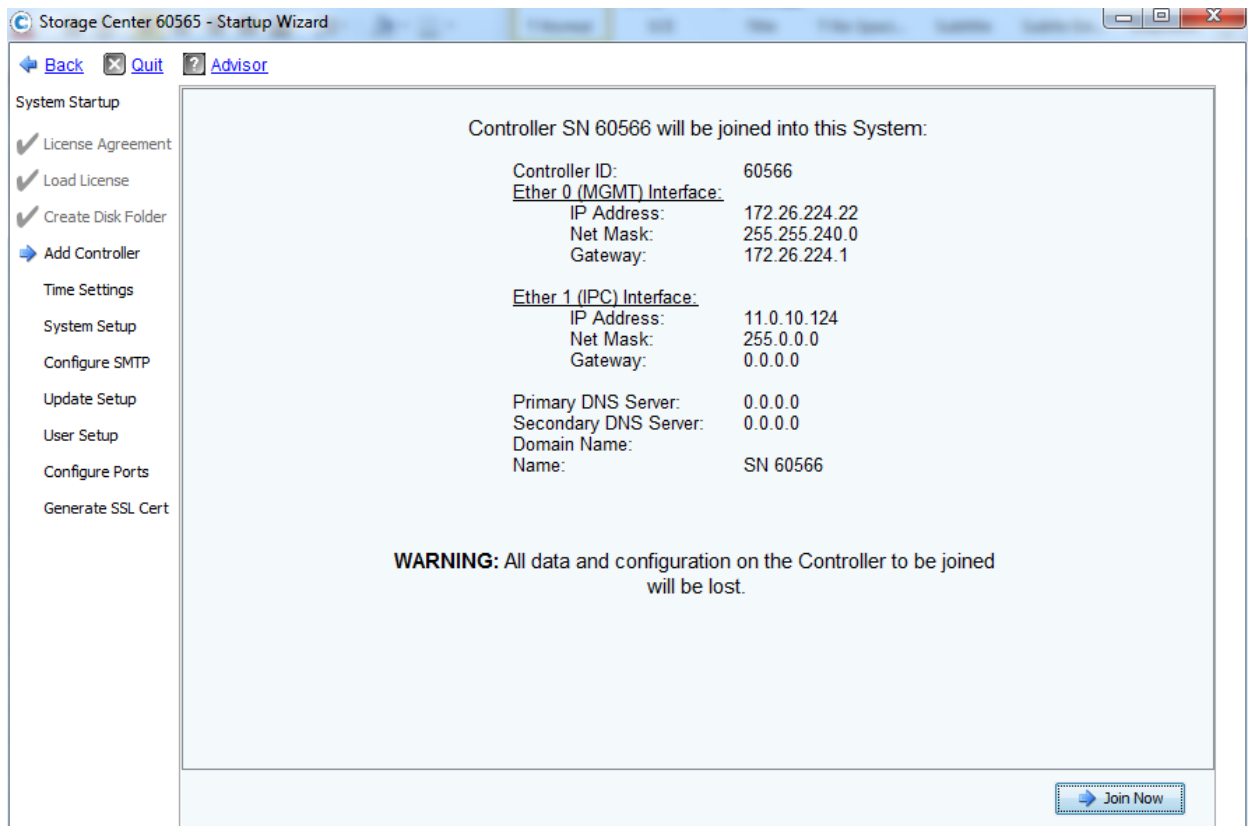
Primary DNS Server: 0.0.0.0

Secondary DNS Server: 0.0.0.0

Domain Name:


[Continue](#)

Provide IP address for "Ether 1(IPC) Interface. This can be found by execute "controller show" on the second controller



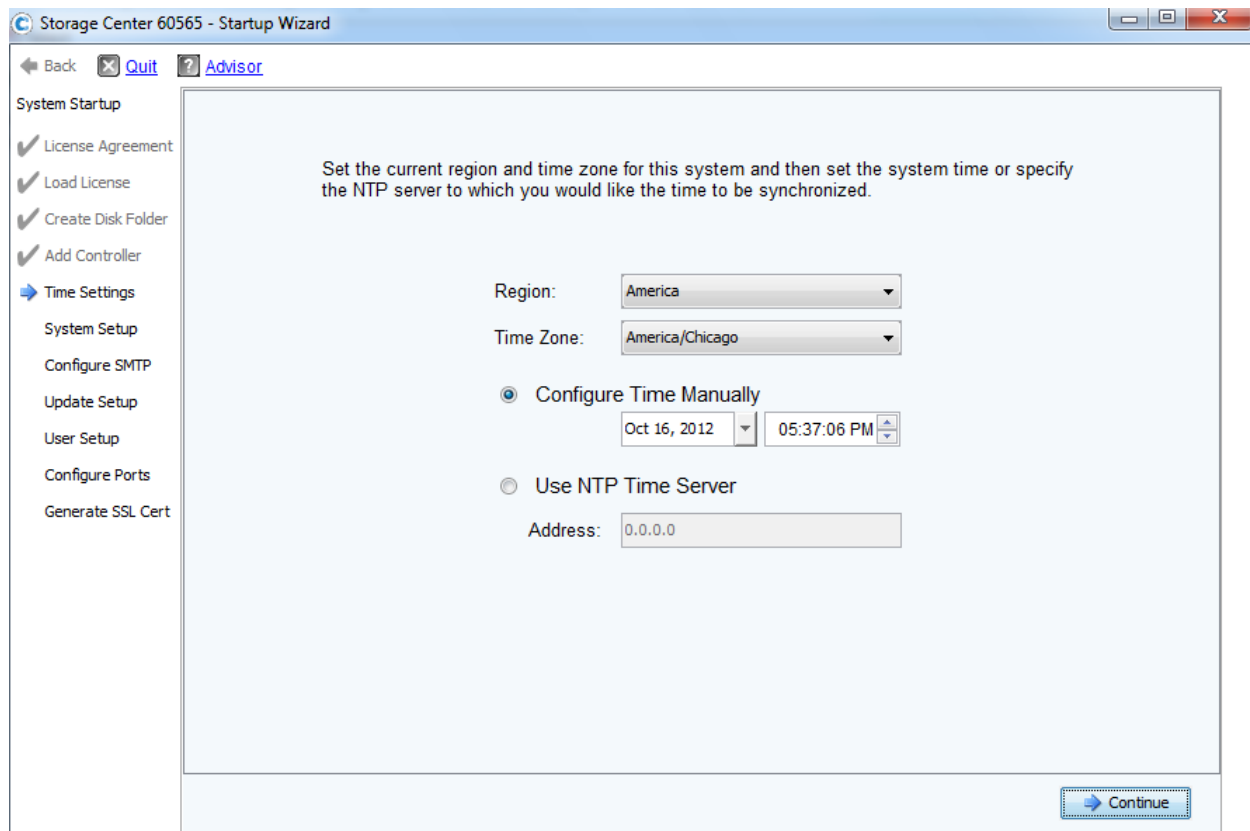
Click "Join Now"

The new added controller will have to be rebooted. This would take a few minutes for a reboot.



## Time Setting

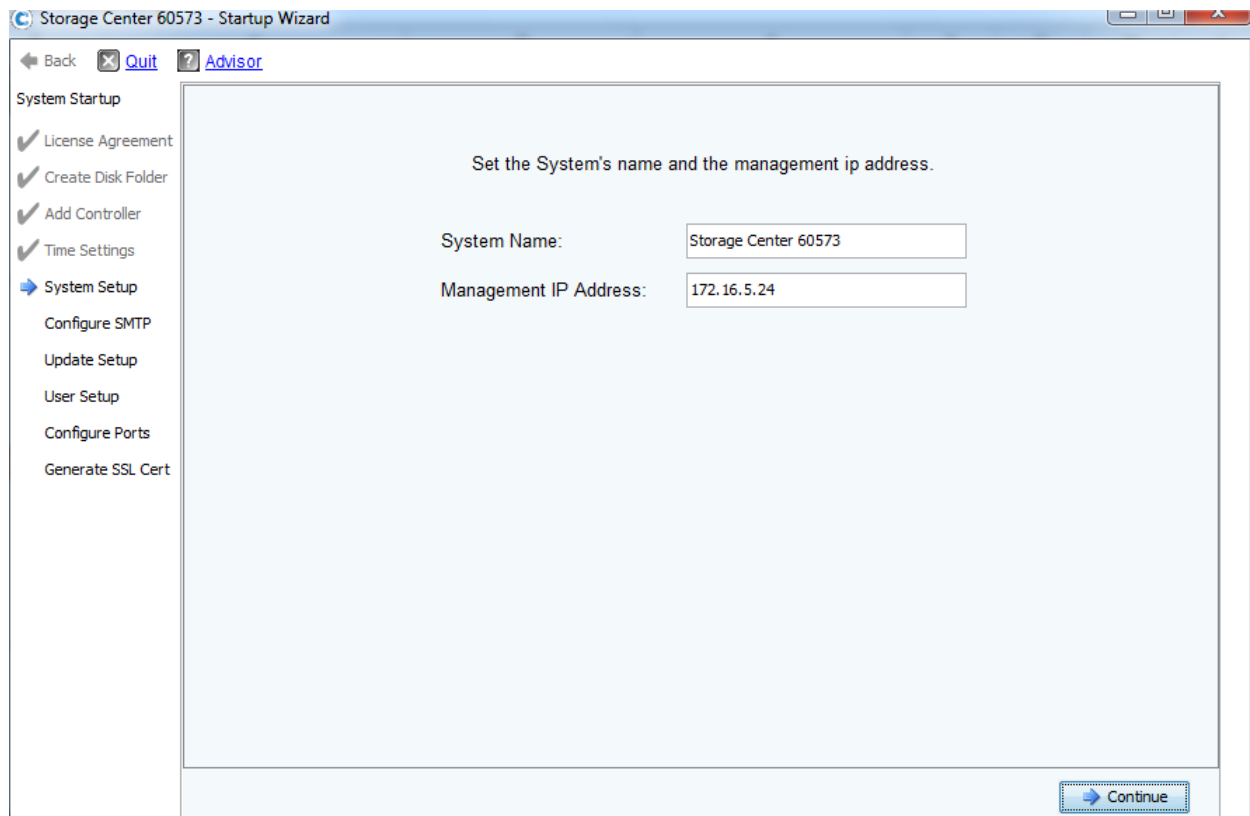
Change the time, its zone, and region accordingly to the destination.



## System Setting

Provide "Management IP address" for the storage array

- For the first storage array: 172.16.5.20
- For the second storage array (if any): 172.16.5.24



Disable read and write cache



## Configure SMTP

Storage Center 60573 - Startup Wizard

Back Quit Advisor

System Startup

- ✓ License Agreement
- ✓ Create Disk Folder
- ✓ Add Controller
- ✓ Time Settings
- ✓ System Setup
- ➔ Configure SMTP
- Update Setup
- User Setup
- Configure Ports
- Generate SSL Cert

Set the IP address of your SMTP mail server and a sender e-mail address to allow e-mailing of alert messages.

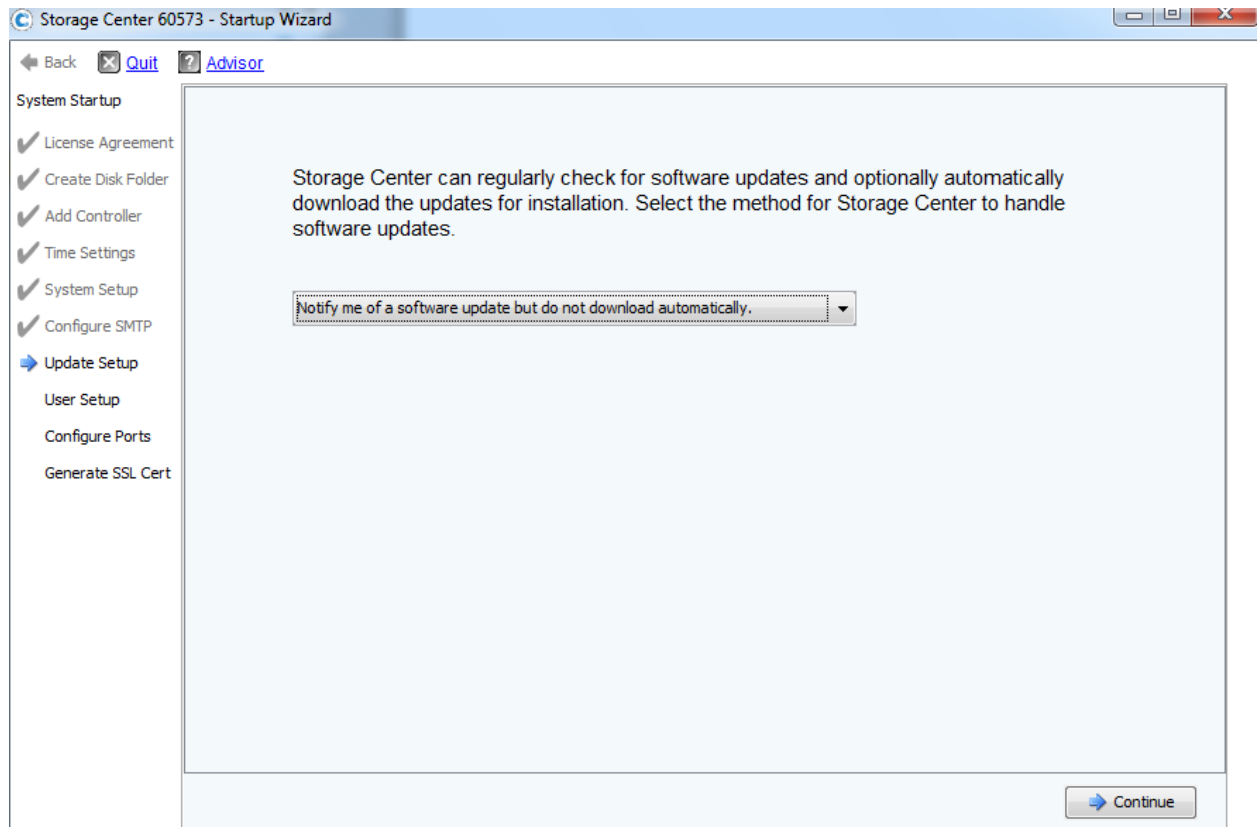
SMTP Mail Server:  ➔ [Test server](#)

Sender E-mail Address:

Most SMTP servers including Microsoft Exchange require a Sender E-mail Address (MAIL FROM) value.

➔ Skip SMTP Configuration   ➔ Advanced   ➔ Continue

## Update Setup



Select "Notify me of a software update but do not download automatically" option. Then Click "Continue" to proceed.

Please check CML administration guide for details on how to configure Compellent Phone Home feature after completing this "Startup Wizard".

## User Setup

The screenshot shows the 'Storage Center 60573 - Startup Wizard' window. The left sidebar lists the following steps: System Startup, License Agreement, Create Disk Folder, Add Controller, Time Settings, System Setup, Configure SMTP, Update Setup, **User Setup** (highlighted with a blue arrow), Configure Ports, and Generate SSL Cert. The main area contains the following text and fields:

Set user Admin's Session Timeout duration as well as the e-mail addresses where alerts are to be sent.

Userid: Admin

Session Timeout: 2 Hours (dropdown menu)

Email:  [Send test e-mail](#)

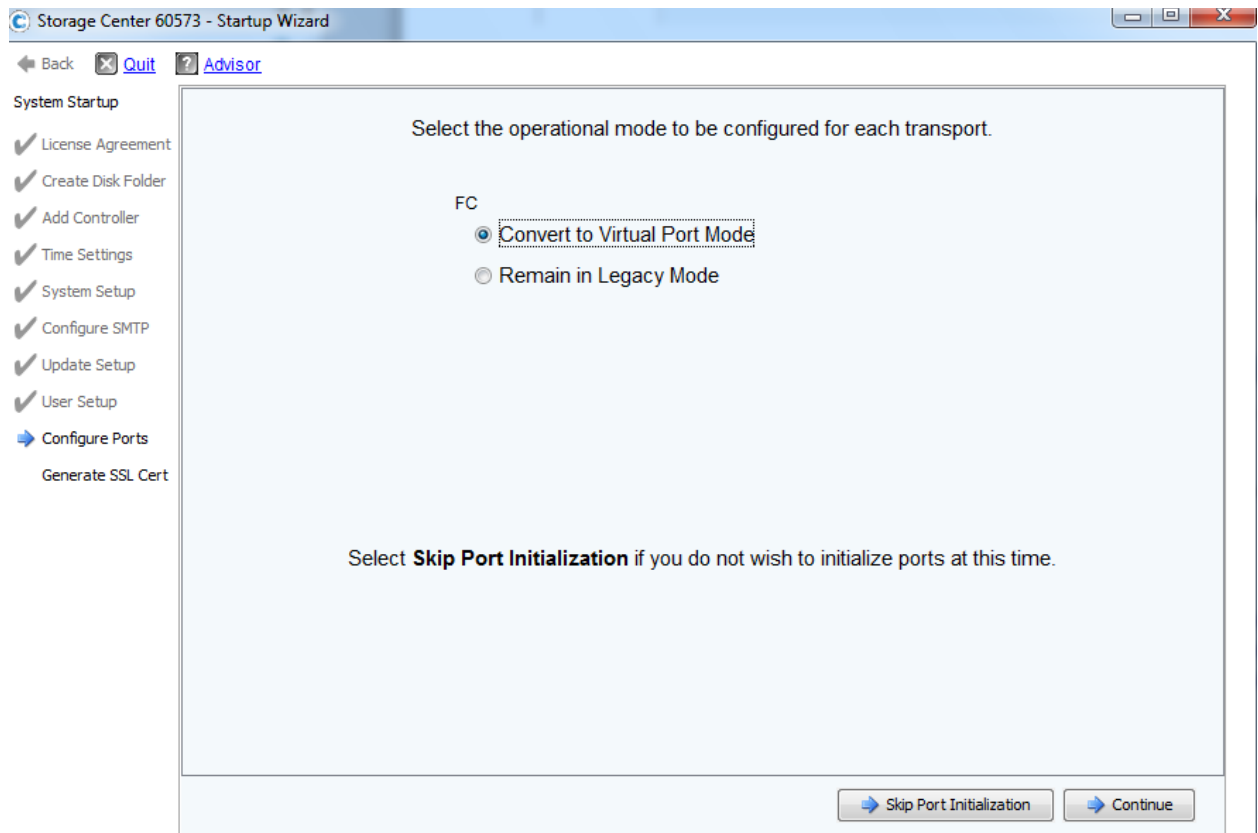
Email 2:  [Send test e-mail](#)

Email 3:  [Send test e-mail](#)

At the bottom right, there is a 'Continue' button with a blue arrow icon.

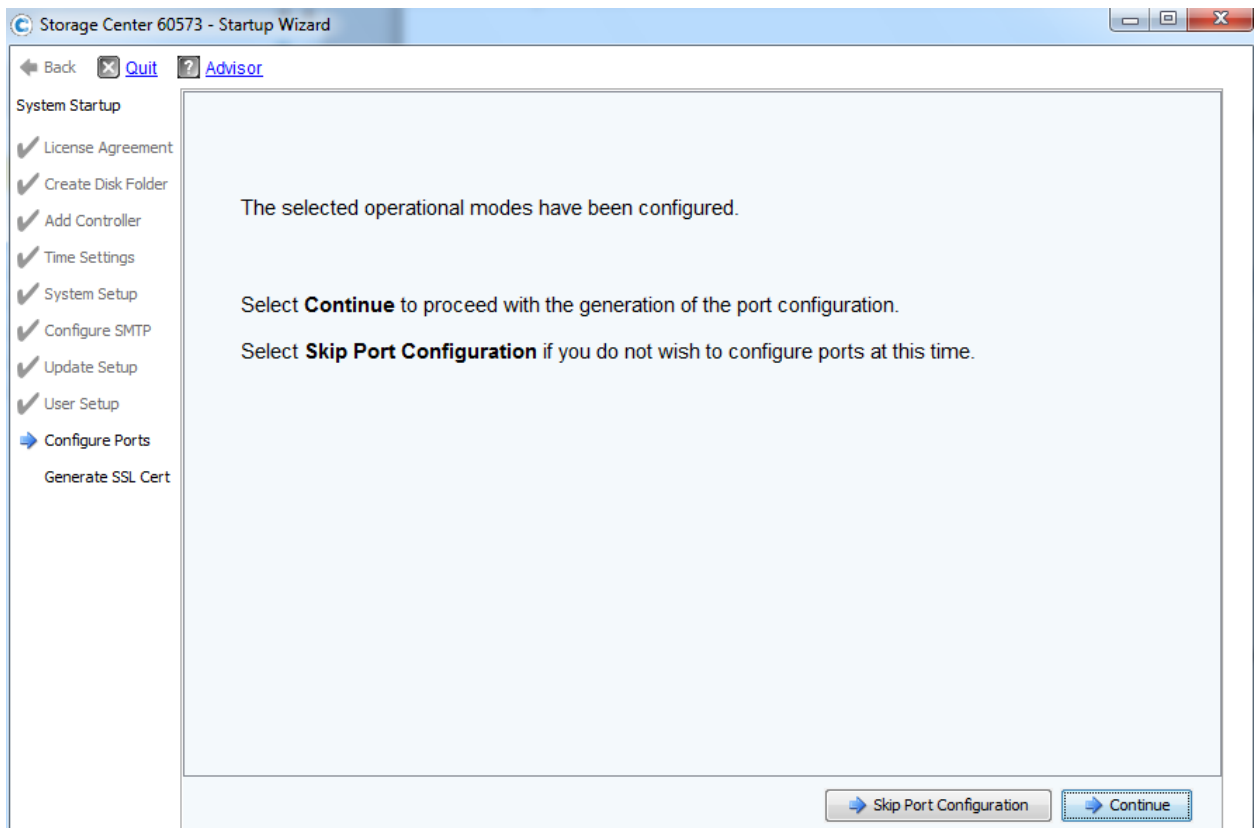
Provide details if applied and click "Continue" to proceed

## Configure Ports

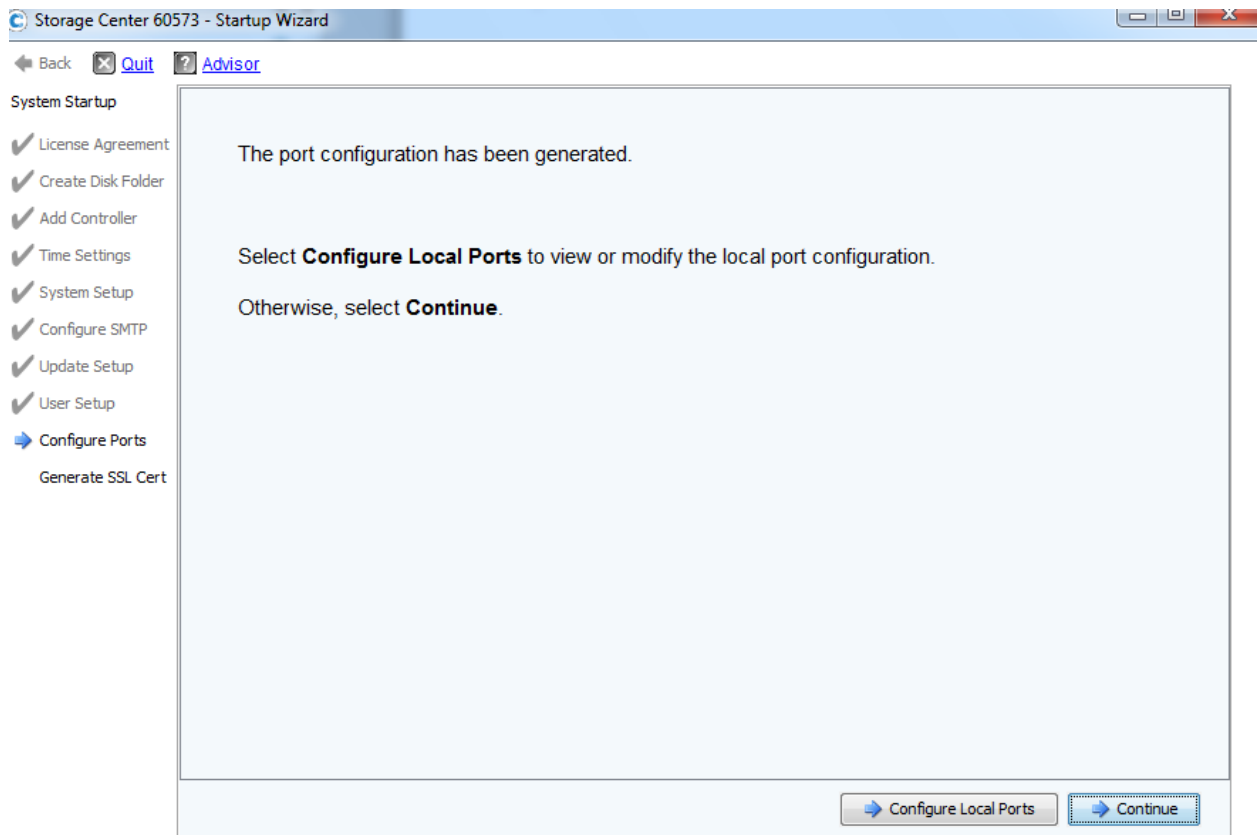


Ensure that all FC ports are connected to FC switches for this step.


Select "Convert Virtual Port Mode" and click "Continue" to proceed



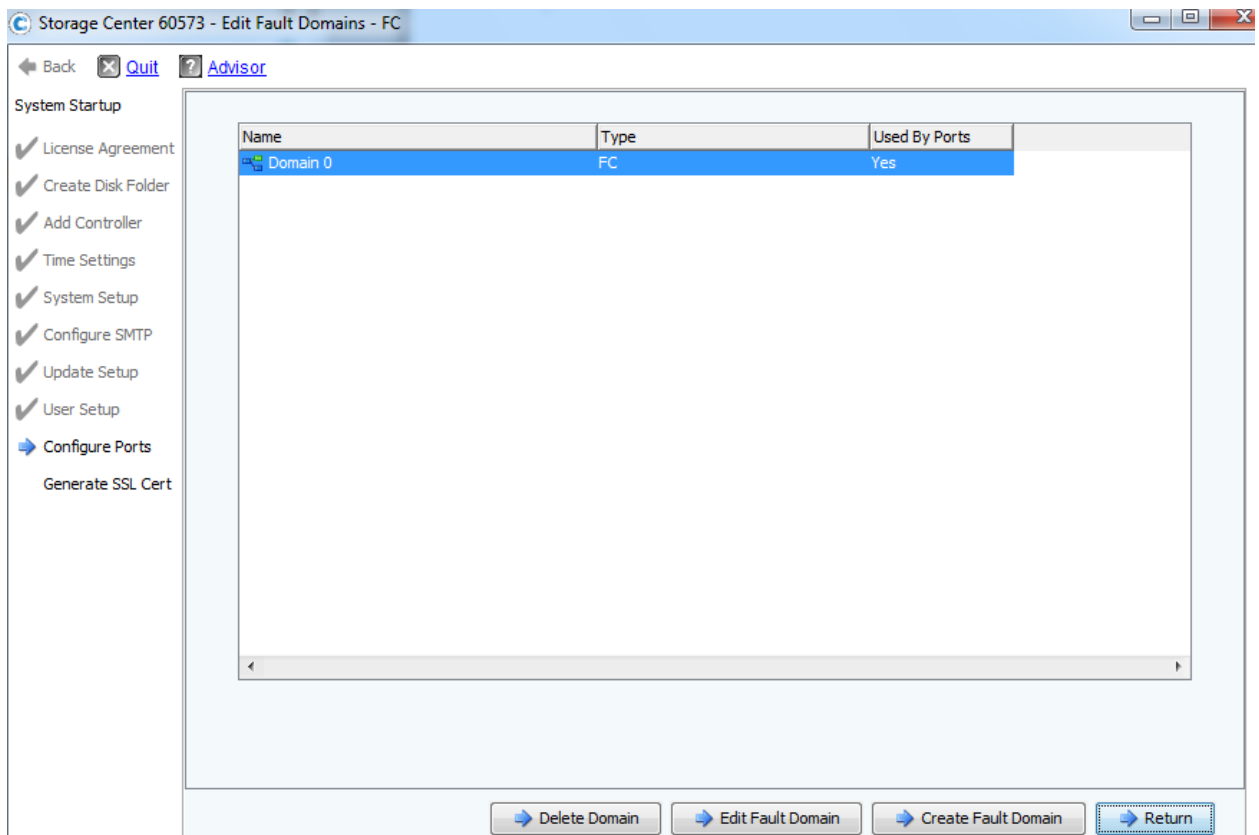
Click Continue to proceed



Click on "Configure Local Ports"



Click on "Edit Fault Domain"




Click on "Edit Fault Domain" to change "Domain 0" fault domain to "Top" fault domain




Storage Center 60573 - Fault Domain Properties - Domain 0

General

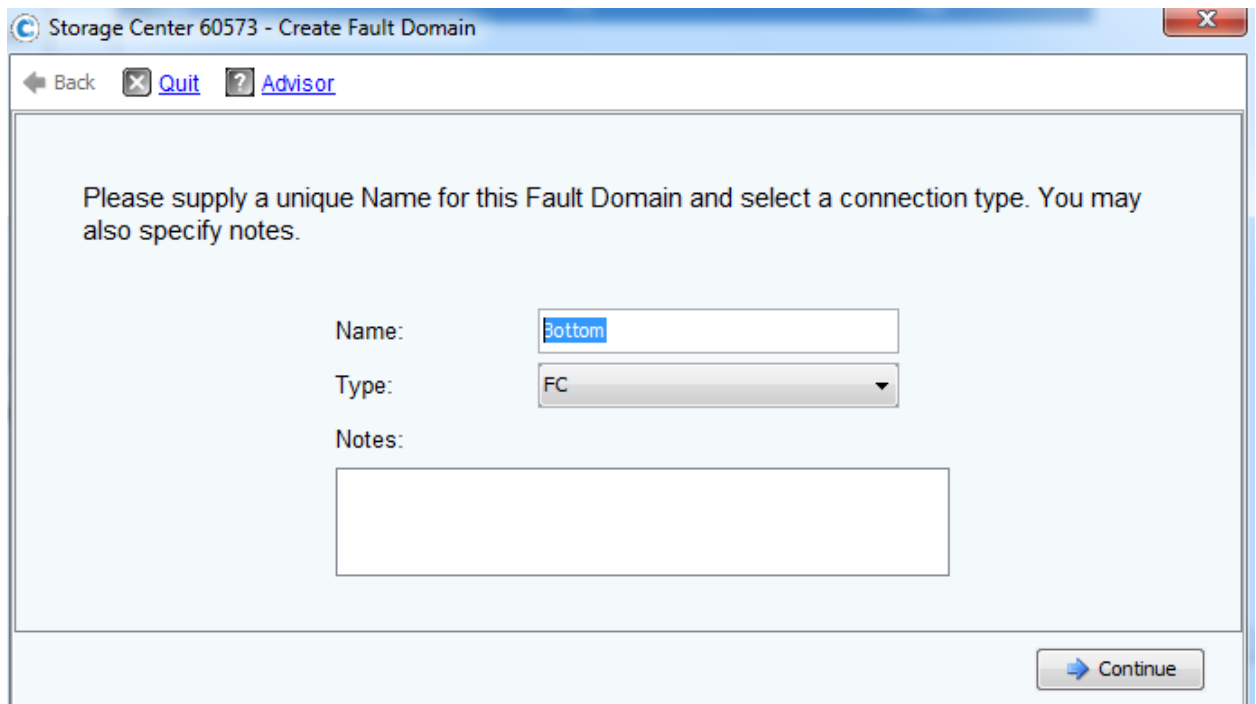


Type: FC

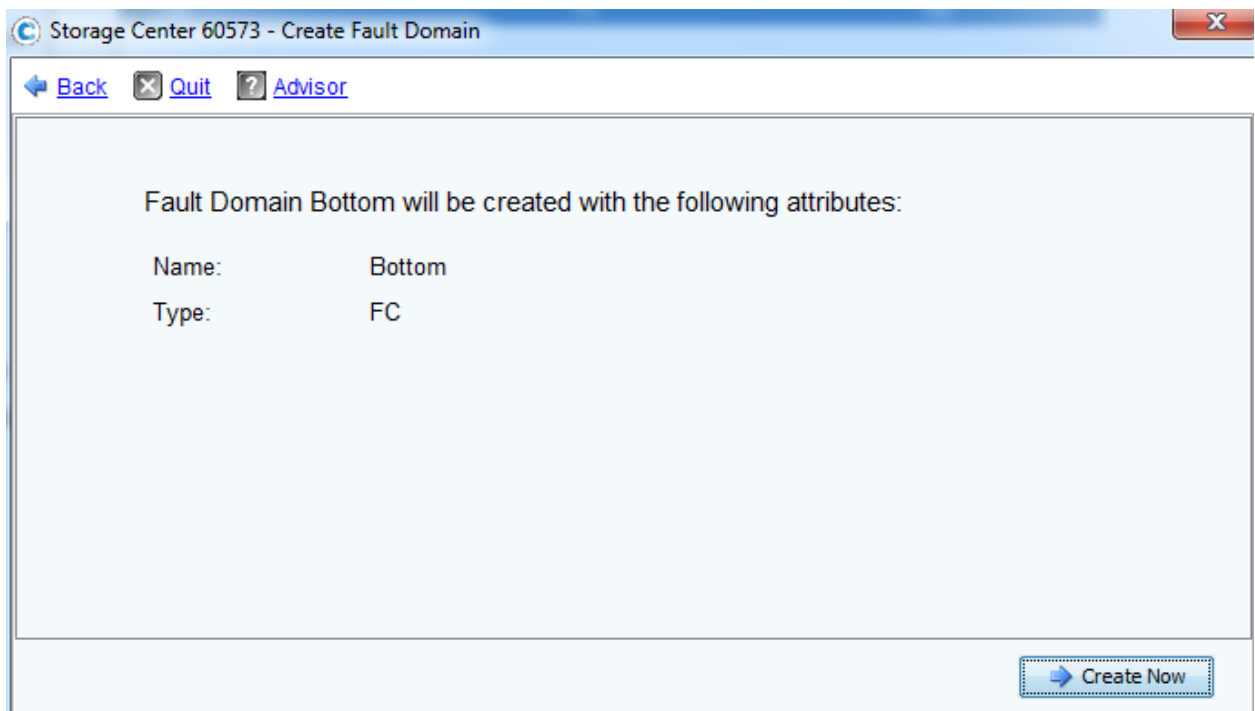
Notes:




Click "Create Fault Domain" to create "Bottom" fault domain




Click "Continue"



Click "Create now"




Click "Return" to go back to "configuration Ports" page



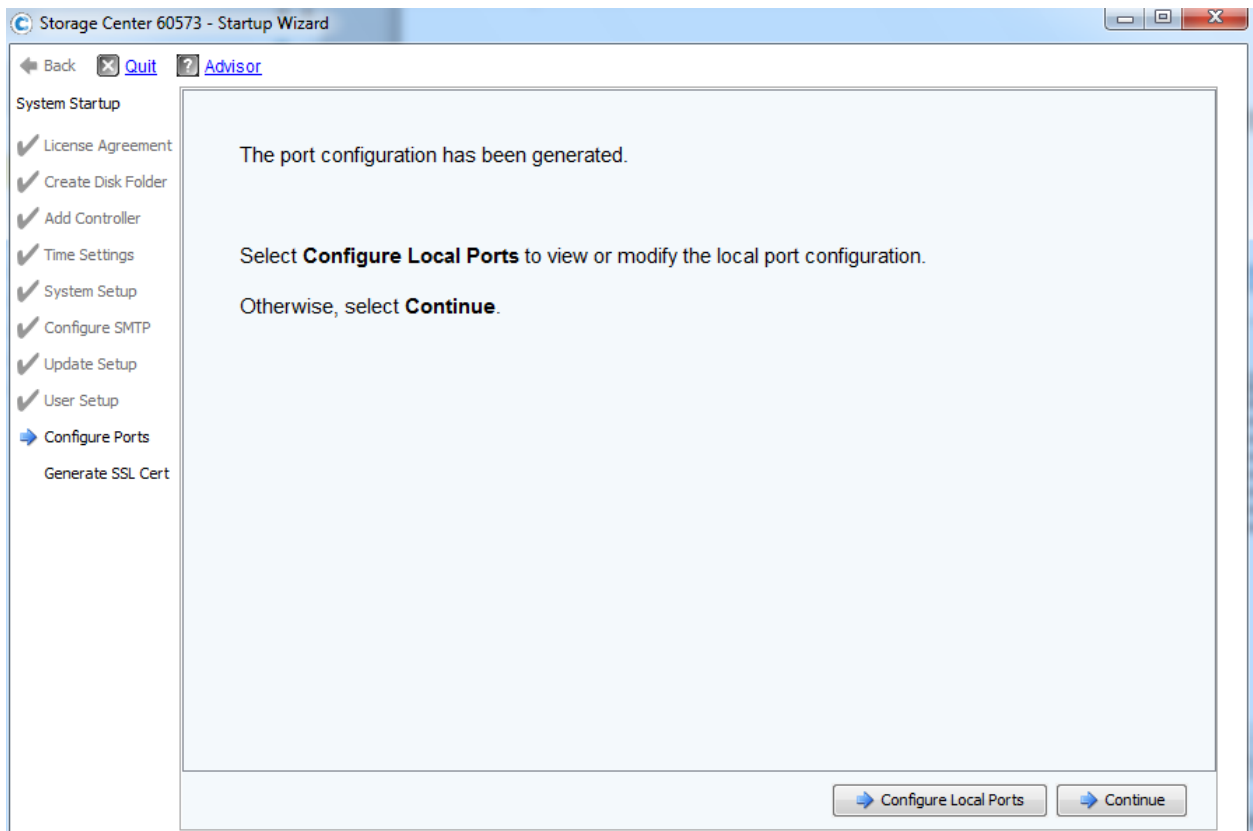
From "Fault Domain" pull-down tab, change a FC port's Fault Domain to a correct "Top" or "Bottom" fault domain as the following.

- 1) Port's Purpose: Set to "Front End"
- 2) Fault Domain:
  - Port-1 is set to "Top" Fault Domain
  - Port-2 is set to "Bottom" Fault Domain

The final settings should look as the following screenshot



Click "Assign Now" to proceed




## Post Settings for Compellent storage array

This provides post setting steps including Storage Profile etc ...

### Enable "Manual Storage Mode"

- 1) Login to the Storage Center
  - Username: Admin
  - Password: "mmm"
- 2) Left click on System → Setup → Enable Manual Storage Mode



Press Continue to proceed

## Configuring Storage Profile

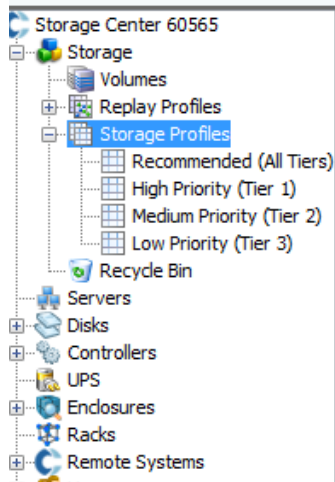
There is storage profile for both data and log volumes profile as the following.

Profile for HANA Volume

1. From a Storage Center, Click on "Create Storage Profile" to start with the process



Create Storage Profile



## Storage Profiles

Name	Type	Created By
<input type="checkbox"/> Recommended (All Tiers)	Storage Profile	System
<input type="checkbox"/> High Priority (Tier 1)	Storage Profile	System
<input type="checkbox"/> Medium Priority (Tier 2)	Storage Profile	System
<input type="checkbox"/> Low Priority (Tier 3)	Storage Profile	System

2. Check and Uncheck dialog boxes of Writeable Data or Replay Data for the storage Profile as shown below.

Storage Center 60565 - Modify Storage Profile

Back Quit Advisor

Select the classes of storage to be used for writable and replay data.  
Changing this Storage Profile will affect all volumes that use it.

### Tier 1 Storage

	RAID Level	Writable Data	Replay Data
	RAID 10	<input checked="" type="checkbox"/>	<input type="checkbox"/>
	RAID 5-5	<input type="checkbox"/>	<input type="checkbox"/>
	RAID 5-9	<input type="checkbox"/>	<input checked="" type="checkbox"/>

### Tier 2 Storage

	RAID Level	Writable Data	Replay Data
	RAID 10	<input type="checkbox"/>	<input type="checkbox"/>
	RAID 5-5	<input type="checkbox"/>	<input type="checkbox"/>
	RAID 5-9	<input type="checkbox"/>	<input type="checkbox"/>

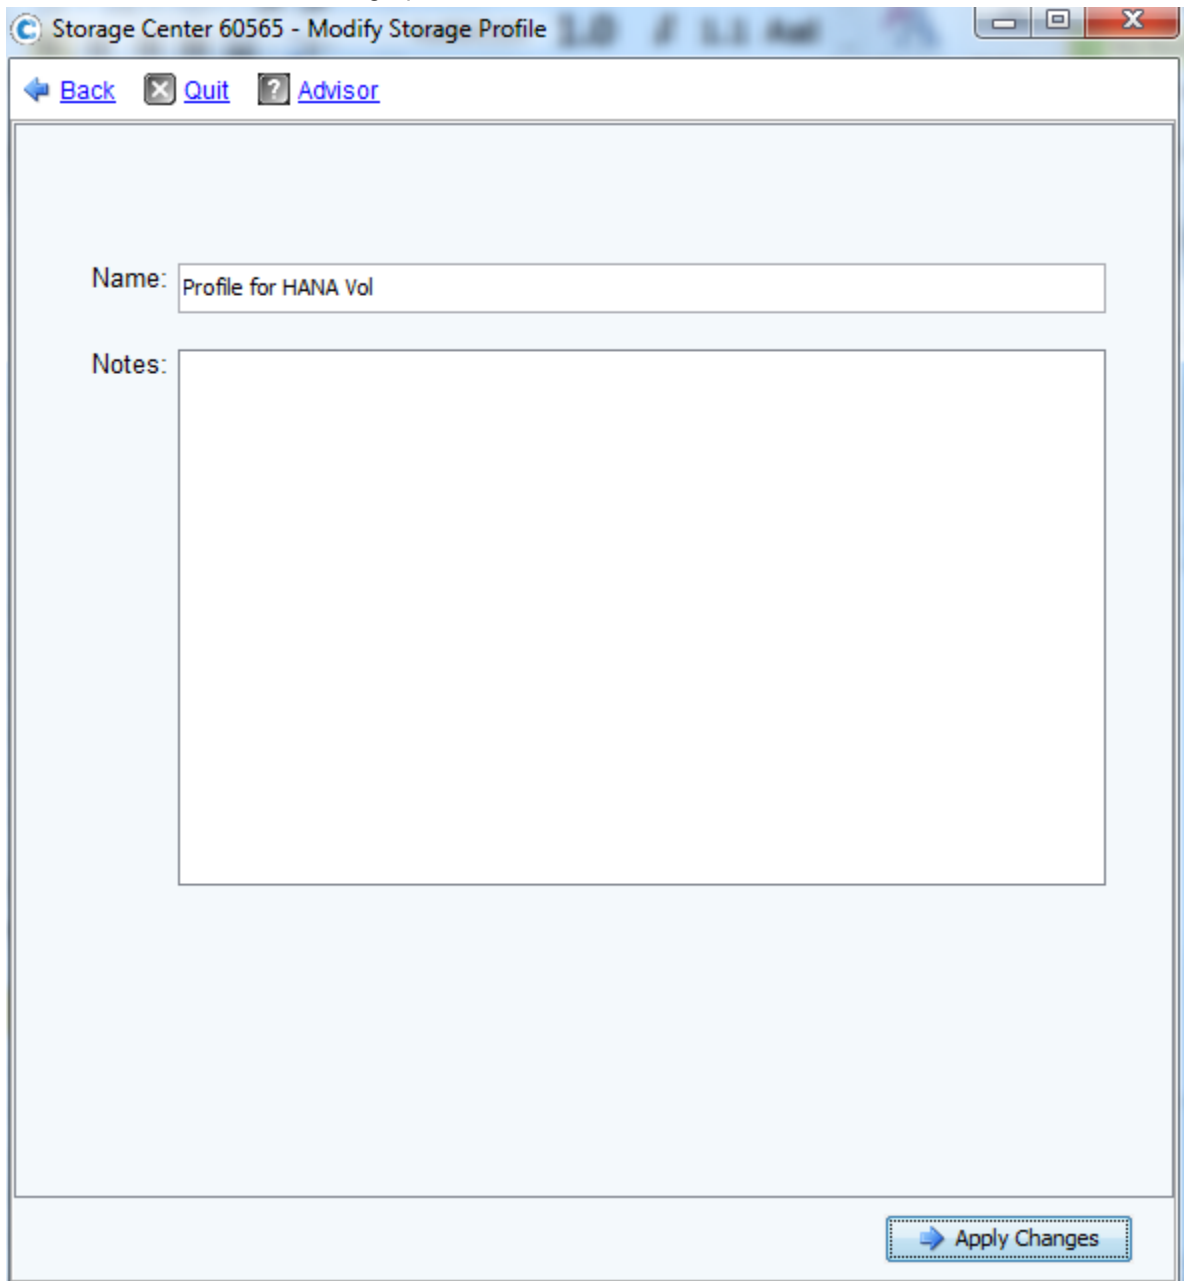
### Tier 3 Storage

	RAID Level	Writable Data	Replay Data
	RAID 10	<input type="checkbox"/>	<input type="checkbox"/>
	RAID 5-5	<input type="checkbox"/>	<input type="checkbox"/>
	RAID 5-9	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

**Notice:** Because both Tier-1 and Tier-3 storage are configured as Single Redundant, the dialog box is shown as above for the initial configuration (as Single Redundant). For later modification the dialog box is shown differently in Single Redundant and Dual Redundant settings. Please ignore a Dual Redundant setting because it is not applicable.

Then Click Continue to proceed.

3. Provide a name for a storage profile as **“Profile for HANA Vol”**.



Storage Center 60565 - Modify Storage Profile

Back Quit Advisor

Name: Profile for HANA Vol

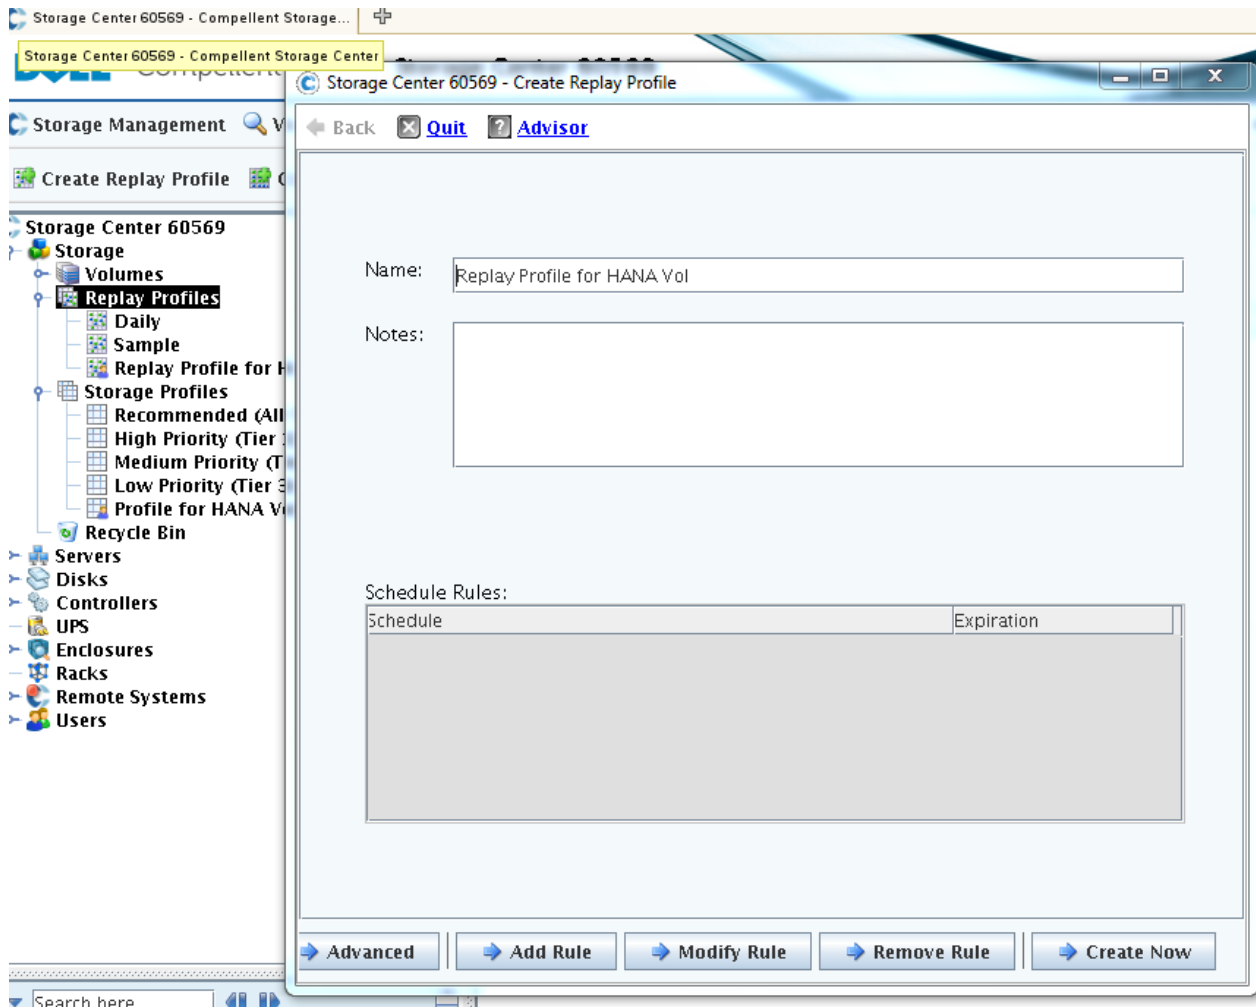
Notes:

Apply Changes

Press “Create Now” to complete the process

## Configuring Replay Profile

- Click on “Create Replay Profile” button to start a Replay Profile for the storage array.



- Provide a name for the replay profile as “Replay Profile for HANA Vol”.
- Click on “Add Rule” button to proceed.

Storage Center 60565 - Modify Replay Profile

[Back](#) [Quit](#) [Advisor](#)

Schedule Type: Daily

☒ at 07 :00 PM

☐ every 15 minutes

☒ between 07 :00 PM and 12 :00 AM

Expiration: 1 days

[Continue](#)

- Select Schedule Type as "Daily"
- Replay started at "7:00PM"
- Expiration: "1 days"
- Press "Continue" to proceed


Storage Center 60569 - Create Replay Profile

Back Quit Advisor

Name:

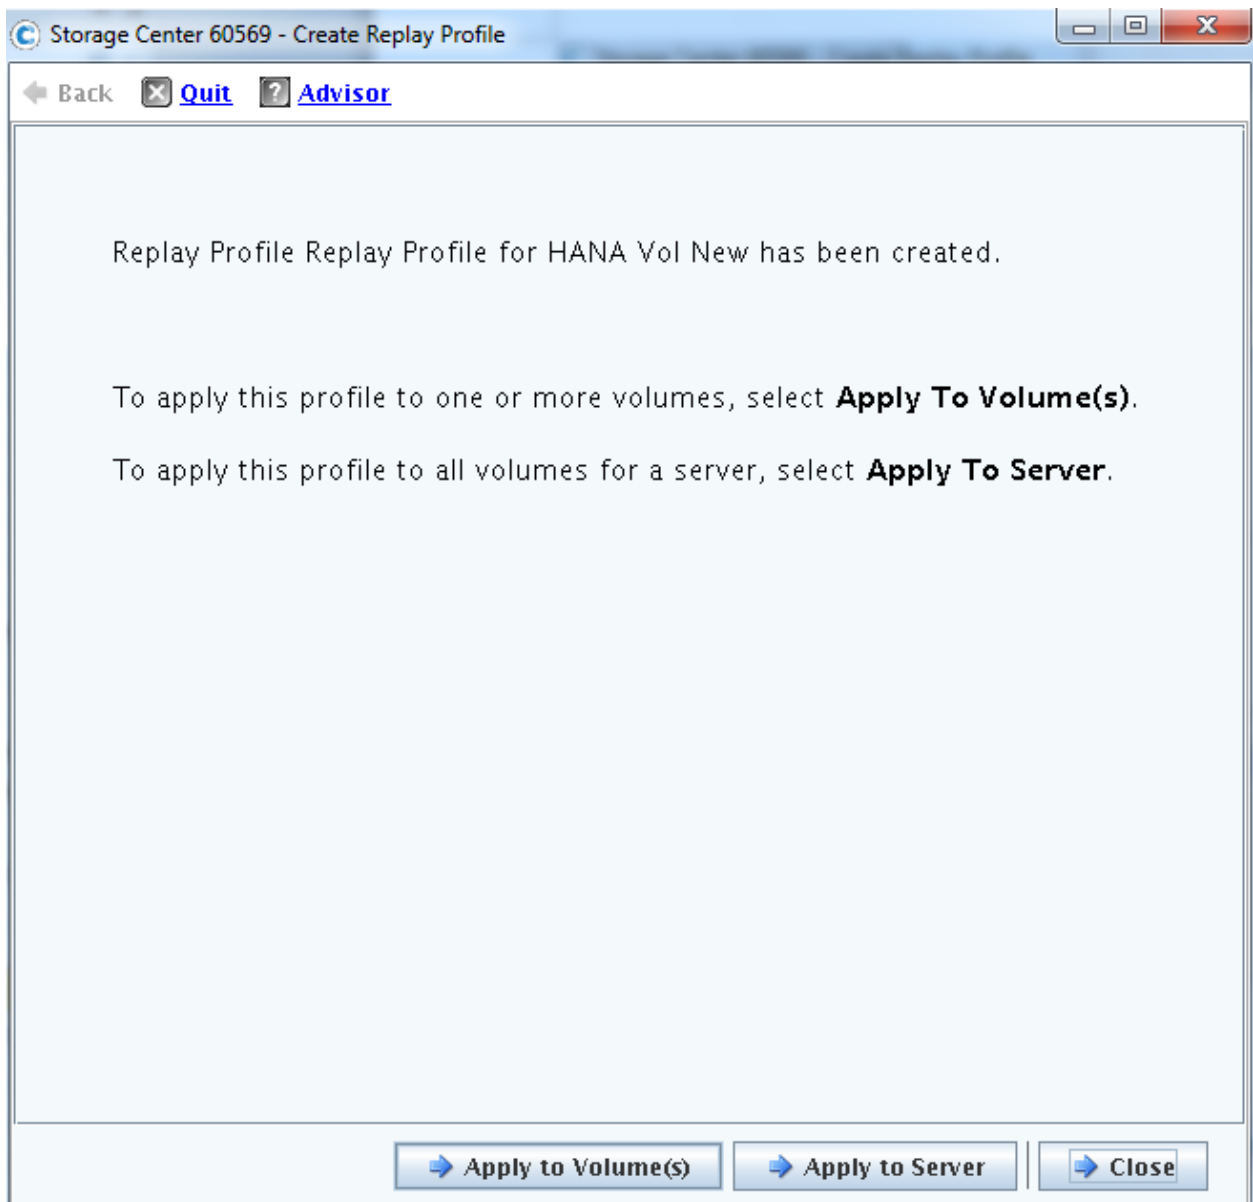
Notes:

Schedule Rules:

Schedule	Expiration
 Daily at 7:00 PM	3 days

Advanced Add Rule Modify Rule Remove Rule Create Now

- Press "Create Now"



- Press "Close" to complete the task.

Repeat these steps for the other storage array if any

## Configure Schedule for Data Progression

**DELL** Compellent **Storage Center 60569**

Storage Management View Refresh Help

Properties Configure My Volume Defaults


**Storage Center 60569**

- Storage
  - Volumes
  - Replay Profiles
    - Daily
    - Sample
    - Custom Replay Profile
    - Custom Replay Profile 2
  - Storage Profiles
  - Recycle Bin
- Servers

**Storage Center 60569**

**Storage Summary** as of Oct 24, 2012

3.88 TB




A horizontal bar chart showing storage usage. The bar is divided into two sections: a dark blue section representing used space and a light gray section representing free space. The value '3.88 TB' is displayed below the bar, indicating the amount of used space.



Storage Center 60569 - System Properties - Storage Center 60569

General Storage Data Progression Cache Disk Sparing Info

 Storage Center 60569

Type: System

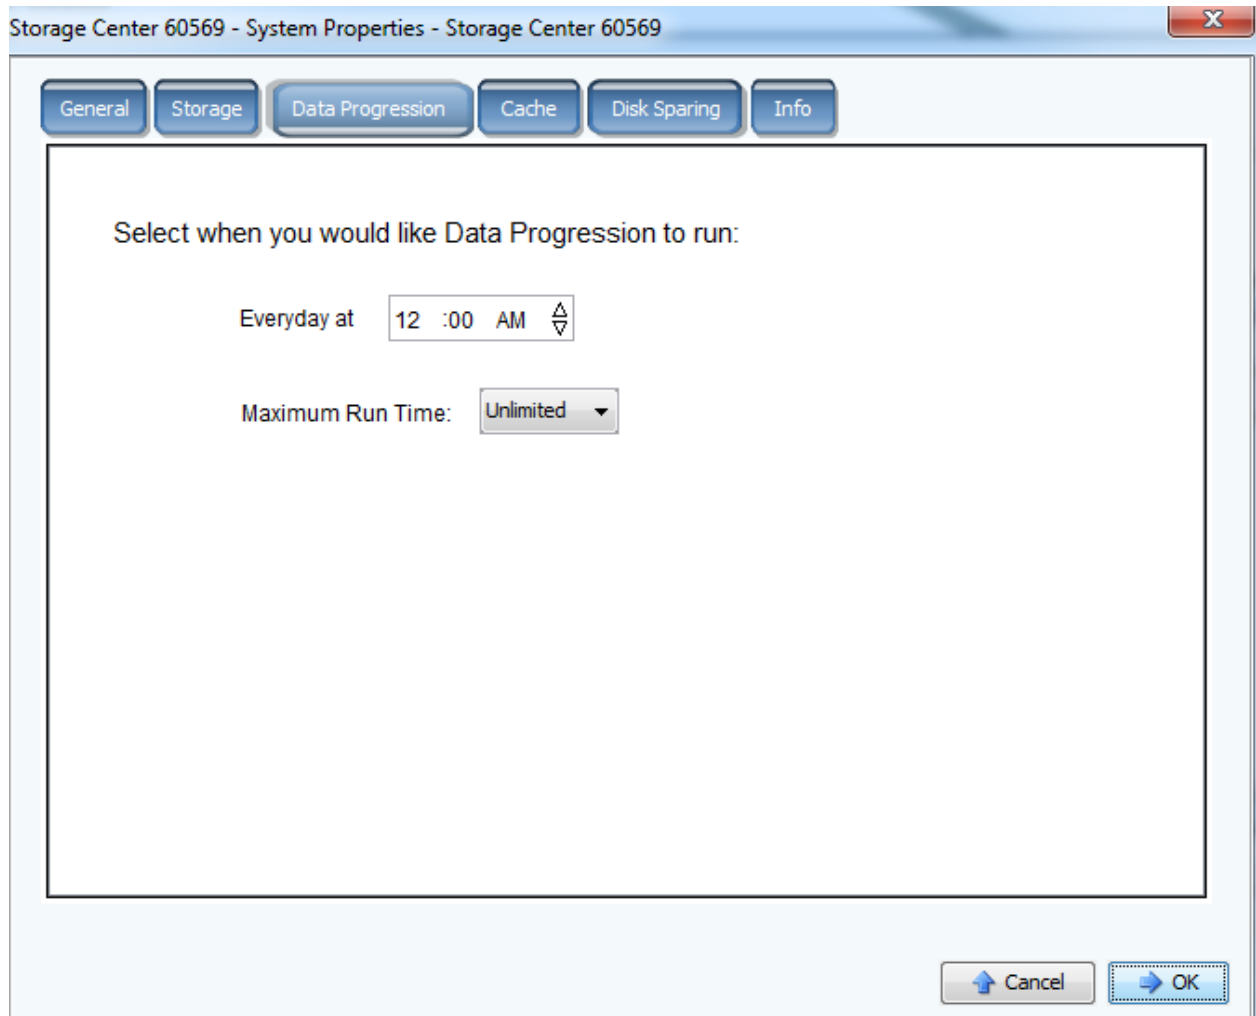
Storage Center ID: 60569

Management IP Address: 172.16.5.20

Operation Mode: Normal

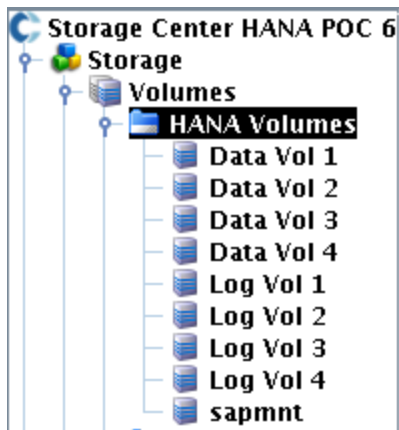
Cancel OK

- Click on "Data Progression" Tab



## Setting up Volumes

1. On each storage array, create up to 8 volumes for DATA (3xRAM/node) and up to 8 (512GB each) volumes for LOG for 8 HANA worker nodes. Also on the first array, create an additional LUN 'sapmnt' which should be sized 512GB for a set of 4 HANA nodes (eg. For 8 nodes, 'sapmnt' should be sized to 1TB) which will be used for /hana/shared as the NFS mount point OR formatted with OCFS2 for SLES deployment and GFS2 for RHEL if not using the FS8610 appliance for NFS option.



## Setting up Hosts

2. create a server cluster comprising of all HANA nodes which will be accessing the volumes on the storage array:

