

Implementing HP Enterprise Virtual Array Solutions

ESG9710SG0306

student guide



Implementing HP Enterprise Virtual Array Solutions

ESG9710SG0306



training

© Copyright 2003 Hewlett-Packard Development Company, L.P.

All other product names mentioned herein may be trademarks of their respective companies.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for HP products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

Implementing HP Enterprise Virtual Array Solutions

Student Guide

June 2003

Course Overview

Introduction	1
Course Objectives	2
Course Prerequisites	3
Module Overview	4
Documentation References	6
Enterprise Virtual Array 5000 Documentation	6
Enterprise Virtual Array 3000 Documentation	6
White Papers	7
Other References	8
Technical Bulletin Board	8
HP Inline	8
HP Internet	9

Module 1 — Product Overview

Overview	1
Objectives	2
Introduction	3
Product Features	4
Key Customer Value	5
Product Positioning	7
Virtualization Technology	8
Functional Layout of the Enterprise Virtual Array Solution	9
Physical Layout of the Enterprise Virtual Array Solution	10
Physical Layout	12
Rack	13
HSV110 Controller Enclosures	18
Disk Drive Enclosures	19
FC-AL Disk Drives	20
Fibre Channel Loop Switches	21
Power Distribution	22
Enclosure Address Bus	23
Virtual Controller Software	24
Command View EVA	25
Storage Management Appliance	26
Storage System Scripting Utility	27
Product Support	28
Operating Systems Support	29
Boot Support	30
Coexisting Product Support	31
Licensing	33

Event Reporting.....	34
Diagnostics	36
Warranties.....	37
Service Offerings.....	38
Installation and Startup Services.....	38
Basic Post Warranty Support	39
CarePaq Priority Services	39
SAN Implementation Service	40
Enterprise Virtual Array 3000 Overview	41
Features	42
Configuration	43
Product Support.....	44
Learning Check	45

Module 2 — HSV Controller

Overview	1
Objectives	2
HSV Controller Block Diagram	3
Built-In Self Test (BIST) Diagnostics.....	4
Module Integrated Self Test (MIST) Diagnostics.....	5
Controller and Management Appliance Communications	6
Controller Features	7
Virtual Controller Software	8
Controller Enclosure.....	10
Controller Front.....	11
Bezel.....	13
Blowers	13
Power Supplies.....	14
Cache Battery Assemblies	15
Operator Control Panel	18
Controller Rear	26
Rear Connections	28
Device Port LEDs	31
Mirror Port	32
Controller Port Numbers	33
Field Replaceable Units.....	34
Controller Replacement Considerations	35
Controller Swapping and Replacement.....	36
Blowers and Cache Batteries	38
Optical Fibre Channel Cables and Transceivers	39
Fibre Channel Copper Cables	42
Enterprise Virtual Array 3000 Controller.....	43
Rear Connections	43
Controller Port Numbering Changes.....	44
Learning Check	45

Module 3 — Disk Drive Enclosure

Overview	1
Objectives	2
Disk Drive Enclosure Block Diagram	3
Drive Enclosure Front	5
Disk Drives	6
Disk Drive Status Icons	7
Disk Drive Power	8
Disk Drive Blank	8
Drive Enclosure Rear	9
Environmental Monitoring Unit	11
EMU Functions	12
EMU LED Status Displays	13
EMU RJ-45 Connectors	14
EMU Alphanumeric Display and Pushbuttons	15
EMU Pushbutton LEDs	16
EMU Display Groups	16
EMU Firmware Update	34
FC-AL I/O Modules	36
Power Supplies and Blowers	43
Field Replaceable Units	46
Learning Check	47

Module 4 — Fibre Channel Loop Switch

Overview	1
Objectives	2
General Description	3
Switch Features	4
Switch Cabling	5
Switch Port Numbering	6
Functional Overview of the Switch	7
Power-On Self Test	8
Switch LEDs	9
System LEDs	9
Port LEDs	10
Field Replaceable Units	11
Learning Check	12

Module 5 — Rack, Cabling, Configuration, and Initial Setup

Overview	1
Objectives	2
Rack Models	3
Switched and Unswitched Models	3
Opal 42U Rack	4
Graphite 41U Rack	5
Rack Power Distribution	6
Dual Zero-U PDU Assembly	7
Wall Receptacles	8
Power Distribution Modules	9
PDM and PDU Connections	10
PDU and PDM Connectivity to Enclosures	11
PDU Failures	11
Power Component Removal and Replacement	12
Configuration Overview	13
Standard Models	13
Expansion Rack Models	13
Expansion Components and Options	14
Nonswitched Configurations	15
Expansion Panel Description	16
Expansion Panel Location	17
Fibre Channel Cable Configurations	18
Cabling Rules for Nonswitched Configurations	19
Cabling Configurations for Nonswitched Models	20
Switched Configurations	32
FC Loop Switches	33
Fibre Channel Cable Configurations	36
Cabling Rules for Switched Configurations	37
Cabling Configurations for Switched Models	39
Cable Management Configurations	47
Cable Management Arms	47
Cable Containment Spools	48
Additional Expansion Configurations	49
Enterprise Virtual Array 3000 Configurations	50
Model 2C2D Features	50
Model 2C2D Connections	52
Cabling for Loop 1	53
Additional Configurations	54
Storage System Setup	57
Getting Started Flowchart	58
Gathering Information	59
Obtaining a License Key	60
Setting Up the Hardware	64
Entering Data into the OCP	69
Configuring the Storage System	70
Learning Check	71

Module 6 — Storage Management Appliance

Overview	1
Objectives	2
Storage Management Appliance Description	3
Storage Management Appliance	5
Hardware Description for G1 Appliances	6
Hardware Description for G2 Appliances	7
SAMS Operating System	8
Appliance-Based Software	9
Storage Management Appliance Software	10
Command View EVA	11
Command View EVA Rules and Limitations	12
Initial System Configuration	13
Passwords	14
ELM Web Server Login	14
SAMS and Windows Console Login	15
RIB Login	15
Appliance Administration	16
Storage Management Appliance Configuration	17
Storage Management Appliance Availability	18
Management Appliance Restore CD	19
Restoring G1	19
Restoring G2	19
Applying Patches	20
Learning Check	21

Module 7 — Concepts and Terminology

Overview	1
Objectives	2
Virtual Storage Terminology	3
Storage System (Cell)	4
Storage System Metadata	4
Disk Group	5
Default Disk Group	6
Disk Group Metadata	6
Disk Group Capacity	6
Disk Group Member Terms	7
Disk Group Configuration	8
Disk Group Creation	8
Virtual Disk	9
Virtual Disk Creation	10
Distributed Virtual RAID Technology	11
Conventional RAID	12
Data Location Independence	13
Virtual Disk Map	15
Virtual Disk Leveling	16

Virtual Disk Expansion	17
Distributed Sparing	18
Examples with Equal-Size Disk Capacities	20
Example with Mixed-Size Disk Capacities	20
Redundant Storage Sets	21
Examples of RSS Calculations	22
Virtual Disk Space Allocation	22
Hosts and LUNs	25
Snapshot Implementation	26
Fully-Allocated Snapshots	26
Demand-Allocated Snapshots	27
Snapclone Implementation	27
Snapshot and Snapclone Comparison	28
Learning Check	29

Module 8 — Storage System Configuration

Overview	1
Objectives	2
Getting Started with Command View EVA	3
Command View EVA Interface Layout	3
Wizards	7
Command View EVA Online Help	8
Command View EVA Home Page	11
Configuration Process Steps	12
Gathering Preliminary Information	13
Develop a Requirements List Per Server	13
Create a List of all FCAs	13
Determine the Host WWN	14
Starting with the Uninitialized Storage System	15
Performing a Controller Code Load	17
Code Load Initial Prompt	18
Performing a Code Load	19
Initializing the Storage System	22
Metadata	22
Initializing an HSV Storage System	24
Performing a Disk Drive Firmware Code Load	30
Creating a Disk Group after Initialization	31
Using the Create a Disk Group Wizard	32
Advanced Options	34
Viewing and Changing Disk Group Properties	35
Creating a Vertical Disk Group	36
Deleting Disk Groups	37
Creating a Host	38
Creating a Host Folder	39
Adding a Host	41
Viewing and Changing Host Properties	45

Adding a Host FCA.....	47
Creating a Virtual Disk.....	48
Creating a Virtual Disk Folder.....	49
Creating a Virtual Disk Family.....	51
Presenting a Virtual Disk to a Host.....	57
Unpresenting a Virtual Disk to a Host.....	58
Creating a Snapshot.....	59
Using the Snapshot Wizard.....	59
Virtual Disk Active Properties Page for Snapshot.....	63
Creating a Snapclone.....	64
Using the Snapclone Wizard.....	65
Snapclone Properties During Background Process.....	68
Virtual Disks with Snapshots and Snapclones.....	69
Learning Check	70

Module 9 — Storage System Management

Overview	1
Objectives	2
Viewing Management Agent Options.....	3
Storage System Password Access.....	4
Licensing Options	6
User Interface Options	8
Page Footer Message Options.....	11
Viewing Hardware Properties	13
Rack Properties	14
Controller Properties.....	15
Disk Enclosure Properties.....	23
Learning Check	36

Module 10 — Error Reporting and Diagnostics

Overview	1
Objectives	2
Event Handling.....	3
Event Logs.....	4
Monitoring and Notification Tools.....	5
Proactive Remote Services.....	5
Event Forwarding Using SNMP Traps	6
System Event Analyzer.....	6
Open Service Event Manager.....	6
Viewing Event Logs.....	7
Management Agent Event Log	11
Controller Event Log.....	16
Controller Termination Event Log.....	21
NT Event Log.....	25
Tools for Translating Controller Events	26

Event Message Formats	31
Interpreting Messages with Physical Addresses	34
Interpreting SCMI Event Codes	36
Analyzing Event Logs	37
Inband Host Event Reporting	41
Event Notification	43
Set Event Notification Options	43
Configure Host Notification Options	50
Notification Using Proactive Remote Services	53
PRS Functional Overview	54
Prerequisites for PRS on the Enterprise Virtual Array	57
Configuring PRS with the Enterprise Virtual Array	58
Notification Using System Event Analyzer	63
System Event Analyzer Architecture	64
Configuring SEA with the Enterprise Virtual Array	64
Disabling SNMP Trap Notification	64
Notification Using Open Service Event Manager	65
Diagnostics Overview	66
BIST and MIST	66
Test Entities and Test Numbers in EIP Type 14	67
Cache Memory Scrubber	69
Controller Temperature Sensing	70
Learning Check	71

Module 11 — Storage System Scripting Utility

Overview	1
Objectives	2
SSSU Description	3
SSSU Architecture	4
Terminology	5
Installing SSSU	6
Starting the SSSU	7
Syntax	7
Command Structure	8
Getting Help	8
File Paths and Naming Conventions	9
SELECT Command	10
SELECT MANAGER	10
SELECT SYSTEM	10
ADD Command	11
ADD SYSTEM	11
ADD COPY	11
ADD FOLDER	12
ADD DISK_GROUP	12
ADD HOST	12
ADD LUN	12

ADD SNAPSHOT	12
ADD VDISK	13
SET Command	14
SET SYSTEM	14
SET DISK	14
SET FOLDER	14
SET DISK_GROUP	14
SET HOST	14
SET MONITOR	14
SET VDISK	15
SET OPTIONS	16
DELETE Command	17
DELETE SYSTEM	17
DELETE FOLDER	17
DELETE DISK_GROUP	17
DELETE HOST	17
DELETE LUN	17
DELETE VDISK	18
CAPTURE CONFIGURATION Command	19
Description of Script Files	20
Restoring Configurations	21
SHOW Command	22
SHOW SYSTEM	22
SHOW DISK	22
SHOW DISK_GROUP	22
SHOW HOST	22
SHOW LUN	22
SHOW MANAGER	22
SHOW MONITOR	22
SHOW OPTIONS	23
SHOW POWER	23
SHOW VDISK	23
SHOW WORLD_WIDE_NAME	23
SHUTDOWN Command	24
RESTART Command	25
FILE Command	25
EXIT Command	26
HELP Command	26
Learning Check	27

Module 12 — Best Practices

Overview	1
Objectives	2
Review of Enterprise Virtual Array Rules and Limitations	3
Rules for Hosts, Connections, and LUNs	3
Other Rules for the Enterprise Virtual Array	4
Available Tools and Documentation	5
Disk Group Configuration Overview	6

Disk Group Configuration Considerations	7
Disk Group Availability	8
Disk Group Performance	11
Disk Group Capacity Utilization.....	13
General Steps for Configuring Disk Groups	15
Disk Group Best Practices Summary	17
Disk Group Sizing	18
Disk Group Sizing Factors	18
Formula for Determining Disk Count.....	21
Disk Group Sizing Procedure Summary	22
Storage Calculation Utilities.....	23
Disk Group Sizing Utility	23
Disk Groups and Virtual Disks Calculator.....	23
Learning Check	24

Module 13 — Troubleshooting

Overview	1
Objectives	2
Troubleshooting Background and Approach.....	3
Developing a Troubleshooting Plan	4
Troubleshooting Aids and Tools	5
Access to Troubleshooting Aids	6
Event and Error Logs	7
Command View EVA GUI	10
SSSU Commands	11
HSV Controller Status Indicators	12
Disk Drive Enclosure Status Indicators	13
FC Loop Switch LEDs	15
FC Switch Commands and Status Indicators.....	16
Fibre Channel Adapter LEDs.....	18
Emulex Configuration Utility	19
LightPulse Utility/NT	20
Storage Management Appliance LEDs.....	22
Troubleshooting References	23
Documents	23
Websites	23
Tips for Analyzing Controller Event Logs.....	24
Replaceable Components	25
Identifying Spare Part Numbers.....	25
Learning Check	26

Appendix A — EMU Error Condition Reports

Overview	1
Condition Report Format.....	2
Navigating the Error Display.....	3

Element Types	4
Drive Conditions	5
Drive Conditions	5
Power Supply Conditions	8
Blower Conditions	10
Temperature Conditions	14
EMU Conditions	17
Transceiver Conditions	23
Voltage Sensor and Current Sensor Conditions	25
Backplane NVRAM	28
I/O Module Conditions	30
Host Interconnect	33

Appendix B — Troubleshooting the Management Appliance

Overview	1
Operational Status	2
System Status LED Indicators	3
Hard Drive LED Indicators	4
RJ-45 Connector Network Status LED Indicators	5
Fibre Channel Adapter Indicators	6
Troubleshooting Open SAN Manager	7
Troubleshooting OSM Login	8
Problems Connecting to Storage Management Appliance	10
OSM Navigation Pane	11
Troubleshooting Appliance Manager	12
Troubleshooting Product Installation	13
Installing or Updating Applications	13
CD Install	14
Network Install	15
Application Status	16
Troubleshooting Resource Managers	17
Element Manager for HSG	17
Troubleshooting Open SAN Manager Applications	20
Storage Allocation Reporter	20

Appendix C — Space Allocation

Overview	1
Formatted Capacity of Disk Drives	2
Formatted Capacities Using Binary Equivalents	2
Disk Group Metadata Overhead	3
Total Storage Space and Available Space	4
Space Available to Create Virtual Disks	6
Virtual Disk Metadata Overhead	7
VRAID Redundancy Overhead	7
Storage Space Used — Example	8
Disk Group Failure Protection Levels	9

Protection Level of Single.....	9
Protection Level of Double	10
Disk Occupancy Levels.....	11
Adding Up the Numbers.....	13
Pulling Out a Disk Group Member Improperly	16
Selecting the Manual Option.....	17
Selecting the Automatic Option.....	18
Disk Failure Protection Space Usage	19
Example Using Spare Capacity.....	20
Disk Sizing Example	25

Appendix D — Learning Check Answers

Introduction	1
Module 1 — Product Overview	2
Module 2 — HSV Controller	3
Module 3 — Disk Drive Enclosure	4
Module 4 — Fibre Channel Loop Switch	5
Module 5 — Rack, Cabling, Configuration, and Initial Setup	6
Module 6 — Storage Management Appliance	7
Module 7 — Concepts and Terminology	8
Module 8 — Storage System Configuration	9
Module 9 — Storage System Management.....	10
Module 10 — Error Reporting and Diagnostics.....	11
Module 11 — Storage System Scripting Utility.....	12
Module 12 — Best Practices	13
Module 13 — Troubleshooting	14

Introduction

This course is designed to provide training for the HP StorageWorks Enterprise Virtual Array family, including the Enterprise Virtual Array 5000 (eva5000) and Enterprise Virtual Array 3000 (eva3000). The material describes Enterprise Virtual Array hardware, how the Enterprise Virtual Array storage system implements virtualization technology, including its use of Virtual Controller Software (VCS), and provides instructions on how to set up, configure, and manage the Enterprise Virtual Array using the storage management appliance and Command View EVA software. A discussion of error reporting mechanisms, best practices, and troubleshooting sources complete the course materials.

As you progress through the course, hands-on laboratories are available. These activities, led by the classroom instructor, are designed to correlate with the lecture materials.

This training release is for the following Enterprise Virtual Array models:

- Enterprise Virtual Array 5000 (eva5000), consisting of the following:
 - Updated software (Enterprise Storage System Software V3.000, VCS 3.00)
 - Updated software (Enterprise Storage System Software V2.000, VCS 2.003)
 - Latest versions of the storage management appliance software (Storage Management Appliance Software V2.0, SP3) and element manager software (Command View EVA V3.0)
- Enterprise Virtual Array 3000 (eva3000) VCS Version 2.004, consisting of the following:
 - Updated software (Enterprise Storage System Software V2.000, VCS 2.004)
 - Latest versions of the storage management appliance software (Storage Management Appliance Software V2.0, SP2) and element manager software (Command View EVA V2.1)

Note

This course discusses the eva5000 unless otherwise noted. Many of the general features of the eva5000 are applicable to the eva3000. Where eva3000 features are different, they are noted at the end of the appropriate module.

Course Objectives

After completing this course, you should be able to:

- Describe the features and functions of the Enterprise Virtual Array hardware components.
- Identify the physical, management, and interconnect features and functions of the HSV controller, disk drive enclosure, FC-AL drives, and FC loop switch.
- Describe the rack models, power cabling, and cabling configurations for all models.
- Describe how to perform initial setup of the Enterprise Virtual Array components.
- Define the new terminology associated with the Enterprise Virtual Array and virtual storage.
- Describe how the Enterprise Virtual Array implements virtualization technology.
- Use Command View EVA to set up, configure, and manage the Enterprise Virtual Array.
- Describe all aspects of Enterprise Virtual Array error reporting and diagnostics generation.
- Describe how to set up and use the Storage System Scripting Utility.
- Describe storage configuration best practices.
- Identify and use troubleshooting sources and tools.
- Identify FRUs and CRUs, and perform removal and replacement procedures.

Course Prerequisites

You should have successfully completed the following courses:

- HP Storage Technologies Web-based Training
- HP StorageWorks Full-Line Product Technical Web-based Training
- HP Storage Software and Solutions Full-Line Technical Web-based Training
- HP Entry-level SAN Solutions

Students should be aware of basic host platform usage and be comfortable navigating through system management utilities associated with the specified platforms.

Module Overview

- **Module 1 — Product Overview**
Describes the major components of the Enterprise Virtual Array storage system.
- **Module 2 — HSV Controller**
Describes in detail the physical, management, and interconnect features of the HSV100 and HSV110 controllers.
- **Module 3 — Disk Drive Enclosure**
Describes in detail the physical, management, and interconnect features of the M5114 and M5214 disk drive enclosures and the FC-AL drives.
- **Module 4 — Fibre Channel Loop Switch**
Describes in detail the physical, management, and interconnect features of the FC loop switches.
- **Module 5 — Rack, Cabling, Configuration, and Initial Setup**
Describes how the Enterprise Virtual Array components are integrated and describes supported configurations and initial storage system setup.
- **Module 6 — Storage Management Appliance**
Describes the major components of the Enterprise Virtual Array storage solution, focusing on the storage management appliance.
- **Module 7 — Concepts and Terminology**
Describes the most important Enterprise Virtual Array terms and concepts.
- **Module 8 — Storage System Configuration**
Describes how to set up and configure the Enterprise Virtual Array using Command View EVA running on the storage management appliance.
- **Module 9 — Storage System Management**
Describes how Command View EVA provides management resources such as defining passwords, and monitoring hardware property displays and event logs.
- **Module 10 — Error Reporting and Diagnostics**
Describes the various event logs and diagnostic sources for the Enterprise Virtual Array.

- **Module 11 — Storage System Scripting Utility**
Describes the Storage System Scripting Utility and outlines the commands available to you.
- **Module 12 — Best Practices**
Describes the suggested best practices to enhance Enterprise Virtual Array storage subsystem performance. Also included are some disk group sizing utilities that provide the data to establish an optimum storage environment.
- **Module 13 — Troubleshooting**
Provides detailed information on troubleshooting resources for the Enterprise Virtual Array and all of its major components. A list of customer replaceable units (CRUs) and field replaceable units (FRUs) is also provided.
- **Appendix A — EMU Error Condition Reports**
Use this appendix in conjunction with Module 3, Disk Drive Enclosure, and Module 13, Troubleshooting. This appendix includes a summary of information taken from the *HP StorageWorks Enterprise Virtual Array Service Manual*. This supplies a single, condensed source for quickly diagnosing EMU errors.
- **Appendix B — Troubleshooting the Management Appliance**
Use this appendix in conjunction with Module 6, Storage Management Appliance, and Module 13, Troubleshooting. This appendix helps you identify the initial steps when troubleshooting the appliance, use standard troubleshooting techniques to identify a problem with the appliance and determine its solution, and identify key troubleshooting areas for applications that reside on the appliance.
- **Appendix C — Space Allocation**
Use this appendix in conjunction with Module 7, Concepts and Terminology, and Module 8, Storage System Configuration. This appendix contains a screen-by-screen view of what happens when disk groups and virtual disks are created, or when a disk drive is removed.
- **Appendix D — NAPP Program Documentation**
Use this appendix in conjunction with Module 10, Error Reporting and Diagnostics. This appendix gives you the necessary background information for running the Not a Post Processor (NAPP) program, a tool that performs bit-to-text translation on both Controller Event logs and Controller Termination Event logs.
- **Appendix E — Learning Check Answers**
This appendix supplies you with all of the answers to the learning checks in each module of this course.

Documentation References

The documents listed in this section are the latest versions of documents that support the release of the eva5000 using VCS V3.0 and the eva3000 using VCS V2.004.

INTERNET

For complete information about all documents pertaining to the Enterprise Virtual Array, see the HP StorageWorks Enterprise Virtual Array Technical Documentation home page, available at <http://h18000.www1.hp.com/products/storageworks/enterprise/documentation.html>

Enterprise Virtual Array 5000 Documentation

The HP StorageWorks Enterprise Virtual Array Read Me First document contains references to the following groups of eva5000 documents:

- Enterprise Virtual Array documents
- HP OpenView Storage Management Appliance documents
- HP StorageWorks Operating System Kit documents

Command View EVA V3.0 documents include the following:

- Getting Started Guide
- Installation Instructions
- Release Notes

Enterprise Virtual Array 3000 Documentation

The HP StorageWorks Enterprise Virtual Array 3000 Catalog of Associated Documentation is a web-based document that contains references to the following groups of eva3000 documents:

- Enterprise Virtual Array 3000 documents
- HP OpenView Storage Management Appliance documents
- HP StorageWorks Operating System Kit documents

Command View EVA V2.1 documents include the following:

- Read Me First
- Getting Started Guide
- Installation Instructions
- Release Notes

White Papers

- *HP StorageWorks SAN Design Reference Guide*,
Part Number: AA-RMPNF-TE
- *HP StorageWorks Enterprise Virtual Array: Performance*,
Part Number: 16ZQ-0702A-WWEN
- *HP StorageWorks Enterprise Virtual Array: Use in HP-UX Environment*,
Part Number: 16ZX-0702A-WWEN
- *Connecting Single HBA Servers to the Enterprise Virtual Array without Secure Path*,
Part Number: 16V4-1002B-WWEN
- *Best Practices To Baseline The Performance Of EVA-Based Configurations*,
November 2002
- *Configuring Disk Groups and Virtual Disks in a StorageWorks Enterprise Virtual Array by Compaq*,
April 2002
- *Sizing Disk Groups in a StorageWorks Enterprise Virtual Array by Compaq*,
April 2002

Other References

Technical Bulletin Board

The Technical Bulletin Board, also called the TechBB or SE Forum, provides a place for participants to have an open exchange of information about storage products. The forum for HSV Storage Systems (Enterprise Virtual Array) is moderated by Storage Engineering and Storage Support.

INTRA NET	The link to the TechBB is http://seconnection.inet.cpqcorp.net:90/pages/static/forums.html
------------------	--

To view the topics of discussion for the Enterprise Virtual Array, access the *Storage* category, then *HSV Storage Systems*.

HP Inline

Electronic documentation is available through the HP intranet by going to the Network Storage Solutions (NSS) product documentation page.

INTRA NET	The website for electronic documentation is http://storage.inet.cpqcorp.net/application/view/menu_products.asp
------------------	--

To access Enterprise Virtual Array documentation, select *Product Families* → *Storage Array Systems* → *Virtual Array Systems* → *hp StorageWorks Enterprise Virtual Array*.

A single source for all technical product support information is available through Cybrary.

INTRA NET	The website for Cybrary is http://cybrary.inet.cpqcorp.net/cybrary.html
------------------	--

For Enterprise Virtual Array information, under *Storage*, select *Solutions and Enclosures*. Select the *Enterprise Array* link, then *EVA5000(HSV110)*.

A source for documentation on the disk drive enclosure is also available.

INTRA NET	The website for the drive enclosure is http://sbuwww.cxo.cpqcorp.net/mvs-hwe/moongazer/mvs-hwe_emu.asp
------------------	---

HP Internet

Electronic documentation is available from the HP Internet.

INTERNET	The website for Enterprise Virtual Array 5000 documentation is http://h18006.www1.hp.com/products/storageworks/enterprise/documentation.html
-----------------	--

Downloadable software and drivers for storage products are available from the HP Internet.

INTERNET	The website for downloadable software and drivers is http://h18007.www1.hp.com/support/files/storage/index.html
-----------------	--

Overview

This module provides an overview of the Enterprise Virtual Array storage system and all of its major hardware and software components. Each component is covered in more detail in later modules.

Note

This material describes the eva5000 unless otherwise noted in the topic. Many of the general features of the eva5000 are applicable to the eva3000. Where eva3000 features are different, they are noted at the end of the module.

Objectives

After completing this module, you should be able to:

- Recognize product features of the Enterprise Virtual Array.
- Describe the functional layout of the Enterprise Virtual Array.
- Describe the physical layout of the Enterprise Virtual Array.
- Recognize the attributes and features of Enterprise Virtual Array hardware, including:
 - Racks and rack models
 - HSV controller enclosures
 - Fibre Channel disk drive enclosures and FC-AL disks
 - FC loop switch
 - Power cables and enclosure address bus connections
- Recognize the attributes and features of Enterprise Virtual Array software, including:
 - Virtual controller software
 - Command View EVA
- Recognize the supported operating systems.
- Identify the current types of warranties and service offerings.

Note

The HSV Element Manager has been rebranded as Command View EVA V3.0 for the eva5000 (VCS V3.0) and Command View EVA V2.1 for the eva3000 (VCS V2.004).

Introduction



The Enterprise Virtual Array is a high-performance, high-capacity, and high-availability RAID storage solution. It complements the HP StorageWorks product line and can coexist on the same Storage Area Network (SAN). The Enterprise Virtual Array is designed for the data center that needs to configure high-capacity systems having application-specific demands for high performance.

Product Features

Features of the Enterprise Virtual Array include:

- **Use of Fibre Channel (FC) technology** — Takes advantage of the benefits of FC in distance, performance, and connectivity. The Enterprise Virtual Array is 2Gb/s end to end.
- **Packaging** — Solutions are built with new storage enclosure packaging.
- **High storage density, OEM wiring benefit** — Fully wired solution, which reduces installation time.
- **No single point of failure** — Redundant architecture and value-added software eliminates single points of failure from server to storage in clustered or single-server configurations.
- **Clustered server and high availability system support** — Dual and multinode cluster support for host-level fault tolerance and high system availability.
- **Multivendor platform** — Support for industry-leading platforms including: Windows NT, Windows 2000, Windows Server 2003, HP Tru64 UNIX, HP OpenVMS, HP-UX, Sun Solaris, IBM AIX, Red Hat and SuSE Linux, and Novell NetWare.
- **Multiserver shared support for storage consolidation** — Heterogeneous and homogeneous host support allows you to share storage between multiple servers.
- **Support for larger capacities** — Support for up to 35TB maximum raw capacity by using an expansion rack.

Key Customer Value

Most customer decision-maker audiences relate to the three information technology (IT) challenges: lack of time, lack of money, and lack of people. These three challenges are translated into Enterprise Virtual Array “value opportunities.”


The key virtualization value opportunities enabled by Enterprise are listed below. All affect cost, and result in a much lower effective price per MB.

- Significantly higher utilization of purchased capacity
 - Up to twice the typical 40% to 50% open systems utilization
 - Importance to business:
 - ◆ Minimizes purchase of unusable capacity
 - ◆ Just-in-time capacity increments
 - ◆ No performance drop
 - ◆ Simplifies storage administration
 - Much lower effective price/MB
- Virtually capacity-free (demand-allocated) snapshot replication:
 - Minimal space reserved up front. Space is only used as the original virtual disk changes data.
 - Importance to business:
 - ◆ Eliminates the need to purchase replication capacity
 - ◆ Catalyzes application development
 - ◆ Enables new backup methodologies
 - ◆ Allows expedited restores
 - Much lower effective price/MB
- Virtually instantaneous business copy replication (snapclone):
 - Data is accessible before data finishes copying
 - Importance to business:
 - ◆ Catalyzes data mining
 - ◆ Mitigates performance contention
 - ◆ Facilitates seamless data migration/sharing
 - ◆ Avoids access scheduling overhead/complexities
 - ◆ Provides clones with full VRAID protection
 - ◆ Provides a foundation for life cycle data management
 - Harmonious mining and exploitation of business data in a time-critical fashion

- Automatic, self-tuning performance
 - Performance no longer limited by bounds of traditional disk array architectures
 - Importance to business:
 - ◆ No data micro-management
 - ◆ Truly scalable capacity and performance
 - ◆ Configuration flexibility
 - ◆ Improved and more consistent user service levels
 - Virtually boundless performance consistently delivered without continuous manual intervention
- Greatly reduced planned and unplanned application downtime
 - Online firmware downloads
 - Battery backup for cache (up to 96 hours)
 - Multiple point-in-time copies with Vsnap
 - ◆ Facilitates more frequent point-in-time backups
 - ◆ Enables expedited restores through reduced forward recovery of transaction logs
 - Importance to business:
 - ◆ Eliminates storage reconfigurations
 - ◆ Minimizes backup windows
 - ◆ Accelerates recovery
 - ◆ Assures business continuance
- Simplified management gives lowest total cost of ownership (TCO)
 - Up to five times storage administration productivity improvements through the use of LiteTouch management
 - Importance to business:
 - ◆ Higher actual capacity utilization
 - ◆ Dynamic LUN and pool expansion
 - ◆ More disks per controller pair
 - ◆ Simplified storage administration
 - ◆ Minimized performance tuning
 - Lowest effective price/MB + highest managed GB/person = lowest TCO

Product Positioning

The Enterprise Virtual Array family is positioned in the upper midrange and lower upper-range part of the HP StorageWorks Online Storage portfolio.



The following describes the midrange products:

- va7110 and va7410
 - Storage and HP-UX servers integrated into a single rack
 - Direct connect (arbitrated loop) or fabric connect
 - No requirement for array-based Continuous Access
- eva3000
 - Upper midrange array
 - Broad operating system support
 - Future plans for continuous access
 - Outstanding performance
 - Up to 8TB capacity
- eva5000
 - Lower upper-range array
 - Capacity growth up to 35TB
 - Maximum performance
 - 168k cache iops, 560MB/s sustained throughput


Virtualization Technology

The Enterprise Virtual Array implements virtualization at the storage system level. Other products use virtualization at different levels. For example, Virtual Replicator uses virtualization at the host level, and the HP OpenView Continuous Access Storage Appliance (CASA) uses virtualization at the fabric level.

Features of storage system level virtualization include:

- Physical storage is abstracted.
- Storage pools are created from physical blocks of storage.
- Virtual disks are created from the storage pool.
- Physical devices and distribution of capacity are transparent to servers and applications.

The following graphic depicts storage virtualization:



Virtualization concepts and terminology are covered in detail in Module 7, Concepts and Terminology.


Functional Layout of the Enterprise Virtual Array Solution

The current trend is to view data storage like a utility, such as electricity or water. Like a utility, storage is centrally managed, always there (except in catastrophic instances), and can be expanded and reconfigured to fit changing needs.

The Enterprise Virtual Array implements this strategy in the following way:

- Centralized management software controls the storage system and lets an administrator configure the storage required by various host computers.
- The management software creates pseudo-disks, called virtual disks, and presents them to a host (or hosts).
- The host sees the virtual disk just as it sees any ordinary disk drive.

The following figure shows the basic functional layout of the Enterprise Virtual Array.



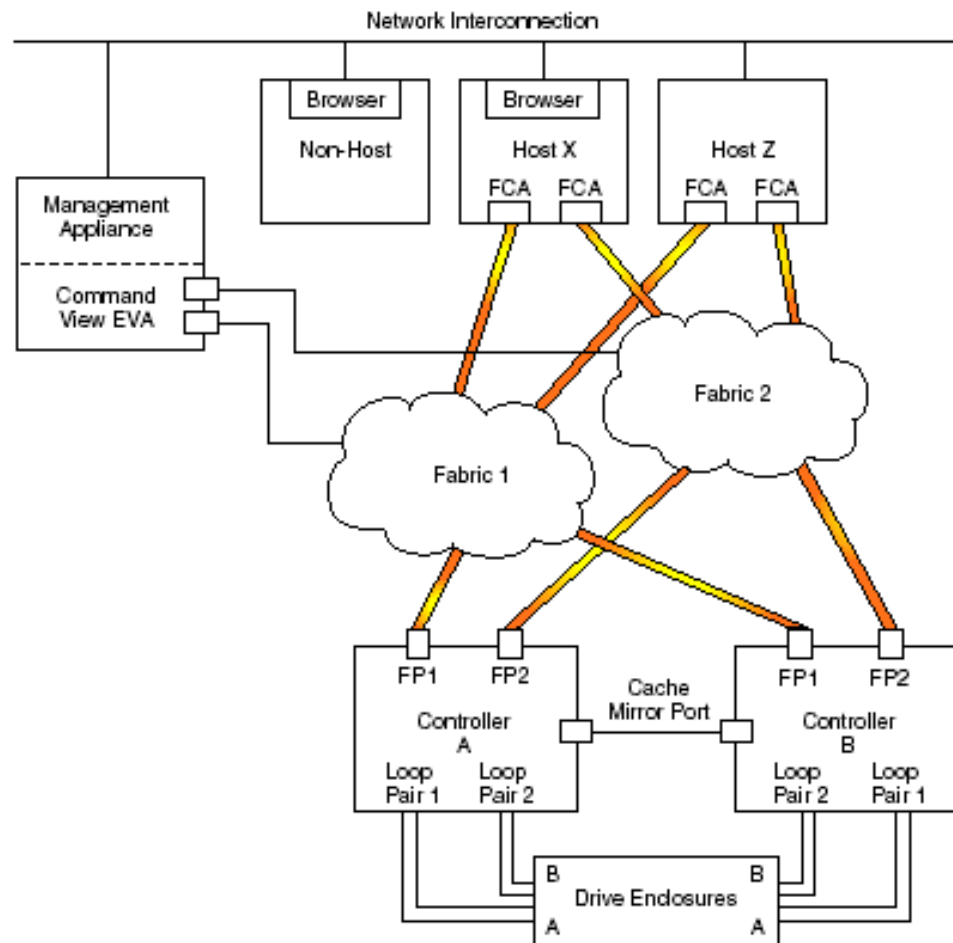
Physical Layout of the Enterprise Virtual Array Solution

Physically, the storage system consists of a pair of controllers and their array of physical disk drives. The physical disk drives plug into disk drive enclosures.

The controller pair is cabled to the disk drive enclosures. A backplane in the disk drive enclosures distributes commands and data to the disk drives.

- The controller pair connects to two Fibre Channel fabrics, to which the hosts also connect.
- The Command View EVA software controls the storage system. It resides on the storage management appliance. The storage management appliance connects into the fabric.
- The controller pair connects to the physical disk array through FC arbitrated loops. The eva5000 switched hardware implements this through four FC loop switches or through nonswitched hardware (expansion panels). The eva3000 uses a configuration without loop switches.
- Two separate loop pairs are used: loop pair 1 and loop pair 2. Each loop pair consists of two loops, each of which run independently and can take over for the other loop in case of failure.

The following block diagram depicts the Enterprise Virtual Array connections.




FP = Fibre (Host) Port
FCA = Fibre Channel Adapter

Note

A connection is supported from the Management Appliance to both fabrics as of Management Appliance Software Update V1.0C, SP4.

Physical Layout



The physical layout of the Enterprise Virtual Array consists of the following primary components:

- Rack
- Pair of HSV110 controllers
- Drive enclosures containing array of physical disk drives
- Loop switches or expansion panels
 - Enclosures connect to FC loop switches in **switched** configurations
 - Enclosures connect to other enclosures in **nonswitched** (expansion panel) configurations

All of these components are described in later modules.

Rack


Two types of racks are available with the Enterprise Virtual Array 5000:

- 42U Cannondale 9000 rack — Available in opal with a depth of 909mm/35.8" with industry-standard 19" mounting rails.
- 41U Millenium rack — Available in graphite with a depth of 993mm/39.1" with industry-standard 19" mounting rails.

Note

1U = 1.75 inches.

Both racks allow the mounting of the standard 19-inch (483 mm) wide, HSV110 controller assembly, Model 5214 drive enclosures, and FC loop switches. The 42U rack can also support expansion panels.



Note

Both racks include four FC loop switches, two above and two below the controller pair. Because the 42U racks of Enterprise Virtual Array V1 contained expansion panels, not FC loop switches, customers can purchase FC loop switch upgrade kits to upgrade their configurations.

Rack Models



42U Opal

41U Graphite

The rack models are the 42U opal and 41U graphite. The 42U opal rack may or may not contain loop switches, whereas the 41U graphite rack does contain loop switches. The front and rear view are shown.


Note

The 42U rack is being phased out mid-2003.

Rack Model Configurations


Three predefined rack models are available in opal and graphite:

- 2C6D — General purpose
- 2C12D — General purpose
- 8C8D — High-performance, scientific



The Model 2C2D is the latest model and has the following features:

- Graphite rack only (41U)
- Field-upgradeable to 2C12D by adding enclosures and loop switches



Rack Capacity

Model	Drive Bays	Maximum Capacity		
		36GB	72GB	146GB
Enterprise 2C2D	28	1.0TB	2.0TB	4.0TB
Enterprise 2C6D	84	3.1TB	6.1TB	12.2TB
Enterprise 2C12D	168	6.1TB	12.2TB	24.4TB
Enterprise 8C8D	112	4.1TB	8.2TB	16.4TB

Note

Maximum capacities are approximations. Enterprise Virtual Array supports 36GB, 72GB, and 146GB 10K rpm, and 36GB and 72GB 15K rpm dual-ported 2Gb/s Fibre Channel disk drives at full rated transfer rates.

Expansion Rack Capacity

Model	Drive Bays	Maximum Capacity		
		36GB	72GB	146GB
Enterprise 0C6D	72	2.6TB	5.2TB	10.4TB
Enterprise 0C12D	144	5.2TB	10.5TB	21TB

Note

The expansion of a Model 2C12D with a 0C6D rack allows a maximum configurable storage of 35TB using (240) 146GB drives. The 0C12D accommodates two Enterprise Virtual Array 2C12D storage systems to give a maximum configurable storage of 70TB using (480) 146GB drives. Configurations are discussed in Module 5, Rack, Cabling, Configuration, and Initial Setup.

Rack Stability



WARNING

To reduce the risk of personal injury or damage to the equipment, be sure that:


- The leveling jacks are extended to the floor.
- The full weight of the rack rests on the leveling jacks.

HSV110 Controller Enclosures

The HSV110 controller enclosure model number is 3220. Each storage system has two controllers, each occupying 1.5U or 2.625 inches. The two controller enclosures occupy 3U, or 5.25 inches.

Note

A rack has seven junction boxes, allowing cabling between all enclosures.



The controller is covered in detail in Module 2, HSV Controller.


Disk Drive Enclosures

The FC disk drive enclosures support the FC-AL connections on both the external (host-controller-to-enclosure) and the internal buses.

Note

The backend connections for the eva5000 will be copper (same as eva3000) post VCS V3.0.

The disk drive enclosure model number is 5214. Each rack configuration has 2 to 12 drive enclosures. Each disk drive enclosure can be populated with up to 14 1-inch FC disks per enclosure. The enclosure supports dual-loop I/O modules.



The disk drive enclosure is covered in detail in Module 3, Disk Drive Enclosure.

FC-AL Disk Drives

The following are characteristics of the disk drives:


- Dual-ported 2Gb/sec FC-AL
- 36GB, 72GB, and 146GB 10K rpm
- 36GB and 72GB 15K rpm
- Up to 120 drives per FC-AL pair



Disk drives are covered in detail in Module 3.

Fibre Channel Loop Switches

Each of the switched models has four FC loop switches, one pair mounted above and one pair mounted below the controller pair. They use a star cable routing scheme (out and back from same enclosure) and operate at 2.125Gb/s. Each FC loop switch acts as a central point of interconnection and establishes a fault-tolerant physical loop topology.



Fibre Channel loop switches are covered in detail in Module 4, Fibre Channel Loop Switch.

Power Distribution

AC power is distributed to the rack through a dual zero-U power distribution unit (PDU) assembly mounted at the bottom rear of the rack. Each PDU is connected to a separate circuit-breaker-protected, 30A AC site power source. The PDUs in the 41U rack are mounted as in the 42U opal rack, but the PDUs are mounted to the side rails of the rack. The 41U graphite rack back door is full length and opens up to allow access to the PDUs. The 42U opal rack door is divided into two sections, with the bottom section used to access the two PDUs attached to the door itself.



Opal




Graphite

Enclosure Address Bus

The enclosure address bus (EAB) is composed of cables and junction boxes (JBs) that interconnect the controller enclosures with the drive enclosures. The functions of the EAB are to:

- Manage and report environmental conditions within the rack.
- Connect drive enclosures and controllers through the EMUs with cables and junction boxes.
- Assign enclosure numbers (addresses) to the EMUs.
- Provide communications within a reporting group.
- Provide a common path for event logging.
- Allow identification of the shelf location in the rack.
- Offer the HSV controller a way to bypass FC disk drives.




The EAB is described in more detail in Module 3 and Module 5.

Virtual Controller Software

The virtual controller software (VCS) is the software engine behind Enterprise Virtual Array virtualization and resides in the controllers. The latest version of VCS is 3.0; however, some customers still run V2.002, V2.003, and V2.004 (eva3000). The following are a few features of VCS:

- Contains software for controllers, EMU, and disk drives
- Allows nondisruptive firmware updates
- Distributed with controller hardware from the factory or through VCS Software kit; not available on PCMCIA cards
- Enabled by software license keys (base, Business Copy, Continuous Access)
- Managed by the storage management appliance and Command View EVA

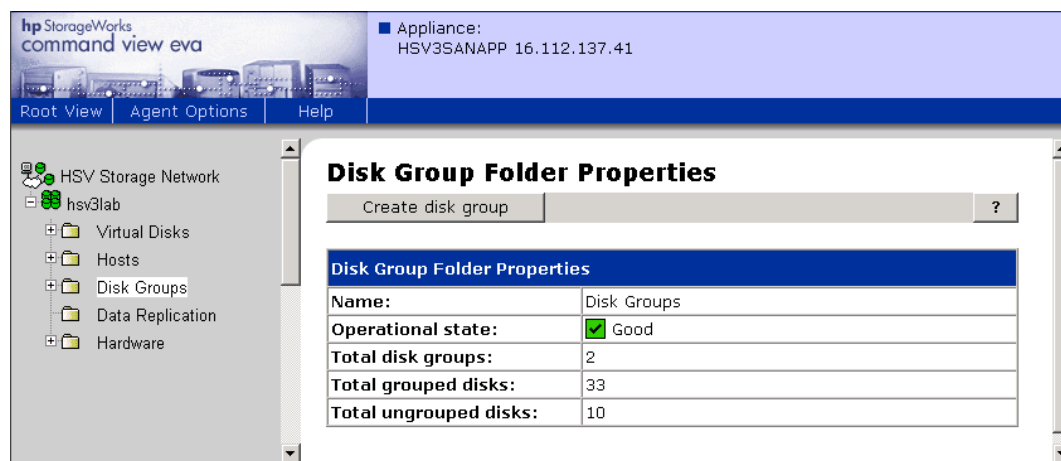
HSV110 Controller Rear



VCS is covered in more detail in Module 2.

Command View EVA

Command View EVA is software that runs on the storage management appliance to allow storage configuration, administration, and management of the Enterprise Virtual Array. Command View EVA is the only graphical user interface for the Enterprise Virtual Array.



Current versions of Command View EVA include the following:

- Command View EVA V3.0, used with VCS V3.0, V2.003
- Command View EVA V2.1, used with VCS V2.004

Details about configuration tasks done through the Command View EVA are provided in Module 8, Storage System Configuration. Details about monitoring and management tasks done through the Command View EVA are provided in Module 9, Storage System Management.

Storage Management Appliance

The storage management appliance is a host-independent appliance (hardware and software) that provides continuous and accurate SAN monitoring for HP StorageWorks products. The appliance provides a central point for SAN management.

The appliance and Command View EVA software are used to set up, configure, and manage the Enterprise Virtual Array storage system.

Versions of the Storage Management Appliance software currently supported are:

- V2.0, Service Pack 3 for the eva5000
- V2.0, Service Pack 1a for the eva3000

Features of the storage management appliance include the following:

- ProLiant DL380 server
- Host platform for enterprise-wide storage management applications
- Rack mounted
- Web-based GUI
- Generation 1 (G1) is 3U with an Intel Pentium III 733MHz processor
- Generation 2 (G2) is 2U with an Intel Pentium III 1.26GHz processor



G1



G2

Details about both appliances are presented in Module 6, Storage Management Appliance.

Storage System Scripting Utility

The Storage System Scripting Utility (SSSU) allows you to do host-based scripting for many of the same functions as Command View EVA. SSSU is the only command-level interface for storage configuration tasks and comes with host solution kits.

The following are two sample SSSU commands:

```
SELECT MANAGER hsvlsanapp user=administrator pass=administrator  
SELECT SYSTEM "HSV Storage System"
```

Details about using the SSSU are presented in Module 11, Storage System Scripting Utility.

Product Support

The most current support matrices are contained in the current product QuickSpecs, Release Notes, and advisories. Areas covered include the following:

- Operating systems
- Coexisting software products
- Switches, FCAs, adapter firmware, adapter drivers
- Disks
- Browsers
- Server types

The following topics summarize operating systems and coexisting software product support.

Operating Systems Support

The following hosts are supported by VCS V3.0. See the product documentation for cluster support.

Windows Platforms

- Windows NT (Intel) V4.0 SP6a
- Windows 2000 (32-bit) V5.0 SP2 and SP3
- Windows Server 2003 Enterprise Edition (32-bit) — No support for Continuous Access

Unix Platforms

- Tru64 Unix V5.1 (BL19), V5.1a (PK4), V5.1b
- Sun Solaris V2.6, V7, V8, V9

Note

Sun Solaris V9 is only available for VCS V3.0, not VCS V2.003.

- HP-UX V11.0 and V11.i
- IBM AIX V4.3.3 and V5.1

OpenVMS

OpenVMS V7.2-2, V7.3 (no CA support), and V7.3-1

Red Hat Linux

Linux Advanced Server 2.1, 2.4.9-e3SMP, 2.4.9-e12SMP

SuSE Linux

SLES 7, 2.4.7-64SMP, 2.4.18-224SMP

Novell NetWare

NetWare V5.1 SP5 and V6.0 SP2

Boot Support

You can boot several platforms and operating systems from Enterprise Virtual Array storage:

- Windows 2000 SP2, Microsoft NT 4.0 SP6a, MSCS
- Tru64 UNIX V5.1 and V5.1a, TruCluster Software products
- OpenVMS V7.2-1H1 with TIMA 300, V7.3, native VMS clusters

There is no Enterprise Virtual Array boot support for the Sun, HP-UX, IBM AIX, Linux, and Novell platforms.

Note

Refer to the *HP StorageWorks SAN Design Reference Guide* for more detail.

Coexisting Product Support

The following describe the most important products that may coexist with the Enterprise Virtual Array:

- SecurePath

You can find the host-specific SecurePath versions that are supported in the standard documentation. In all cases, SecurePath is required to support multiple paths to a single LUN for a host.

Enterprise VCS 2.X and above also offer single path support in which SecurePath is unnecessary. For configuration rules and support matrices, see the white paper, *Connecting Single HBA Servers to the Enterprise Virtual Array without Secure Path*.

- Continuous Access

The VCS V3.0 remote data replication functionality supports various operating systems. You can find the host-specific Continuous Access versions that are supported in the standard documentation.

- Business Copy

The Enterprise Virtual Array supports HP StorageWorks Business Copy (BC) in the following versions:

- BC V2.1 supports VCS V2.X.
- BC V2.1A supports VCS V3.0.

Support for specific operating systems is documented in the standard areas.

Note

Differences exist between traditional HP Business Copy snapshot technology and what is currently called EVA Business Copy technology. These differences are covered in Module 7.

Business Copy allows you to create two types of snapshots.

- Fully-allocated snapshot operations are similar to the HSG80 controller-based snapshot. They require that you reserve and set aside the capacity amount equal to that of the original volume for the snapshot.
- Demand-allocated (Virtually Capacity-Free) snapshots have minimal space reserved up front. Space is only used as the original virtual disk changes data.

Business Copy allows you to create snapclones (Virtually Instantaneous snapclones), an improved type of data cloning. A snapclone is similar to a fully-allocated snapshot in that duplicate space is reserved. However, a complete copy of the original virtual disk is created, instead of only copying metadata. Two identical copies of the data are created with a snapclone. Unlike traditional clones performed by other controllers, you can access the Enterprise Virtual Array snapclone data immediately instead of waiting until the copy is complete.

Licensing

Licenses for the Enterprise Virtual Array include:

- Base
- Business Copy EVA, separate, optional license
 - Snapshots and snapclones
 - Licensed by capacity point based on the disk storage capacity connected to the controller pair
- Continuous Access for EVA (VCS V3.0 and above), separate, optional license
 - Real-time remote data replication
 - Licensed by capacity point based on the disk storage capacity connected to the controller pair

Licenses are entered through the storage management appliance. More detail is provided in Module 5.





Event Reporting

The Enterprise Virtual Array generates events based on such things as state changes, failures, and job status. It uses four event log types to relay error information. These include:

- Management Agent Event Log
- Controller Event Log
- Controller Termination Event Log
- NT Event Log

The following is an example of a Controller Event Log.

Controller Events (Initialized system)

OK	Get log file	Get parse file	Send parse file	?
Display Range:	1 - 400	Previous group	Next group	
Controllers:	Controller B: 5005-08b4-0001-4529-0000-0000-0000-0000 Controller A: 5005-08b4-0001-4523-0000-0000-0000-0000			
Date/Time Controller	Severity	Event Code	Sequence #	Description
17:43:23:368 15-Apr-2003 Controller A		06014a08	#1965	A Fibre Channel port on the HSV110 controller has failed to respond. ✖ Corrective action code: 4a More details
17:42:53:806 15-Apr-2003 Controller A		09070005	#1964	The state of the Fibre Channel port identified in the attribute.value.str field and located on the rear panel of the HSV110 controller identified in the handle field has transitioned to the NORMAL state. ✖ Corrective action code: 00 More details
17:42:11:21 15-Apr-2003 Controller A		0df00011	#1963	The status has changed on one or more of the drive enclosures. ✖ Corrective action code: 00 More details
17:41:48:742 15-Apr-2003 Controller A		060f4013	#1962	The Drive Enclosure Environmental Monitoring Unit is able to communicate with a physical disk drive but this HSV110 controller is unable to communicate with that physical disk drive on the Fibre Channel bus. ✖ Corrective action code: 40 More details

All logs are viewable from Command View EVA except the NT Event Log. All events, except controller termination events, can be forwarded using Simple Network Management Protocol (SNMP) to any given receiver.

Download the log files through the **Get log file** button in the Command View EVA GUI. Management Agent events are saved in an ASCII file, whereas controller events are saved as binary files and must be decoded by using any of three tools:

- Not a Post Processor (NAPP)
- Event View EVA
- HSV Event Translation (HSVET)

The Enterprise Virtual Array also allows in-band host event reporting, including the following features:

- Events can be correlated with host system events.
- Controller events are sent to host system's event log.
- Uses controller host fabric ports.
- Uses SCSI sense data, unique to host.

Details about event logging are presented in Module 10, Error Reporting and Diagnostics.

Diagnostics

Several diagnostics are performed, either internally or externally, through the Enterprise Virtual Array, including the following:

- Built-in self test (BIST) diagnostics
 - First software to run during initialization
 - Verifies controller functions well enough for functional software
- Module integrated self test (MIST) diagnostics
 - Software that runs during initialization of functional image
 - Verifies everything not checked through BIST
- Test entities and test numbers in Event Information Packet (EIP) type 14
 - Diagnostic failures reported as controller events using EIP type 14
 - Uses test entity numbers and test numbers
- Periodic diagnostics — Battery, cache memory, idle resources
 - Cache memory — Scrubs correctable ECC errors from cache
 - Controller temperature sensing — Checks temperature of controller every 10 seconds
- Other diagnostics
 - Disk scrub — System uses idle cycles to periodically check the readability of unused or unaccessed storage
 - Metadata check

Details about diagnostics are presented in Module 10.

Warranties

Warranties are available for Enterprise Virtual Array software and hardware. These offerings are included under the general titles of HP-branded StorageWorks hardware and software, and HP-branded SANworks software.

INTERNET

For more information, see the HP Storage Services website
<http://www.hp.com/hps/storage/>.

The following describes the included warranties:

- Software is warranted against media defects for a period of 90 days from purchase.
- Hardware warranty features include a 3-year Limited Warranty that includes 3-year parts exchange, 3-year duration onsite labor, 24x7 coverage, and 4-hour on-site response.

Service Offerings

The following describe the available service offerings with the Enterprise Virtual Array:

- **Installation and Startup Services**
Ensure top performance of hardware and software, including the Enterprise Virtual Array, right from the start.
- **Basic Post Warranty Support**
Available for hardware and software.
- **CarePaq Priority Services**
Extended hardware and software warranties are available as CarePaq Priority Services (CPS).
- **SAN Implementation Service (SIS)**
Customers who need a comprehensive approach to the implementation of a complex enterprise SAN solution.

Installation and Startup Services

HP offers an Installation and Startup Service to provide a way to ensure top performance of hardware and software, including the Enterprise Virtual Array, right from the start. This service is usually bundled with the hardware (sales bundle).

Major categories of work include:

- Preinstallation planning and consulting
- Equipment unpacking and rack mounting
- VRAID array configuration initialization
- StorageWorks/SANworks/OpenView product installation, configuration, and verification
- Service tools installation, configuration, and verification
- Up to eight hours of basic product usage orientation for the customer

Basic Post Warranty Support

Basic Post Warranty Support provides an option to extend the product warranty. This service bridges the gap between the basic warranty and CarePaq Priority Services.

- The post warranty support for hardware is 1 year (payable monthly), ranging from 9x5, next business day to 24x7, 4-hour.
- The post warranty support for software is 1 year (payable monthly), ranging from 9x5 to 24x7 with 2-hour phone support and a media and documentation subscription service.

As with software, upgrades to the basic warranty are available as CarePaq Priority Services.

CarePaq Priority Services

Extended hardware and software warranties are available as CarePaq Priority Services (CPS). Of the six priority services normally available for CarePaq, the Enterprise Virtual Array has the following four, each carrying a term of three years:

- **Priority 24** — Basic support round-the-clock: 24x7 hardware/software support, 4-hour hardware response, 2-hour software phone response, and a named hardware engineer; Software Product Subscription for VCS base, platform kit, and Secure Path kit.
- **Priority Silver** — 24x7 hardware support with named engineer, 4-hour on-site response; 24x7 software phone support with named account representative, 1-hour response Monday through Friday, 8 AM to 5 PM local time, 2-hour response remaining hours; Silver Support Team led by assigned Technical Account Manager; quarterly service activity review; technical newsletter; proactive revision management; Software Product Subscription for VCS base, one platform kit, and one Secure Path kit.
- **Priority Gold** — Same coverage as Priority Silver but with 30-minute callback on critical software calls, 1-hour callback on noncritical software calls; monthly service activity review; annual customer site visit; critical problem management; and 10 hours of upgrade impact planning.
- **Priority Gold Executive** — Same coverage as Priority Gold but with 24x7, 2-hour response for on-site hardware support.

SAN Implementation Service

Several value-added services are available with the Enterprise Virtual Array, including Migration Services, Managed Services, Availability Services, and the SAN Implementation Service.

The SAN Implementation Service (SIS) is provided for customers who need a comprehensive approach to the implementation of a complex enterprise SAN solution. A predefined set of services is provided to ensure expert and smooth implementation of the SAN. The result is a powerful storage solution that the customer can use immediately.

Key benefits of the SIS include:

- Ensures the right configuration and performance to deliver the results customers need from their enterprise SAN—right from day one
- Creates a cohesive storage management entity that puts customers in control of this complex technology solution
- Reduces overall cost of ownership

The SIS provides the following:


- Assigned HP Project Manager
- Preinstallation consulting and planning
- Logical and physical disk configuration design
- Installation, configuration, and verification of hardware/software/tools
- Installation verification testing
- Customer orientation
- Full documentation

INTRA NET	All of the essential information for the SIS is available through the SIS Service Delivery Guide (SDG), available through the website http://sdg.inet.cpqcorp.net/release/sdg/sis .
------------------	---

This link supplies you with the essential standards, procedures, templates, and checklists, including the following:

- SIS Implementation Document
- SIS Naming and Configuration Standards
- SIS Customer Orientation Checklist
- SIS Installation Verification Test

Enterprise Virtual Array 3000 Overview



Enterprise Virtual Array 3000 (eva3000) is a replacement for the MA/EMA product line and an alternative to the HP VA product line for lower midrange storage system requirements. The eva3000 is part of the Enterprise Virtual Array family and has nearly identical features and support as Enterprise Virtual Array 5000. Serviceability requirements are identical.


Features

Features of Enterprise Virtual Array 3000 include the following:

- 42U Cannondale 10000 rack
 - Graphite, not opal
 - Same dimensions as the eva5000 42U rack
- No loop switches
- HSV100 controller (based on the HSV110)
 - Model 3200
 - Dual power supply (DPS)
 - HSV100 firmware (VCS V2.004)
 - No support for Continuous Access functionality
 - Two covered back-end ports (For now, sheet metal covers the ports; in the future, the internal hardware will be removed.)
- Disk drive enclosure
 - Model 5114
 - No SFP transceivers
 - Copper interconnects between controllers and enclosures
- Command View EVA V2.1 or V3.0

Configuration

The configuration is shown in the diagram.



The following are the main points regarding the configuration:

- Base configuration is 2C2D
 - Expandable to 2C4D
 - ◆ Maximum of 2TB using 36GB drives
 - ◆ Maximum of 4.1TB using 72GB drives
 - ◆ Maximum of 8.2TB using 146GB drives
 - Supports 8 to 56 drives
- In 42U rack, supports a maximum of 23TB
 - Can stack two 2C4Ds and one 2C3D
 - Can use 154 drives
- Controller upgrade to HSV110 is available

Product Support

The eva3000 supports all eva5000 products, but has the following exceptions:

- Supports VCS V2.004
- Supports Storage Management Appliance Software V2.0, SP1A
- Supports Command View EVA V2.1 and V3.0
- Operating systems
 - See the product documentation for the eva3000
- Coexisting products
 - See SecurePath versions in the product documentation
 - No support for Continuous Access
 - Business Copy V2.1
 - ◆ Basic support on all operating systems
 - ◆ See the product documentation for the eva3000

Learning Check

1. List the operating systems that are currently supported on the Enterprise Virtual Array.
.....
.....
2. What is the maximum storage capacity of one Enterprise Virtual Array rack with 72GB disks?
 - a. 3.1TB
 - b. 6.1TB
 - c. 12.2TB
3. Choose the best answer to complete this sentence: The HSV controller pair connects to the physical disk drive array through the:
 - a. Fibre Channel fabric.
 - b. Fibre Channel arbitrated loops.
 - c. Storage management appliance.
4. List the components within a rack that are used to distribute power. Include the total number of each power distribution component.
.....
.....
5. Complete this sentence: The enclosure address bus consists of
.....
6. List the interfaces that connect on the rear of the controller enclosure.
.....
.....

7. How many dual-loop I/O modules are there in the disk drive enclosure?
 - a. Four
 - b. One
 - c. Two
 - d. Three

8. How many EMUs can communicate directly with the host at one time?

.....

Overview

This module describes, in detail, the physical, interconnect, and management features of the HSV controller. Each Enterprise Virtual Array storage system has a pair of HSV controllers. This next-generation, high-performance array controller is easily migrated between arrays with varying function and performance.

This module concentrates primarily on single power supply controllers; however, one topic describes the dual power supply controller differences.

Note

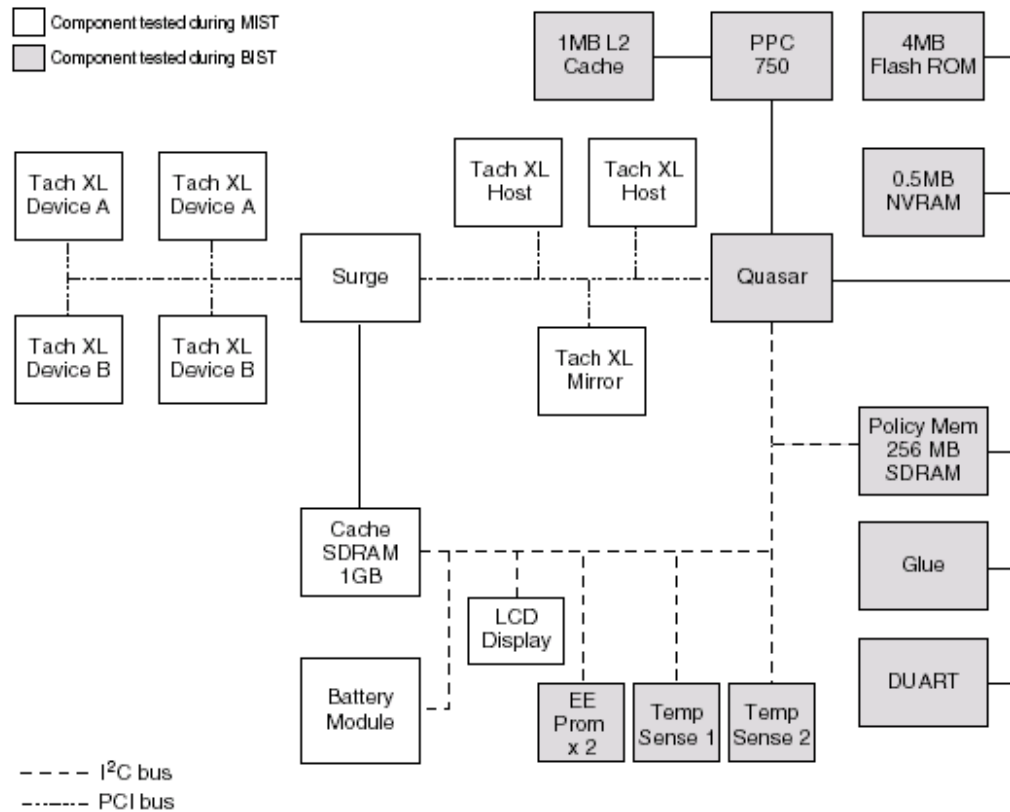
In this module, the term HSV controller implies the HSV series because the Enterprise Virtual Array 3000 (eva3000) uses the HSV100 controller. Differences for the HSV100 are noted at the end of the module.

Objectives

After completing this module, you should be able to:

- Recognize the main features of the HSV controller.
- Recognize the components and functions of the front controller elements.
- Recognize the components and functions of the rear controller elements.
- Identify the differences between the single and dual power supply controllers.
- Identify all field replaceable units (FRUs) for the controller.
- Identify the differences between the HSV110 and the HSV100 controllers.

HSV Controller Block Diagram



Note

You see most of these items (for example, Quasar, Glue, PPC, Surge Code) when you view system software and firmware versions from the controller Operator Control Panel.

Built-In Self Test (BIST) Diagnostics

BIST diagnostics are the first software to run when the Enterprise Storage System hardware initializes. The goal of the BIST is to verify that the controller functions well enough to support the execution of the functional software. To accomplish this purpose, BIST tests all components attached directly to the PowerPC and those accessible through the local bus on the Quasar chip.

Components are tested in the following order:

1. Processor initialization — The BIST configures the PowerPC processor for use with the Enterprise Storage System design.
2. Quasar configuration — The Quasar chip is configured for access to the components on its bus, such as the DUART.
3. Policy memory — The DIMM for policy memory is interrogated, using the Quasar I²C controller to determine size and configuration requirements. After the Quasar chip has been configured for this DIMM, the memory is tested for address and data integrity.
4. Program card — The BIST image located on the program card is verified to ensure that the code being executed is valid.
5. Glue chip — The Glue chip (G3 to signify the third glue chip) is an FPGA chip containing much of the miscellaneous control functionality that was implemented using more discrete components in previous controller designs. This component contains features that include PCI bus arbiters, interrupt controllers, bus watchdog timers, and general-purpose I/O pins (diagnostic registers). The Glue chip is read to verify initial values, written where possible to verify address and data integrity, and configured for use.
6. L2 cache — The PowerPC uses an L2 cache. This cache is verified to ensure functionality before it is enabled for use by the processor.
7. NVRAM and TOY clock — The NVRAM component of the Enterprise Virtual Array is composed of a battery backed-up SRAM chip. The NVRAM battery controller contains a Time of Year (TOY) clock. The battery controller, the battery, and the memory are checked for functionality, and the TOY clock is initialized if needed. The functionality of the NVRAM system is not required to begin execution of the functional image, but it is tested as a BIST component.
8. Functional image — The functional image is moved from the program card to the required locations in Policy memory. During this step, the functional image is decompressed and a checksum is calculated over the length of the image. The image checksum is verified as it is moved from the card to memory. After the functional image is placed in memory, the BIST code cleans up and begins executing the functional image.

Module Integrated Self Test (MIST) Diagnostics

The diagnostics that run after initialization of the functional image are referred to as the Module Integrated Self Test (MIST). The MIST diagnostics verify all the devices and interconnecting buses that were not tested during BIST.

The Diagnostic Operations Generator (DOG) executes the MIST diagnostics. It also collects and reports error information as required. In VCS V2.X and higher, events generated by the DOG are reported in the Controller Event Log using EIP type. Events using EIP type 14 identify the component being tested and test being performed (when multiple tests are run).

Components are tested in the following order:

1. Near PCI bus configuration — The devices on the PCI bus between the Quasar and the Surge (including the Surge chip) are configured. Configurations are also verified. Component revisions are compared during configuration (or during component diagnostic) to ensure compatibility between hardware and firmware.
2. Surge chip — The Surge chip is configured for Enterprise Storage System use.
3. Far PCI bus configuration — The devices on the primary PCI bus of the Surge chip are configured.
4. Cache memory — The cache DIMM modules are interrogated for size and setup information, and the Surge chip is configured accordingly. The contents of cache memory are checked. Data integrity of the cache memory, the XOR engine, the DMA engines, and the ECC functionality of the Surge chip are verified.
5. Cache battery — The initial status of the cache battery module is checked and passed to the cache manager and Command View EVA.
6. Tachyon XL2 chips — This diagnostic tests the Tachyon XL2 chip that drives each Fibre Channel port. It also tests the write-sensitive data area of the Quasar, used for port interrupts to the processor. This diagnostic does not transmit over the external cables.

Note

For the eva3000, one set of Device A and Device B ports is not required or involved in the diagnostics.

7. Glue chip — Additional functions of the Glue chip that could not be tested during BIST diagnostic phase or previous MIST diagnostics are tested here. These functions include a test of the bus timeout and illegal state detection logic in the Glue Chip.
8. Module load test — The goal of this diagnostic is to ensure that the hardware functions correctly under load. All devices that can move data (Tachyon XL2 chips, Surge DMA engines, and the processor) are tested, using cache, cache XOR, and policy memory as both source and destination.

Controller and Management Appliance Communications

The controller communicates with the storage management appliance using StorageCell Management Interface (SCMI) commands. The following are features of this communication:

- In-band FC communications
- Object-oriented interface
 - Objects are disks, hosts, and so forth
 - Object have attributes such as size and VRAID type
- Defines commands to:
 - Get or set object attributes
 - Create and destroy objects
- Commands are called SCMI packets, sent over SAN using FC SCSI Send and Receive diagnostic commands
- Processed only by the Storage Array Controller Device (SACD), LUN 0

Information on interpreting SCMI commands and return codes in event logs is discussed in Module 10, Error Reporting and Diagnostics. For a complete list of SCMI commands and return codes, see Appendix I of the Service Manual.

Controller Features

The following are features of the HSV controller:

- PowerPC microprocessor
- Operator control panel (OCP) for easy operation
- Two 2Gb/s FC-Switch fabric host ports
- Four 2Gb/s FC-AL device ports
- 2Gb/s FC cache mirroring port
- 1GB cache per controller, mirrored, with battery backup
- FRUs
 - Cache battery assembly
 - Blowers (single power supply controllers)
 - Transceivers
 - Cables
 - Power supply/blower (dual power supply controllers only)

The redundancy of the controller pairs ensures that the failure of a controller element does not disable the system.

- The complete data redundancy includes two loop A and two loop B data paths.
- Even the failure of a controller power supply would not disable the system.
- A single controller can fully support an entire system until the defective controller, or controller element, is repaired.

Virtual Controller Software

Virtual Controller Software (VCS) contains the virtualization engine for the Enterprise Virtual Array. The latest versions for the eva5000 are VCS V2.003 and V3.0. VCS V2.004 is the version compatible with the eva3000.

The following are some features of VCS:

- Stored in flash memory.
- Distributed with controller hardware from the factory.
- Updates are distributed on CD-ROMs and on the Web.
- Code is enabled (base, Business Copy EVA, and Continuous Access EVA) by software license keys.
- License keys are managed through Command View EVA.

The upgrade paths for upgrading the storage system software from V2.003 to V3.0 are listed in the following table. Only the eva5000 applies.

Current VCS Version	Upgrade Path
2.003 (recommended minimum)	Upgrade to V3.0
2.002 (mandatory minimum)	Upgrade to V2.003. Upgrade to V3.0.
2.001 or 2.000 (supported versions)	Upgrade to V2.003. Upgrade to V3.0.
1.021 (obsolete and unsupported version)	Upgrade to V2.002. Upgrade to V3.0.

When upgrading from VCS V2.000 to a higher version and the host server is running Windows 2000, no rolling upgrade is allowed, that is, you must take the storage system offline so that it is not servicing host commands.

Note

A software and firmware compatibility matrix, and details for upgrading VCS versions, are provided in the Enterprise Virtual Array Instructions for Upgrading to V3.0.

The following shows version compatibility with Command View EVA:

VCS Version	HSV Element Manager V2.0A	Command View EVA V2.1	Command View EVA V3.0
2.002 (eva5000)	Yes	Yes	Yes
2.003 (eva5000)	Yes	Yes	Yes
3.0 (eva5000)	No	No	Yes
2.004 (eva3000)	No	Yes	Yes

VCS also provides the following capabilities:

- Support for up to 240 disk drives per controller pair
- Management of up to 512 virtual disks per disk pool, ranging in size from 1GB to 2TB per virtual disk
- Dynamic capacity expansion and virtual disk data load leveling
- Distributed sparing of disk capacity
- Demand-allocated (Virtually Capacity-Free) Snapshot (Vsnap)
- Virtually Instantaneous Snapclone
- Dual redundant controller operation for increased fault tolerance
- Multiple path failover support
- Virtual RAID arrays (VRAID0, VRAID1, VRAID5)
- Nondisruptive software upgrade capability
- Connects up to 256 hosts
- Multivendor platform support
- Controller password protection for configuration control
- Selective storage presentation
- GUI for management and monitoring

Controller Enclosure

The HSV controller is model 3220, and is 1.5U high, or 3U as a pair. The HSV is the interface between Command View EVA and the storage system.



Two models of HSV110 controllers with differing physical dimensions exist.

- Controllers with a single power supply (SPS controllers). SPS controllers are less deep than dual power supply controllers. They can be mounted using either short (rack V1-style rack) or long (V2-style rack) rails.
- Controllers with dual power supplies (DPS controllers). DPS controllers are deeper than SPS controllers and require long (rack V2-style rack) rails. The part number of the dual power supply controller is 70-40927-01.

The original single power supply (SPS) controller assembly has one power supply per assembly, with each power supply having the following characteristics:

- Each power supply has a design spec of 500,000 hours—highly reliable.
- Each power supply is powered from different AC sources.
- If a power supply fails, one controller is inoperable until the controller assembly is replaced.
- If an AC source is lost, one controller is inoperable until AC is restored.

The new dual power supply (DPS) controller assembly has been introduced to increase availability and has the following characteristics:

- Each controller assembly has two power supply/blower hot-plug FRUs.
- Each controller assembly has two AC sources, one per power supply.
- If a power supply/blower assembly fails, the controller remains up, and the assembly can be removed/replaced without stopping the controller.


If an AC source is lost, that is, one power supply/blower is without power, the controller remains operable.

Note

If you are accessing the system remotely, you can use Command View EVA to determine which type of controller is installed in the system by selecting the controller, clicking the Enclosure tab, and observing the FRU type in the Blowers segment of the display.

Controller Front

The following diagrams show the essential differences between the SPS and DPS front (without bezel).




The major elements found on the front of the HSV controller are as follows:


- A **removable bezel** that provides access to the blower or blower/power supply compartment and the cache battery assembly compartment.
- The **blower compartment** that contains two blowers or two blower/power supplies to cool the controller.
- A **cache battery compartment** that contains **two lead-acid battery assemblies**. The battery assemblies are located in a compartment behind the OCP.
- The OCP that includes:
 - Four status light emitting diodes (LEDs)
 - A liquid crystal display (LCD)
 - A four-position pushbutton switch

By using the OCP pushbuttons and the LCD, you can display information about the controller, and, in some cases, enter controller operational data. The use of the displays and pushbuttons is described in more detail later in this module.

The following front view shows the removable bezel and front panel of an SPS controller.




The following front view shows the doors and replaceable components in an SPS controller.



Bezel

The plastic bezel allows access to FRUs (blowers, cache batteries) and pops off when pulled.



Blowers

For the SPS controller, the two hot-replaceable blowers (shown in the photograph) are physically located in the left front of the controller (behind the bezel).




Note

The amber light on the blower face panel indicates the blower is operating at slow speed or has failed.

For the DPS controller, two hot-replaceable blower/power supply assemblies are located at the left front of the controller.

Dual Redundant Power Supply/Blower FRUs

Access door



The amber light on blower/power supply face panel indicates the following:

- The blower is operating at slow speed or has failed and requires replacement.
- The power supply has failed and requires replacement.

Power Supplies


Each power supply provides the DC voltages required for controller operation. Nominal input voltage is from 202 to 240 VAC, 50 Hz to 60 Hz. Each power supply generates the necessary 3.3 VDC, 5.0 VDC, and 12.0 VDC system voltages. The two power supply/blower units in a DPS controller are identical.

A controller power supply failure causes the HSV Element Manager to report that the controller enclosure is no longer part of the system. Furthermore, power supply failure indications and recovery actions are different on SPS and DPS controllers as follows:

- DPS controller
 - A failure in either component of the power supply/blower assembly causes the LED on the front of the assembly to become amber.
 - Hot replacement procedures are available as described in a later topic.
- SPS controller: You must replace the entire controller assembly.

Cache Battery Assemblies

Each HSV series controller has two lead-acid cache battery assemblies that provide power to the cache memory dual inline memory modules (DIMMs). Each battery assembly has three lead-acid, nonspillable cells. When both battery assemblies are fully charged, they can provide power to the DIMMs for up to 96 hours. Batteries have a four-year expected lifetime.



Battery Operation

The following are some features of the batteries:

- Batteries are tested every month (by draining a battery cell), and the results give you expected future battery life.
 - The batteries are discharged to test them one cell at a time.
 - The amber LED on the battery/HSV front panel will not be turned on.
 - It will take six months to test both bricks, because each brick has three cells.
 - It is not user selectable, and the diagnostic does **not** light the fault LED.
- On charge-up, charging order is cell 1, 2, 3, and so on. It takes about four to eight hours for new batteries to fully charge.
- When any one cell fails, the controller assumes that all cells have failed on the controller, and you must replace both batteries.
- If a battery system is bad, and if the battery state on the other controller is better, the units (LUNs) are moved to the other controller.
- Never remove batteries from the system while it is powered off.
- Use the HSV power-off option to flush the cache and disable the batteries—do not just turn off the switch.
- Never install a battery that was previously failed by any controller.

Cache Battery Assembly Status LED

The battery cells are mounted in a plastic case, which is the battery assembly. The orientation label is located in the upper left corner at the rear of the case. The battery assembly amber status LED is located in the lower right corner at the rear of the case.



Caution

You should shut down the controllers in a graceful way, that is, through Command View EVA or controller OCP, to avoid draining the battery.

- When the amber LED next to the icon (shown below) is ON (lit solid) the battery assembly is discharged or may need to be replaced.



When the LED is OFF (not lit) the battery assembly is charged, and operational.



Caution

If you are replacing one battery and the other battery fault LED illuminates, you must replace both batteries.

- For newly installed batteries, the LED may remain on for up to two minutes.

Cache Options

The cache options consist of the following:

- Write-back modes
 - Mirrored
 - Nonmirrored (When using nonmirrored, performance goes up but it may not provide failover.)
- Read modes
 - Read ahead cache/adaptive cache
 - No Read cache

Batteries protect half of the controller cache memory.

- The protected portion is where write-back data is stored.
- Read data is not protected.
- Virtual disk cache mirroring is enabled/disabled on a per unit basis.

Mirrored Cache Battery Policy

When the battery system on one controller is no longer good, the HSV will move the units to the other controller if its battery state is better. The following table describes the writeback and writethrough policies for controller master and slave controller battery states.

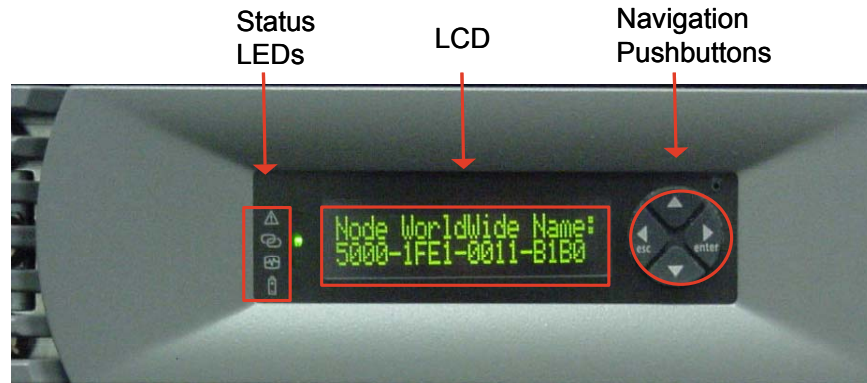
	Storagecell Slave Battery System Good	Storagecell Slave Battery System Low	Storagecell Slave Battery System Bad
Storagecell Master Battery System Good	All units write back on both Storagecell Master and Slave	All units write back on Storagecell Master	All units write back on Storagecell Master
Storagecell Master Battery System Low	All units write back on Storagecell Slave	All units write through on both Storagecell Master and Slave	All units write through on Storagecell Master
Storagecell Master Battery System Bad	All units write back on Storagecell Slave	All units write through on Storagecell Slave	No unit presentation except Storage Array Controller Device (SACD)

Note

The SACD, located at LUN 0, is a special LUN defined by the SCSI-3 specification. The SACD is a portal to the controller, similar to the Command Console LUN (CCL) used in previous products.

Operator Control Panel

The OCP on the front panel of the HSV controller displays information and is used to enter data to set up an HSV controller pair, or to check controller status. The following illustration shows the elements of the OCP.



Navigation Pushbuttons

The information displayed by the OCP, the data you can enter, and the navigation pushbutton functions depend on the alphanumeric display mode. To avoid confusion, the pushbutton reference names are *left*, *right*, *top*, and *bottom*.

Using the pushbuttons, you can navigate through menus and perform the following functions:

- Uninitialize the system.
- Enter the WWN and checksum (at initialization).
- Display the storage system menu tree.
- Set a password.
- Display system information.
- Shut down and restart the controller.
- View revision levels.

Storage System Menu Tree

The Storage System Menu Tree lets you select information to be displayed (for example, System Information, Fault Management) or procedures to implement (for example, Shutdown System, System Password). The following describes the tree:

- **Information**
 - *Port Config*
 - *Versions*
- **Fault Management**
 - *Last Fault*
 - ◆ Term Code
TC: 04240960 IDX: 04 (termination code)
 - ◆ Done
 - ◆ LTEA[0] (last termination event array)
 - *Detail View*
 - ◆ Done
 - ◆ LTEA[0]
- **Shutdown Options**
 - *Restart*
 - ◆ Restart System
No
 - *Power Off*
 - ◆ Power Off System
No
 - *Uninitialize System*
 - ◆ Really Uninit System
No
- **System Password**
 - *Change Password*
 - *Clear Password*
 - *Current Password*

The Port Config option is available beginning with VCS V2.002. It allows you to enable a failed port device loop before a five-minute retry period. This is helpful when troubleshooting loop problems because you do not have to wait for the port to be enabled.

Using the Storage System Menu Tree

The following are some general guidelines when using the tree:

- To display the Storage System Menu Tree, press any of the pushbuttons when either of the default displays (Storage System Name, Node WorldWide Name) is active.
- Use the up and down pushbuttons to sequence up or down through the tree.
- Use the right pushbutton to select an active display, or use the left pushbutton to select the default display.

LED Status Displays

The status LEDs indicate the internal status of the controller; however, during initial setup, the status LEDs may not be fully operational.

- **Fault LED** — When the amber LED to the right of this icon is ON (flashing) there is a controller problem. Check the Command View EVA GUI or the LCD Fault Management displays for a definition of the problem and recommended corrective action.



Note

You do not see a fault LED with a removal.

- **Host link LED** — When the green LED next to this icon is ON, there is a link between the storage system and a host or the appliance. If OFF, there is no link.



Note

The host link LED remains on after the last host disconnects. It is not an indication of fabric login.

- **Controller heartbeat LED** — When the green LED next to this icon is flashing slowly, the controller heartbeat is operating normally.



When this LED is not flashing, there is a problem.


- **Cache battery assembly LED** — When the amber LED next to this icon is OFF, the battery assembly is charged.



When this LED is ON, the battery assembly is less than 50% charged or failed.

LCD

Each LCD row can display up to 20 alphanumeric characters. The LCD alternates between displaying the *Storage System Name* and the *Node Worldwide Name*. These are the default displays. When there is no activity for approximately 15 seconds, the display automatically returns to the default display. The following figure shows an example of the default LCD displays.




Important

An active (flashing) display, an error condition message, or a user entry (pressing a pushbutton) overrides the default display. When none of these conditions exist, the default display is active after approximately 15 seconds.

The displays typically seen on the HSV110 controller LCD are as follows:

- **Startup display** — Resetting or applying power to the storage system controller pair creates this display. The first line defines the controller model number. The second line is a series of asterisks for VCS V1.x, or a series of test entity numbers for VCS V2.x and above.



```
HSV110 Startup
*****
```

Note the following correct controller power-on procedure:

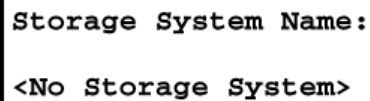
1. Ensure that the controller power switches are off.
2. Power on the drive shelves.
3. Wait two minutes for drive spinup.
4. Power on the controllers.

An asterisk is displayed when one of the 20 startup procedures is **successfully** completed. After completing all the startup procedures, the display changes to the *Initial Setup Display*.

Note

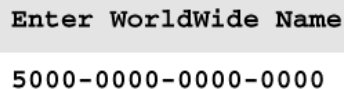
The Startup display may not appear for up to seven minutes after initial application of power.

- **Initial setup display** — The following shows the initial setup display when active, after the initial startup sequence is complete.



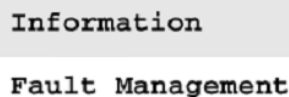
Storage System Name:
<No Storage System>

- **Worldwide name and WWN checksum display** — Using the pushbuttons on the OCP after initialization activates this display. The display changes as you enter each character of the WWN and checksum. The following is an example of the display with the first character entered. Before characters are entered, the display for each character position is “0.” Details on how to enter the WWN and checksum are discussed later in this module.




Enter WorldWide Name
5000-0000-0000-0000

- **Storage system menu tree display** — To activate this display, press any of the arrow pushbuttons on the OCP when either of the default displays (*Storage System Name* or *Node Worldwide Name*) is active. The following appears.




Information
Fault Management

- **Removing the power or uninitializing the system** — Navigate to this display using the arrow pushbuttons. Shutting down the system removes power from the controller pair, flushes I/Os to the user disks, and removes the storage system from the Command View EVA view.



Uninitializing the system removes the controller configuration and causes data loss. Shutdown does not remove the configuration or data from the system.



Entering the Worldwide Name and Checksum

Entering the WWN and the checksum are part of the initial setup of the HSV110 controller. Enterprise Virtual Array operations require that each **controller pair** have a unique Node WWN. This 16-character alphanumeric name uniquely identifies the system.

HP assigns the WWN to each storage system before shipment. The *Enterprise Virtual Array World Wide Name Label* document defines the WWN for each storage system. The Node WWN labels, similar to the one shown here, specify the system-specific WWN and checksum.



Setting a Password at Initial Setup

When you enter a storage system in Command View EVA (Management Options), you can optionally enter the password. This binds the controller pair with Command View EVA. The password must be a unique, **eight-character, alpha password** using the letters A through Z, uppercase or lowercase.

Note

The password **must** be eight characters exactly, nothing less and nothing more.

The default password is AAAAAAAAAA.


Controller Rear

The major elements found on the rear of the HSV controller are as follows:


- Power connections (switch)
- Data connections (with link status LEDs)

The data connections are the interface to the FC loop switches (switched hardware) or disk drive enclosures (nonswitched hardware), the other controller, and the fabric (host) or Command View EVA on the storage management appliance.

This photograph shows a rear view of the HSV110 controller (uncovered) with an SPS.



This photograph shows a rear view of the HSV110 controller (uncovered) with a DPS.



Rear Connections

On the rear of the HSV controller are the data and power connections. Data connections are the interfaces that connect from the HSV controller to the:


- Disk drive enclosures on nonswitched hardware models, or to FC loop switches on switched hardware models.
- Other controller.
- Fabric (host).

Data Connections


The data connections consist of the following:

- Two 2Gb/s FC-Switch fabric host ports
- Four 2Gb/s FC-AL device ports
- 2Gb/s FC cache mirroring port (device ports as backups)

The following shows all data and power connections for the SPS controller.



The following shows a DPS controller.



The rear connections consist of the following:

- **VCS card** — The VCS card (4MB) holds the VCS firmware. New firmware upgrades do not require new VCS cards.
- **Serial line and UPS line** — These DB9 serial ports are used for Engineering only. There is no support currently for UPS.
- **Host fabric 1** — Used to connect the HSV controller to the fabric.
- **Host fabric 2** — Used to connect the HSV controller to the fabric.
- **Mirror port** — Used for communication between the pair of controllers.
- **Loop pair 1 (A and B)**
 - In a 2C12D configuration, connects to an FC loop switch for switched hardware, which in turn connects to drive enclosures 1 to 6
 - In a 2C6D configuration, connects to an FC loop switch for switched hardware, which in turn connects to drive enclosures 1 to 3
 - In an 8C8D configuration, connects to the top drive enclosure of each of the four enclosure pairs
- **Loop pair 2 (A and B)**
 - In a 2C12D configuration, connects to an FC loop switch for switched hardware, which in turn connects to drive enclosures 9 to 14 (switched) or 8 to 13 (nonswitched)
 - In a 2C6D configuration, connects to an FC loop switch for switched hardware, which in turn connects to drive enclosures 9 to 11 (switched) or 8 to 10 (nonswitched)
 - In an 8C8D configuration, connects to the bottom drive enclosure of each of the four enclosure pairs

Note

In a 2C12D + 0C6D configuration, the loop pairs are extended through either an expansion panel (nonswitched hardware) or through FC loop switches (switched hardware) to connect to drive enclosures 15 through 20 of the expansion rack. Loop pair 1 extends to drive enclosures 15 through 17, and loop pair 2 extends to drive enclosures 18 through 20.

Power Connections

The newest controllers include two power supplies per HSV controller. If a power supply failure occurs, the remaining power supply provides the required power to the controller.



Dual Power
Connectors

PDMs 2 ,3 ,6, and 7 are used for the dual power supply connections. PDMs 2 and 3 are used for the single power supply connections.




Caution

A controller power supply failure causes the Command View EVA GUI to report that the controller enclosure is no longer part of the system. If the power supply fails in an SPS controller, you must replace the controller assembly with an SPS controller. If the power supply fails in a DPS controller, you must replace the failed blower/power supply assembly.

Device Port LEDs

To the right of each data connector (FP1, FP2, MP, 1B, 2B, 1A, 2A) is a two-colored LED that defines the link status. These LEDs do not indicate whether there is communication on the link. They indicate only whether the link is attached and can transmit and receive data.



The following describe the LED states:

- Amber LED on — Port failed, usually unseated cable
- Green LED on — Port is fully functional, logged into something on the other end of the cable
 - Host port — Logged into fabric
 - Device port — Logged into a loop
 - Mirror port — Logged into the other HSV controller
- Off — Port failed, unable to complete loop initialization
- Flashing green — Port failed, able to complete loop initialization but nonfunctional
 - The controller has determined that a device loop is failing and has disabled the port.
 - Use the *Port Config* option in the menu tree to troubleshoot.

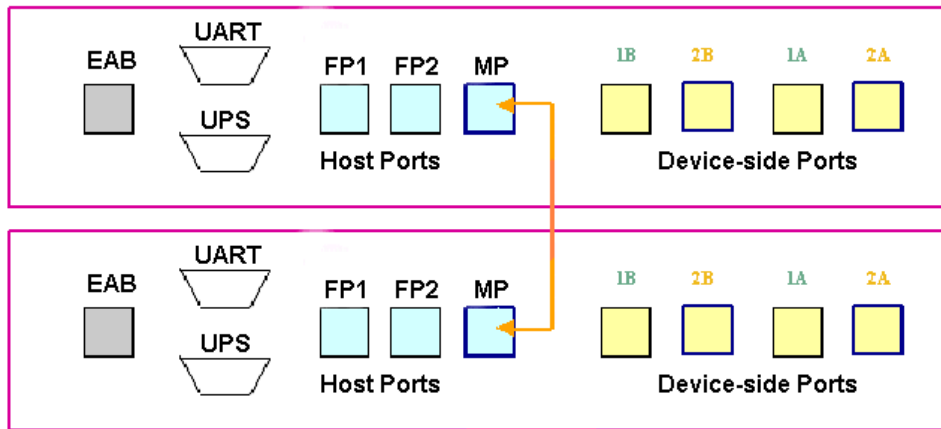
The Port Config option was added to the LCD panel as of VCS V2.002. It allows enabling a failed port device loop before a five-minute retry period. This is helpful when troubleshooting loop problems because you do not have to wait for the port to be enabled.

Mirror Port

The following are characteristics of the mirror port:

- Dedicated communications link between the two HSV controllers
- Point-to-point FC connection
- Primary heartbeat path between the two controllers
- Primary path for HSV cache mirroring
- If the mirror port fails:
 - The HSVs use one of the four FC-AL loops as the mirror link.
 - Caching performance is impacted slightly.


Each device side port transmits both directions on the cable.



Controller Port Numbers

When reporting events, ports are presented in two different formats:


- `cerp_id` — Controller enclosure rear panel identifier, for example, DP-1B
- Port identified by an internal number, for example, 1



These issues are covered in detail in Module 10, Event Logs and Diagnostics.

Diagnostic failures are reported as controller events of event information packet (EIP) type 14, which have the following characteristics:

- Contain a test entity number (`TE_num`), for example, TE 17.(0x11) is Port 4 (MP) test
- Viewable through the OCP on controller startup



Field Replaceable Units

The only FRUs in the single power supply HSV110 controller are the controller itself and the following:

- Blowers (hot-pluggable)
- Batteries (hot-pluggable)
- Cables (fibre and copper)
- Transceivers

For the dual power supply controller, the power supply/blower is also a hot-pluggable FRU.



Caution

The failure of any other controller elements requires replacing the controller enclosure.

Removing and replacing the FRUs are discussed in the troubleshooting lab. This module describes these elements and discusses where they are located in the HSV110 controller. Refer to the Service Manual for FRU repair and replacement procedures.

Controller Replacement Considerations

The following are some important considerations for replacing single power supply controllers with dual power supply controllers:

- You **cannot** mix different controller assemblies.
- Replacement requires VCS V2.002.
- **Original** controller assemblies are spared with **original** controller assemblies.
- **New** controller assemblies are spared with **new** controller assemblies.
- This is **not** an upgrade; it is a replacement, that is, the original controller assembly **cannot** be modified/upgraded to a dual power supply configuration.
- A Field Retrofit Order (FRO) **is not** generated.
- There is only one kit, and that kit includes long rails to replace units with short rails.
 - The latest models (with FC loop switches) have long rails. Nonswitched hardware versions Model 2C6D and 2C12D have short rails.
 - Hot replacement is supported on models with long rails only.
 - The Service Manual contains instructions for using the replacement kit for racks with long or short rails.



Important

A replacement controller assembly pair will be offered for sale to customers wanting to replace the original controller assembly; however, this expensive and complicated replacement is not being promoted.

Controller Swapping and Replacement

Both HSV controllers in a storage system must be running the same version of VCS. The operational HSV controller in the pair (that is, the HSV controller that is operational at the time another controller is added) dictates the VCS version used by the pair.

When you swap an HSV controller, the operational controller inspects the version of the replacement controller. If the VCS version on the new controller does not match the VCS version of the operational controller, the operational controller upgrades (or downgrades) the VCS of the replacement controller. When both controllers are new in the system, the **lower** software revision is upgraded to the **higher** revision.



Caution

To avoid an unintended upgrade or downgrade, HP highly recommends that you replace one HSV controller at a time in an operational environment. Simultaneous swapping of the controllers could result in a forced upgrade of the Enterprise Virtual Array, or even an unbootable system.

As an example, consider the following scenarios:

- Suppose you are replacing one controller of a pair running V1.02. The replacement controller (call it Controller B) from the pool of spares contains V2.0. If Controller A (running V1.02) is operational in the system when Controller B (with V2.0) is added, the operational Controller A loads its software version (V1.02) onto Controller B. As intended, both controllers are now running the customer's software version (V1.02).
- Suppose, however, that the system is shut down when Controller B is added. When power is re-applied, both controllers come up simultaneously. Because neither was operational when the other was added, the higher revision (V2.0 from Controller B) is loaded onto Controller A. EMU software is also upgraded. An unintentional upgrade has now occurred. Because no path to downgrade is available, you will have to upgrade other software components (such as Command View EVA) and purchase the appropriate V2.0 licenses.

It is important to consider compatibility interactions when replacing controllers. For example, if you must replace both controllers of a pair simultaneously (strongly not recommended), you should ensure that the VCS version in the replacement controllers is at least as high as the version in the former controllers. If two simultaneous replacement controllers contain a VCS version lower than the metadata version on the storage, the replacement controllers will be unable to locate the storage system name and will not boot. Because the controllers will not boot, the VCS version cannot be upgraded and the system is unusable.

Note

The Service Manual contains procedures for replacing a controller, and special procedures for replacing one controller of a pair running VCS V2.x or higher with a controller spare loaded with VCS V1.021.


Blowers and Cache Batteries

The blowers and cache batteries are located behind the bezel. The blowers are on the left side of the controller, while the cache batteries are located on the right side.

Use the following conventions when referring to displays:


- Blower 1 is blower A, the blower farthest from the blower door hinge.
Blower 2 is Blower B, the blower closest to the blower door hinge.
- Battery module 1 is battery module A, the battery assembly closest to the cache battery door hinge.
Battery module 2 is battery module B, the battery assembly farthest from the cache battery door hinge.

The following is a rear view of the SPS controller.



**SPS
Controller
Front View**

For the SPS controller, to access the blower or cache batteries, remove the front bezel and blower or battery door. The DPS controller has levers to pull each unit.




- ❶ Bezel
- ❷ Blower Door

Optical Fibre Channel Cables and Transceivers

Fibre Channel optical FC cable connections are used in the eva5000 for all host, device, and mirror port connections. The following are the three primary classifications:

- FC cables used to connect older 1Gb/s HBAs and older 1Gb/s SAN switches. These are Siemens Connector (SC) to SC cables and use SC connectors on both ends.
- FC cables used to connect HSV110 controllers, 2Gb/s HBAs, 2Gb/s SAN switches, and FC loop switches. These are Lucent Connector (LC) to LC cables and use LC connectors on both ends.
- FC cables used to connect HSV110 controllers to 1Gb/s HBAs, connect 1Gb/s HBAs to 2Gb/s SAN switches, or connect 2Gb/s HBAs to 1Gb/s SAN switches. These are LC to SC cables and use one transceiver type on each end.



The cables for switched hardware use LC connectors, also called small form-factor pluggable (SFP) transceivers (shown on the left below), on both ends to connect the controller with the FC switches, and to connect the controller with a 2Gb/s SAN switch.




Important

To remove a transceiver, pull it down slightly and then towards you.




Caution

When replacing optical Fibre Channel cables, you should replace one at a time. Replacing more than one cable at the same time can result in data loss.



The transceiver cannot be removed while the fiber optic cable is installed. When you grasp the optical Fibre Channel cable connector, you may also grasp the transceiver tab. If you pull on both the connector and the tab, you may break off the tab.

This diagram shows top, side, and bottom views of the transceiver. The side view reveals the potential breakage problems that might occur when pulling the fibre optic cable connector and transceiver tab.



Protecting Fiber Optic Connections

To ensure optimum operation, the fiber optic components (transceivers, fiber optic cable connectors, and fiber within the cable) require protection from contamination and mechanical hazards. Failure to provide this protection can reduce the amount of light passing through a component, thereby degrading operation.

The connections can be contaminated or damaged by:

- Being touched
- Dust
- Debris
- Being dropped



Caution

To protect the connectors against contamination or damage, always install the dust covers or dust caps whenever a transceiver or a fiber cable is disconnected. Remove, but do not discard, the dust covers or dust caps from transceiver or fiber cable connectors only when they are connected.

Use the guidelines in the following topics to minimize the risk of contamination or damage.

Dust Covers

Remove **and retain** the dust covers and dust caps when installing an I/O module, a transceiver, or a cable. Install the dust covers when disconnecting a transceiver or cable.

Note

The customer receives a plastic bag full of dust caps with the Enterprise Virtual Array.


Connectors

Any time you think a connector may be contaminated, or a connector has not been protected by a dust cover for an extended period, clean the connector. To clean a connector:

1. Wipe the connector with a lint-free tissue soaked with 91% isopropyl alcohol.
2. Wipe the connector with a dry, lint free tissue.
3. Dry the connector with moisture-free compressed air.

Fibre Channel Copper Cables

Fibre Channel copper cable connections are used in the eva3000 for device and mirror port connections, but not host connections. The copper cables are 2Gb (2.125 Gbps) duplex with two integrated SFP transceivers and are available in multiple lengths. A standalone SFP is not used.



For the eva3000, transceivers are only present in host connections.



Caution

To avoid deforming or possibly breaking the copper cable, do not allow it to bend with a radius smaller than 60 mm (2.4 in).

Enterprise Virtual Array 3000 Controller


The features of the eva3000 controller include the following:

- Controller model is HSV100 (based on HSV110).
 - Bezel on the controller indicates HSV100.
 - EEPROM contains model HSV100.
- Firmware is VCS V2.004.
- Firmware does not allow the eva5000 firmware to run on the eva3000 (enforced by EEPROM entry).
- Firmware programmatically disables two back-end ports that are not used (2A, 2B).
- Controller events do not reflect missing ports.
- Firmware restricts the number of disks to 56.

Rear Connections

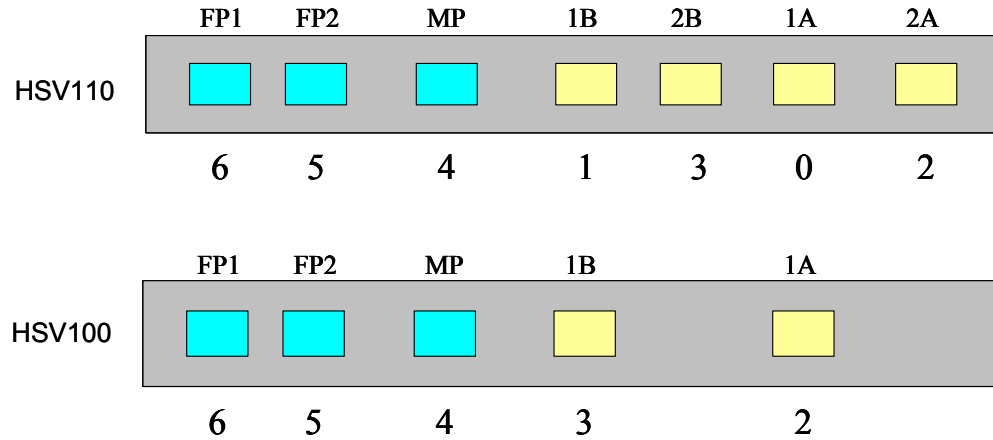
The following are features of the HSV100 rear connections:

- No back-end ports for data connections 2A and 2B
- Dual power supply only (diagram shows only one)



Controller Port Numbering Changes

The design of the HSV100 port numbering required that the mirror and host ports retain the same numbering as the HSV110. The main difference is the device port numbers. For the HSV100, the numbering for data connections 2A and 2B was not needed. Connections 1A and 1B took on those port numbers (2 and 3).



Learning Check

1. List four major elements located on the front of the HSV controller.
.....
.....
.....
.....
2. List two HSV controller features that provide redundancy to ensure high availability of the Enterprise Storage Array.
.....
.....
3. Which controller LED indicates a controller problem? Include the LED name and color.
.....
.....
4. What information displays on the Operator Control Panel LCD default display?
.....
5. Which three conditions override the default information that displays on the Operator Control Panel LCD?
.....
.....
.....
6. What displays on the Operator Control Panel LCD after each startup procedure (diagnostic) successfully completes?
.....
.....
7. Which controller features can you use to activate LCD displays, navigate through the menu tree, and enter values?
.....
.....
8. What are the password requirements (number and type of characters)?
.....

9. List the field-replaceable units (FRUs) in the HSV controller.

.....

.....

.....

10. Where does each of the data connections located on the rear of the HSV connect?

.....

.....

.....

Overview

This module describes, in detail, the physical, interconnect, and management features of the M5214 disk drive enclosure. The Model 5214 disk drive enclosure is a 2.125-Gb, dual-loop, 14-disk drive enclosure. The disk drive enclosures used in the Enterprise Virtual Array storage system support the FC-AL connections on both the external (host-controller-to-enclosure) and the internal buses.

In this module, the M5214 front and rear layout is described, as well as the components of the drive enclosure, including the disk drives, I/O modules, environmental monitoring unit (EMU), power supplies, and connections. This module also covers the displays and controls of the disk drive enclosure and its elements.

Note

The M5214 disk drive enclosure is used with the eva5000 and the M5114 is used with the eva3000. They are identical except that the I/O modules in the M5214 accept SFP transceivers and those in the M5114 do not. Instead, the M5114 accepts copper cables that have two integrated SFP transceivers.

Objectives

After completing this module, you should be able to:

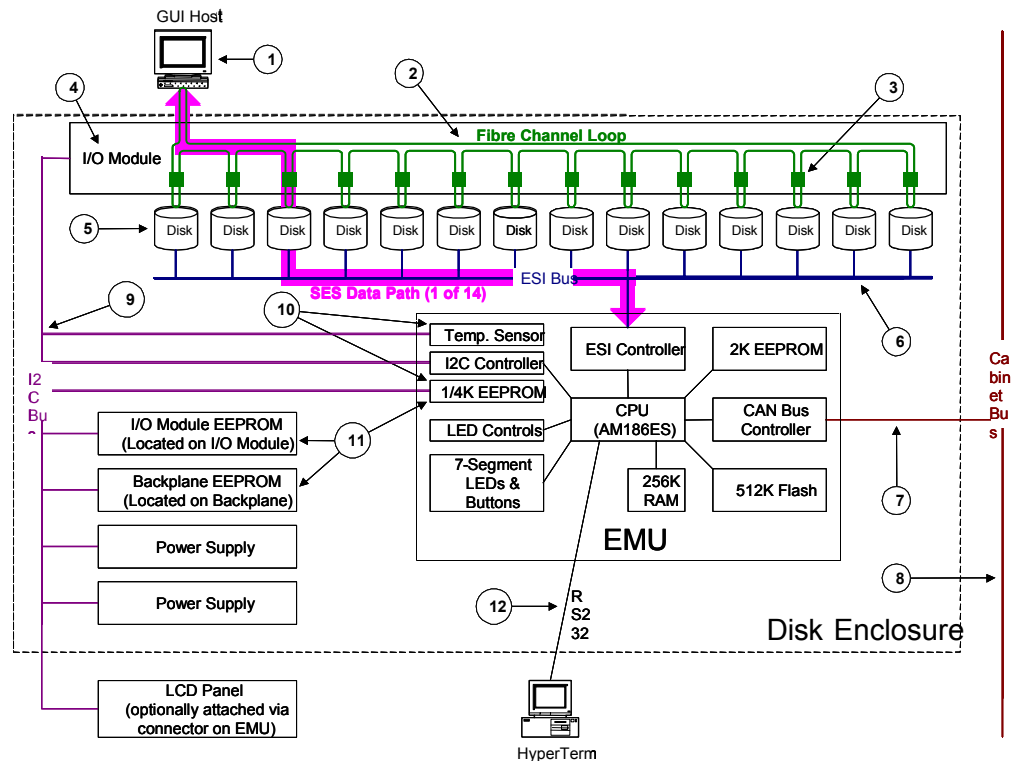
- Describe the features of the M5214 disk drive enclosure.
- Describe the front and rear layouts of the disk drive enclosure.
- Describe the components and functions of the EMU and I/O modules of the disk drive enclosure.
- Describe the components and functions of the blowers and power supplies of the disk drive enclosure.
- Identify the features of the enclosure address bus.
- Identify and describe each of the LEDs used for the disk drive enclosure.
- Recognize the LED display of I/O modules A and B.

Disk Drive Enclosure Block Diagram

SCSI-3 Enclosure Services (SES) is part of the SCSI-3 specification and provides a set of commands (SCSI-3 Primary Commands (SPC)) to be used by the storage management appliance (Command View EVA host) to communicate to a disk drive or EMU.

The following diagram shows the two distinct paths to the EMU:

1. Through the enclosure address bus (EAB)
2. Through the FC-AL using disk drive to Enclosure Services Interface (ESI) to EMU




Key to the previous diagram:

1. Storage management appliance running the Command View EVA.
2. FC-AL — Only one shown, but enclosure has two loops.
3. Bypass circuit for each drive to allow the EMU or I/O modules to bypass a problematic disk drive on the FC-AL.
4. I/O module — Only one shown but each enclosure has two.
5. Fibre Channel disk drive elements — Dual-loop disk drives can communicate through loop A or loop B.
6. Enclosure Services Interface (ESI) — Provides a second path to the EMU if the EAB is defective or unavailable.
7. Control area network (CAN) bus controller.
8. EAB.
9. I2C bus — Inter-IC bus. Internal communication path between EMU and enclosure elements.
10. Temperature sensor.
11. 1/4K EEPROM, I/O module EEPROM, and backplane EEPROM — Store revision information, WWN, serial names, and so on.
12. RS232 — Use only to correct EMU firmware download faults.

Drive Enclosure Front

The Fibre Channel disk drives are the only components of the drive enclosure that mount in the front of the disk drive enclosure. They mount in the bays that are numbered from the left bay to the right.

The common method of referring to a drive is by the bay number. In the following diagram, the drive in bay 1 is drive 1, the drive in bay 3 is drive 3, and so forth.




The enclosure status LEDs are located on the enclosure backplane and illuminate the status icons on the lower right front corner of the enclosure. These icons are described in the EMU section.

The disk drive enclosure status icons are also located on the front of the disk drive.

Disk Drives

The 14-disk drive enclosure supports only 1-inch, 3.5-inch form factor drives that consume less than 24 W of power. The drives are dual-ported 2Gb/s FC-AL, providing redundancy and load balancing. Up to 120 drives are supported per FC-AL pair, either 36GB, 72GB, or 146GB. A maximum of 126 arbitrated loop physical addresses (AL-PAs) are supported per loop, two for controllers and 120 for drives.



Disk Drive Status Icons

Three status icons on the front of each drive define the drive operational status. The LEDs that illuminate these icons are located on the enclosure backplane.

To determine the drive status, you must observe these icons. In some configurations, the host controller can control the status LEDs.

The drive status icons are:

- **Drive activity icon** — The green LED for this icon is controlled by the disk drive. Neither the HSV110 controller nor the EMU can control the LED. This prevents this LED from flashing when the locate function is active.




This LED is ON when there is no bus activity (opposite of traditional use). A LED may flash either slowly or rapidly to assist in identifying a defective drive.



Important

The drive activity and the drive on-line LEDs are controlled by the drive, and cannot be overridden by the Enterprise Virtual Array controller. The status of the LED is either off, flashing, or on.

- **Drive online icon** — The green LED for this icon is controlled by the disk drive. Neither the HSV controller nor the EMU can control the LED.



- **Drive failure icon** — The orange LED illuminates steadily with the other two LEDs in response to the EMU Locate command.



The LED for this icon flashes when the controller detects an error condition on the disk drive.

Middle and Bottom LED Status

The following describe the states of the middle and bottom LEDs:

- Middle flashing slowly, bottom on solid
 - Locate command issued
- Middle flashing fast, bottom off
 - Enabled as ILF disk
- Middle off, bottom flashing quickly
 - Drive failed, ready for removal
- Middle on, bottom off
 - Normal
- Middle on, bottom on
 - Drive is setting bottom LED because it is not receiving a signal. This is common for the first drive in the shelf when the controllers are off.

Disk Drive Power

The drive voltage backplane sensors (+5.1 VDC and +12.1 VDC) can detect a drive overcurrent condition.

In a drive overcurrent condition, the sensor disconnects the voltage from the drive to prevent writing data to the drive. The drive is disabled until one of the following conditions occurs:

- The defective drive is replaced.
- The overcurrent condition no longer exists.

Disk Drive Blank

To maintain the proper enclosure airflow, a drive or a drive blank must be installed in each drive bay. The function of the drive blanks is to control airflow within a bay.

Replacing a disk drive is covered in detail in the troubleshooting lab; however, keep in mind the following caution.




Caution

HP recommends that you never remove more than one drive at a time because this can cause the enclosure to overheat. To prevent overheating and ensure proper operation, a drive blank should be installed as soon as possible.

Drive Enclosure Rear

This illustration shows the components that mount in the rear of the disk drive enclosure. The power and the data connections are also located on the rear of the drive enclosure.



Key to the diagram:

- 1 — Environmental monitoring unit (EMU)
- 2 — FC-AL I/O module B
- 3 — Power supply 1 and blower 1
- 4 — Power supply 2 and blower 2
- 5 — FC-AL I/O module A

The following is a description of each element:


- Dual-loop I/O modules
 - Each disk drive enclosure is connected to two independent, external FC-AL buses, or loops through I/O module A and I/O module B. The HSV controller establishes these loops.
 - The I/O modules establish the internal enclosure buses, and all the drives connect to both loop A and loop B. I/O module A connects to loop A, and I/O module B connects to loop B.
- Power supplies — Two each
 - The FC-AL enclosure uses a redundant power supply configuration (two power supplies). This prevents the failure of a single power supply from disabling the enclosure.
 - A single power supply can support all enclosure operations.

- Blowers — Two each
 - All enclosures have two blowers. If one blower fails, the other blower automatically operates at a higher speed. Therefore, the enclosure is operational as long as a single blower is operational.
- Environmental monitoring unit (EMU)
 - The EMU is fully SCSI-3 Enclosure Services (SES) compliant and mounts in the left rear bay of a disk enclosure. Each EMU contains an Enclosure Services Processor (ESP) that controls the Enclosure Services Interface (ESI).
 - All EMUs can transmit information to the host. However, only one EMU in each reporting group or on a loop can communicate directly with the host at a time.

Environmental Monitoring Unit

The EMU collects the data for the associated enclosure. The EMUs fully comply with the SES standards, and they mount in the left rear bay of a disk enclosure. In addition:

- Each EMU contains an ESP that controls the ESI.
- The EMU uses a combination of LEDs, icons, the two-character alphanumeric display, and an audible alarm to notify you of the operational status of the enclosure and the enclosure elements.
- The EMU is a hot pluggable element (you can remove and replace without stopping data transfers or removing power).



EMU Functions

The EMU performs the following functions:

- Displays the enclosure number, based upon the enclosure address bus (EAB) location
- Displays the assigned bay 1 loop ID
- Detects, reports, records, and displays errors
 - Blowers/power supplies
 - Drives (insertion, removal, failure, bypass)
 - EMU/enclosure
 - I/O modules and transceivers
- Implements automatic corrective actions for some errors
- Provides enclosure status data to a GUI and the controller



Important


Although the EMU can determine the AL_PA of a drive, the EMU cannot display or change this information. The Command View EVA GUI can display the addresses from the EMU-supplied status information.

EMU LED Status Displays

The EMU status LEDs are arranged horizontally above the alphanumeric display. These icons are the same as those on the front, lower-right corner of the disk drive enclosure. When the EMU and the enclosure are operational, the LEDs, from left to right, are flashing, on, and off.

The EMU LED displays include:

- EMU heartbeat status (1)
 - When **not** flashing green, there is a problem with the EMU.
- Enclosure power status (2)
 - On when levels are OK
 - Off when there is a power problem
- Enclosure fault status (3)
 - Flashes amber when there is any enclosure problem



EMU "Heartbeat" (green LED)	Power (green LED)	Enclosure Fault (orange LED)


EMU RJ-45 Connectors

The following are the RJ-45 connectors on the EMU:

- RS-232 (1)
 - Allows EMU firmware updates
- Cabinet LCD (2)
 - Unused
- EAB (3)
 - Connects to the EAB

Note


All three connectors are keyed differently.



EMU Alphanumeric Display and Pushbuttons

The following are characteristics of the EMU alphanumeric display and pushbuttons:

- Seven-segment, two-character display (1)
- Pushbuttons control the data displayed or entered. They are display group dependent. The default display is the enclosure number, a decimal number in the range 00 through 14 for a single rack.
 - Function Select button (2)
 - Group Select button (3)
- Using the pushbuttons
 - Push and hold
 - ♦ Push the button and hold until the interface changes.
 - ♦ Release the button.
 - Push and release



EMU Pushbutton LEDs

The EMU pushbutton LEDs define error conditions and the state of the audible alarm. The LEDs in the pushbuttons may be on when:

- An error occurs (the top pushbutton LED is on (lit solid)).
 - When a single error occurs, the LED is on until you view the error.
 - When multiple errors occur, the LED is on until you view the last error.
- The audible alarm is muted or disabled (the bottom pushbutton LED is on (lit solid)).


EMU Display Groups

Display group symbol descriptions are as follows:

- **En (enclosure number)** — The enclosure number is the default display and is a decimal number in the range 01 through 14 (for a single rack) that is a subset of the enclosure number display group.
- **Li (bay 1 loop ID)** — This display group has a single sublevel display that defines the enclosure bay 1 loop ID. Valid loop IDs are in the range 00 through 7F. The setting 7F indicates to **not** request a bay 1 loop ID. Do not change this setting because it is the only supported setting.
- **rG (reporting group)** — This display group has two 2-digit displays that define the reporting group number in the range 0000 through 4095.
- **Au (audible alarm)** — This display group provides control over the audible alarm. The sublevel displays are audible alarm enabled (**on**) or audible alarm disabled (**oF**).
- **Fr (firmware revision)** — This display group defines the EMU firmware version.
- **Er (error report)** — This display group has precedence over all other display groups and is only active when an error occurs. The report defines the element type, the element number, and the error code.

Regardless of what is being displayed, anytime you press and release the bottom pushbutton, the display changes to **En**, **Li**, **rG**, **Au**, **Fr**, or **Er**.

A flashing alphanumeric display indicates that you can edit an address or state, or display an error report.



Note

The display above would indicate Er in place of En if an error occurred at the time.

You can cycle through all the display groups rapidly by pressing and holding the bottom pushbutton. When you press and hold the pushbutton, only hold it for five to eight seconds. If you hold too long, the display goes into change mode instead of view mode.

Enclosure Number (En)

In a single-rack configuration, the enclosure number is a decimal number in the range 00 through 14, which is automatically assigned by the EAB. By default, the two-character alphanumeric display shows this number.

- Pressing the bottom pushbutton changes the display to **En**, (enclosure number display mode).
- When the display is **En**, pressing and releasing the top pushbutton displays the enclosure number.

Note

If your storage system uses an expansion rack, enclosure numbers may be in the range 00 to 24.

A display of **00** indicates that the enclosure is not connected to the EAB. When this condition exists, there is no EMU-to-EMU communication over the EAB.

A display of **01 through 14** indicates that the enclosure is physically connected to the EAB and can exchange information with other enclosures on the EAB.

The decimal number indicates the physical position of the enclosure in relation to the bottom of the rack.


- **01** is the address of the enclosure connected to the bottom connector in the first (bottom) junction box.
- **14** is the address of the enclosure closest to end of the bus, the top connector in the last (upper) junction box.

Enclosure Address Bus Components

The EAB is composed of cables and junction boxes (JBs) that interconnect the controller enclosures with the drive enclosures.

The following are some facts about EAB connections:

- EAB junction boxes are placed at 6U increments in the left rear rail of the rack.
- The junction boxes fit between the enclosure mounting rails.
- Racks come fully wired with seven junction boxes per rack whether or not the rack is fully populated with device enclosures.



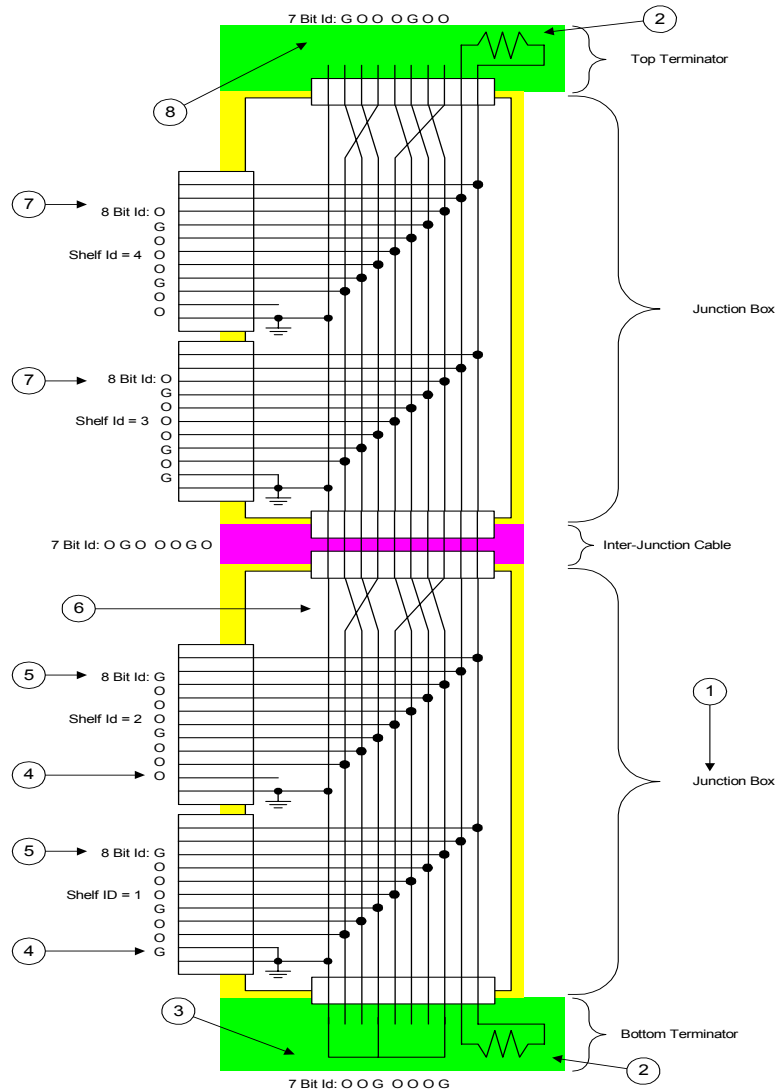
The functions of the EAB are to:

- Allow management and reporting of environmental conditions within the rack.
 - Use environmental monitoring units (EMUs) to collect data.
 - One EMU for each drive enclosure.
- Connect drive enclosures and controllers through the EMUs through cables and junction boxes.
- Assign enclosure number/address to the EMUs.
 - Allows identification of shelf location in the rack.
- Offer the HSV controller a way to bypass FC disk drives.

EAB Cable

The EAB cable uses a ground wire to establish enclosure addresses. The cable consists of:


- Controller Area Network (CAN) bus (two wires), which carries cabinet-specific data.
- Address bus (seven wires), used to establish enclosure numbers.
- A ground wire (one wire), used to create unique addresses.



File: Cab Cable Overview Diagram.vsd
Modified: March 6, 2000, 02:05

EAB Cabling

The following diagram shows the enclosure numbering (1 – 14) in an eva5000 41U rack **with** Fibre Channel loop switches.



The following are characteristics of the 41U rack with switches:

- Controllers are assigned enclosure number 7.
- 3U space for enclosure 8 is skipped by the EAB.
- JB 4 connects to the controller pair with a Y cable.


Note

The top connector of JB 4 is not used, the top connector of JB 7 is used, and the enclosure number of the highest drive enclosure is 14.

This scheme also applies to:

- Model 2C2D, which comes in a 41U rack without switches.
- A 42U rack that has had switches added.

The following diagram shows the enclosure numbering (1 – 14) in an eva5000 42U rack **without** Fibre Channel loop switches.




The following are characteristics of the 42U rack without switches:

- Controllers are assigned enclosure number 7.
- JB 4 connects to the controller pair with a Y cable and to the drive enclosure in the next highest 3U space (enclosure 8).

Note

The top connector of JB 4 is used, the top connector of JB 7 is not used, and the enclosure number of the highest drive enclosure is 13.

The following diagram shows the enclosure numbering (1 – 5) in an eva3000 42U rack with a 2C4D configuration.



The following are characteristics of the eva3000 rack:

- Controllers are assigned enclosure number 3.
- JB 2 connects to the controller pair with a Y cable and to no other enclosure.

Loop ID (Li)

The Loop ID is a hexadecimal number that identifies the assigned AL_PA for enclosure bay 1. Loop IDs are in the range 00 through 7F.

- Each loop ID corresponds to an AL_PA address (the Service Manual has a table).
 - Loop IDs 7C and 7D (AL_PA 01 and 02) are reserved for the controllers.
 - Loop ID 7E (AL_PA 00) is reserved for the fabric port.
 - Loop ID 7F has a special meaning: no specific AL_PA is recommended, and the device or drive should choose the AL_PA. Therefore, the bay 1 loop ID will be determined by negotiation.
- The loop ID is a recommendation to assign a specific AL_PA to a specific bay.



Caution

Use the Li function for viewing the bay 1 loop ID only. Do not try to manually set or alter the loop ID because you may create problems on the loop.

Physical Addressing

Fibre Channel eliminates the need to manually set any address by allowing the interconnecting topology or the associated ports to assign the addresses. An arbitrated loop can have a maximum of 126 devices (although using fewer devices tends to improve performance because each port on the loop causes a signal delay). A loop must have a minimum of two devices, a transmitter and a receiver.

Characteristics of AL_PAs include the following:

- A 1-byte (8-bit) value that identifies a port in the arbitrated loop topology.
- Device address used by the Fibre Channel loop.
- Have a range of consecutive, but not contiguous, values between 00 and EF.
- Assignments are made automatically by the arbitrated loop protocol.
- Assignment is temporary and is negotiated each time the loop is initialized.

VCS versions before V2.002 use soft addressing, in which:

- No attempt is made to set the AL_PAs.
- Initially, each drive has the loop ID set to 7F to indicate that the EMU is not recommending an AL_PA.
- The drives negotiate for the first available AL_PA starting with bay 1 in each enclosure, based upon the sequence in which the drives spin up.
- It is **not** recommended to use drives 13 and 14 on enclosures 17 and 20.

VCS V2.002 and higher uses hard addressing, in which software attempts to force each drive enclosure to use hard addressing when obtaining AL_PAs on each loop.

The following conditions must be met to use hard addressing:

- No hardware faults or system misconfigurations exist.
- Two enclosures with the recommended same AL_PA range are not on the same loop.
- The configuration contains only the drive enclosure numbers listed in the following table. For configurations containing drive enclosure numbers **not** listed in the table, refer to the Service Manual.

Drive Enclosure Numbers	AL_PA Range (in hex)	Loop ID (in hex)
1, 8, and 14 ¹	36 through 26 ²	62
2 and 9	51 through 39	54
3 and 10	6A through 52	46
4 and 11	80 through 6B	38
5 and 12	A5 through 81	2A
6 and 13	B6 through A6	1C
15 and 18 ³	D1 through B9	0E
16 and 19 ³	EF through D2	00
17 and 20 ³	25 through 01 ⁴	70

Hard Addressing of Disk Drives in VCS V2.002 and Higher

The superscripts in the table refer to the following:

1. In VCS V2.002, drive enclosure 14 does not have an assigned loop ID and AL_PA range.
2. The range contains more than 14 numbers because AL_PAs skip some sequential numbers.
3. Starting in VCS V2.002, drive enclosures 18, 19, and 20 do not have an assigned loop ID and AL_PA range.
4. Drive bays 13 and 14 of any drive enclosure with loop ID 70 overlap with AL_PAs 02 and 01 on the controllers. This overlap rarely occurs because, under most circumstances, only seldom-used drive enclosures (17) have loop ID 70. However, if nonstandard addressing takes place, a drive enclosure may occasionally receive loop ID 70. If this condition should occur, the disk drives in bays 13 and 14 of the enclosure with loop ID 70 should be moved to other bays.

Reporting Group (rG)

A *reporting group* is an HSV controller pair and the associated Enterprise Virtual Array storage system disk enclosures. The HSV controller pair assigns a unique four-digit reporting group number (RGN) to each FC-AL loop. Therefore, each Enterprise Virtual Array has two FC-AL loop pairs.

Each of the drive enclosures on a loop pair is in one reporting group.

- All of the drive enclosures on loop pair 1, both loop 1A and loop 1B, share a unique reporting group number.
- All of the drive enclosures on loop pair 2, both loop 2A and loop 2B, share a unique reporting group number.

Each EMU collects environmental information from the associated enclosure and broadcasts this information to reporting group members using the EAB.

Information from enclosures in other reporting groups is ignored. The reporting group feature provides a way to segregate, identify, categorize, isolate, and separate controller enclosure and drive enclosure environmental information. The feature does not use the FC-AL bus, or inband communications, but instead uses out-of-band communications.

Note

The EMU can request a different RGN for an enclosure. Implementing this option in an Enterprise storage system may cause problems. Therefore, procedures for requesting a different RGN are not documented.

Reporting Group Numbers

The reporting group number (RGN) range is 0000 through 4099, decimal.

- 0000 is reserved for enclosures that are not part of any reporting group, a JBOD configuration.
- 0001 through 0015 are RGNs reserved for use by the EMU. When entered by the user, they default to 0000.
- 0016 through 4095 are valid RGNs.
- 4096 through 4099 are invalid RGNs. When entered by the user, they default to 4095.


The reporting group numbers are displayed on the EMU alphanumeric display as a pair of two-digit displays. These two displays are identified as **rH** and **rL**.

- Valid **rH** displays are in the range 00 through 40, and represent the high-order (most significant) two digits of the RGN.
- Valid **rL** displays are in the range 00 through 99, and represent the low-order (least significant) two digits of the RGN.


Viewing a Reporting Group Number

The following is the procedure to view a reporting group number, such as 2553. All the examples for navigating the alphanumeric display are for **En 06** and a bay 1 loop ID of 18.

1. Press and release the bottom pushbutton until the alphanumeric display is **rG**.



2. To display the two most significant Reporting Group numbers press and hold the top pushbutton until the display is **rH**.



3. Press and release the top pushbutton to display the first two digits of the RGN, 25.
4. Press and release the top pushbutton until the alphanumeric display is **rH**.
5. Press and hold the top pushbutton until the alphanumeric display is **rL**.
6. Press and release the top pushbutton to display the last two digits of the RGN, 53.
7. To exit the display, press and release the bottom pushbutton until the alphanumeric display is **rG**.

Audible Alarm (Au)


Whenever an error occurs, the EMU audible alarm automatically sounds until all errors are corrected. You can either mute or disable the alarm. Using these options establishes the following conditions:

- Muted alarm — Sounds whenever another error occurs.
- Disabled alarm — **Cannot sound** when another error occurs.

Audible Alarm Patterns

The duration and number of times the audible alarm sounds (the sound pattern) is a function of the error severity.

The following diagram depicts the duration and the approximate relationship of these alarms. The most severe active error controls the alarm pattern.



Controlling the Audible Alarm

Control (mute and disable) the alarm with the bottom pushbutton on the EMU.

When an error occurs, the alphanumeric display is **Er**, the alarm sounds, and you can do any of the following:

- **Correct all errors**, thereby silencing the alarm.
- **Mute or temporarily disable the alarm** by pressing and holding the bottom pushbutton. The alarm remains off until another error occurs, or until you enable (unmute) the alarm. When a **new** error occurs, the alarm sounds and the bottom pushbutton LED is off (not lit).

Using this feature ensures you are aware of any new errors so that you can evaluate them and implement the necessary corrective action.

- **Disable or turn off the alarm.** A disabled alarm cannot sound for any new error, regardless of severity.

The bottom EMU pushbutton LED defines the state of the audible alarm.

- When the alarm is enabled (**on**) or **unmuted**, the bottom pushbutton LED is off.
- When alarm is disabled (**oF**), or **muted**, the bottom pushbutton LED is on.

If an error does not require immediate corrective action, you can turn the alarm off by muting it.

**Important**

Muting the alarm is the preferred method for turning the alarm off. This process reduces the noise level while ensuring that you are aware of any new errors.

Error Reports (Er)

The EMU reports errors to you by changing the alphanumeric display to **Er** and sounding the audible alarm. An error report has precedence over all other displays and activates the error menu. An error always generates a condition report; however, not all condition reports are generated by errors.

The following are some other characteristics of a condition report:

- Contains the element type, element number, and error code for the EMU-detected condition
- Appears on the alphanumeric display
- Sounds the audible alarm
- Is stored in the error queue
- Is provided to the HSV110 controller for processing and display through the Command View EVA

For a full list of error codes for all element types, see Appendix A.

Analyzing an Error Report

The following shows a sample condition report.

e.t.nn.er	S	Error Description
0.1.**.01	C	Drive configuration error; unsupported link rate, replace drive.
0.1.**.02	I	Drive missing; install a drive or blank.
0.1.**.03	I	Drive removed when locked; clear software lock before removal.

Note that e.t. is the element type, nn is the element number, and er is the error code.

Analyzing an error report involves:

- Identifying the element.
- Determining the major problem.
- Defining additional problem information.

Before you can do this, you must understand the error condition format and how to navigate the error display.

Error Condition Report Format

Each condition report identifies the element affected (type and number), and the primary problem (the error code). The full format of an error condition is ***e.t.en.ec***, followed by the severity level. The following defines each item:

- **Element type** — This two-digit **hexadecimal** display defines the element type reporting the problem. The format for this display is *e.t.* with a period after each character. Valid element types are *0.1.* through *F.F.*
- **Element number** — The second display is a two-digit **decimal number** that defines the element reporting the problem. The format for this display is *en.* with a period after the second character.
- **Error code** — The third display is a two-digit **decimal number** that defines the specific error code. The format for this display is *ec* without any periods.

The severity levels, in order of precedence, are:

- **UNRECOVERABLE** — Most severe condition. Take corrective action immediately.

This most severe error is active when one or more elements have failed and disabled some enclosure functions. The enclosure may be incapable of correcting or bypassing the failure, which requires you to take action to correct the error. Other elements may be able to continue normal operations. This type of error requires immediate corrective action by you.

An unrecoverable error establishes the following conditions:

- This error has precedence over critical and noncritical errors.
- This error has precedence over informational messages.
- The alarm sounds constantly.
- **CRITICAL** — Less severe than UNRECOVERABLE condition. Take prompt corrective action to prevent system degradation.
- **NONCRITICAL** — Less severe than CRITICAL condition. Take early corrective action to prevent system degradation.
- **INFORMATION** — A condition exists that does not reduce the capability of an element. Prepare to implement corrective action if necessary.

Example

0.3.02.01 N Fan speed alert; replace fan soon

0.3 is a blower condition.

0.2 is blower number two.

01 indicates a blower speed alert to replace the blower soon.

N indicates NONCRITICAL.

For a full list of error codes for all element types, see Appendix A.

EMU Error Queue

The EMU stores each condition report in the error queue until the problem is corrected, or for at least 15 seconds after the error is reported. This ensures that when there are multiple or recurring errors, each is displayed. You can recall each report using a combination of the top and bottom pushbuttons. When you correct an error, the condition report is removed from the queue. Furthermore, to remove information from the queue requires either clearing the error or replacing the EMU.

The EMU displays the error queue information based on the severity of the error and the timestamp of the error. The hierarchical display ensures that the most severe errors have precedence over less severe errors. Within the error type, earlier errors have precedence.

For example, if there are errors at all levels, the EMU displays them as follows:

- UNRECOVERABLE errors in the sequence they occurred
- CRITICAL errors in the sequence they occurred
- NONCRITICAL errors in the sequence they occurred
- INFORMATION errors in the sequence they occurred

Navigating the Error Display

To navigate the error display to collect error codes, use the top pushbutton as indicated:


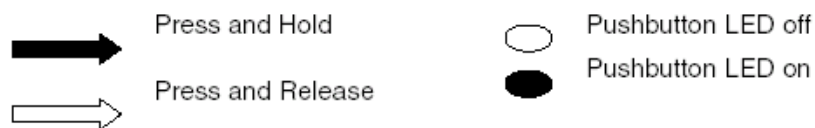
- Move from the **error display** to the **element type display** — Press and hold the pushbutton. Record the element type.
- Move from the **element type display** to the **element number display** — Press and release the pushbutton. Record the element number.
- Move from the **element number display** to the **error code display** — Press and release the pushbutton. Record the error code.
- Move from the **error code** back to the **element type** — Press and release the pushbutton.
- Move from one **error entry** to the **next** — Press and hold the pushbutton.

The bottom pushbutton function allows you to move from any sublevel directly to the **Er** display. To do this, press and release the pushbutton.

Sample Error Display

The following shows a high-temperature CRITICAL condition for the disk drive in bay 9 (element number 12).

LEGEND



The element type is 0.4 (temperature sensor), the element number is 12 (disk drive in bay 9), and the error code is 02 (critical high temperature). All of these codes are available in Appendix A of this student guide or in the Service Manual.

EMU Firmware Update

Two different types of firmware run in the EMU:

- *Boot code* initializes the EMU internal microprocessor and other electronic components. It then jumps to the application code and begins execution. Boot code contains the program code that downloads new versions of application code. Boot code is factory loaded using a special fixture and cannot be loaded in the field.
- *Application code* performs the various EMU functions: monitoring the disk enclosure and its components, communicating with other enclosures and devices, and so forth. Unlike boot code, application code can be downloaded both automatically by the controller and manually by authorized service personnel.

Loading Firmware

Firmware is loaded in different ways depending on type.

For boot code:

- If boot code becomes corrupt (-b in display)
 - No field service repair
 - Must be returned to repair facility

For application code:

- Loaded after shipment in two ways
 - Downloaded automatically by the controller
 - ◆ Uses superfile
 - ◆ Occurs if EMU firmware mismatch
 - ◆ Begins as soon as the controller becomes aware of the EMU
 - Downloaded manually by HP Services
 - ◆ Needed if automatic download fails
 - ◆ Indicated by -A in display

An automatic download may fail due to a firmware update interruption before completion. This could be caused by a power interruption event such as:

- Hot swapping the EMU.
- Powering down the drive enclosure.
- Powering down the rack by cycling the PDU.

If the EMU becomes nonfunctional, you have two choices:

- Return the unit to a repair facility.
- Download the firmware manually.
 - Uses an RJ-45 connector on the EMU
 - Requires use of an EMU-to-PC cable
 - Requires built-in Xmodem software in the EMU and HyperTerminal in the PC

Common alphanumeric displays during update include the following:

- Ld — Code load in progress or incomplete
- A — EMU firmware application code is missing
- b — EMU firmware boot code is corrupt
- — Both pushbuttons pushed in and held while inserting the EMU
- 8.8. — EMU is not seated properly

Note


The Service Manual contains several pages of detailed instructions on how to do a manual download.

FC-AL I/O Modules

The two I/O modules (I/O module A and I/O module B) are the interfaces between the disk drive enclosure elements and the host controller. They route data to and from the drives using a dual-loop configuration, loop A and loop B.

The Enterprise Virtual Array only supports dual-controller, dual-loop operation. Each controller, by way of FC loop switches (switched hardware), is connected to both I/O module A and I/O module B in each disk drive enclosure.

I/O
Module B




Note

The photo shows the Enterprise Virtual Array nonswitched hardware configuration. The switched hardware configuration, using FC loop switches, does not use a connection to the bottom transceiver.


I/O Modules

I/O modules A and B are similar in appearance and function in the same manner, but they **are not interchangeable**. The physical differences between the I/O modules ensure that you can only install them as follows:

- The A module at the right end of the enclosure, behind drive bay 1.
- The B module at the left end of the enclosure, next to the EMU.



The following graphics show how the modules are physically different, and identify the three green LEDs that display the I/O module status.



I/O Module Revisions

Two revisions of I/O modules are current.

- Revision D2
 - These and earlier I/O modules require SFP transceivers in the lower ports of the module, even when the configuration contains FC loop switches (lower port is not used).
- Revision Ex
 - Revision E2 and higher modules can be identified by the port-colored locking tab that holds the module in place.
 - Revision E1 modules do not have port-colored locking tabs. A small number of these modules were shipped.
 - Revision Ex modules do not require SFP transceivers in the lower ports of the module when the configuration contains FC loop switches. A protective plug is installed in the lower port.

Note

Field spares can use revision D2 or later in all situations.

Revision E modules have different packaging compared to lower-revision modules. A protective cover is installed over some components on the printed circuit board to protect these components from scraping and damage during module insertion into and removal from the drive enclosure.

All supported revisions of I/O modules can coexist in a drive enclosure.

I/O Module Status LEDs

The three green LEDs display the I/O module status. The following describes each LED:

- Top port status LED — Monitors the receive function of transceiver element 2 or 4
- Power status LED — Monitors the presence of power
- Bottom port status LED — Monitors the receive function of transceiver element 1 or 3

I/O Module Communication Paths

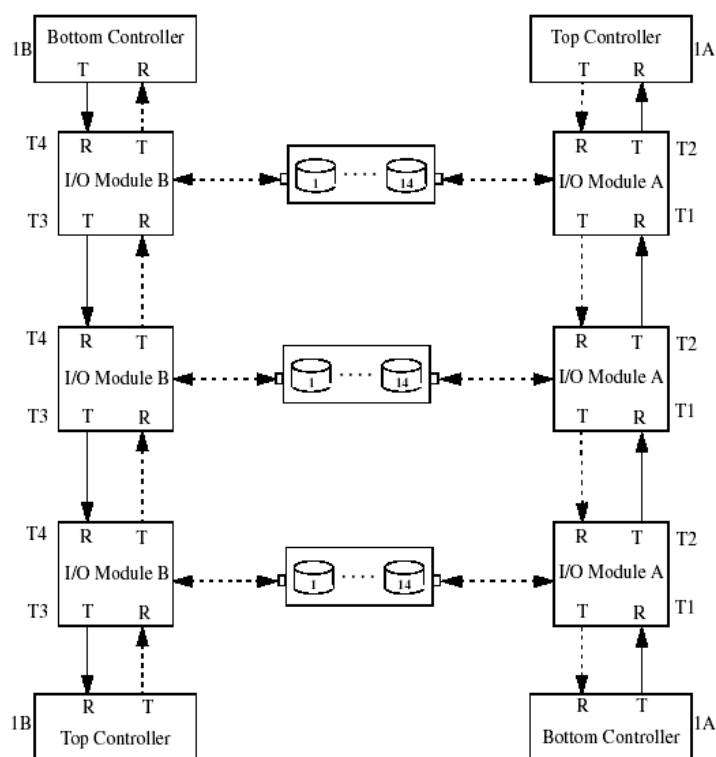
The I/O modules are major components in the Fibre Channel loop. Each module has two ports that can transmit and receive data.

To activate a port, you must connect a transceiver to the port. The FC-AL enclosures use short wavelength, fiber optic transceivers with 50 micron (μm) diameter, multimode fiber cables. The port function depends upon the loop.

Different communication paths are taken, depending on whether the configuration is switched or nonswitched.

Nonswitched

In a nonswitched configuration, each FC loop (1A, 1B, 2A, 2B) has a bidirectional communication path that flows through the top and bottom transceivers and includes every I/O module.

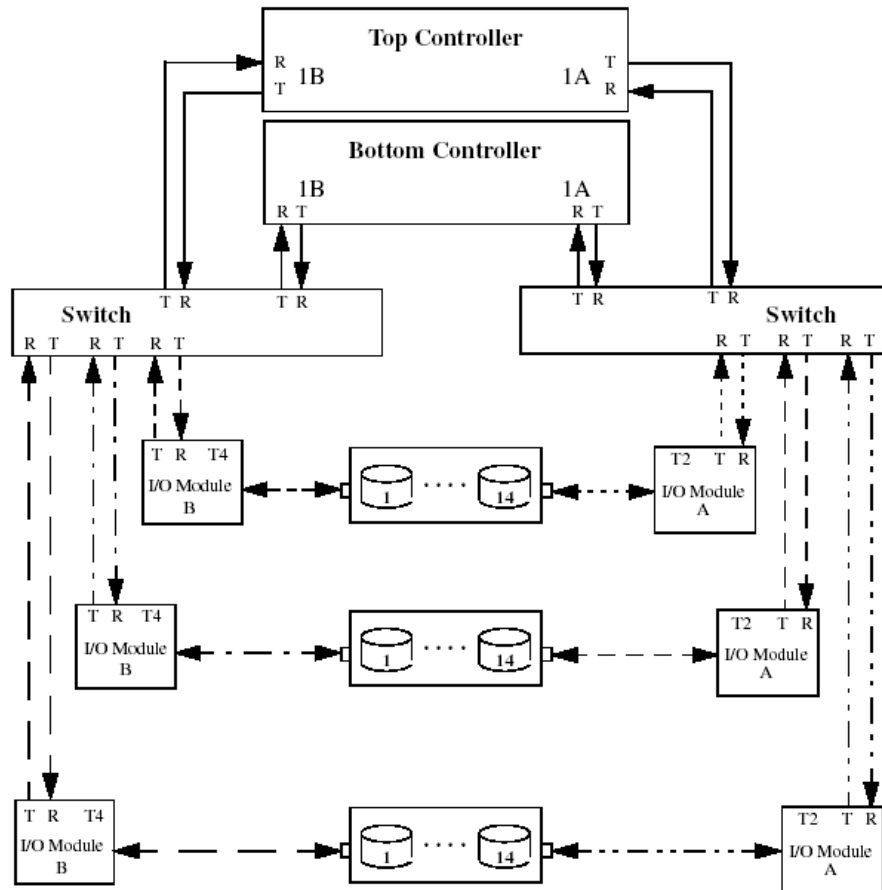


Note the following features of this communication design:

- Each transceiver (T1 through T4) in each I/O module and controller has both a transmit (T) and a receive (R) function.
- The signal flows from the receive function of an upper I/O module transceiver (T2 in I/O module A or T4 in I/O module B) through the I/O module and out the transmit function of the bottom transceiver (T1 in I/O module A or T3 in I/O module B).
- The signal flows from the receive function of a lower I/O module transceiver (T1 in I/O module A or T3 in I/O module B) through the I/O module and out the transmit function of the top transceiver (T2 in I/O module A or T4 in I/O module B).
- The signal traverses the drive enclosure between the receive function and the transmit function of each I/O module. Communication within the drive enclosure starts with bay 1 and proceeds sequentially through bay 14. (The graphic of the drive enclosures depicts the dual-ported design of the drives, not the access order.)
- The signal flows through the loop in an indeterminate order. Each I/O module independently determines whether to link its enclosure drive chain on either the upward traversal or the downward traversal shown in the graphic. The determination of the I/O module is related to a light race that takes place on each initialization. An established drive linkage order remains in effect until an I/O module loses and then regains laser signal, which can happen on a reboot or a cable change.
- The port status LEDs on the I/O module connect to the receive function of each transceiver. Thus, a port LED may be lit due to a problem on that transceiver or I/O module, or on the downstream or upstream transceiver or I/O module. This design has implications for error reporting related to transceivers, I/O modules, and other elements of the FC loop.

Switched

In a switched configuration, the communication path flows through the controllers and switches to the top transceivers of each I/O module to form each FC loop.



Note the following features of this bidirectional communication:

- Each transceiver in each I/O module, switch, and controller has both a transmit and a receive function.
- Each I/O module actively uses only one transceiver located in the top port (transceiver element number 2 in I/O module A and transceiver element number 4 in I/O module B).

Note

In I/O module Rev D2 and lower, a transceiver must be present in the lower port of each I/O module, even though it is not used. In I/O module rev E and higher, the lower port does not require a transceiver; a dust cover is installed instead.

- All drive enclosures within the storage system share the same address space, so each disk drive needs a unique address.


Note

Communication within the drive enclosure starts with bay 1 and proceeds sequentially through bay 14. The graphic of the drive enclosures depicts the dual-ported design of the drives, not the access order.

- Intelligence in the switches allows data transfers to and from the controllers to be confined to the individual I/O modules, indicated in the figure by a unique line pattern from the switch to the I/O module and its drive enclosure.
- The port status LEDs on the I/O module connect to the receive function of each transceiver. Thus, a port LED may be lit due to a problem on that transceiver or I/O module, or on the downstream or upstream transceiver or I/O module. This design has implications for error reporting related to transceivers, I/O modules, and other elements of the FC loop.

Power Supplies and Blowers

Each enclosure has two operational power supplies and blowers.



Power Supplies

The two power supplies mount in the rear of the enclosure. The supplies are auto-ranging and operate on a country-specific AC input voltage of 202 to 240 VAC $\pm 10\%$, 50Hz to 60Hz, $\pm 5\%$, (188 to 264 VAC, 47Hz to 63Hz).

The DC outputs of this power supply are:

- +5.1 VDC for the EMU, I/O module, backplane, and drives
- +12.1 VDC for the drives
- +12.5 VDC for the blower

The nominal output of each power supply is 499 W, with a peak output of 681 W. A single power supply can support an enclosure with a full complement of disks, blowers, EMU, and I/O module.

The power supply circuitry provides protection against:

- Overloads
- Short circuits
- Overheating

Power supply status and diagnostic information is reported to the EMU with voltage, current, and temperature signals.

Blowers

The blower mounts on the rear of the power supply. A power supply connector is the interface between the blower and the enclosure for:

- Blower speed control to the blower.
- Blower speed to the EMU through the power supply.
- Power supply high speed enable.
- Blower operating voltage.

The power supply-mounted blowers cool the enclosure by circulating air through the enclosure and the elements. These blowers, under the control of the EMU or the associated power supply, can operate at multiple speeds to ensure that when the enclosure temperature changes, the blowers can automatically adjust the airflow. Should a blower operate too slowly or stop (a blower failure), internal circuitry automatically causes the operational blower to run at a higher speed.

Simultaneously, the error condition is reported in several ways:

- Power supply LED
- Audible alarm
- Enclosure fault LEDs
- EMU alphanumeric display

Should both blowers fail, the power supplies automatically shut down in seven minutes.



Caution


If a blower is removed, the shelf powers off after about seven minutes. Do not remove a blower before you have a replacement available.

Note

The failure of a power supply +12.5 VDC circuit disables the associated blower.

Power Supply and Blower Assembly Components

The following diagram shows all of the components in the assembly.



The EMU shuts down the drive enclosure for the following reasons:

- A power supply is missing for more than seven minutes.
- Both blowers are missing or fail for more than seven minutes.
- Two of the three temperature sensor groups (EMU, disk, power supply) exceed critical levels for more than seven minutes.

Field Replaceable Units

The hot pluggable FRUs that do not require halting Fibre Channel data transfers include:

- Blowers
- Power supplies
- EMU

The following FRUs are pluggable:

- I/O module
- Transceiver
- Fiber and copper cables

Replacing these FRUs always interrupts data transfers on the Fibre Channel loop.

Replacement of FRUs and CRUs is covered in the troubleshooting lab.

Note

Disk drives and disk drive blanks are the only CRUs for the Enterprise Virtual Array. They are hot pluggable.

Learning Check

1. List three elements of the disk drive enclosure that are accessible from the rear of the enclosure.
.....
.....
.....
2. Discuss the drive bay numbering scheme starting at number one.
.....
.....
3. At which end of the disk drive enclosure does the EMU reside?
.....
4. List the EMU status LED displays.
.....
.....
5. How many characters can display in the alphanumeric display of the EMU?
 - a. One
 - b. Four
 - c. Two
 - d. Seven
6. What does an **Er** display represent on the EMU?
.....
7. At which location in the disk drive enclosure is I/O module B installed?
.....
8. How many status LEDs are there on the I/O modules?
 - a. Three
 - b. Two
 - c. One
 - d. Four

9. Which EMU display group has precedence over all other display groups and is only active when an error occurs?

.....

10. What can you check first to determine whether the enclosure elements are functioning properly?

.....

Overview

This module describes the physical, interconnect, and management features of the Fibre Channel loop switch. The FC loop switches, available only with Enterprise Virtual Array switched hardware configurations, are the key interface between the controllers and the disk drive enclosures. They act as a central point of interconnection and establish a fault-tolerant physical loop topology. This module describes switch features and functions, and defines switch components that allow serviceability.

Objectives

After completing this module, you should be able to:


- Describe the general features of the FC loop switch.
- Identify the components of the front and rear layouts of the FC loop switch.
- Recognize the major functional elements of the FC loop switch.
- Recognize the LED display of the FC loop switch.
- Identify the field replaceable units of the FC loop switch.

General Description



Fibre Channel Loop Switch — Rear View

Enterprise Virtual Array switched hardware implements FC loop switches in the original 42U rack and the new 41U rack. There are four switches in a rack, a pair of switches between the top controller and the M5214 enclosure above it, and a pair of switches between the bottom controller and the M5214 enclosure below it. You can upgrade nonswitched racks using switch upgrade kits. The presence of the switches eliminates the need for an expansion panel in the rack.




FC Loop Switch Pairs — Rear View

Switch Features

The major features of the FC loop switch are the following:

- 2.125Gb/s operating speed
- Half-width, 1U size
- Each bezel covers two FC loop switches in a space of 1U


Front Bezel



Other features include:

- Twelve ports
- System and port status LED indicators
- Universal power supply that operates between 100 to 250 VAC (or 50Hz to 60Hz)
- Small form-factor pluggable (SFP) transceivers

Rear View




The loop switch technology enables a router to properly use the switch core in sending data from one port to another. This process allows multiple, simultaneous conversations between ports—effectively multiplying bandwidth.

Switch Cabling


The simplified cabling in the rack improves serviceability by:

- Simplifying addition or removal of shelves and expansion racks.
- Allowing detection and repair of defective/intermittent components.

The photograph shows the four loop switches installed above and below the controllers, with management guides and cables in place.




As shown in the following diagram, cabling goes from controller device ports to the loop switch ports, and in turn to disk drive enclosure I/O modules.



Switch Port Numbering

The port numbering is important when discussing connections from the controllers and to the disk enclosures, or for referencing connections when upgrading nonswitched configurations to switched.



FC Loop Switch — Rear View with Port Numbering

Note

Connections to ports are dependent on switch placement within the rack, that is, whether connecting to loop 1 or 2.

Functional Overview of the Switch

Each FC loop switch acts as a central point of interconnection and establishes a fault-tolerant physical loop topology. The following describes the communication flow of the switch in the Enterprise Virtual Array:

- Each of the four switches handles a loop (A or B) within loop pair 1 or loop pair 2.
- Communication begins from a controller to the switch, which in turn communicates to the target transceiver on the appropriate I/O module A or B port.
- Communication returns through the same transceiver, closing the loop. The M5214 enclosures are not connected serially from top to bottom or bottom to top using the bottom transceivers.

Any loop node that is missing or is inoperative is detected by the FC loop switch. The data is automatically routed to the next operational port and attached node in the loop. LED indicators provide status information to service personnel to indicate whether the port is active or bypassed.



Important

Loop switches are not recognized by the storage management appliance; therefore, no specific errors are trapped for loop switches. Loop switch problems may cause problems to be reported by controller or disk enclosures that connect with the loop switches. So, in general, do not assume that because a specific component is reporting errors, that component is bad. You must look at all connections to and from the suspect component.

Using FC loop switches also eliminates the need for expansion panels in the racks. This is detailed in Module 5.

Power-On Self Test

When you power on the Fibre Channel switch, the switch performs a power-on self test (POST). A POST verifies that the switch is functioning properly. During a POST, all of the LEDs turn on for approximately two seconds. Then, all of the LEDs, except the power LED, turn off.



Important


The loop switch acts as a hub during loop initialization, and so it may propagate errors.

If the port bypass LEDs are blinking at a constant rate, and the POST fault LED is on, the switch detected a fault during the POST. When a POST detects a fault, customers should contact an HP Authorized Service Provider.

Switch LEDs


The FC loop switch uses two sets of LEDs to indicate switch and port status:

- System LEDs — Four separate LEDs that indicate the status of the switch.
- Port LEDs — Two LEDs per switch port that indicate status of that specific port.



System LEDs

Four LEDs indicate the status of the switch, independent of the port LEDs.




The following indicates the status of the system LEDs for a correctly functioning switch:

- Power — Green LED should be **ON**.
- Loop operational — Green LED should be **ON**.
- POST fault — Amber LED should be **OFF**.
- Over temp — Amber LED should be **OFF**.

Port LEDs

Two port LEDs indicate the status of the port.



The following indicates the status of the port LEDs for a switch:

- Both LEDs **OFF** — Normal status of operation for ports in which SFPs are not installed.
- SFP status **ON** and port bypassed **OFF** — Normal operation. Port and device are fully operational.
- Both LEDs **ON** — The port is nonoperational due to loss of signal, poor signal integrity, or the attached node is sending LIP(F8,xx).
- SFP status **OFF** and port bypassed **ON** — Tx Fault. The port is nonoperational due to an SFP transmitter fault, improperly seated SFP, or another failed device.

Field Replaceable Units

The FC loop switch has two FRUs:

- Loop switch itself (hot-pluggable)
- Loop switch transceiver

The Service Manual contains detailed procedures for doing the following:

- Removing a failed switch
- Installing a replacement switch
- Replacing a FC loop switch transceiver

Learning Check

1. The presence of FC loop switches eliminates the need for what component within the rack?
.....
2. List three features of the FC loop switches.
.....
.....
.....
3. Describe the cabling of the switches within the rack.
.....
.....
.....
4. Which LEDs do you check to determine whether the FC loop switch is functioning properly?
.....
5. Name the FRUs of the FC loop switch.
.....
.....
.....

rack, cabling, configuration, and initial setup

module 5

Overview

This module describes rack models, the power cabling that is common across all rack models, configurations, and initial setup of the Enterprise Virtual Array. This module provides the rules and diagrams for attaching Fibre Channel cables between the disk drive enclosures and the controller pair of the supported configurations. Enterprise Virtual Array 5000 cabling is described for nonswitched configurations with expansion panels (V1 hardware plus Model 2C2D) and those with FC loop switches (V2 hardware). Enterprise Virtual Array 3000 cabling is described for its nonswitched configurations.

Note

This material describes the eva5000 unless otherwise noted in the topic. Many of the general features of the eva5000, such as power cabling, are applicable to the eva3000. Where eva3000 features are different, they are noted at the end of the configuration section.

The initial setup of the Enterprise Virtual Array hardware is also covered in detail, including a flowchart for getting started; information you must gather; and procedures for inspecting the storage system, applying power to the system, and attaching the controllers to the fabric. Information on configuring the storage system is also presented.

Objectives

After completing this module, you should be able to:

- Identify the rack models and characteristics of each.
- Describe the power components and cabling in the rack.
- Identify the rules for cabling all configurations, including expansion racks.
- Describe the characteristics of switched and nonswitched configurations.
- Identify the key components and processes for cabling the 2-, 6-, 12-, and 8-drive enclosure configurations.
- Identify the key components and processes for cabling the expansion options.
- Identify the steps used to set up a storage system.
- Identify the process used to obtain a license key.

Rack Models

- Graphite 41U — Contain FC loop switches (switched)
- Opal 42U — May or may not (nonswitched) contain FC loop switches
 - Version 1
 - ◆ Uses Modula-style horizontal rails.
 - ◆ Rails have studs that extend through the holes in the rear vertical rails.
 - Version 2
 - ◆ Uses Millenium-style horizontal rails.
 - ◆ Rails connect between two inner vertical members.

The rack has four feet and four casters. Raising the adjustable feet places the rack weight on the casters so you can easily move the rack. Lowering the feet places the rack weight on the feet and prevents you from moving the rack. Procedures are listed in the Service Manual for this as well as the following:

- Removing and installing rack side panels
- Installing rails

Physical specifications, shipping dimensions, and AC power specifications of the racks are listed in Appendix A of the Service Manual.

Switched and Unswitched Models


The following table summarizes the current switched and unswitched models.

Opal Models (Nonswitched)	Opal Models (Switched)	Graphite Models (Switched)
		2C2D-B (nonswitched)
2C6D	2C6D-A	2C6D-B
2C12D	2C12D-A	2C12D-B
8C8D	8C8D (nonswitched)	8C8D-B (nonswitched)
0C6D ¹	0C6D-A ²	0C6D-B ³
0C12D ¹	0C12D-A ²	0C12D-B ³
¹ Wired to connect to nonswitched models.		
² Wired to connect to switched (-A) models.		
³ Wired to connect to switched (-B) models.		

Opal 42U Rack

The following are characteristics of the opal 42U rack:


- Can be nonswitched (original V1 hardware) or switched.
- The nonswitched hardware has 3U spare space.
- The switched hardware has 1U spare space.



Graphite 41U Rack

The following are characteristics of the graphite 41U rack:

- Racks are all switched except models 2C2D and 8C8D.
- Racks have no spare space.



Rack Power Distribution

AC power is distributed to the rack through a dual zero-U power distribution unit (PDU) assembly mounted at the bottom rear of the rack. Each PDU is connected to a separate circuit-breaker-protected, 30A AC site power source. The PDUs in the 41U rack are mounted like the 42U opal rack, but the PDUs are mounted to the side rails of the rack. The 41U graphite rack back door is full length and opens to allow access to the PDUs. The 42U opal rack door is divided into two sections, with the bottom section used for access to the two PDUs attached to the door itself.



Opal



Graphite

This configuration provides complete power redundancy and eliminates all single points of failure for AC and DC power distribution. Requirements for the fully redundant power configuration of an Enterprise Virtual Array storage system rack are as follows:


- Two separate circuit-breaker-protected, 30 Amp site power sources with a compatible wall receptacle
- One dual PDU assembly (each PDU connects to a different wall receptacle)
- Eight PDMs (AC power strips) per rack
 - Four PDMs mount vertically on each side of the rack.
 - Each set of four PDMs connects to a different PDU.
 - The drive enclosure power supplies on the left (PS 1) connect to the PDMs on the left with a gray, 26-inch (66 cm) power cord.
 - The drive enclosure power supplies on the right (PS 2) connect to the PDMs on the right with a black, 26-inch (66 cm) power cord.
 - The upper controller connects to a PDM on the left with gray, 60-inch (152 mm) power cord.
 - The lower controller connects to a PDM on the right with a black, 26-inch (66 cm) power cord.

Dual Zero-U PDU Assembly

Each Enterprise Virtual Array storage system rack has a 50Hz or 60Hz dual zero-U PDU mounted at the bottom rear of the rack. This assembly contains two PDUs: PDU 1 and PDU 2. Each PDU has:

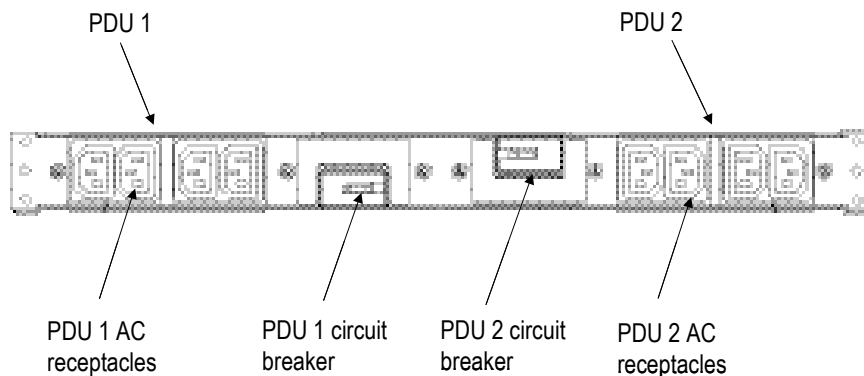
- A 250 VAC, 30A circuit breaker.
- Four IEC 320-C13 AC output power receptacles.

The circuit breakers and AC receptacles are accessible when the PDU assembly is in the upright position. You can fold down the PDU by 1) unscrewing the spring-loaded screws, and 2) folding the PDU away from the rack towards the floor.



Each of the two PDU power cables has an AC power-source-specific connector. The circuit-breaker-controlled PDU outputs are routed to a group of four AC receptacles.


The voltages are then routed to PDMs (AC power strips) mounted on the two vertical rails in the rear of the rack.




Wall Receptacles

Type of wall power receptacles include the following:

- North America — Three-wire NEMA No. L6-30R, 30 amp (208V to 240V 60Hz 30A)
- Europe — Three-wire, two-pole IEC 309, 30 amp, (220V to 240V 50Hz 32A), single phase



NEMA




IEC

Power Distribution Modules

The PDMs distribute the AC power from the PDUs to the enclosures. Each of the thermally protected PDMs connects to one of the four AC receptacles on a PDU. Each PDM has six AC receptacles and a thermal circuit breaker.

The PDMs on the right of the rack connect to the right PDU. The PDMs on the left of the rack connect to the left PDU. Individual power cords connect the controller and drive enclosure power supplies to these modules.



Eight PDMs are mounted in the rear of each rack.

- The four mounted on the left vertical rail connect to PDU 1.
- The four mounted on the right vertical rail connected to PDU 2.

Each controller pair and drive enclosure has two power sources. Even when a single PDU fails, the system remains operational.

PDM and PDU Connections


The power distribution in an Enterprise Virtual Array rack is the same for all variants. The site AC input voltage is routed to the dual PDU assembly mounted in the rack lower rear. Each PDU distributes AC to a maximum of four PDMs mounted on the left and right vertical rails.

- PDMs 1 through 4 connect to receptacles A through D on PDU 1. **Grey** power cords connect these PDMs to the number 1 drive enclosure power supplies and to the **upper** controller enclosure.
- PDMs 5 through 8 connect to receptacles A through D on PDU 2. **Black** power cords connect these PDMs to the number 2 drive enclosure power supplies and to the **lower** controller enclosure.



Important

The locations of the PDUs and the PDMs are the same in all Enterprise Virtual Array racks.



Rack AC Power Distribution

PDU and PDM Connectivity to Enclosures

The following table shows how the PDUs and PDMs are connected to the drive enclosures. This table is also found in the Service Manual.

PDU	PDM	Enclosures
1	1 (bottom, left side)	Power supply 1 on drive enclosures 01, 02, 03
1	2 (next up, left side)	Power supply 1 on drive enclosures 04, 05, 06, and bottom controller left power supply (on hardware with dual power supply controllers)
1	3 (next up, left side)	Power supply 1 on drive enclosures 08, 09, 10, and top controller left power supply (on hardware with dual power supply controllers)
1	4 (top, left side)	Power supply 1 on drive enclosures 11, 12, 13
2	5 (bottom, right side)	Power supply 2 on drive enclosures 01, 02, 03
2	6 (next up, right side)	Power supply 2 on drive enclosures 04, 05, 06 and bottom controller
2	7 (next up, right side)	Power supply 2 on drive enclosures 08, 09, 10 and top controller
2	8 (top, right side)	Power supply 2 on drive enclosures 11, 12, 13

PDU Failures

PDU 1 connects to AC power distribution source 1. A failure with PDU 1 would:

- Remove power from PDMs 1, 2, 3, and 4, which would eliminate power from:
 - PS1 in the drive enclosures.
 - The upper controller power supply.

PDU 2 connects to AC power distribution source 2. A failure with PDU 2 would:

- Remove power from PDMs 5, 6, 7, and 8, which would eliminate power from:
 - PS 2 in the drive enclosures.
 - The lower controller power supply.

Power Component Removal and Replacement

The Service Manual provides procedures for power component removal and replacement. The following summarizes what is available for the different assemblies:

- PDU assemblies
 - Different assemblies in 41U and 42U racks.
 - Procedures are available for replacing PDUs in each.
- PDM assemblies
 - Same PDM used in the 41U and 42U racks.
 - Mounting is different.
 - Procedures are available for replacing PDMs in each.

Configuration Overview

The two primary hardware classifications for the Enterprise Virtual Array 5000 configurations are nonswitched and switched. These are implemented in standard models; however, provisions are made to allow expansion into expansion models or other configurations using expansion components.

Standard Models

The standard model classifications are described by the following:

- Nonswitched models do not contain FC loop switches, but may or may not contain expansion panels.
 - Models that support expansion panels include the 2C6D and 2C12D.
 - Other models include the 2C2D and 8C8D.
- Switched models contain FC loop switches.
 - Models that support FC loop switches include the 2C6D and 2C12D.

Expansion Rack Models

The Enterprise Virtual Array has two standard expansion racks used to expand 2C12D configurations. They include the following:

- Model 0C6D
 - Provides up to 17.5TB of storage when combined with one master rack (2C12D)
 - Used with 2C12D + 0C6D configuration
- Model 0C12D
 - Provides up to 35TB of storage when combined with two master racks (2C12D)
 - Used with 2 X 2C12D + 0C12D configuration
- Either of the models are usable with switched or nonswitched versions
 - Nonswitched models use expansion panels
 - Switched models use FC loop switches

Note

The expansion racks supply 12 disks per enclosure, not 14. So, the 0C6D allows you to add 72 disks to one master rack (240 total disks), and the 0C12D allows you to add 144 disks in combinations with two master racks (480 total disks).

Expansion Components and Options

The Enterprise Virtual Array 5000 has several expansion components that allow several expansion options.

- Expansion components
 - M5214 drive enclosure kits (nonswitched or switched)
 - Expansion panel kits (nonswitched only)
 - FC loop switch upgrade kit
- Numerous configuration options
 - Drive enclosures allow expansion of 2C2D, 2C6D, or 0C6D models—for example, 2C2D + 2D, 2C6D + 2D, and 2C6D + 4D.
 - Expansion panel kits allow expansion from 2C6D to 2C12D and from 0C6D to 0C12D.
 - FC loop switch upgrade kits allow you to upgrade a nonswitched configuration to a switched configuration.

Note

Procedures for performing the expansion options are detailed in the Service Manual.

Nonswitched Configurations

The nonswitched configurations are divided between those that may or may not use expansion panels. The panels allow the extension of FC loops in a master rack to an expansion rack or racks by way of an expansion panel or panels in the expansion rack.


Configurations **with** expansion panels include the following:

- Model 2C12D master rack contains one panel.
This allows connection to an expansion rack to create a 2C12D + 0C6D configuration.
- Model 0C6D expansion rack contains one panel.
The panel in this expansion rack allows a single 2C12D to expand to a 2C12D + 0C6D configuration.
- Model 0C12D expansion rack contains two panels.
These panels allow two 2C12Ds to expand to a three-rack configuration, that is, 2 X 2C12D + 0C12D.

Configurations **without** expansion panels include the following:

- Model 2C2D
If these models expand, it will be through loop switches.
- Model 8C8D
No expansion panels or loop switches are required for this configuration.

Expansion Panel Description




The expansion panel is a thin panel mounted on the rear of the rack. It contains eight ports. Each port allows the controller pair to connect to the I/O ports on each drive enclosure. The eight ports on the expansion panel correspond to the four Fibre Channel loops contained in a rack: 1A, 1B, 2A, and 2B. Each loop has an IN and an OUT port. Fibre Channel cables returning from the drive enclosures enter the IN port. Fibre Channel cables going to the drive enclosure enter the OUT port.



Expansion Panel


Expansion Panel Location

The master rack contains a single expansion panel enabling its controllers to connect to expansion rack FC cables. The 0C6D expansion rack contains one expansion panel, and the 0C12D expansion rack contains two expansion panels.



Fibre Channel Cable Configurations

The Enterprise Virtual Array has many possible configurations and cable routing schemes. To ensure proper operation, HP recommends identifying the transceivers using the numbering system shown in the following figure.



Important

Enterprise Virtual Array configurations with FC loop switches use only the top transceivers (04 and 02). A black rubber plug is installed in the unused SFP connectors. Command View EVA marks the unused SFP transceiver as FAILED in the I/O Module page of the GUI.

The HP supported configuration is a dual-loop with 6 or 12 drive enclosures per dual loop, expandable. This configuration eliminates single points of failure. Limiting the total number of enclosures reduces the transmission delays inherent in longer loops.


With expansion racks, drive enclosure kits, and expansion panel options, configurations support 6, 8, 10, 12, or 18 enclosures per dual loop.



Important

Keeping the loops completely separate ensures that the failure of a single component (controller, I/O module, transceiver, or cable) does not stop data transmissions to all the active drives. Therefore, only connect the loop A cables to an I/O module A and the loop B cables to an I/O module B.

Cabling Rules for Nonswitched Configurations



Follow these rules when configuring the nonswitched Enterprise Virtual Array:

- Loop pair 1 enclosures are below the bottom controller pair.
- Loop pair 2 enclosures are above the top controller pair.
- The A loops are connected from the topmost drive enclosure downward.
- The B loops are connected from the bottommost drive enclosure upward.

Also note:

- The A loop attachments are on the right side of the disk drive enclosures (looking from the back).
- The B loop attachments are on the left side of the disk drive enclosures (looking from the back).
- If an expansion panel is present, the loops should be daisy-chained through the panel to make future expansion easier.

Cabling Configurations for Nonswitched Models

Enterprise Virtual Array nonswitched configurations that are supported include:

- Model 2C6D (single rack)
- Model 2C12D (single rack)
- Model 8C8D (single rack)
- Expanded configuration 2C12D + 0C6D (one Master rack plus one Expansion rack)
- Expanded configuration 2 X 2C12D + 0C12D (two Master racks plus one Expansion rack)


Cabling Configuration for Nonswitched Model 2C6D

The nonswitched Model 2C6D contains two controllers and six drive enclosures. A Fibre Channel loop is formed when the FC drive enclosures and the HSV controller pair are connected by FC cables.

The nonswitched Model 2C6D uses an expansion panel to complete the loop configuration. Each dual-signal FC cable is run from one controller, through the drive stack, to the other controller, and back again.

Cabling Diagram

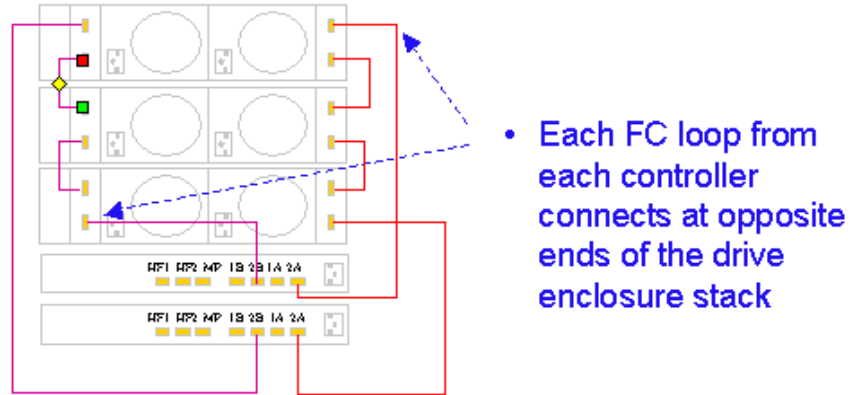
The following diagram shows the four redundant 2Gb/s FC-AL device loops for the Enterprise Virtual Array nonswitched Model 2C6D. Although it is not shown here, an expansion panel can be implemented with this configuration for future use with expansion racks.



Loop 1B is at the bottom because it is part of loop pair 1. The loop is connected from the bottommost enclosure upward because it is a B loop. Each loop of the redundant pairs (1A-1B and 2A-2B) is also connected at opposite ends of the stack for each controller.

Fibre Channel Cable Routing

The following illustration shows the FC cabling for loop pairs 2A and 2B.



Each FC cable is a dual-signal cable. Notice that each loop runs from one controller, through the drive stack, to the other controller, and back again. As explained in Cabling Rules, notice that loop 2B is at the bottom because it is part of loop pair 2. The loop is connected from the bottommost enclosure upward because it is a B loop.


Each loop of the redundant pair is connected at opposite ends of the stack for each controller. This provides the additional redundancy. Therefore, if a loop is interrupted, each controller can (potentially) access the stack from both ends. Each drive enclosure can reroute a loop, or shorten the loop, should the connection to the next enclosure fail.

For example, should one FC cable fail in the middle of the loop, each enclosure can reroute the loop. One loop essentially becomes two smaller loops, with each controller at one end of one loop, and the other controller at the end of the second loop.

Cabling Configuration for Nonswitched Model 2C12D

The Enterprise Virtual Array nonswitched Model 2C12D contains two controllers and 12 drive enclosures. A Fibre Channel loop is formed when the FC drive enclosures and the HSV controller pair are connected by FC cables.

The following diagram shows the four redundant 2Gb/s FC-AL device loops for the Enterprise Virtual Array nonswitched Model 2C12D. This configuration comes with an expansion panel for potential use with expansion racks. Jumper cables can be used when it is not connected to an expansion rack configuration. The loops are arranged in redundant pairs—1A-1B and 2A-2B.




Note

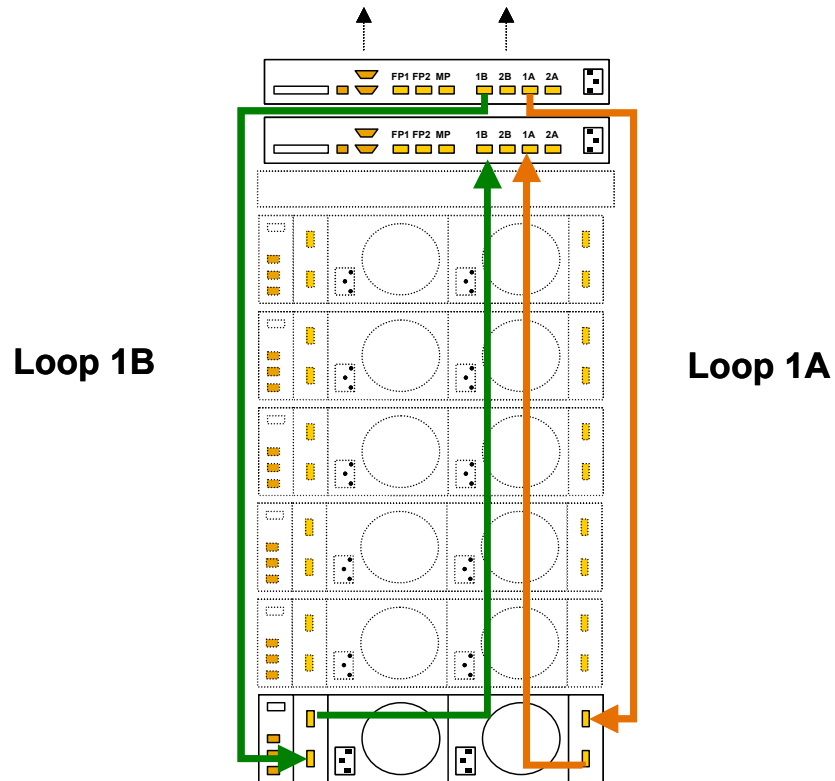
When no expansion rack is connected to a 2C12D configuration, four 0.5m jumper cables are used to cable the loops within the master rack expansion panel. If an expansion rack is used later, it contains an expansion panel to cross-couple the cables from the master rack to the expansion rack.

Cabling Configuration for Nonswitched Model 2C2D

The Enterprise Virtual Array nonswitched Model 2C2D contains two controllers and two drive enclosures. A Fibre Channel loop is formed when the FC drive enclosures and the HSV controller pair are connected by FC cables. Cabling is similar to the Model 2C6D.



The following diagram shows the loop 1 cable routing for the Enterprise Virtual Array nonswitched Model 2C2D. This configuration does not use an expansion panel, but can later implement FC loop switches for use with expansion racks.

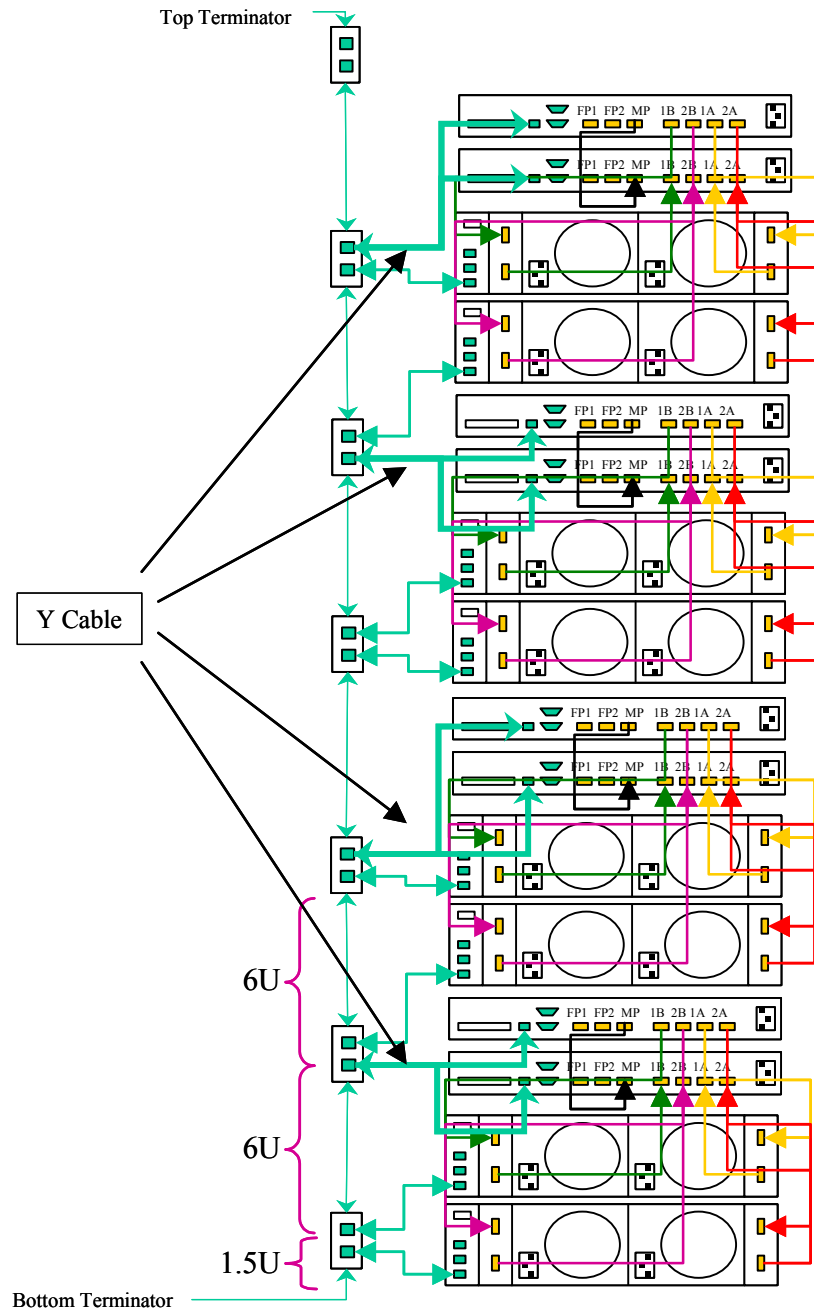


Cabling Configuration for Nonswitched Model 8C8D

The Enterprise Virtual Array Model 8C8D contains four storage systems, each containing two controllers and two drive enclosures. No expansion configurations are available, nor are any FC loop switches involved in the cabling.


Cabling Diagram

The following diagram shows the cabling of each storage system.



Fibre Channel Cable Routing

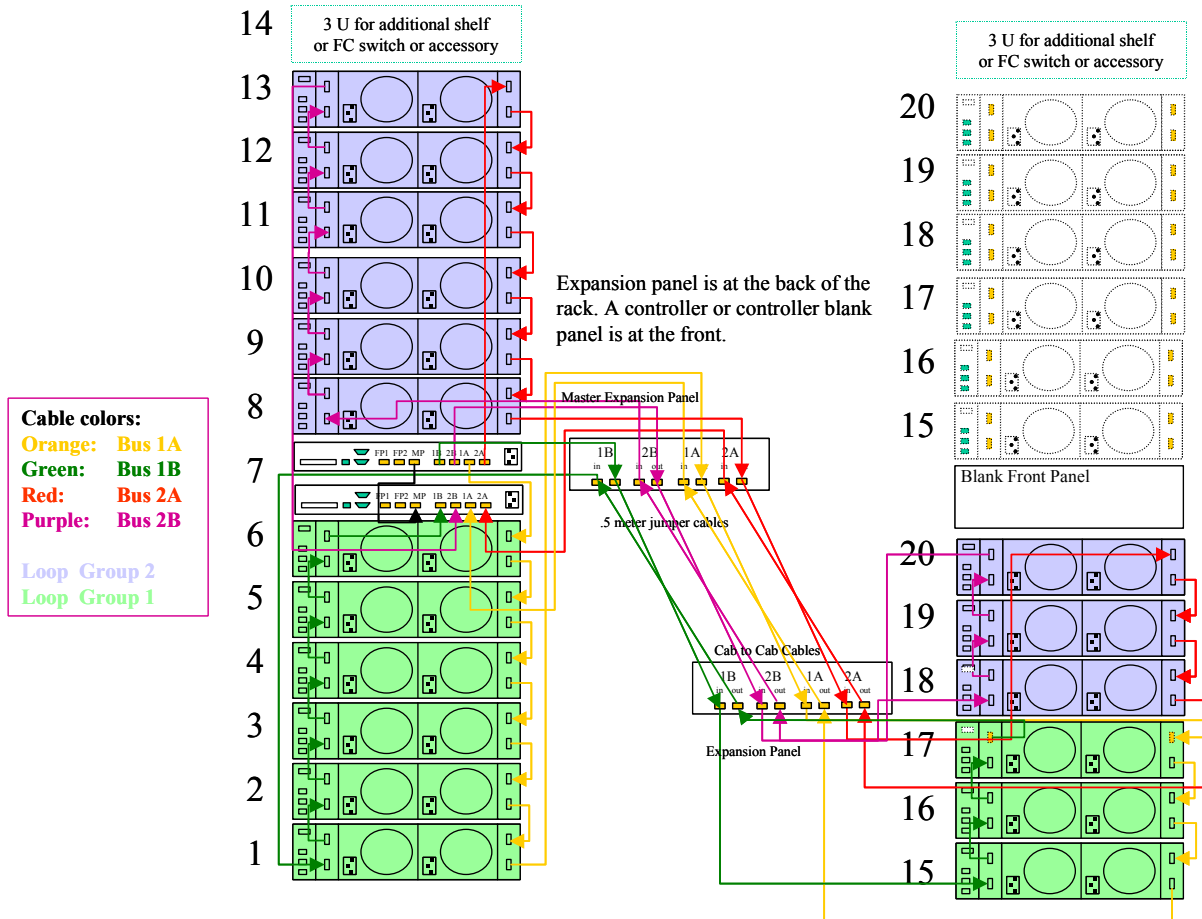
The following figure shows the cabling for one Enterprise Virtual Array Model 8C8D storage system, that is, the controller-to-enclosure cabling for a storage system with two disk enclosures.



This configuration has a total of 16 Fibre Channel loops. Because the controller pair has only two enclosures, the cabling is a little different. Rather than each loop beginning at the master controller, going through both enclosures, and back to the slave, each loop goes through one enclosure. The same level of redundancy and high availability is achieved.

Cabling Configuration for Nonswitched Model 2C12D + 0C6D

The following diagram shows the four redundant 2Gb/s FC-AL device loops for the Enterprise Virtual Array nonswitched Model 2C12D with a six-drive enclosure expansion rack. The nonswitched Model 0C6D rack contains a single expansion panel and connects to the 2C12D rack with an enclosure address bus (EAB) and Fibre Channel cables.



When an expansion rack is connected to a 2C12D configuration, the four jumper cables on the master rack expansion panel are removed, and eight FC interconnect cables are wired between the two racks. This alleviates an otherwise confusing and difficult rewiring of the 2C12D configuration to add the other six disk drive enclosures into the four FC device loops.

Note

The enclosures are numbered 15 to 20 for the 0C6D expansion rack.


Detailed diagrams for each loop are included in the *Enterprise Virtual Array Hardware Configuration Reference Guide*.

EAB Configuration for Nonswitched Model 2C12D + 0C6D

Each rack contains EAB junction boxes at 6U increments in the left-rear rail of the rack.

- The 0C6D rack contains six EAB junction boxes; however, the M5214 drive enclosures use only three of the boxes.
- The 2C12D rack contains seven EAB junction boxes.
- The M5214 drive enclosures and controller pair in the 2C12D rack use all seven EAB junction boxes.

As shown in the following figure, the two racks are connected by one EAB cable, which connects to the **bottom** EAB junction box on each rack. The EAB cable between the racks is polarized. The **P1** end of the cable connects to the **bottom** EAB junction box on the 2C12D rack. The **P2** end of the cable connects to the **bottom** EAB junction box on the 0C6D rack.

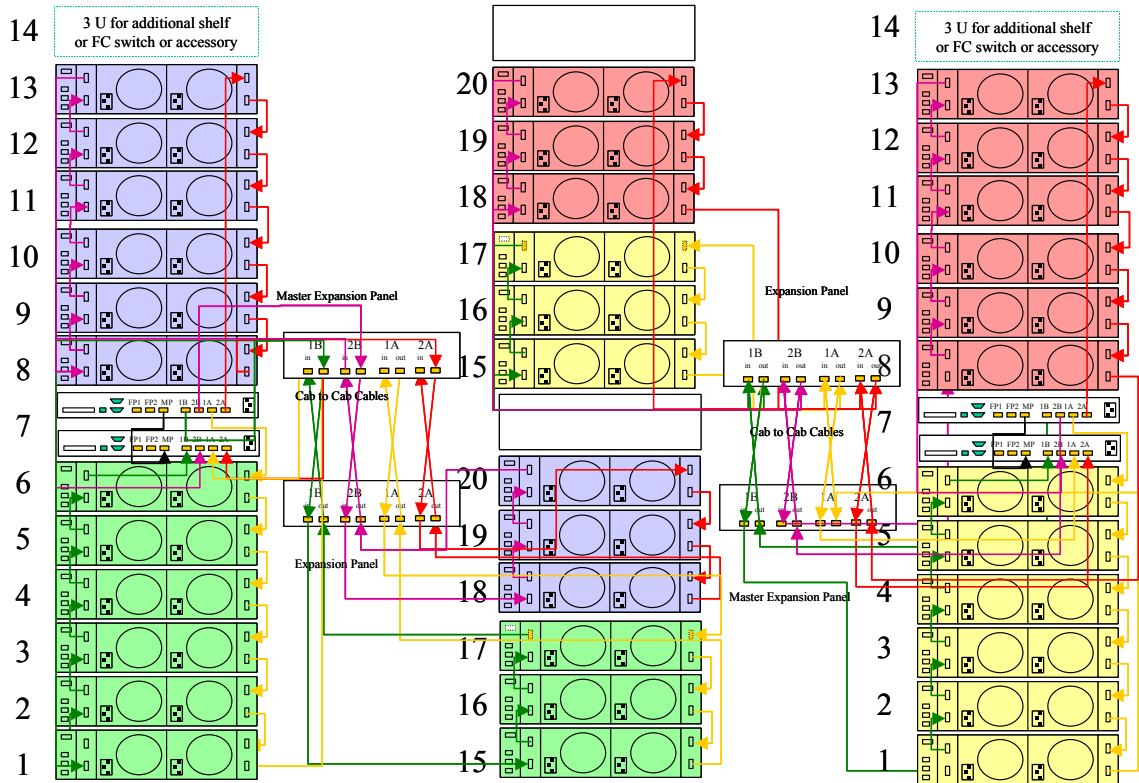


Note

For complete procedures and cabling diagrams for expanding racks, see the Hardware Configuration Guide and the Service Manual.

Cabling Configuration for Nonswitched Model 2 X 2C12D + 0C12D

The following diagram shows the four redundant 2Gb/s FC-AL device loops for two Enterprise Virtual Array nonswitched Model 2C12Ds with a 12-drive enclosure expansion rack. The nonswitched Model 0C12D rack contains two expansion panels and connects to the 2C12D rack with an EAB and Fibre Channel cables.



For detailed diagrams for each loop, see the Hardware Configuration Guide.


EAB Configuration for Nonswitched Model 2 X 2C12D + 0C12D

Each rack contains EAB junction boxes at 6U increments in the left-rear rail of the rack.

- The 0C12D rack contains six EAB junction boxes. The M5214 drive enclosures in the 0C12D rack use all six boxes.
- The 2C12D rack contains seven EAB junction boxes. The M5214 drive enclosures and controller pair in the 2C12D rack use all seven boxes.

The three racks are connected by two EAB cables. Each master rack is connected to the expansion rack by one EAB cable. As shown in the following figure, one master rack runs an EAB cable to the **bottom** EAB junction box on the expansion rack; the other master rack runs an EAB cable to the **fourth** EAB junction box (from the bottom) on the expansion rack.

The EAB cables are polarized. As shown in the figure, the **P1** end of the cable connects to the **bottom** EAB junction box on the 2C12D rack. The **P2** end of the cable connects to the EAB junction box on the expansion rack.



- ❶ Terminator
- ❷ Enclosure Address Bus Junction Box
- ❸ 5-meter Rack-to-Rack Enclosure Address Bus Cable

Note

For complete procedures and cabling diagrams for expanding racks, see the Hardware Configuration Guide and the Service Manual.

Switched Configurations

The switched configurations (those that contain FC loop switches) include:

- Model 2C6D (single rack)
- Model 2C12D (single rack)
 - Direct connection to an expansion rack 0C6D creates the expanded configuration 2C12D + 0C6D (one master rack plus one expansion rack).
 - Connection of two 2C12Ds to an expansion rack 0C12D allows expansion to a three-rack configuration 2 X 2C12D + 0C12D.


The expanded configurations use direct FC switch connections in the master and expansion racks. You can also expand a rack by adding drive enclosures. See the Hardware Configuration Guide for these procedures.

FC Loop Switches


Enterprise Virtual Array switched Model 2C6D and 2C12D configurations use four FC loop switches in the master rack to connect controllers with disk enclosures, allowing a fault-tolerant physical loop topology. The switches also function as a central point of interconnection by allowing the extension of FC loops to an expansion rack through switch-ready Models 0C6D and 0C12D expansion racks. All single-rack and expansion configurations are discussed in later topics.

Switch Description

The FC loop switches are mounted in pairs in 1U enclosures above and below the controller pair. The use of the 12 ports varies according to placement of the switch in the loop configuration.



Loop Switch — Rear View




Loop Switch — Rear View with Port Numbering

Switch Location

Fibre Channel switch and loop location, viewed from the rear of the rack:

- Loop 1A and switch are located at the lower-right side of the rack.
- Loop 1B and switch are located at the lower-left side of the rack.
- Loop 2A and switch are located at the upper-right side of the rack.
- Loop 2B and switch are located at the upper-left side of the rack.

These loops, but no cable connections, are shown in this picture.



Loop Switch and Loop Location

FC Loop Switch Upgrade

You can upgrade an Enterprise Virtual Array nonswitched configuration containing an expansion panel to a switched hardware configuration containing FC loop switches. This provides a high-performance, fault-tolerant configuration for the rack and allows you to add the switched expansion racks through switch cabling rather than expansion panels.

Three key operations are involved in the upgrade:

- Removing expansion panels
- Installing FC loop switches
- Connecting FC cabling to complete the switched configuration


Detailed procedures for performing this upgrade are included in the Service Manual.

Fibre Channel Cable Configurations

Switched configurations use only the top transceivers (04 and 02) of the disk drive enclosure. A black rubber plug is installed in the unused SFP connectors.

Transceiver 04


Transceiver 02



Important

Enterprise Virtual Array configurations with FC loop switches use only the top transceivers (04 and 02). A black rubber plug is installed in the unused SFP connectors. Command View EVA marks the unused SFP transceiver as FAILED in the I/O Module page of the GUI.

Cabling Rules for Switched Configurations



Follow these rules when configuring the switched Enterprise Virtual Array:

- Loop pair 1 enclosures are below the bottom FC loop switch pair.
- Loop pair 2 enclosures are above the top FC loop switch pair.
- Loop 1A cables connect in order from the bottommost drive enclosure, up to the left-most switch port in the bottom row of switch ports, and proceed from left to right.

- Loop 2A cables connect in order from the bottommost drive enclosure, down to the right-most switch port in the top row of switch ports, and proceed from right to left.
- Loop 1B cables connect in order from the bottommost drive enclosure, up to the right-most switch port in the bottom row of switch ports, and proceed from right to left.
- Loop 2B cables connect in order from the bottommost drive enclosure, down to the left-most switch port in the top row of switch ports, and proceed from left to right.

Also note:

- The A loop attachments are on the right side of the disk drive enclosures (looking from the back).
- The B loop attachments are on the left side of the disk drive enclosures (looking from the back).
- No expansion panel is present. Closed loops are formed between each switch port and enclosure.

Cabling Configurations for Switched Models

Supported Enterprise Virtual Array switched configurations include:

- Model 2C6D (single rack)
- Model 2C12D (single rack)
- Expanded configuration 2C12D + 0C6D (one Master rack plus one Expansion rack)
- Expanded configuration 2 X 2C12D + 0C12D (two Master racks plus one Expansion rack)


Cabling Configuration for Switched Model 2C6D

The Enterprise Virtual Array switched Model 2C6D contains two controllers and six drive enclosures. A Fibre Channel loop is formed when the FC drive enclosures and the HSV controller pair are connected by FC cables.

The switched Model 2C6D uses four FC loop switches to complete the loop configuration. Each dual-signal FC cable directly connects each FC drive enclosure to the associated FC switch.

Cabling Diagram


The following diagram shows the four redundant 2Gb/s FC-AL device loops for the Enterprise Virtual Array switched Model 2C6D that uses four FC loop switches.



Fibre Channel Cable Routing

The following figure shows the switch-to-controller cabling required for the switched Model 2C6D configuration. The connections use 0.5-meter cables.


- For loop 2, ports 2 and 4 of the switches are used.
- For loop 1, ports 1 and 3 of the switches are used.



Notice that if a controller fails, or an individual cable fails (for example, 2A), the other controller takes over to complete the loop. If a switch fails, access to the loop (loop 1 or loop 2) is still available through the other switch.

Switch-to-Enclosure Cabling

The following illustration shows the FC cabling between the switches and drive enclosures for loop pairs 2A and 2B. Two-meter FC cables are used.



Note that the cabling uses different ports for the switch pairs of loop 1 and loop 2. Loop 2 uses the odd-numbered ports (1, 3, 5, and so on), while loop 1 uses the even numbered ports (2, 4, 6, and so on). Also notice that the cabling to the drive enclosures only uses the top transceiver and is a closed loop. Should a loop connection fail (for example, in loop 2B), the controller provides access to loop 2A to access data.


Cabling Configuration for Switched Model 2C12D

The Enterprise Virtual Array switched Model 2C12D contains two controllers and 12 drive enclosures. A Fibre Channel loop is formed when the FC drive enclosures and the HSV controller pair are connected by FC cables.

The Model 2C12D uses four FC loop switches to complete the loop configuration. Each dual-signal FC cable directly connects each FC drive enclosure to the associated FC switch.

Cabling Diagram

The following diagram shows the four redundant 2Gb/s FC-AL device loops for the Enterprise Virtual Array switched Model 2C12D that uses four FC loop switches.




The cabling is similar to the switched Model 2C6D; however, all six ports on the top row (odd numbered ports) of the loop 2 switches and all six ports in the bottom row (even-numbered ports) of the loop 1 switches are connected to disk enclosures.

- Cabling for the loop 2 switches proceeds from the outside port (port 1 on the A side, port 11 on the B side) to the lowest (nearest) disk enclosure and works up.
- Cabling for the loop 1 switches proceeds from the outside port (port 2 on the A side, port 12 on the B side) to the highest (nearest) disk enclosure and works down.

Cabling Configuration for Switched Model 2C12D + 0C6D

The following diagram shows the four redundant 2Gb/s FC-AL device loops for the Enterprise Virtual Array switched Model 2C12D with a six-drive enclosure expansion rack. The switched Model 0C6D rack connects directly to the 2C12D rack with an EAB and Fibre Channel cables that connect directly to the switches in the master rack. For clarity, the diagram does not show switch-to-enclosure cabling.



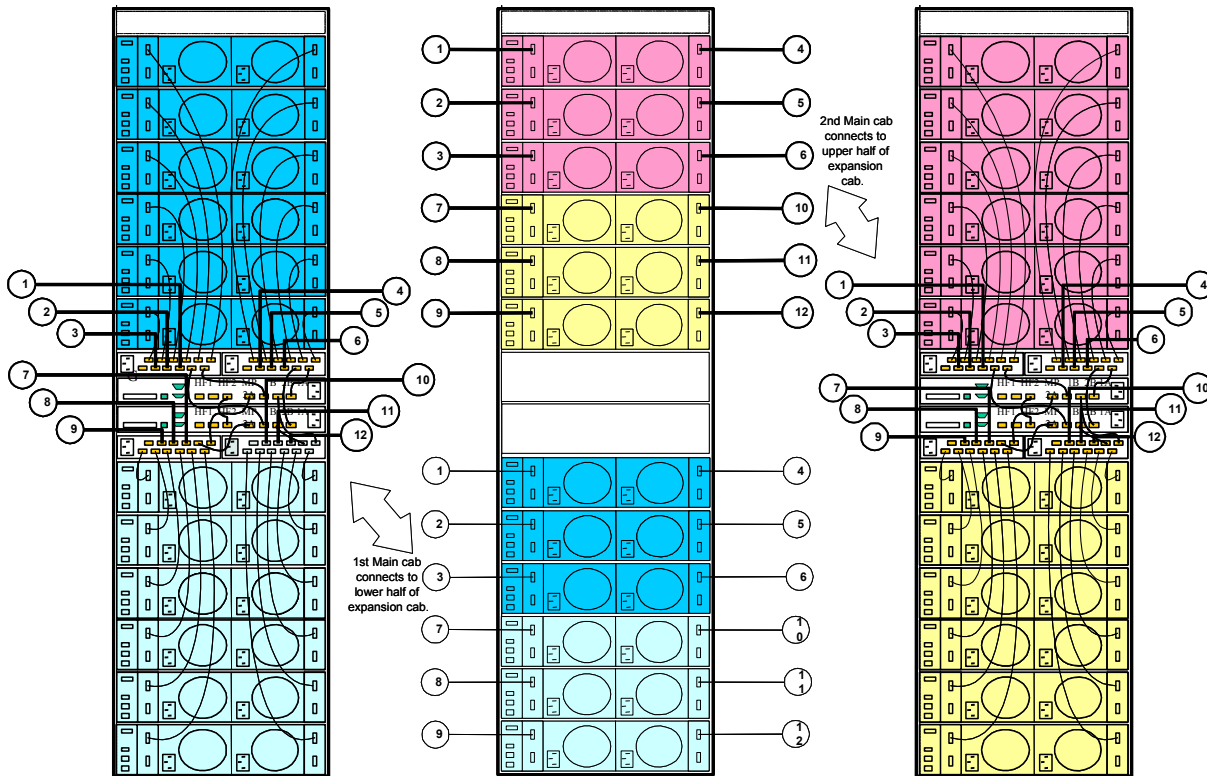
When a switched Model 0C6D expansion rack is connected to a switched Model 2C12D configuration, three of the four unused FC loop switch ports are attached between racks with 5-meter or longer cables, as shown in the preceding diagram. For loop 2, ports 6, 8, and 10 are used; for loop 1, ports 5, 7, and 9 are used.

Notice that the leftmost port of the top row for the loop 1 switches and the leftmost port of the bottom row for the loop 2 switches are not used. EAB connections are the same as those for nonswitched configurations.

Detailed diagrams for each loop are included in the Hardware Configuration Guide. Procedures for adding the expansion rack are included in the Service Manual.

Cabling Configuration for Switched Model 2 X 2C12D + 0C12D

The following diagram shows the four redundant 2Gb/s FC-AL device loops for two Enterprise Virtual Array switched Model 2C12Ds with a 12-drive enclosure expansion rack. The switched Model 0C12D rack connects to the switched Model 2C12D rack with EAB and Fibre Channel cables that connect directly to the switches in the master rack.



When a switched Model 0C12D expansion rack is connected to two switched Model 2C12D configurations, cabling is performed twice (as for the switched Model 0C6D). EAB connections are the same as those for nonswitched configurations.

Detailed diagrams for each loop are included in the Hardware Configuration Guide. Procedures for adding the expansion Rack are included in the Service Manual.

Cable Management Configurations


A configuration has either of two cable management configurations:

- With FC loop switches, the configuration uses a combination of cable containment spools and a cable management arm to organize the FC cables.
- With expansion panels, the configuration uses cable containment spools only.

Cable Management Arms

When FC loop switches are used, there are two cable management arms. All of the FC cables in the configuration pass through one of the arms. Each cable management arm can hold several radio clips, each clip holding a maximum of seven FC cables.

Each rack contains four flumes that are placed next to the cable management arms on the left and right sides of the rack. Each flume guides the FC cables from the cable containment spools to the cable management arms.



Cable Containment Spools

The cable containment spools hold the FC cables. The spools gather up extra cable length and guide the FC cables from the FC drive enclosures to the flumes located on the sides of the rack.

The number of cable containment spools varies by rack configuration; however, in general, racks with more FC drive enclosures have more spools.

The number of cable containment spools and management arms for each configuration are as follows:

- 2C12 — 12 cable containment spools and two cable management arms
- 2C6D — Six cable containment spools and two cable management arms
- 8C8D — Eight cable containment spools and no cable management arms

Additional Expansion Configurations

An authorized service provider can expand the 2C6D and 0C6D racks by adding M5214 drive enclosures, expansion panels, Fibre Channel cables, and EAB cables to the existing racks. In general, the M5214 drive enclosures are added in pairs.

The 2C6D rack can be expanded to the following:

- 2C6D + 2D (expands the 2C6D rack to 112 disks)
- 2C6D + 4D (expands the 2C6D rack to 140 disks)
- 2C6D + 6D (expands the 2C6D rack to 168 disks)

The 0C6D rack can be expanded to the 0C12D rack.

The procedures for expanding a rack are included in the Hardware Configuration Guide.

Enterprise Virtual Array 3000 Configurations


The HP StorageWorks Enterprise Virtual Array 3000 (eva3000) is the latest addition to the Enterprise Virtual Array family. The eva3000 comes in a single 2C2D configuration, but is expandable. The following topics describe features of the eva3000 and cabling configuration differences where appropriate.

Model 2C2D Features


The eva3000 comes in a Model 2C2D with the following features:

- Graphite 42U (10000) rack
- Two HSV100 controllers
- Two 14-drive bay FC drive enclosures containing up to (28) 36GB, 72GB, or 146GB drives, maximum of 4TB
- Six internal copper cables
- Three 2-port EAB junction boxes
- Eight AC strips
- Two 0U PDUs

The diagram shows the essential components.




As you can see in the following diagram, 9U of rack is used, leaving 33U of blanks available for additional enclosures or for stacking configurations.



Model 2C2D Connections

The Model 2C2D uses the same EAB as the eva5000; however, it has only three EAB junction boxes (JB). Each drive enclosure is connected to a JB, and a Y cable is used for the controller pair at the top of JB 2.


The following diagram shows the connections.



Cabling for Loop 1

The eva3000 has only a loop 1. A copper cable is used to connect controller device ports 1A and 1B to the drive enclosures. The cabling follows the nonswitched pattern of the Enterprise Virtual Array 2C2D, 2C6D, or 2C12D.


The following two diagrams show the loop 1 cabling.



Additional Configurations

You can expand the configuration by adding drive enclosures to the rack. Options include the following:


- 2C3D — Up to 6.1TB using (42) 146GB drives
- 2C4D — Up to 8.2TB using (56) 146GB drives
- These options are added by Field Service only.



Model 2C4D Connections

The expanded Model 2C4D uses the same EAB as the Model 2C2D; however, all three EAB junction boxes are used. Each drive enclosure is connected to a JB, and a Y cable is used for the controller pair at the top of JB 2. The top two enclosures connect to JB 3.


The following diagram shows the connections.



Model 2C4D Cabling for Loop 1

The expanded configuration 2C4D cabling is very different from any eva5000 cabling. Instead of enclosure number 2 returning to the top controller, it connects to enclosure 3. Enclosure 4 returns to the top controller.

The following two diagrams show the loop 1 cabling.



Storage System Setup

The initial setup of the Enterprise Virtual Array recommends that you reference the Getting Started flowchart to review the process. The following steps in the process are highlighted in this module:

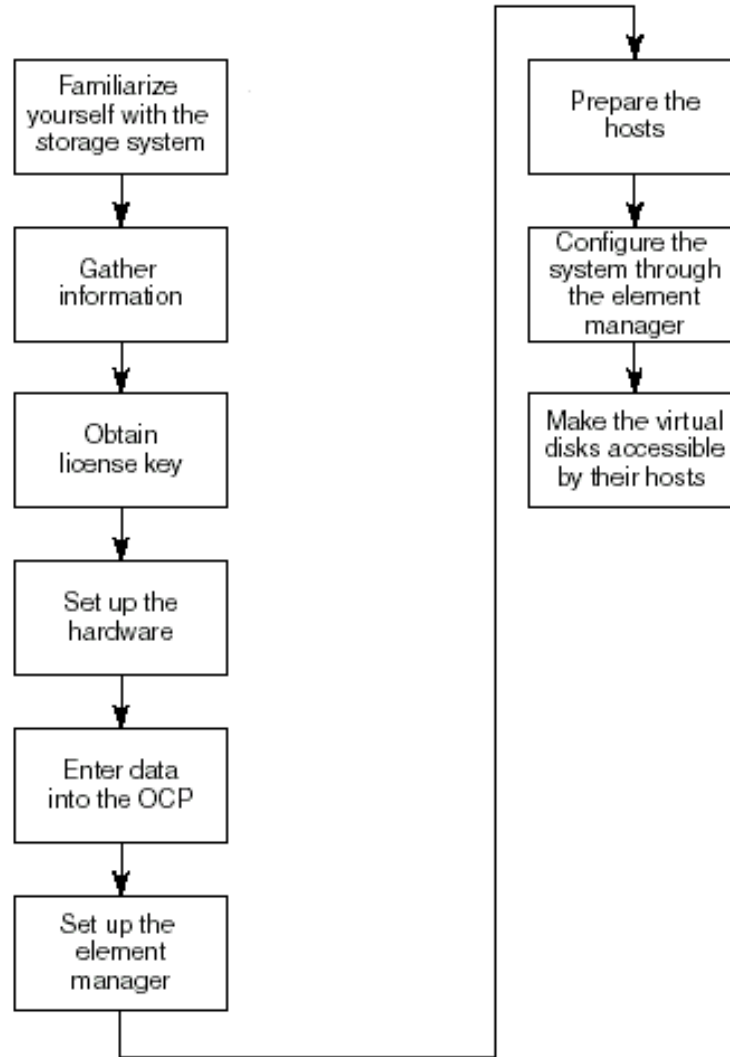
- Gathering all the necessary information
- Obtaining a license key
- Setting up the hardware, including:
 - Inspecting the storage system
 - Applying power to the system
 - Attaching the controllers to the fabric
- Entering data into the operator control panel (OCP)

Setting up Command View EVA is covered in Module 8.

The User Guide and Service Manual have more information about the entire process.

Getting Started Flowchart

The following diagram guides you through the steps of setting up and configuring the Enterprise Virtual Array.



Gathering Information

Retrieve the items described here, and perform the procedures described in each. After you complete these procedures, you can initialize and operate the Enterprise Virtual Array.

Locate the following items to install the storage system:

- *HP StorageWorks Enterprise Virtual Array World Wide Name (WWN) Label*, which is a separate sheet of paper shipped with the system.
- *HP StorageWorks Enterprise Virtual Array Read Me First* document.
- *HP StorageWorks VCS License Key Retrieval Instruction Sheet*, which is a separate sheet of paper that ships in the *HP OpenView VCS* kit. This sheet provides instructions for obtaining a Basic license for the storage system software.
- If you bought a license for the snapshot feature, you should find another license key retrieval instruction sheet. This is a separate sheet of paper that ships in the kit called *HP OpenView VCS Snapshot V2.0 for Dual HSV Controllers*.
- The latest *HP OpenView Storage Management Appliance Update*, which consists of the Management Appliance Update CD and its associated documentation.

You can determine the latest update version available by checking your Release Notes or contacting your HP Authorized Service Representative.

- The boxed kit for the operating system of the host computer. If you have hosts running different operating systems, you need a boxed kit for each operating system. This kit ships separately from the storage system.
- The boxed kit that contains the hardware documentation and ships with the system (the box that contained the *Enterprise Virtual Array User Guide*).

Locate these items and keep them handy.

Make a list of information for each computer (host) to be used in the storage system. The information you need for each host is as follows:

- LAN name of the host
- A list of World Wide Names of the Fibre Channel adapters, also called host bus adapters, through which the host connects to the fabric on which the storage system resides
- IP addresses of each host
- Operating system
- Available LUN numbers

Obtaining a License Key

You need a license key to unlock the VCS that runs on both controllers in a storage system. One license covers both controllers in a storage system.

Obtaining a license key requires the following documents:

- *HP StorageWorks Enterprise Virtual Array World Wide Name (WWN) Label.*
- *HP OpenView License Key Retrieval Instruction Sheet*, which includes the License Authorization ID.
- If you purchased the Business Copy EVA or Continuous Access EVA license, you also need the snapshot *License Key Retrieval Instruction Sheet*, which includes the License Authorization ID.

VCS License Types and World Wide Names

There are three types of VCS licenses:

- Basic
- Business Copy EVA (ordered separately)
- Continuous Access EVA (ordered separately) — For VCS V3.0 only

These licenses are sold as part of a software kit. Business Copy and Continuous Access licenses are sold based on the capacity the customer requires. There is one WWN per controller pair on a storage system.



Important

Both the basic license and the WWN are required to create an initialized storage system.

Basic License

The basic license provides the right to use VCS to create and operate an Enterprise storage system whose Fibre Channel identity is the WWN that was generated by manufacturing and attached to the hardware. A single VCS software kit provides the basic license that covers both controllers in the storage system.

The license is bound to the WWN. The WWN is a characteristic of the entire storage system as an entity. The license and the WWN do not belong to any specific component of the storage system. Individual components can be removed and replaced—even the controllers—but the storage system remains and retains its WWN identity.

The WWN belongs to the storage system. It is entered into the OCP of one of the controllers, but it does not belong to the controllers. The controllers simply keep a record of the WWN until the storage system is initialized. At initialization, the WWN becomes tied to the storage system.

The basic license must be tied to the WWN, which you do through the License Fulfillment Process. You receive an Authorization ID on a License Key Retrieval Instruction Sheet, and you provide the Authorization ID and the WWN to the license fulfillment website. The license fulfillment website then generates and sends back a license key. This license key ties the WWN to the license. When you enter this license key into Command View EVA, it unlocks VCS on both controllers, and you can begin using your storage system.

Business Copy EVA

Business Copy EVA can create nearly instant copies of any active volume on the Enterprise Virtual Array for use by other applications or systems for batch processing, backup, or testing, without interrupting current workflow. This replication eliminates data movement between the host and storage device, allowing you to create business copies for online backup or application testing environments.

Continuous Access EVA

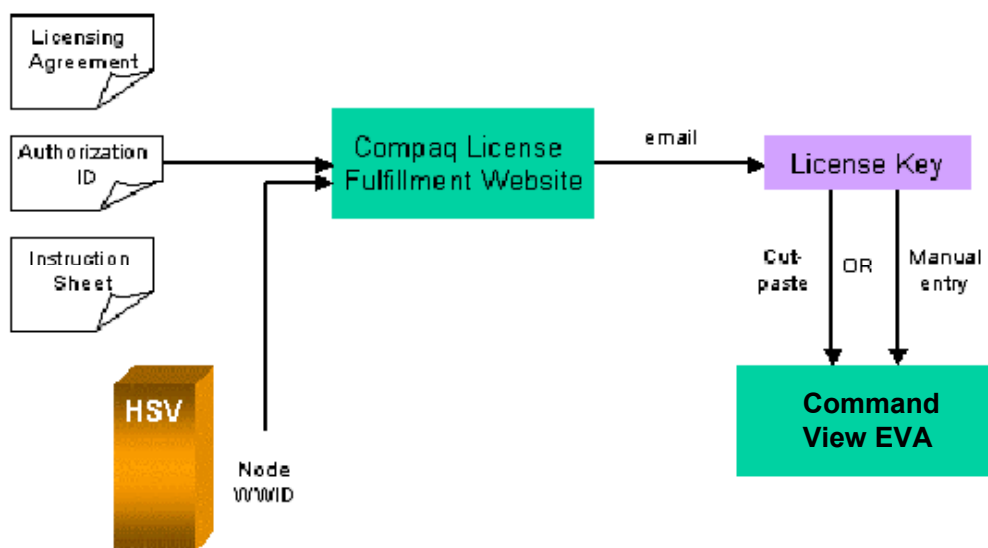
Continuous Access EVA provides replication functionality on the EVA and requires a separate license for an HSV controller pair. The license is ordered separately from the required basic VCS license. These licenses are supported on VCS V3.0 and above.

License Fulfillment Process

The shipment of the Enterprise Virtual Array storage system includes the following:

- Licensing agreement
- Authorization ID
- Instruction sheet

You can complete the license fulfillment process manually through email or fax, or by visiting the HP website. The following diagram describes the process.



Basic Licensing Process

Follow the instructions on the License Key Retrieval Instruction Sheet to obtain a license key at the license fulfillment website. If you do not have Web access, the License Key Retrieval Instruction Sheet describes how to obtain a license key manually through email or fax.

The license fulfillment website prompts you to enter the storage system WWN. This number is on the sheet called *HP StorageWorks Enterprise Virtual Array World Wide Name (WWN) Label*. This number may also be on labels on both sides of the controllers.



Caution

Enter the WWN exactly as it is on the label before you submit the licensing information. The WWN is case sensitive, and the hyphens must be entered. Once the WWN is entered into the licensing database, the license key is irrevocably locked to that WWN.

Do not enter the checksum that appears on the label. Doing so invalidates the license key.

When you obtain the license key, copy it or print it. You must enter the license key into Command View EVA when you initialize the storage system. The license key is made up of alphanumeric characters and may have as many as four lines of text. Enter the license key exactly as it is displayed.

Add a license

Add license Cancel ?

Enter a license key and click the **Add License** button to activate special features on your storage system.

```
1F1C377E0B3B
HOSTID=HSV\WWN=5000-1FE1-0013-A220 NOTICE="Authorization = \
D101DADDIEC021364684, Qty 1, QM-ENTRP-RI.SE - Enterprise \ '
Controller Software - SNAPSHOT, Quickspec 78*90*12 ck=129|
```

The license keys use checksum and are invalidated if they are tampered with. They persist on the appliance hard drive and on each Management Logical Disk (MLD). Loss of the appliance does not lose the licensing state.

Licensing Additional Storage Systems

If you have more than one license or more than one storage system, you can obtain up to 10 license keys at the same time. Each license key is locked to the WWN of one storage system, so be sure to keep records that show the association of license keys with WWNs.

Setting Up the Hardware

To install storage system hardware, contact an HP Authorized Service Representative. The service representative must ensure that the following steps are completed:

1. Complete the prerequisites — Complete the procedures printed on the shipping carton to unpack the rack and remove it from the pallet.
2. Inventory the system — Ensure that there is no physical damage and the system is complete.
3. Verify the site requirements — Ensure that the site satisfies safety and environmental requirements.
4. Move and stabilize the rack — Ensure that proper procedures are performed while moving and stabilizing the rack.
5. Inspect the storage system — Ensure that every element of the storage system, such as disk drives, cables, and PDUs, has been inspected.
6. Install the controller cache batteries — Ensure that the controller cache batteries are properly installed in both controllers.
7. Apply power — Ensure the correct procedures are used for applying power.
8. Attach controllers to the fabric — Ensure that controllers are properly connected to Fibre Channel switches for the particular configuration.


This module only describes the steps involved with inspecting the storage system, applying power, and attaching controllers to the fabric.

The Service Manual contains details regarding setting up site hardware.

Inspecting the Storage System

After the rack has been installed and stabilized with the help of an HP Authorized Service Representative, ensure that the following tasks have been completed:

1. All disk drives are fully seated.
2. All internal data cables are connected and fully seated.
3. The PDU power cable connectors and the wall receptacles are compatible.
4. The AC power distribution modules are connected to PDUs and fully seated.
5. The controller and drive enclosure power cords are connected and fully seated.
6. The power supplies, blowers, I/O modules, and EMUs are fully seated.
7. All data cables are properly connected and fully seated.
8. All panels, cable clamps, wire ties, and so forth are fastened securely.
9. The WWN labels are installed properly on either side of the controllers.



If the labels are not in the required location, refer to the label sheet, *HP StorageWorks Enterprise Virtual Array World Wide Name (WWN) Label*, and install the labels. The sheet contains labels and instructions for installing them.

Applying Power

Before applying power, you must install cache battery pairs in both controllers.




Caution

Before applying power to the system, be sure that the temperature has stabilized and is in the range +10°C to +35°C (+50°F to +95°F). If the system temperature is outside of this range, an error condition occurs when you apply power.

To apply power to the storage system, complete the following steps:

1. Set the circuit breakers on both PDUs to the Off position.

The PDUs are located in the rear of the rack. The following figure shows a close-up view of the PDU circuit breakers.



2. Set the controller enclosure power switches to the Off position.
3. Connect the PDU power cable to the wall receptacle.

At this time, you can apply power to the rack enclosures.

4. Set both PDU circuit breakers to the On position.

The drive enclosure power supplies start operating. The audible alarm sounds, but should turn off as soon as both controllers are fully powered up.

5. Verify that the power supply and blower status LEDs (green) are On.

The middle drive status LED should be On; the other two drive status LEDs should be Off. The I/O module status LED (middle) should be On.

The drives start to spin up, during which time the top drive status LED blinks. When the two top drive status LEDs are On and the bottom drive status LEDs are Off, the drive has spun up. Wait until the drives have spun up before turning on the controllers.

! Important

If you see any error indications, refer to the *Enterprise Virtual Array User Guide* for possible causes and corrective actions.

6. When all the drives are spun up, set the power switches of both controllers to the On position.

The LCD on the controller enclosure OCPs should show HSV110 Startup. When one of the controller OCPs displays the message Enter WorldWide Name, the controllers are running their software and are operational.

When both controllers are powered up, the audible alarm should turn off.

Attaching the Controllers to the Fabric


The controller pair attaches to its hosts through a Fibre Channel fabric. For redundancy, the controllers should be connected to two fabrics. If one fabric becomes nonoperational, the hosts can still access the data through the other fabric.



Important

Some operating systems have specific requirements for how switches are configured. Consult the Installation and Reference Guide for your host operating systems. The specific operating system constraints always take precedence over this generic description.

HP recommends using a cabling scheme that is easy for you to remember, and be sure to label both ends of the cables.



Entering Data into the OCP

You must enter into the OCP two pieces of data critical to the initial storage system setup:

- WWN and checksum — Mandatory.
- Storage system password — Optional. A password is a security interlock that allows only specific instances of Command View EVA to access the storage system.

This data is required to set up the HSV controller pair, which allows your storage system to communicate with Command View EVA through VCS.

Refer to the *Enterprise Virtual Array User Guide* for more detailed information about OCP operation.

Setting Up an HSV-Series Controller Pair Using the OCP

Enterprise Virtual Array storage system operations require that each **controller pair** have a unique Node WWN. This 16-character alphanumeric name identifies the system on the storage system. HP assigns the WWN to each storage system before shipment. The *Enterprise Virtual Array World Wide Name Label* document defines the WWN for each storage system. The Node WWN labels, similar to the one shown, specify the system-specific WWN and checksum.



Sample Node WWN Label



Important

The controller pair WWN is unique to a controller pair, and cannot be used for any other controller pair or device anywhere on the network. It is the only WWN applicable to any controller installed in a specific physical location, even a replacement controller. Once a WWN is assigned to a controller, it cannot be changed as long as that controller is part of the same storage system.

Entering the WWN, WWN Checksum, and Password

The *Enterprise Virtual Array User Guide* includes the detailed instructions for entering the 16-character WWN, the two-character checksum, and optional eight-character password by way of the OCP.

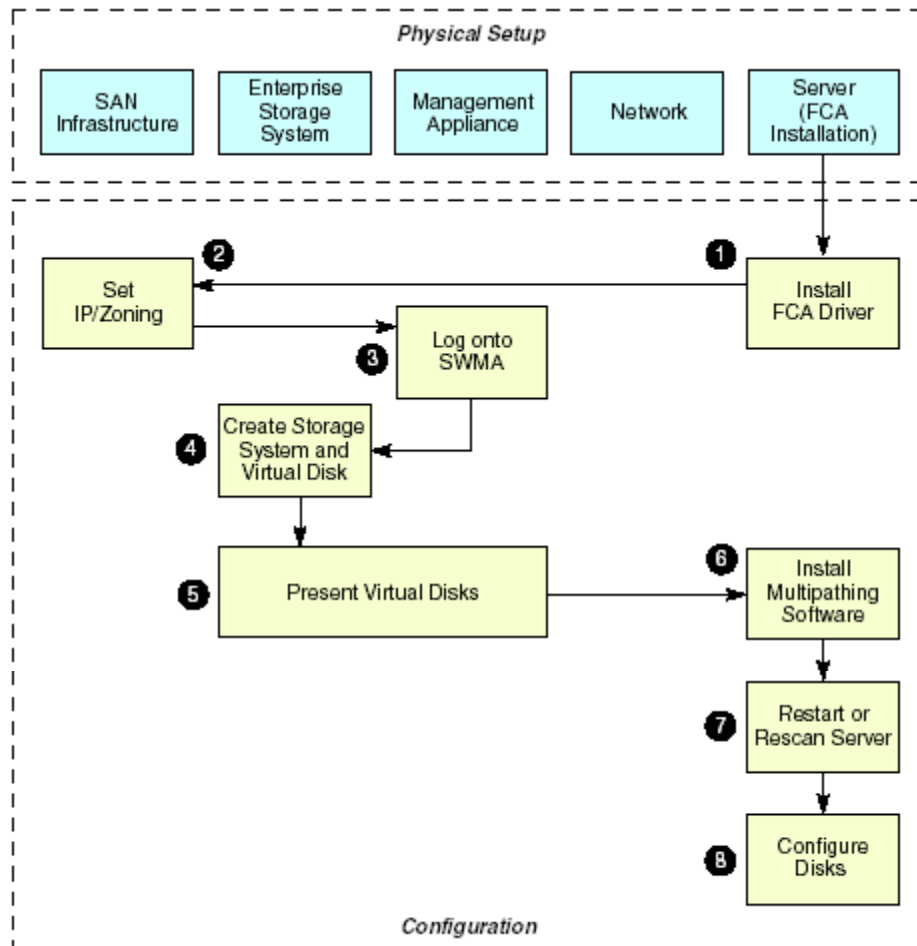
Configuring the Storage System

After you have completed the initial setup of physical components, you can configure the Enterprise Virtual Array. The following diagram shows the configuration process for all operating systems.

The order in which you perform the steps to configure the storage system is important. All of these steps are necessary, and each is dependent on the previous step being completed.

Note

Step 6, Installing Multipathing Software, is only required for certain operating systems. Multipathing software is embedded within OpenVMS and Tru64 UNIX.




Note

Refer to Module 8 for detailed information on configuring the system.

Learning Check

1. The two rack models are the _____ and the _____.
.....
2. Power is distributed through the rack by the _____, and power is distributed to the enclosures by the _____.
.....
3. The two hardware configuration classifications are _____ and _____.
.....
4. Identify each transceiver in the diagram in terms of its loop connections.



① =

② =

③ =

④ =

5. Complete the sentence. When connecting cables, connect _____.
.....
 - a. Loop A cables to an I/O module A
 - b. Loop A cables to an I/O module B
 - c. Loop B cables to an I/O module A
 - d. Loop A cables to either an I/O module A or an I/O module B

6. When configuring the Enterprise Virtual Array, which of the following cabling rules is correct?
- a. Loop pair 2 enclosures are at the bottom of the rack, below the controller pair.
 - b. Loop pair 2 enclosures are at the top of the rack, above the controller pair.
 - c. Loop pair 1 enclosures are at the top of the rack, above the controller pair.
 - d. The A loops are connected from the bottom-most drive enclosure upward.
7. How is an expansion rack connected to a 2C12D configuration?
-
-
8. Complete the sentence. When powering up the storage system, before you can apply power to the controllers you must ensure that _____.
-
-
9. What is the correct sequence of steps for setting up a storage system?
- 1) Obtain a license key.
 - 2) Gather all necessary information.
 - 3) Enter data into the OCP.
 - 4) Set up the hardware.
- a. 2, 1, 3, 4
 - b. 2, 1, 4, 3
 - c. 4, 2, 1, 3
 - d. 2, 3, 4, 1
10. The _____ and the _____ are required to create an initialized storage system.
-

Overview

This module introduces the HP OpenView storage management appliance and its essential role in a SAN with the Enterprise Virtual Array storage solution. It describes appliance hardware and software components, including the storage management appliance software and Command View EVA. You are provided guidelines for configuring the appliance to operate with the Enterprise Virtual Array, and the limitations for configuring the Enterprise Virtual Array in a mixed storage system environment.

Note

Troubleshooting information for the appliance is contained in Appendix B.

Objectives

After completing this module, you should be able to:

- Specify the hardware and software components of the HP OpenView storage management appliance.
- Describe the storage management appliance software compatible with the Enterprise Virtual Array.
- Describe the tasks that you can perform with the storage management appliance software.
- List and discuss the features, rules, and limitations of Command View EVA software.
- Identify configuration tasks for the storage management appliance.
- Describe the purpose and features of the Restore CDs.

Storage Management Appliance Description

The storage management appliance is an element in the HP SAN management strategy. The appliance is a server on the SAN fabric, and it contains SANworks software products that allow you to perform specific tasks for management of hardware and software members of the SAN.


The appliance uses standardized, web-based GUIs accessed from supported web browsers. One of these GUIs, Command View EVA, resides on the storage management appliance to control and monitor the Enterprise Virtual Array.

The appliance is configured at the factory for compatibility with the Enterprise Virtual Array. You should not load any additional software or tools on the appliance except what is required and qualified by HP.

The following are key features and benefits of the appliance:

- **Graphical storage configuration and monitoring** — Provides easy installation, configuration, and monitoring of the HSV110 storage systems.
- **Automated failure notification** — Provides timely broadcast notification of FRU-level failures through email, web browser, and SNMP traps.
- **Host-independent fault monitoring** — Supports serviceability of storage and switch products without the involvement of host computers.
- **FRU-level callout analysis** — Provides analysis of events to facilitate faster mean-time-to-repair and identification of specific repair actions.
- **Improvement of overall system performance** — Operates independently of host servers, preserving precious application processing cycles. Also, the appliance is not located in the SAN data path, so data transfers proceed without impact from SAN management.
- **Rack-mountable** — Enables ease of installation and administration.
- **Improves security** — Provides SAN monitoring and management, independent of host servers.
- **SAN-scale monitoring** — Monitors all supported switches and HP storage arrays, lowering the cost of ownership for the customer through a single point of monitoring across the SAN. The appliance supports the Open SAN concept.
- **Simplifies SAN management** — Provides highly scalable, host-independent management point SANs.
- **Provides for SAN redundancy** — Supports multiple appliances on the same SAN. This requires zoning with each appliance in a separate zone.

The following diagram shows the logical connectivity between the storage management appliance and the HSV110 controller.



Note

The Fibre Channel connection relies on StorageCell Management Interface (SCMI) commands, the means of communication between the appliance and HSV controller. SCMI was covered in Module 2, HSV Controller.

Storage Management Appliance

The storage management appliance is built on a ProLiant DL 380 G1 or G2 server chassis. The appliance provides a host platform for enterprise wide storage management applications. The following are some other characteristics:

- Uses a web-based GUI
- Operating system is a customer version of Windows 2000 Advanced Server
- Supports SNTP time server sync
- Has virus and backup software support
- Can be mounted in the Enterprise Virtual Array rack if space is available

The following pictures show the G1 and G2 appliances.



Management Appliance — G1




Management Appliance — G2

With a web-based GUI, the appliance provides a centralized point for monitoring SAN elements. This simplifies management tasks and helps to reduce the costs associated with managing a SAN.

Note

The SANworks Management Appliance DL380 G1 is end of life. Any new appliances ordered will be G2.

Hardware Description for G1 Appliances




1. Low-profile 24x IDE CD-ROM
2. 1.44-MB Diskette Drive
3. Protected Power Switch
4. Front Panel Status LEDs
5. One full height or two half height removable media bays
6. Four 1-inch Wide Ultra2/Ultra3 SCSI hot plug drive bays

The hardware specifications for Management Appliance Generation 1 (G1) include the following.

Hardware Feature	Attributes
Form factor	3U
Processor	733MHz
Memory	512MB
CD-ROM drive	24x Max
Floppy drive	1.44MB
PCI slots	4
Network controller	One 100Mb/s Fast Ethernet network interface card (NIC)
Fibre Channel host adapter	2 KGPSA-CB
Modem interface *	Compaq 56K PCI fax modem
Remote Insight Board	Y cable used for keyboard, mouse (serviceability only)
Bays	5 or 6
Internal storage	2 X 18GB Ultra3 hard drives mirrored via RAID On Chip (ROC)
Factory Installed Software (FIS)	OSM V1.0
Restore CD	Restores to OSM V1.0C**, Services Only CD
System Name	SWMAxxxxxx, where xxxxxx = last 6 digits of appliance serial number
* The modem is solely for pager notification and should not be used for RAS or any other purpose.	
** A new version of the G1 Restore CD restores to OSM V1.0C.	

Hardware Description for G2 Appliances



The G2 appliance was designed to replace the Compaq ProLiant DL380 line. The appliance has a faster processor and increased memory to position it for use, in the near future, within platforms that use StorageWorks software.

The G2 and G1 appliance differences are shown in the following table.

Hardware	G2 Enhanced or Changed Attributes
Form factor	2U
Processor	1.26GHz
Memory	1GB
PCI slots	3
Network controller	Two 100Mb/s Fast Ethernet NIC
Modem	None
Remote Insight Board	Ribbon cable, rather than Y cable
Bays	8
Internal storage	2 X 18GB Ultra3 hard drives mirrored via SMART Array 5i
Factory Installed Software (FIS)	January 2002 Update
Restore CD	Restores to January 2002 Update, OSM V1.0C, Customer CD
System Name	SMAxxxxxxxxxxx, where xxxxxxxxxxxx = full serial number of appliance

Previous system naming conventions, using only six digits, created duplicates of appliance names; therefore, the full appliance serial number is required.

! Important

You cannot use the G2 Restore CD on a G1 appliance. The G1 appliance still uses a Services-only CD, while customers receive a Restore CD with G2.

The process for migrating from G1 to G2 does not currently give a complete update; that is, you cannot duplicate your G1 image by backing up the G1 configuration and restoring it to the G2.

SAMS Operating System

The operating system for the appliance is an HP proprietary version of Windows 2000 Server. The SAN Application and Management Software (SAMS) allows the hardware to communicate with the Fibre Channel storage network and with the Ethernet network. It supports the collection of data and allows the data to be accessed by the user and by the SANworks applications that reside on the appliance.

The SAMS operating system supports the following:

- Peripheral support (NIC, modem, FCA)
- Microsoft Data Engine (MSDE) (SQL version)
- Emulex Fibre Channel combo driver
- Remote Insight Board (RIB) software
- ELM and Internet Information Server (IIS) servers (web servers)

Note that IIS is Microsoft's web server that runs on Windows NT platforms.



Caution

At no time should you reload the Windows 2000 Server operating system. This causes the appliance to stop functioning.

Appliance-Based Software

The appliance provides a host platform for enterprisewide storage management applications, including:

- Enterprise Virtual Array compatible software
 - Open SAN Manager (OSM) V1.0C or Storage Management Appliance Software V2.0, SP1A and SP3
 - HSG Element Manager V1.0E
 - Command View EVA V2.1 or V3.0 shipped with Enterprise Virtual Array storage system or VCS V2.X or V3 software kits
- Additional supported applications
 - Business Copy EVA V2.1, V2.1A
 - Network View V2.0B
 - Storage Allocation Reporter (SAR) V4.0A
 - Virtual Replicator V3.0
 - See the release notes for additional information

The storage management appliance only runs applications that are specifically designed for it. Do not install other applications.

The software used in conjunction with the Enterprise Virtual Array storage system includes:

- **Command View EVA** — Used to set up and manage the HSV110 controllers.
- **API** — Interface to Command View EVA used by the Storage System Scripting Utility (SSSU) to issue its commands. SSSU is a user interface with a command prompt interactive mode.

Storage Management Appliance Software

The Storage Management Appliance Software V2.0 is the latest version of what was formerly called Open SAN Manager (OSM) V1.0C.

This software interface performs the following:

- Centralizes the monitoring and management interfaces
- Configures appliance parameters such as IP network address, system name, date and time, and administrator password
- Installs and sets up SAN applications
- Upgrades system components such as operating system and drivers (through service packs)
- Starts and restarts the Command View EVA service
- Launches installed StorageWorks applications such as Command View EVA

The storage management appliance can be monitored by Insight Manager.

Command View EVA

The Enterprise Virtual Array storage system implementation includes Command View EVA. Command View EVA is the GUI to the Enterprise Virtual Array storage system. You use this GUI to set up and manage the Enterprise Virtual Array with its physical and virtual disks. The GUI is accessible through a standard Internet browser running on various operating systems.

For each Enterprise Virtual Array storage system, Command View EVA assists you with the following:

- Initializing the storage system
- Adding hosts to a storage system
- Creating disk groups, virtual disks, snapshots, and snapclones
- Configuring distributed spare capacity
- Monitoring existing virtual disk, controller, physical disk, and host properties
- Making changes to these properties for different configurations

Command View EVA Rules and Limitations

You can manage both HSG and HSV subsystems on the same storage management appliance. Basic Command View EVA rules and limitations include the following:

- One instance of Command View EVA software per appliance, and, if required, one instance of HSG Element Manager software on the same appliance.
- A maximum of 16 Enterprise Virtual Array storage systems per Command View EVA in a homogeneous (HSV controller only) environment.
 - For homogeneous environments:
 - HSG controllers, a maximum of 25 HSG subsystems are supported.
 - HSV controllers, a maximum of 16 subsystems are supported.
 - For heterogeneous environments, a maximum of 25 HSG subsystems and eight HSV subsystems are supported.

Note

Performance degrades well before you reach the maximum subsystem limit. Remember that you can add more subsystems by adding more appliances and proper zoning.

- Limit of one active Storage System Scripting Utility (SSSU) session per Command View EVA installation
- Command View EVA and SSSU session not running at the same time
- Multiple simultaneous user interface sessions
- One Command View EVA actively managing the Enterprise Virtual Array

Note

If you have multiple storage management appliances, you should configure the system so that all of the Enterprise storage systems are managed by a single storage management appliance at any given time.

Initial System Configuration

The appliance is configured at the manufacturing site with the SAMS operating system installed and active. The system is configured using the onboard NIC and the DHCP protocol to access the network. You need to set static IP addresses for storage management appliance networking components.

System names are configured as described in the following:

- For G1 appliances, the system name is SWMAxxxxxx, where xxxxxx is the last six digits of the serial number of the appliance.
- For G2 appliances, the system name is SMAxxxxxxxxxxx, where xxxxxxxxxxxx is the full serial number of the appliance.

Note

The serial number is located on the front of the appliance enclosure, under the bezel.

The appliance is designed to plug and play, and, under normal conditions, only requires a standard installation. It is designed to work remotely and allow the administrator to interact with the SAN and its components.

Passwords

The appliance password for the administrator (and Administrator) user determines that user's access to the SWMA. Those passwords enable access to the:

- Web agent (ELM server)
- SAMS operating system and Windows console
- RIB board

The following table lists the default user name and passwords.

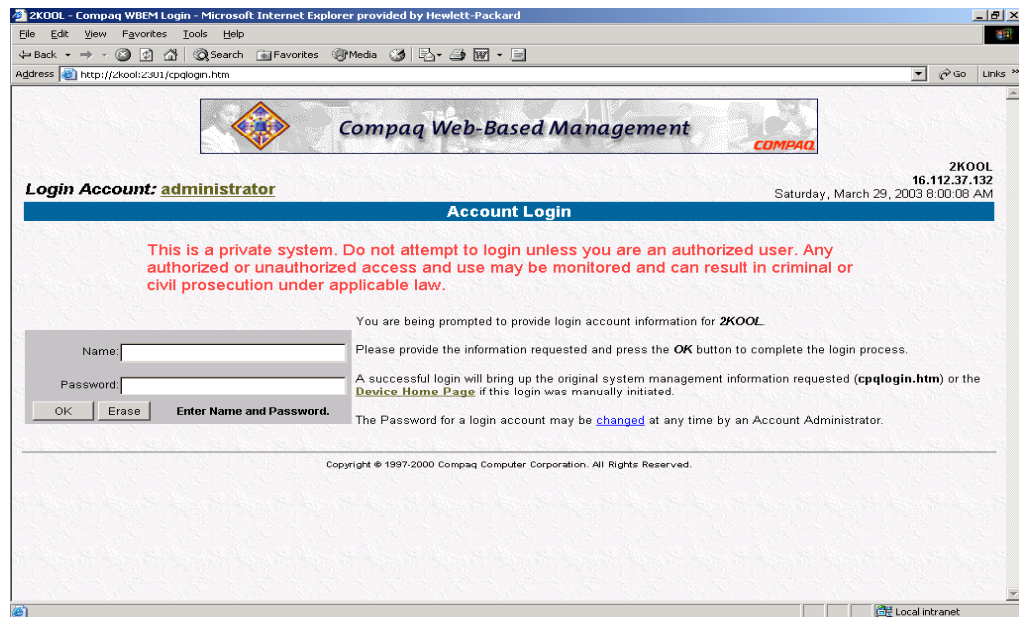
Area	User Name	Password*
Web agent initial logon	administrator*	administrator
SAMS operating system	administrator	admin followed by the last six digits of the serial number reversed, for example: admin100K13 Alpha character is capitalized.
Remote Insight Board	Administrator*	Per attached tag on the SWMA, for example: 2VH5

Note

* These are case sensitive.

ELM Web Server Login

The following is a sample screen for a login to the ELM server.



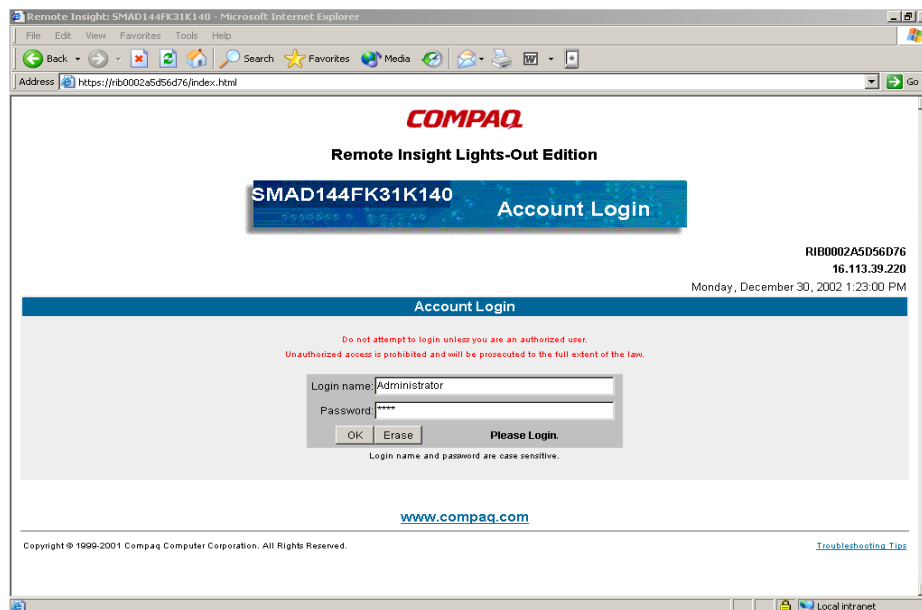
SAMS and Windows Console Login

The following is a sample screen for a login to SAMS and the Windows console.



RIB Login

The following is a sample screen for a login to the RIB.



Appliance Administration

Always use the storage management appliance web-based software remotely to carry out operations on the appliance. Do not use Windows system tools to perform operations, because this can lead to unpredictable results.

You should always browse to the appliance through the following URLs:

- `http://appliancename:2301`
- `http://ipaddress:2301`

Storage Management Appliance Configuration

Most of the appliance configuration is done at the factory. The SAMS software is already loaded. You need to perform a few procedures and checks to ensure communication and discovery of the switches by the appliance.

The following are some configuration steps that may be required for the appliance (exact order is not implied):

- Establish communication with the appliance by performing an IP host name resolution using the Fully Qualified Domain Name (FQDN). You should be able to communicate using this name if the DNS server is functioning.
- Configure static IP addresses for the network interface card (NIC), and one for the RIB board. If an appliance uses Dynamic Host Configuration Protocol (DHCP), DHCP might assign a different IP address on restarts. HP recommends that you define static IP addresses.
- Set the time zone (Mountain time is the default).
- Add to a domain (if necessary).
- Reconnect to the appliance.
- Configure the switches (zoning).
- Change the appliance name (if absolutely necessary).

Some of these steps are covered in the lab exercises.




Caution

Renaming the appliance can affect the installed applications and should be handled carefully. The best time to rename the system is before installing any of the value-added HP storage management applications. See the release notes for additional information.

Storage Management Appliance Availability

You have dual-redundant SAN support with one appliance. The active/passive configuration support enables quick recovery in case of a fatal appliance outage.

The following diagram shows the dual-redundant support.



Dual Redundant SAN with HSV and HSG Controllers in Multibus Mode

Management Appliance Restore CD

If required, you can restore the storage management appliance operating system to original default state. You should use these options only after all other attempts to restore the operating system and storage management software have failed.

Two appliance versions are available for a restore, the G1 and G2.

- The G1 appliance uses a V1 Restore CD or a V2 Restore CD.
- The G2 appliance uses its own Restore CD (Management Appliance II V2.0 Restore CD). You cannot use the G2 Restore CD on a G1 appliance.

All of the procedures for using the Restore CDs are listed in the Module 6 lab.

Restoring G1

For restoring the G1 appliance, you have two options:

1. Use the V1 Restore CD. This has a default of OSM V1.0.
 - Apply the June 2001, January 2002, and December 2002 patches.
 - This CD is only available through Customer Services.
2. Use the V2 Restore CD. This is already patched to the MA January 2002 update (OSM V1.0C).
 - Apply the December 2002 patch.
 - The customer can order this CD from hp.com.

Restoring G2

For restoring the G2 appliance:

1. Use the Management Appliance II QuickRestore CD V2.0 (shipped with the appliance). The default state is OSM Version 1.0C (January 2002).
2. Apply the December 2002 patch (Storage Management Appliance Software V2.0).

Applying Patches

You will need to apply several patches to bring all software up to date to support VCS V3.0. These upgrades and patches are contained in the HP StorageWorks Storage Management Appliance Software V2.0 Service Pack 3 Install CD and include the following:

- Storage Management Appliance Software V2.0 SP1A
This is a mandatory service pack that addresses a communication problem between Network View V2.0B in a large SAN environment with multiple Enterprise Virtual Array storage systems.
- Microsoft SQL Server Virus correction
This is a mandatory security patch to offset the **W32.Slammer** virus.
- Microsoft Windows 2000 Security Patch — **DisableWebDAVPatch.SWP**
- HP StorageWorks Command View EVA V3.0
- Fibre Channel HBA Performance Optimization patch

At the end of the restore and software update process, verify the end state of the storage management appliance software:

- Storage Management Appliance Software V2.0, SP3
- Command View EVA V3.0
- HSG Element Manager 1.0E

Note

These updates should be applied regardless of hardware version.

Learning Check

1. List five major hardware components of the G2 storage management appliance.
.....
.....
2. What storage management appliance software is currently compatible with the Enterprise Virtual Array?
.....
3. List five functions you can perform using the Storage Management Appliance Software V2.0.
.....
.....
.....
4. Which of the following Command View EVA rules is correct? (Choose one.)
 - a. Command View EVA permits only one user interface session at one time.
 - b. You can only install one Command View EVA instance on one SMA.
 - c. With one Command View EVA instance, you can manage up to 20 Enterprise Virtual Array storage systems in a SAN.
 - d. You can have more than one Command View EVA actively managing an Enterprise Virtual Array at one time.
5. List the steps to license Command View EVA.
.....
.....
.....
6. What are the required configuration tasks to set up the storage management appliance to operate with the Enterprise Virtual Array storage system?
.....
.....
.....
7. Describe the purpose of the Restore CD.
.....
.....

Overview

This module describes concepts and terminology essential to understanding how the Enterprise Virtual Array virtualization software operates. It defines the terms that you need to know and then explores more deeply the concepts behind the terms.

The module discusses the following virtual storage concepts: storage system, disk group, virtual disk, distributed sparing, redundant storage set (RSS), host, and LUN. This module also describes how the HSV controller provides snapshot and clone data replication methods, specifically for creating demand-allocated snapshots, fully-allocated snapshots, and snapclones.

Note

Refer to Appendix C for details and issues regarding disk allocation, metadata, disk failure, and occupancy level.

Objectives

After completing this module, you should be able to:

- Define the virtual storage terms applicable to the Enterprise Virtual Array.
- Name the primary features of a disk group and default disk group.
- Describe the ways in which disk groups are configured.
- Name the primary features of virtual disks.
- Differentiate between conventional RAID and distributed virtual RAID (VRAID) technology.
- Describe virtual disk leveling.
- Describe distributed sparing and disk failure protection levels.
- Define the replication methods used with the Enterprise Virtual Array.

Virtual Storage Terminology

You should be familiar with the following terms as they pertain to the Enterprise Virtual Array storage solution:

- Storage system (cell) — An initialized pair of HSV controllers with a minimum of eight physical disk drives.
- Disk group — A group of physical disks, from which you can create virtual disks.
- Virtual disk — A logical disk with certain characteristics, residing in a disk group.
- Virtual disk leveling — Distribution of all user data within the virtual disks in a disk group across all physical disks within the disk group.
- Distributed sparing — Allocated space per disk group to recover from physical disk failure in that disk group.
- Redundant storage set (RSS) — A subgrouping of drives, usually 6 to 11, within a disk group, to allow failure separation.
- Host — A collection of host bus adapters that reside in the same (virtual) server.
- Logical unit number (LUN) — A virtual disk presented to one or multiple hosts.

Storage System (Cell)

A storage system, also referred to as a *cell*, is an initialized pair of HSV controllers with a minimum of eight physical disk drives. The storage system requires a name up to 24 characters long, in which:

- Special characters are not allowed (for example, ?, “, /, \).
- Two consecutive spaces are not allowed.
- The name is changeable.

A storage system is created at system initialization. This does the following:

- Makes the storage system ready for use
- Binds the controllers together as an operational pair
- Creates the first disk group
- Establishes preliminary data structures on the disk array (metadata)

Note

The term *cell* is used primarily by Engineering, however, as shown in the updated SSSU commands, it is being replaced by *storage system* or *system*.

Storage System Metadata

System-level metadata is stored on quorum disks and contains:

- Controller information
- WWN and storage system name
- Character map of disk groups and virtual disk members

Disk Group

A disk group is a group of physical disks from which you can create virtual disks. Disk group features include the following:

- A maximum of 16 disk groups are allowed per storage system.
 - The minimum number of physical disk drives in a disk group is 8.
 - The maximum number of physical disk drives is the number present in the system up to 240.
 - Unassigned or new disks can be added to an existing group.
 - Mixed drive sizes within a group are allowed.
- A disk group has a disk failure protection level.
 - The level can account for no, single, or double disk failures (not concurrent failure).
 - The default protection level is no disk failures (None).
- Space is allocated in 2MB segments (PSEGs).

Note

A *PSEG* is the smallest physical disk space allocated on a disk drive. A virtual disk is built of multiple PSEGs from all the disk drives in a disk group. A PSEG is not related to the selected redundancy level of the virtual disk.

- The chunk size is 256 blocks (128KB), fixed.

Note

A *chunk* is the amount of space on a virtual disk that is read or written to any one physical disk before the next physical disk is accessed. For example, in a VRAID5 virtual disk, when a large write (or read) is issued, the system writes the first chunk of data to one disk drive, then writes the next chunk to the next disk drive; after the fourth disk is written, the fifth disk contains the parity data.

Data that is smaller than the chunk size is held in write cache and consolidated with more data, when available, and then written to disk.

- Redundant Storage Sets (RSSs) are subgroupings of drives, usually from 6 to 11, within a disk group to provide failure separation. For example, if a disk group has two RSSs, a drive can fail in both RSSs and we will still have VRAID 5 and VRAID 1 availability.

RSSs are covered in detail in a later topic.

Default Disk Group

The default disk group is the first disk group created when a storage system is initialized. A default disk group:

- Is used like any other disk group.
- Can be deleted after another disk group is created to protect metadata.
- Can be renamed.

Disk Group Metadata

When a storage system is created, there are five quorum disks in the default disk group. There is always a maximum of 16 quorum disks (one per disk group).

When a new disk group is created:

- One quorum disk is created on it.
- One quorum disk is removed from the default disk group until just one quorum disk remains.
- Metadata overhead is $\approx 0.2\%$ of total disk group capacity.

Disk Group Capacity

A disk group has two types of capacity: assigned and unassigned. Assigned capacity is used for virtual disks, snapshots and clones, and spare space.

The system draws on unassigned capacity as needed for:

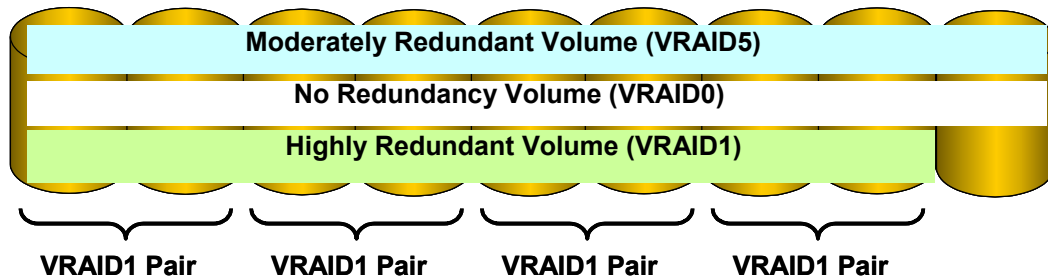
- New virtual disk creation, snapshots and clones.
- Freeing a physical disk for removal or reassignment.
- Data reconstruction after disk failure.

Unassigned capacity increases when virtual disks are deleted or new physical disks are added to the disk group.

Disk Group Member Terms

The following disk group terms are important when discussing VRAID1 disk allocation and drive failures:

- Unpaired — The disk that does not contain VRAID1 (mirrored) data on it because the disk group was created with an odd number of drives or because it lost its mirror partner because of a failure.
- Paired — The disk that was previously unpaired. A disk was added to the group, which now contains an even number of disks.



These terms may be more meaningful during the discussion on virtual disk leveling later in this module.

Disk Group Configuration


When configuring disk groups, the controller chooses the best drives, using a vertical organization and the most efficient parity group organization. For example, if you have eight shelves in the configuration and create a 16-member disk group (two 8-member RSSs), the two leftmost drive banks will be taken up vertically. If another 16-member disk group is created, it will use the two rightmost drive banks vertically. After creating the default disk group, you can group drives by physical location.

The following table describes the pros and cons of a vertical or horizontal configuration.

Configuration	Description	Pros	Cons
Vertical	One disk drive per shelf, ideally in a system with 8 or more shelves	100% guaranteed to survive a shelf meltdown failure	When drives fail or are added, may no longer be in a non-stop configuration
Horizontal	Multiple disk drives per shelf	Performance is only limited by the number of disk drives	Not guaranteed to survive a shelf meltdown failure

Disk Group Creation

Enterprise Virtual Array objects are the basic structures that comprise a storage system. During disk group creation, physical stores (Pstores), the raw FC-AL disks, combine with metadata to create volumes. A grouping of these physical disks creates a Logical Disk Allocation Domain (LDAD), what is commonly called a disk group.



Note

These objects are used extensively in controller event logs. This is discussed more in Module 10.

Virtual Disk

A virtual disk is a logical disk with certain characteristics, residing in a disk group. Virtual disk features include the following:

- Virtual disks can be created in 1GB increments with a minimum size of 1GB and a maximum size of 1.999TB (2TB less 1GB), except for Sun Solaris, which has a 1TB maximum.
- Redundancy levels on virtual disks offer the same protection level as RAID0, RAID1, and RAID5, but I/O per second, parity distribution, and other characteristics are implemented differently than with the HSG80. The redundancy levels for the HSV virtual disks are categorized as follows:
 - **None** (VRAID0, or striping) — Data is striped across all physical disks in the disk group. This provides about 100% of disk group capacity.
 - **Moderate** (VRAID5, or striping with parity) — Data is striped with parity across all physical disks in the disk group. Always five (4+1) physical disks per stripe are used. This provides about 80% of disk group capacity.
 - **High** (VRAID1, or mirroring) — Data is striped mirrored across all physical disks in the disk group. Two physical disks per mirror are used. This provides about 50% of disk group capacity.

Virtual disk leveling enables high spindle utilization.


Virtual Disk Creation

The HSV controller tracks which groups of blocks are written on a virtual disk. This is done so that snapshot and snapclone operations that move virtual disk data can limit their actions to only those blocks that contain data. Because HSV keeps track of this information, the first write to a block in a new LUN requires a metadata update. This is called *first write overhead*, or *first written bit*.

Subsequent writes touch only the block, not metadata. For normal uses with high I/O rates, overhead is not an issue. However, for performance testing, where write performance is critical, this overhead can be critical. Therefore, to eliminate first-write overhead, write all data blocks once before conducting a test.

The following abbreviations are used in the following diagram, which shows the relationship between Enterprise Virtual Array objects during virtual disk creation:

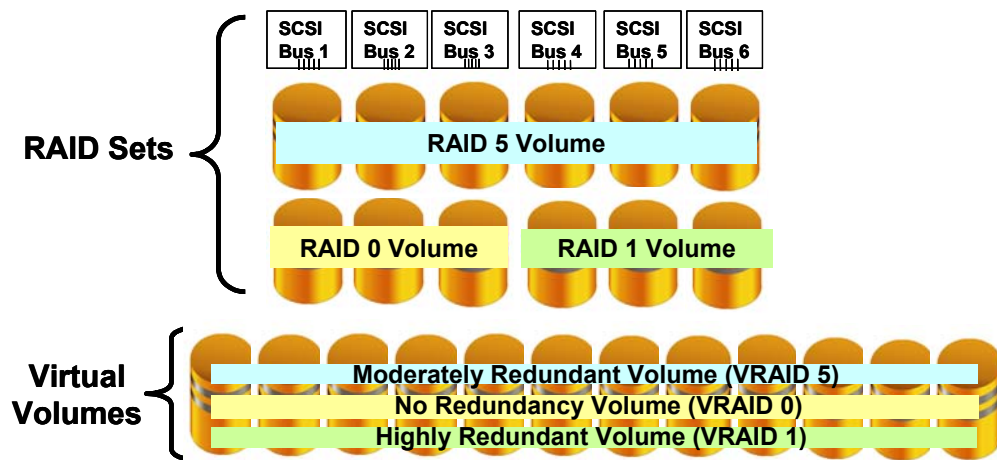
- Ldisks: Logical Disks (containers)
- SCVDs: Storage Cell Virtual Disk
- Dunits: Derived Units
- Punits: Presented Units



Distributed Virtual RAID Technology

In the Enterprise Virtual Array solution, the Command View EVA software manages the Enterprise Virtual Array. The HSV controller and VCS handle the disk virtualization with the implementation of distributed virtual raid (DVR) technology.

As shown in the diagram, conventional RAID sets use a fixed number of disks that have a single RAID level. Multiple RAID levels must be created on exclusive disks. The Enterprise Virtual Array uses virtualized RAID sets that allow you to create virtual disks and dynamically expand storage capacity. Each RAID set is distributed across the disks assigned to the storage pool. By using virtual disks that span a group of physical disks (spindles), all the spindles in the pool should carry the same workload.



Note

VRAID1 uses an even number of disks.

The Enterprise Virtual Array can support multiple virtual disks of varying capacity and RAID types. Virtualization allows data to be redistributed across physical disks within a disk group if an activity occurs that causes a change to the virtual disk data or disk group membership.


Conventional RAID

The following table compares HSG80 conventional RAID technology with HSV distributed RAID technology.

Conventional RAID	Distributed Virtual RAID
Performance is limited by the number of disk drives in the storageset.	Performance is limited by the number of disk drives in the disk group.
Possible to find customer data if you know the logical block number (LBN) and chunk size.	Not possible to determine on which physical disk customer data resides.
Load balancing is required of the application and databases over available back-end (SCSI) buses.	HSV controller eliminates load balancing procedure for applications and databases.
I/Os are balanced across the storageset.	I/Os are balanced across the disk group.


The features of each type of RAID become clearer by examining what would happen if a:

- RAID 5 drive failed without sparing?
- RAID 0 drive failed?
- RAID 1 drive failed without sparing?



Of course, with RAID 0, you would have a failure. For RAID 5 or RAID 1, you would run reduced.

For distributed virtual RAID, it is interesting to note what happens when you have disk drive group fail without sparing for VRAID0, VRAID1, or VRAID5.



Of course, with VRAID 0, you would be inoperative. For VRAID 5 or VRAID 1, you would run reduced.


Data Location Independence

Many of the benefits of virtualization in the Enterprise Virtual Array arise from the fact that it doesn't matter where data is located. But to realize those benefits, storage administrators must change their way of thinking and recognize the following:


- The backend disk access pattern is no longer correlated with the host access pattern.
- Configuration should not be used as an organizing tool—instead, use Enterprise Virtual Array capabilities.
- There is no need to micromanage the array.
- There is no need to be concerned with using every byte.

Block Mapping

The HSV controller assigns the logical blocks of a virtual disk to specific physical disk blocks through a dynamic map managed by the controller. This is similar to page tables in a virtual memory operating system, with minimal overhead in the I/O path.



The map is not block for block but rather 2MB segments (4096 blocks) of physical disk. For example, an 8MB Rstore requires 8MB VRAID0, 10MB VRAID5 (see diagram), or 16MB of VRAID1 storage.



Note


A redundant store (RStore) is 8MB of user-addressable space. This is important to know when Redundant Storage Sets (RSSs) are discussed. A physical segment (PSEG) is the smallest physical disk space allocated on a disk drive. A virtual disk is built of multiple PSEGs from all the disk drives in a disk group. A PSEG is not related to the selected redundancy level of the virtual disk.

Physical segments for a given virtual disk can come from any of the disks in the disk group and:

- Can and will change over time.
- May be reserved but not assigned until later.
- May be assigned when a unit is created or when needed.

Conclusions Regarding Data Location

In the Enterprise Virtual Array, the most that can be known about the location of data is that “the data for **this** SCSI LUN is somewhere in **that** disk group.” So, given that the only control over data location is which disk group a virtual disk is to reside in, and which physical disks comprise that disk group, disk group configuration is the key to configuring an Enterprise Virtual Array subsystem.



Virtual Disk Map

A virtual disk map, contained in the disk group metadata, contains pointers that identify which virtual disks use which chunks of capacity. Capacity chunks within a disk group are allocated at virtual disk creation time. The allocated size equals the size of the virtual disk.

The allocated chunks are not contiguous. Chunks are within stripes (moderate or high redundancy levels) but may not be between stripes.

This map is stored on the disk drives in the default disk group. During reboot, the map is read into the HSV controllers. The map is written to when the state of a disk group changes, such as when one of the following conditions occur:

- Disk drive failure
- Disk drives added or removed
- Traditional snapshot created
- Snapclone fully copied
- Virtual disk added or deleted
- Virtual disk enlarged
- Each write I/O to a virtual disk that has demand-allocated snapshot

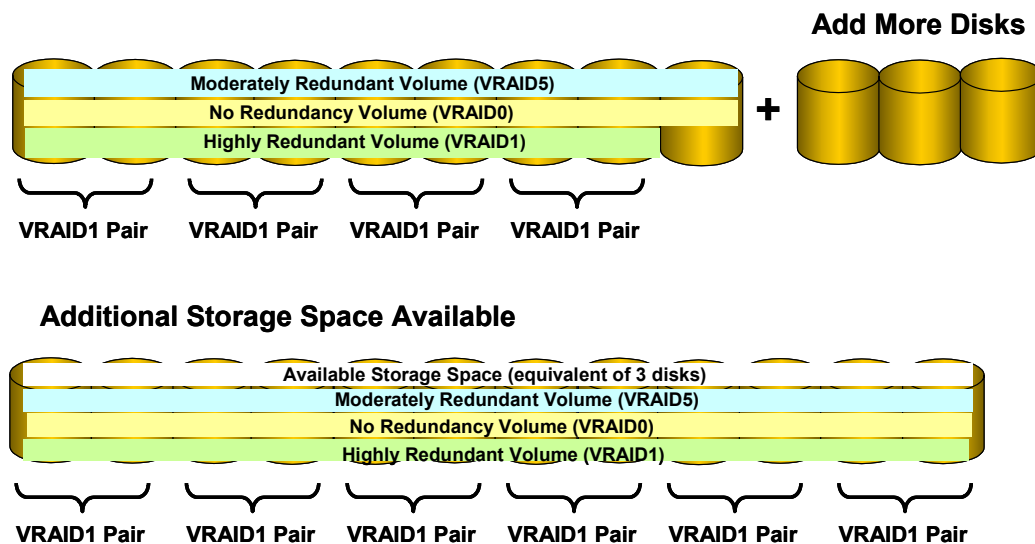
Virtual Disk Leveling

Leveling is the distribution of all data within virtual disks across all physical disks within the disk group. The Enterprise Virtual Array uses a data-leveling algorithm to perform leveling and to balance performance and capacity across all disk drives within the disk group. This process is transparent and does not affect ongoing client workloads. The process redistributes each virtual disk's blocks proportionally across all spindles in that disk group.

The leveling process activates whenever the Enterprise Virtual Array detects an opportunity to improve utilization, such as a change in the number of disks in the pool. In particular, leveling takes place when:

- Disk devices are added to a disk group. You can add them one or more at a time.
- Disk devices are ungrouped from a disk group. You can only do this one disk at a time.
- Disk devices within a disk group fail (assuming VRAID survival).

The following depicts the addition of three disks.



Leveling is done proportionally to drive capacity. When group members vary in capacity (36GB, 72GB, and so on), proportional amounts of drive capacities are used during virtual disk creation.

Example 1

Disk group = 10 drives, all 36GB

If a virtual disk is created using all of the drives, each 36GB disk contains an equal share (10%) of the capacity.

Example 2

Disk group = 10 drives, 5*36GB, 5*72GB

If a virtual disk is created using all of the drives:

- Each 36GB and 72GB disk contains a proportional share of the capacity.
- The proportion maintained in this example must be 2:1—that is, twice as much 72GB drive capacity must be used as 36GB drive capacity (360GB of 72GB drive capacity out of a total of 540GB, and 180GB of 36GB drive capacity out of a total of 540GB).

Proportional amounts of disk capacity are used whatever virtual disk capacity is created.

Note

With regard to example 2:

If a virtual disk of 270GB is created, the proportion of 2:1 must be maintained. The capacity used for all 36GB must still be one third of the total (90GB) and the capacity of all 72GB drives must be two-thirds of the total (180). So 18GB of each 36GB drive is used and 36GB of each 72GB drive is used.

Virtual Disk Expansion

You should be aware of the following with regard to virtual disk expansion:

- Once a virtual disk is created, you can increase the size, but not decrease it.
- It is supported on Sun Solaris with the `growfs` command.
- It is supported on Windows 2000.
- It is not supported on all other operating systems but there are workarounds. See the installation and configurations guides for additional information.

Distributed Sparing

Distributed sparing refers to the allocation of space per disk group to recover from physical disk group member failure. The Enterprise Virtual Array distributes capacity across all members of a disk group. There is no spare spindle; you have either assigned capacity or unassigned capacity. The unassigned capacity is spare capacity.

The system draws on unassigned capacity as needed for items such as:

- Writes after Virtually Instantaneous Snapclones (Vsnap)
- New virtual disk creation
- Freeing a physical disk for removal or reassignment
- Data reconstruction after disk failure

Unassigned capacity increases when virtual disks are deleted or new physical disks are added to the disk group.

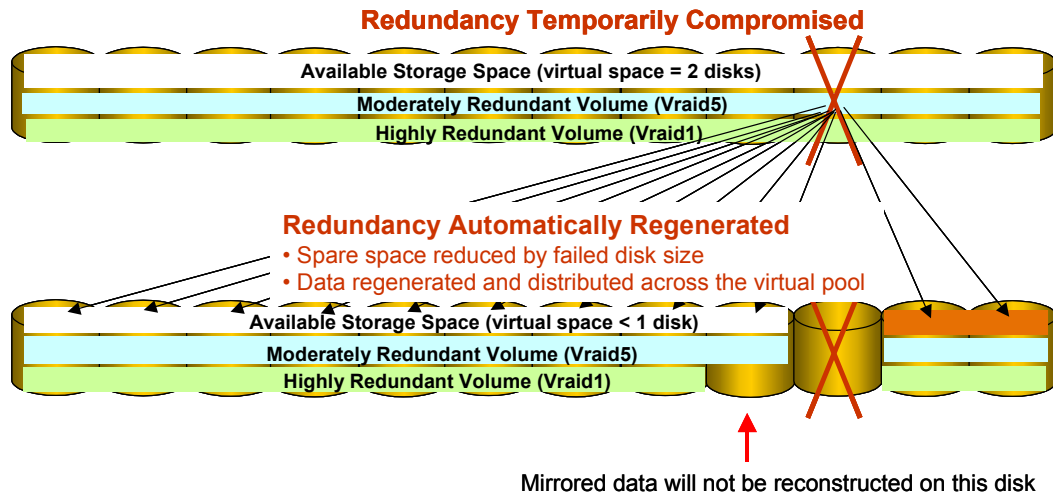
You reserve spare capacity by selecting a disk protection level for the disk group. The disk group protection levels are:

- None, which reserves no physical disk capacity
- Single, which reserves one single disk for spare capacity, and requires the use of two spindles to save enough space to recover a VRAID1 pair
- Double, which reserves two disks for spare capacity, and requires the use of four spindles to save enough space to recover two VRAID1 pairs

The protection levels do not assign specific capacity and do not change the amount of unassigned capacity, but ensure that new virtual disk creations and Vsnap writes leave sufficient free space for reconstruction. The protection levels are not structurally important if you already have sufficient free capacity.

The capacity chunks are allocated on all disk drives to survive one or two disk drive failures, depending on whether you selected moderate or high redundancy levels.

VCS uses unassigned space first, then spare space. For VRAID1 data, data is moved off the surviving disk, creating an unpaired disk (this is shown in the diagram). The unpaired disk can pair up with a newly-grouped disk.



How a disk group uses spare capacity depends on the following factors:

- The number of members of the disk group
- Whether the disk group contains disks with varying capacities
- The types, sizes, and combinations of virtual disks that are carved from the disk group
- Whether there is adequate free space in the disk group to use in place of the spare space

Examples with Equal-Size Disk Capacities

Computing maximum virtual disk sizes with same-sized disk capacities is straightforward; multiply the number of disks by the capacity, then subtract the number of spindles required for the protection level.

Note

The capacities given in the examples are close to real life; they account for disk formatting, disk group metadata overhead, and virtual disk metadata overhead. For example, a 36GB disk, when formatted, yields approximately 33.91GB of usable capacity. Even less is available after creating disk groups and virtual disks.

Example 1

A disk group of 10 physical disks at 36GB each, with a protection level of None, has a maximum virtual disk size of 338GB.

Example 2

A disk group of 10 physical disks at 36GB each, with a protection level of Single, has a maximum virtual disk size of 270GB (338 – 68).

Example 3

A disk group of 10 physical disks at 36GB each, with a protection level of Double, has a maximum virtual disk size of 202GB (338 – 136).

Example with Mixed-Size Disk Capacities

When computing maximum virtual disk sizes for disk drives of mixed sizes, you must use the following more specialized rules:

- Single protection level for a disk group requires two times the **largest** drive capacity of **any** drive in the group.
- Double protection level for a disk group requires four times the **largest** drive capacity of any drive in the group.

Example

If a disk group has ten 36GB disks and one 72GB disk, and you need double protection, the space of four 72GB drives are set aside for spare capacity. The maximum virtual disk size would be 134GB (406GB – 272GB).

Note

Module 12, Best Practices, has details on how to determine disk counts based on redundancy and protection level requirements.

Redundant Storage Sets

An RSS is a subgrouping of drives within a disk group to allow failure separation. User data is segmented into 8MB redundant stores (RStores). Physical data is allocated for the redundant store based on the VRAID type—that is, 8MB for VRAID0, 10MB for VRAID5, and 16MB for VRAID1. The physical storage for a given redundant store is fully contained within a single RSS. Relocation of data may occur but it is always within a single RSS. Data may be completely reallocated or copied to a different RSS in the disk group.

The purpose of an RSS is to reduce the risk that a multidisk failure will result in the loss of customer data, a condition commonly called meltdown but identified in Command View EVA as an Unresolved Disk Group Hardware Problem. This condition occurs when an excessive number of disks have either failed or been removed from a disk group without using Command View EVA, which performs data migration prior to disk removal.

An inoperative disk group affects all the virtual disks of a given VRAID type within a disk group when drive loss within the disk group occurs as shown in the following table.

VRAID Type	Description	Inoperative Disk Group Criterion
VRAID0	Striping with no redundancy	Loss of one disk in the disk group
VRAID5	Striped parity (4+1 data/ parity model)	Loss of two disks in the same RSS
VRAID1	Mirrored data	Loss of two disks in the same RSS index pair. An index pair consists of disks with RSS indices of n and $n+1$, where n is zero or even (for example, indices would be 0 and 1, 2 and 3, 4 and 5, and so on).

Note

Each disk group is divided into one or more RSSs. Virtual disks are composed of segments that are allocated from RSSs to help reduce the risk of data loss. Disks within each RSS are identified by a zero-based number called an RSS index. The RSS and RSS index are displayed by Command View EVA V3 and higher.

The normal range for an RSS is 6 to 11 physical disks, but the desired number is 8 disks. An RSS always has an allocation of an even number of disks except for the remaining disks in the disk group. This supports VRAID1 pairing. If an RSS falls to fewer than six members, an attempt is made to merge it with another RSS.

The rules for RSS creation are as follows:

- Create 8-member RSSs until there are less than 16 drives.
- If the remaining drives are greater than 11, create two evenly split RSSs if possible. If the remaining drives are less than 11, create one RSS.

Examples of RSS Calculations

The following are some examples or RSS calculations.

Example 1

A 38-disk disk group would consist of three 8-member RSSs and two 7-member RSSs ($8+8+8+7+7=38$).

Example 2

A 50-member disk group would contain five 8-member RSSs and one 10-member RSS ($8+8+8+8+8+10$).

- This would allow up to six drive failures before affecting a VRAID5 virtual disk
- Could allow up to 25 drive failures before affecting a VRAID1 virtual disk, but chances are that metadata would have failed.

Note

If a disk group member is removed (or fails), the RSSs are adjusted to ensure that each RSS has at least six members. This adjustment results in a re-leveling of the virtual disk's data.

Virtual Disk Space Allocation

For all VRAID levels, space is allocated across all members of the disk group. The following describes behaviors for each VRAID type:

- VRAID0
 - Failure of any disk group member results in a virtual disk failure.
 - Spare space allocated to the disk group does **not** help you.
- VRAID1
 - Within each RSS, each disk is married to one other disk. Disks of different sizes are not married together.
 - An RSS with an odd number of members (for example, a seven-member RSS) would have one disk unpaired. This means one disk would **not** have any data for the VRAID1 virtual disk written to it.
 - HP recommends that disk groups have an even number of members.
- VRAID5
 - Within each RSS, data is written in a 4+1 parity scheme.
 - Odd or even numbered RSS memberships do not affect the writing scheme.

Refer to Appendix C for details and issues regarding disk allocation, metadata, disk failure, and occupancy level.

Disk Group Member Failures with VRAID1

If a disk group member failure occurs with a VRAID1 virtual disk, and assuming the disk group has spare space allocated **or** the disk group has enough available free space:

- Two copies of the failed disk are reconstructed into the free space of the disk group, if it is available.
 - The Enterprise Virtual Array uses the spare space if there is insufficient free space.
- Once the data has been reconstructed, RSSs for the disk group are recalculated, and the data is re-leveled across the remaining members of the disk group.
 - For odd numbered disk groups, one disk is not used.
- The re-leveling may take many hours to complete.
- If during the reconstruct the unpaired disk were to fail, the VRAID1 virtual disk would be lost.
- If, during the re-leveling, another disk fails (not the unpaired disk):
 - With adequate disk group free space, or another allocated spare space—the above process would repeat itself.
 - Without adequate disk group free space, or another allocated spare space—the process outlined next would take place.

Assuming no spare space is allocated **or** no free disk group space is available:

- A failed disk group member cannot be reconstructed (no place to put the reconstructed data).
- RSSs are **not** recalculated.
- Failure of another disk group member:
 - In another RSS, does not disable the VRAID1.
 - In the same RSS, only destroys the VRAID1 if it happens to be the unpaired disk of the original lost disk.

When one or more disk group members are added back into the disk group:

- Data is mirrored from data on the unpaired disk to new disk.
- RSSs are recalculated.
- Data is re-leveled across the disk group.

Disk Group Member Failures with VRAID5

If a disk group member failure occurs with a VRAID5 virtual disk, and assuming the disk group has spare space allocated **or** the disk group has enough available free space:

- Data from the failed disk is reconstructed into the free space.
 - The Enterprise Virtual Array uses the spare space if there is insufficient free space.
 - The Enterprise Virtual Array chooses where to put the reconstructed data based on what will maintain the highest level of redundancy.
- Once the data has been reconstructed, RSSs for the disk group are recalculated, and the data is re-leveled across all remaining members of the disk group.
 - The leveling may take many hours to complete.
- If during the reconstruct another member of the same RSS were to fail, in most circumstances, the VRAID5 virtual disk would be lost.
- If, during the re-leveling, another member of the same RSS were to fail:
 - With adequate disk group free space, or another allocated spare space—the above process would repeat itself.
 - Without adequate disk group free space, or another allocated spare space—the process outlined next would take place.

Assuming no spare space is allocated **or** no free disk group space is available:

- A failed disk group member cannot be reconstructed (no place to put the reconstructed data).
- RSSs are **not** recalculated.
- Failure of another disk group member in another RSS does not disable the VRAID5.
- Failure of another disk group member within the same RSS will most likely destroy the VRAID5.

When one or more disk group members are added back into the disk group:

- Data is reconstructed from remaining RSS members.
- RSSs are recalculated.
- Data is re-leveled across the disk group.

Hosts and LUNs

A host is a collection of host bus adapters that reside in the same (virtual) server. A LUN is a virtual disk presented to one or multiple hosts.

The following are Enterprise Virtual Array rules for hosts, connections, and LUNs:

- A maximum of 256 hosts
- A maximum of 1,024 Fibre Channel adapter connections
- A maximum of 256 LUNs on any one Fibre Channel adapter
- A maximum of 8,192 presentations of LUNs to hosts

Examples

1 LUN presented to 1 host = 1 presentation

1 LUN presented to 256 hosts = 256 presentations

256 LUNs presented to 1 host each = 256 presentations

Snapshot Implementation

Snapshot and snapclone technology is implemented through an optional license for Business Copy.

Note

There is a difference between HP snapshots and EVA snapshots.

Snapshots are created from a virtual disk on demand—that is, as changes are made to the original volume. Up to seven snapshots can be created from the virtual disk. You cannot create a snapshot from a snapshot.




Important

All snapshots within a virtual disk family must have the same allocation policy; for example, if one snapshot is created using allocation on demand, all snapshots must use allocation on demand.

Fully-Allocated Snapshots

Fully-allocated (traditional) snapshot operations are similar to the HSG80 controller-based snapshot. At the moment of snapshot creation, chunks are reserved in the disk group. The snapshot remains available if the disk group becomes full because the snapshot capacity is reserved.



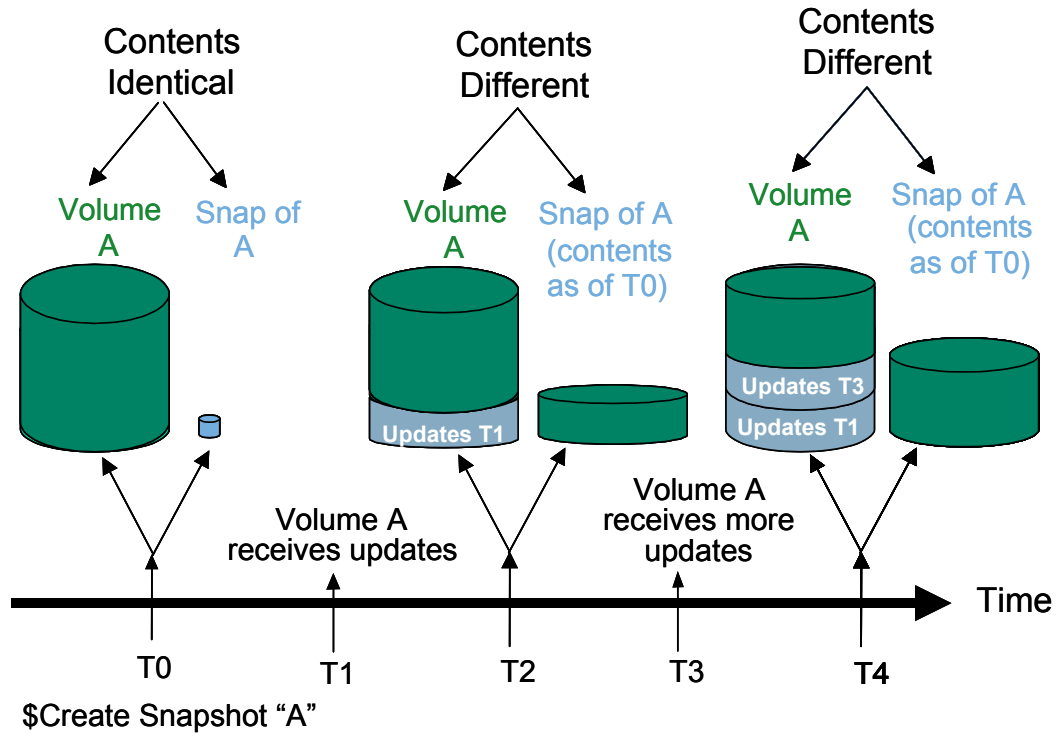
Fully-Allocated Snapshot Creation and Utilization

Note

VCS supports up to seven snapshots/Vsnaps per LUN of a VRAID type and supports re-snap out of sequence. You cannot do snapshots of snapshots.

Demand-Allocated Snapshots

Demand-allocated (Virtually Capacity-Free) snapshot operations allocate chunks in the disk group on demand. If the disk group becomes full, the snapshot is marked as overcommitted and put into an inoperative state. A warning is sent as capacity reaches its threshold (occupancy level).



Demand-Allocated Snapshot Creation and Utilization

Snapclone Implementation

Snapclone (Virtually Instantaneous snapclone) operation performs a full background copy from the original virtual disk. Chunks are allocated in the disk group at the moment of snapclone creation. Once the background copy is complete, it severs the sharing relationship with the virtual disk and becomes its own entity. The clone can then be set to a different preferred controller than the original virtual disk setting after the cloning process is complete.

At creation time, the snapclone is available as a fully-allocated snapshot. Once the background process is finished, it becomes a clone.

Note

VCS supports snapclones across disk groups. A wizard in Command View EVA allows you to choose in which group to place the snapclone.

Snapshot and Snapclone Comparison

The following table compares the snapshot and snapclone operations.

Operation	Description	Pros	Cons
Demand-allocated snapshot	<ul style="list-style-type: none">■ Pointer-based copy before write■ Allocate space on demand	Space efficient (allocated on demand)	Overcommit problem
Fully-allocated snapshot	<ul style="list-style-type: none">■ Pointer-based copy before write■ Reserve space on creation	Overcommit problem is reduced	Space inefficient (reserved right away)
Snapclone	Same as the fully-allocated snapshot but with a background copy process to separate the virtual disk	<ul style="list-style-type: none">■ Overcommit problem is reduced■ Repeatable, separate virtual disks	<ul style="list-style-type: none">■ Space inefficient■ Consumes some background process time

Snapshot and snapclone creation are discussed in Module 8 of this course.

Learning Check

1. A _____ is an initialized pair of HSV controllers with a minimum of eight physical drives.
.....
2. For a disk group, space is allocated in _____ segments and the chunk size is _____ blocks.
.....
3. By default, the Enterprise Virtual Array configures a disk group in a _____ configuration.
.....
4. Describe the virtual disk redundancy levels: None, Moderate, and High. Include VRAID level and amount of disk group capacity consumed in percent.
.....
.....
.....
5. Label each of the following characteristics as Conventional RAID or Virtual RAID.
 - a. Performance is limited by the number of disk drives in a disk group.
 - b. I/Os are balanced across the storageset.
 - c. Controller eliminates the load balancing.
 - d. Possible to find customer data if you know the logical block number (LBN) and chunk size.
6. Complete the sentence. The Enterprise Virtual Array balances performance and capacity across all disk drives by a process called _____.
.....
7. Describe the disk failure protection levels.
.....
.....
.....
8. Describe the three types of data replication used by the Enterprise Virtual Array.
.....
.....
.....

Overview

This module gets you started with Command View EVA by describing the user interface and explaining how to access online help. It then details Command View EVA functions that enable you to configure the Enterprise Virtual Array to present virtual disks to one or more hosts in a SAN environment.

This module includes a description of storage system initialization and code loads, followed by a discussion of folders, and the creation and modification of disk groups, virtual disks, and hosts. You are shown how to present and unpresent LUNs to a host.

Finally, this module explains how to copy active virtual disks to create fully-allocated snapshots, demand-allocated snapshots, and snapclones.

Note

This module applies to Command View EVA V3.0, found in the HP OpenView Storage Management Appliance Software V2.0 Service Pack 3 Install CD.

Objectives

After completing this module, you should be able to:

- Identify the Navigation, Session, and Content panes in Command View EVA.
- Describe how to access system, page, and field help.
- Describe how to locate the Command View EVA version number and home page.
- List the steps in the configuration process.
- Describe how to use Command View EVA to initiate and complete the configuration process steps.
- Identify the properties of disk groups, virtual disks, and hosts.
- Describe how to create disk groups, virtual disks, hosts, snapshots, and snapclones.

Getting Started with Command View EVA

Command View EVA is:

- The GUI through which you can control and monitor the storage system.
- The software that supports the GUI.

Each installation of Command View EVA on the storage management appliance is a storage management agent. The client for the agent is a standard Internet browser.


Command View EVA Interface Layout

Command View EVA is comprised of three main panes:

- Session pane
- Navigation pane
- Content pane

Session Pane

The Session pane (top pane) provides information about the appliance name and IP address, and provides menu options for Command View EVA functions.




The Session pane provides the following menu options:

- **Root View** — Displays the HSV Storage Network Properties page
- **Agent Options** — Selects interface and management agent options
- **Help** — Displays online help in a new window

Navigation Pane

The Navigation pane (left pane) provides a directory tree structure for access to virtual disks, hosts, disk groups, data replication entities (new for VCS V3.0), and hardware in the Enterprise Virtual Array storage system. Click on the top-level folder to expand the tree structure to view folders inside it.



You use folders to organize objects for easy retrieval. Each folder or object has a corresponding Properties page. When you select the folder or object in the Navigation pane, the Properties page displays.

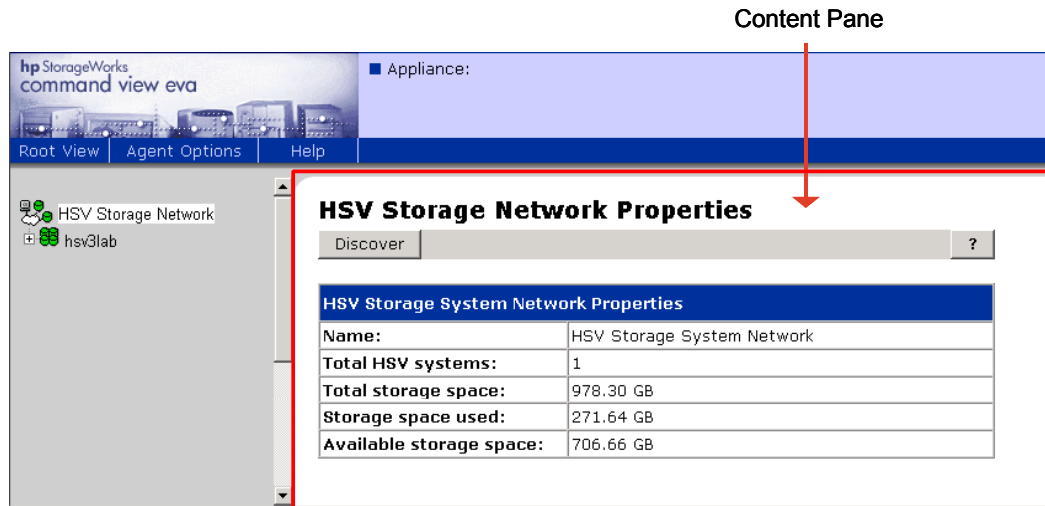
To customize the organization of components, you can create folders within the Virtual Disks and Hosts folders. You can drill down into any folder to get specific information.

Note

Use the “+” symbol next to a folder to expand it without updating the Content pane.

Content Pane

The Content pane (right pane) is the main body of the page where you do most of your work. The type of information that displays in the content pane depends upon what you select in the Navigation pane. The Content pane displays the properties of such objects as disk groups, physical disk drive, virtual disks, or hosts.



The Content pane also contains the question mark (?) button, which displays help for that page. Additional question mark (?) buttons provide field-level help.

Data Entry

The Command View EVA GUI provides several options for entering data. These include the following:

- **Text boxes** — Enter text in a standard format.
- **Drop-down lists** — Select an item from a list of choices.
- **Selection boxes** — Click a box to make a selection.
- **Comment boxes** — Enter text (boxes) up to 128 characters. You can enter useful comments that you want to keep track of in a central place; for example:
 - Hardware locations
 - Installation and configuration dates
 - Past failure information
 - Departmental ownership

Buttons

Some buttons you see frequently in the Content pane when performing tasks include the following:

- **Finish** — Completes a task.
- **Advanced options** — Presents more characteristics for you to specify. Most of the time, these are optional. If not specified, default values are used.
- **Next step** — Displays the next step/next page.
- **Previous step** — Displays the previous step/last page.
- **Save changes** — Saves the changes made on the current page.
- **Cancel** — Cancels an action.

Wizards

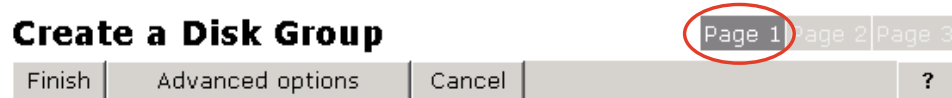
Command View EVA provides wizards to guide you through many processes, such as initializing your storage system, creating disk groups, and creating virtual disk families. Launch a wizard by clicking a button in the Properties page of a folder or object.

The wizard provides step-by-step instructions organized in pages. In the right corner of the Content pane, the wizard displays the number of pages and highlights the page currently displayed at the top right.

Example

This example is the first page of the Create a Disk Group wizard. Note the page numbers at the top right.

Create a Disk Group



Complete this step and click **Finish** to create a disk group in the simplest way possible. If you'd like more control over the creation of your disk group, complete the step and click **Adv Options** instead.

STEP 1: Enter a Name

Enter a name for your disk group.



Access online help from the wizard to obtain further instructions.

Note

It is a user interface option whether to use wizards for virtual disk creation. See Module 9 for an example of how this is done.

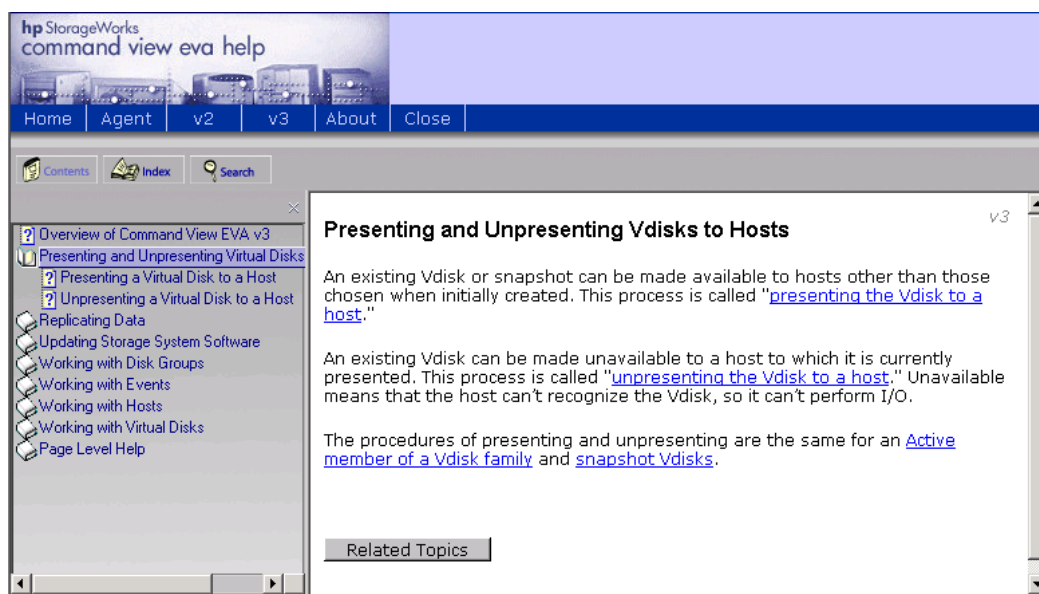
Command View EVA Online Help

Command View EVA has an extensive online help system that includes the following:

- Application help
- Page help
- Field help

Application Help

Application help lets you navigate by a table of contents, index, and keyword search for system wide information. You can access system help by clicking *Help* on the Session pane and selecting the Command View version you are using. You can also determine the Command View EVA version by selecting *About*.

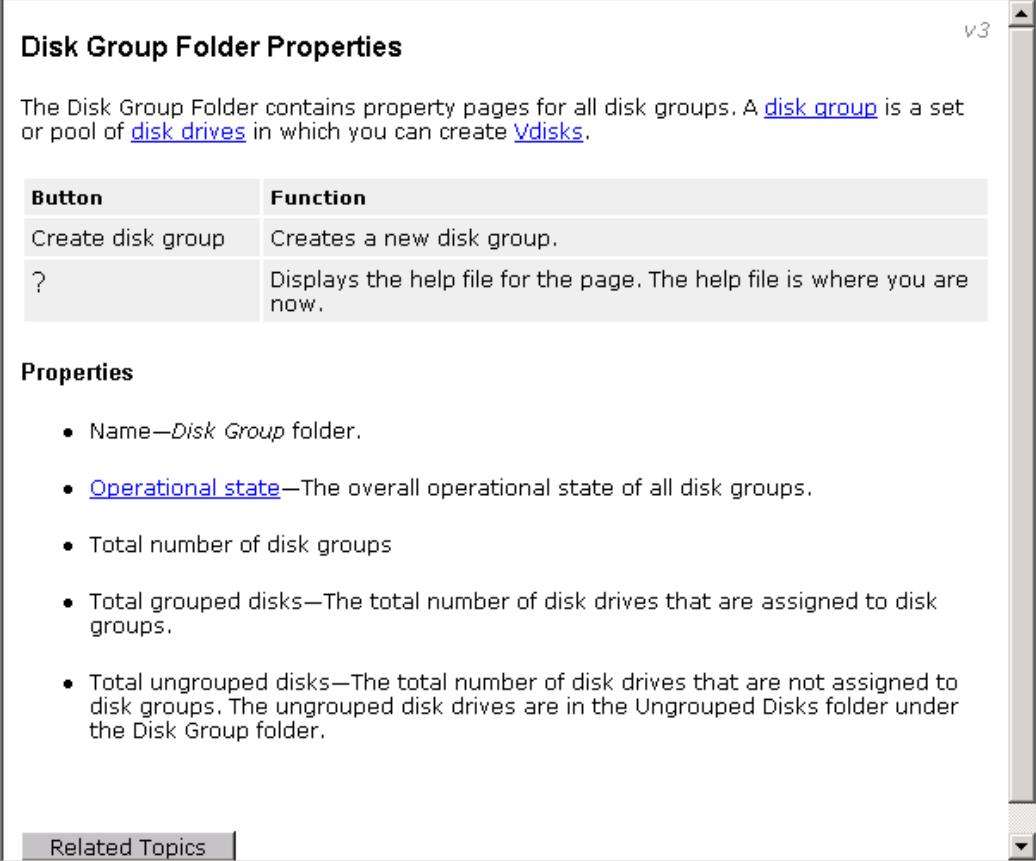


Page Help

Page help displays help for the page displayed in the Content pane. Access this online help by clicking the question mark (?) in the Content pane.

Example

Here, the question mark was pressed on the Disk Group Folder Properties page.



The screenshot shows a help window titled "Disk Group Folder Properties" with a version number "v3" in the top right corner. The window contains a description of the Disk Group Folder, a table of buttons and their functions, a list of properties, and a "Related Topics" section at the bottom.

Disk Group Folder Properties v3

The Disk Group Folder contains property pages for all disk groups. A [disk group](#) is a set or pool of [disk drives](#) in which you can create [vdisk](#)s.

Button	Function
Create disk group	Creates a new disk group.
?	Displays the help file for the page. The help file is where you are now.

Properties

- Name—*Disk Group* folder.
- [Operational state](#)—The overall operational state of all disk groups.
- Total number of disk groups
- Total grouped disks—The total number of disk drives that are assigned to disk groups.
- Total ungrouped disks—The total number of disk drives that are not assigned to disk groups. The ungrouped disk drives are in the Ungrouped Disks folder under the Disk Group folder.

Related Topics

Field Help

Field help provides field descriptions and describes acceptable data entry formats. Access this online help by clicking the question mark (?) next to the field.

Example

This example shows field help for the Enter a Name field, which is Step 1 of the Create a Virtual Disk wizard.

Create a Vdisk Family

Page 1 Page 2 Page 3 Page 4 Page 5

Finish

Advanced options

Cancel

?

Complete these steps and click **Finish** to create your Vdisk family. For more control, complete the steps and click **Adv Options** instead.

STEP 1: Enter a Name

Vdisk001	
----------	---

The following is a portion of the field level help.

Virtual Disk Name	v3
The Vdisk name is assigned to a Vdisk or to a Vdisk family. The maximum length is 32 characters. Names can contain any characters, except for the following:	
<ul style="list-style-type: none">• ? (question mark)• " (double quotes)• / (slash)	

Command View EVA Home Page

When you click the *command view eva* button from the **hp OpenView storage management appliance**, **Devices** page, the HSV Storage Network Properties page displays. This is the same page as you get if you click *Root View*. This page shows what the Command View EVA management agent has discovered for HSV storage systems on the SAN.

Example

This example shows the discovery of three HSV storage systems. All storage systems initialized.

The screenshot shows the 'hp StorageWorks command view eva' interface. The top navigation bar includes 'Root View', 'Agent Options', and 'Help'. The main content area is titled 'HSV Storage Network Properties' and features a 'Discover' button and a help icon. A table displays the following properties:

HSV Storage System Network Properties	
Name:	HSV Storage System Network
Total HSV systems:	1
Total storage space:	978.30 GB
Storage space used:	271.64 GB
Available storage space:	706.66 GB

Clicking the *Discover* button updates the display with any new discoveries.

Configuration Process Steps

Configuring the storage system and host requires a sequence of steps:

1. Gather preliminary information
 - a. Develop a requirements list per server.
 - b. Create a list of all Fibre Channel adapters (FCAs).
 - c. Determine the Host WWN.
2. Perform a software code load (if necessary).
3. Initialize the storage system.
4. Perform a disk drive firmware code load (if necessary).
5. Create additional disk groups as desired.
6. Create the hosts.
7. Create the virtual disks.
8. Create snapshots and snapclones.

These process steps are discussed in more detail later in this module. Depending on your configuration, there may also be a requirement to perform data replication tasks. These tasks and screens are described in the Continuous Access EVA documentation and in the Services training course for Continuous Access EVA.

Gathering Preliminary Information

The setup procedure described in Module 6 details the specific documents you need to get started with the Enterprise Virtual Array. The following topics describe other considerations.

Develop a Requirements List Per Server

For each server determine the following virtual disk characteristics:

- Size (in GB)
- Virtual RAID level (VRAID0, VRAID5, and VRAID1)
- Operating system LUN identifier (used for OpenVMS and Tru64 Unix)

After you have identified the virtual disk characteristics, determine which virtual disks can share a disk group.

Create a List of all FCAs

To add a host to a storage system through Command View EVA, you must enter certain information to identify the server and the communication path to the controller. The information includes the following:

- WWID of FCAs in server
- Which FC switch port
- Operating system on server

Determine the Host WWN

Use the following to determine the host WWN name for each operating system:

- OpenVMS (V7.21H1 or 7.3)
 - \$ Analyze/system
 - SDA> FC SHOW DEVICE FGA0 (FGB0, and so on)
 - Alpha SRM console
 - ◆ SHOW DEVICE
- Windows NT/2000
 - LPUTILNT (LightPulse Utility/NT)
- Tru64 Unix
 - emxmgr (select 1)
 - /var/adm/messages
 - Alpha SRM console
 - ◆ SHOW DEVICE
- SUN Solaris
 - /var/adm/messages (only if adapter driver is loaded)
- IBM AIX
 - lsconfig
- HP-UX
 - fcmsutil or tdutil

Starting with the Uninitialized Storage System

If you select *Uninitialized Storage System* in the Navigation pane under HSV Storage Network, the Uninitialized HSV Storage System Properties page displays.

Example

This example shows the properties of an uninitialized storage system. Note the buttons at the top of the Content pane.

Uninitialized HSV Storage System Properties																																									
Initialize Set time View events Code load Shut down ?																																									
<table border="1"> <tr> <th colspan="2">Identification</th> <th colspan="2">Condition/State</th> </tr> <tr> <td>Name:</td> <td>Uninitialized Storage System</td> <td>Operational state:</td> <td><input checked="" type="checkbox"/> Attention (Uninitialized)</td> </tr> <tr> <td>Node World Wide Name:</td> <td>5000-1FE1-5000-2CD0</td> <td colspan="2">System</td> </tr> <tr> <td>UUID:</td> <td>5005-08b4-0001-4523-0000-0000-0000-0000</td> <td>Type:</td> <td>HSV110</td> </tr> <tr> <td colspan="2">Licensed features</td> <td>Version:</td> <td>3000</td> </tr> <tr> <td>Basic:</td> <td>Yes</td> <td>Time:</td> <td>19 Apr 2003 09:01:10</td> </tr> <tr> <td>Snapshot:</td> <td>Yes</td> <td colspan="2"></td> </tr> <tr> <td>Data replication:</td> <td>Yes</td> <td colspan="2"></td> </tr> <tr> <td colspan="4">Comments</td> </tr> <tr> <td colspan="4"> <p>***** THIS SYSTEM IS UNINITIALIZED *****</p> <p>Click the Initialize button to prepare the system for data storage. If you have not yet entered a license key for this system, you will be prompted to enter it.</p> </td> </tr> </table>		Identification		Condition/State		Name:	Uninitialized Storage System	Operational state:	<input checked="" type="checkbox"/> Attention (Uninitialized)	Node World Wide Name:	5000-1FE1-5000-2CD0	System		UUID:	5005-08b4-0001-4523-0000-0000-0000-0000	Type:	HSV110	Licensed features		Version:	3000	Basic:	Yes	Time:	19 Apr 2003 09:01:10	Snapshot:	Yes			Data replication:	Yes			Comments				<p>***** THIS SYSTEM IS UNINITIALIZED *****</p> <p>Click the Initialize button to prepare the system for data storage. If you have not yet entered a license key for this system, you will be prompted to enter it.</p>			
Identification		Condition/State																																							
Name:	Uninitialized Storage System	Operational state:	<input checked="" type="checkbox"/> Attention (Uninitialized)																																						
Node World Wide Name:	5000-1FE1-5000-2CD0	System																																							
UUID:	5005-08b4-0001-4523-0000-0000-0000-0000	Type:	HSV110																																						
Licensed features		Version:	3000																																						
Basic:	Yes	Time:	19 Apr 2003 09:01:10																																						
Snapshot:	Yes																																								
Data replication:	Yes																																								
Comments																																									
<p>***** THIS SYSTEM IS UNINITIALIZED *****</p> <p>Click the Initialize button to prepare the system for data storage. If you have not yet entered a license key for this system, you will be prompted to enter it.</p>																																									

The Uninitialized HSV Storage System Properties page displays the following information:

- **System Node World Wide Name** — The node identifier for the system. A unique, 16-character hexadecimal number assigned to the system.
- **System UUID** (universal identifier) — Unique universal identifier for each object in the storage system, whether hardware element or software structure.
- **Licensed features** — Status of features that are enabled including basic, snapshot, and data replication.
- **Operational state** — The operational state of the HSV storage system. The possibilities include good, attention, failed, and operation in progress.
- **Version** — System firmware version.
- **Controller time** — Usually set to the same time as the storage management appliance.

The buttons in the Content pane enable you to launch the following functions:

- **Initialize** — Initialize your Enterprise Virtual Array.
- **Set time** — Set the system time base.
- **View events** — View event logs that include errors. For an uninitialized storage system, you can only view Management Agent Events.
- **Code load** — Download a SuperImage file to your Enterprise Virtual Array.
- **Shut down** — Shut down the storage system.

Performing a Controller Code Load

You can perform a code load on an uninitialized storage system, and on an initialized storage system (for updates). The code load operation loads SuperImage files (.sss) to the appliance and then downloads the files to the Enterprise Virtual Array controllers and EMU. The files are loaded from a browser on a client machine.

When performing a code load you should be aware of the following:

- The controllers typically perform a fast reboot once VCS is updated.
- The storage system firmware file must be a properly formatted, supported firmware release file available on the CD shipped with the Enterprise Virtual Array.
- You can only download the SuperImage file (.sss) to a single subsystem. If you have multiple subsystems, you must repeat the code load operation for each one.
- Any error incurred with the SuperImage load results in an abort of the entire SuperImage processing.
- Code cannot be removed from any component once it is installed.
- The SuperImage load requires a valid Basic license key matching the VCS functional version number.

You can also perform a code load to update firmware after the storage system is initialized. Command View EVA provides a Code Load wizard to guide you through the operation.


Refer to Service documentation for information on controller compatibility and how to upgrade HSV controllers to the current version.

Code Load Initial Prompt


To initiate the Code Load wizard, click *Code load* in the Uninitialized HSV Storage System Properties page. If the storage system is already initialized, you can start the Code Load wizard from various Storage System Properties pages.

After clicking *Code load* you are prompted with a message as shown.

Uninitialized HSV Storage System Properties

Initialize	Set time	View events	Code load	Shut down	?												
Identification Name: Uninitialized Storage			Condition/State Operational state:  Attention (Uninitialized)														
Licensed features <table border="1"> <tr> <td>Basic:</td> <td>Yes</td> </tr> <tr> <td>Snapshot:</td> <td>Yes</td> </tr> <tr> <td>Data replication:</td> <td>Yes</td> </tr> </table>			Basic:	Yes	Snapshot:	Yes	Data replication:	Yes	<table border="1"> <tr> <td>n</td> <td>HSV110</td> </tr> <tr> <td>n:</td> <td>3000</td> </tr> <tr> <td></td> <td>09 Apr 2003 07:12:54</td> </tr> </table>			n	HSV110	n:	3000		09 Apr 2003 07:12:54
Basic:	Yes																
Snapshot:	Yes																
Data replication:	Yes																
n	HSV110																
n:	3000																
	09 Apr 2003 07:12:54																
Comments ***** THIS SYSTEM IS UNINITIALIZED ***** Click the Initialize button to prepare the system for data storage. If you have not yet entered a license key for this system, you will be prompted to enter it.																	

Microsoft Internet Explorer

 This action will update the firmware in BOTH of your controllers and your other major system components. IT WILL ALSO RESTART BOTH CONTROLLERS. Are you sure you wish to continue?

OK Cancel

By clicking on *OK*, you can begin the controller code load.

Performing a Code Load

To code load the controllers, you must first select a firmware image file by manually entering the entire path or browsing to it. If you manually enter the path, it **must be in Windows format**.

Example

This is page 1 of the Code Load wizard, titled Code Load Storage System. The path was entered by browsing to the image.

Code Load Storage System Page 1 Page 2 Page 3

Next step Cancel ?

Complete the following steps to update the operating code in your storage system.

STEP 1: Select a firmware image
Enter the complete path to your storage-system firmware image file.

cs\wcs_v3_w030414.sss Browse... ?

STEP 2: Upload your firmware image file

Click the **Next Step** button to upload the firmware image file to your management appliance.

The buttons in the Content pane let you continue with this operation or cancel it as described in the following choices:

- **Next step** — Downloads the firmware image and moves to the next step (page 2) of the download procedure. This uploads the firmware image into the SMA on which the management agent is running, but does **not** install the code into the storage system.
- **Cancel** — Cancels the download process.

The second page of the wizard gives the following information about what the code load does and recommended reading.

Code Load Storage System

Page 1 Page 2 Page 3

Previous step

Next step

Cancel

?

STEP 3: Read the pre-update application notes

Read the pre-update application notes associated with your firmware image. The notes offer important cautions applicable to code update process. If you wish to cancel the code load operation after reading the notes, click the **Cancel** button.

STEP 4: Code Load Your Storage System Firmware

Click the **Next Step** button to update your storage system's firmware using the image file you uploaded to your management appliance.

Pre-Update Application Notes

This code load procedure will bring all components in the HP StorageWorks Enterprise Virtual Array 5000 to Version 3.000.

In order to assure reliable operation of rolling upgrade to version 3.000, HP strongly recommends upgrading to version 2.003 prior to upgrading to v3.000

Please read all HP StorageWorks documentation pertaining to code load before proceeding with this operation.

When you press *Next step*, you get the following prompt before continuing with the code load.

Code Load Storage System

Page 1 Page 2 Page 3

Previous step


Next step

Cancel

?

STEP 3: Read the pre-update application notes

Read the pre-update application notes associated with your firmware image. The notes offer important cautions applicable to code update process. If you wish to cancel the code load operation after reading the notes, click the **Cancel** button.



firmware using the

This code load procedure will bring all components in the HP StorageWorks Enterprise Virtual Array 5000 to Version 3.000.

In order to assure reliable operation of rolling upgrade to version 3.000, HP strongly recommends upgrading to version 2.003 prior to upgrading to v3.000

Please read all HP StorageWorks documentation pertaining to code load before proceeding with this operation.

When you select *OK* to continue, you are taken to Page 3 of the Code Load wizard as shown in the next example.

Code Load Storage System		Page 1	Page 2	Page 3
Finish	Restarting controllers. Time remaining: 0 seconds			



Your system firmware has been successfully updated, and your controllers are restarting.

Step 5: Read the post-update application notes

Read the post-update application notes associated with your firmware image for tips and cautions on the use of your new firmware.

Post-Update Application Notes

All components of the hp Enterprise Virtual Array have been updated to Version 3.000.

The disk enclosure EMUs will continue to be updated until the LD indication is no longer displayed on their LCDs. At that point, their Firmware Version will read 02020070 for each enclosure.

Until then, it is recommended that no operations be attempted that affect EMUs, particularly powering them off with Element Manager or by the PDU.

The **Time remaining** field will steadily decrement until the code load is complete. Controllers will restart when the time remaining is 0 seconds.

Initializing the Storage System

The initialization process makes the Enterprise Virtual Array known to Command View EVA. Initialization binds the controllers together as an operational pair, establishing the preliminary data structures (metadata) on the disk array, and creates the default disk group.

You only need to perform initialization the first time you configure the storage system or if the storage systems is uninitialized.

Metadata

The Enterprise Virtual Array has three levels of metadata:

- Storage system level metadata on quorum disks
- Disk group metadata
- Command View EVA metadata on the Management Logical Disk (MLD)

Quorum Disks

Quorum disk information includes:

- WWN
- Storage system name
- Character map of disk groups and virtual disk members, but not virtual disk chunk mappings

The following describe requirements of the quorum disks:

- A minimum of two quorum disks to boot HSV controllers.
- A minimum of five quorum disks.
- A maximum of 16 quorum disks (one per disk group).
- Storage requirements are pre-allocated at 0.03% on each disk.

Disk Group Metadata

Each disk group has disk group-specific metadata, and each virtual disk has virtual disk-specific metadata. The disk group metadata holds the following disk-level information:

- Disk group characteristics
 - Number of spindles
 - Spare space allocation
- Virtual disk chunk mapping

Command View EVA Metadata

The Management Logical Disk (MLD) provides metadata for Command View EVA. The following are characteristics of the MLD:

- Devoted to the storage management appliance
- Spread across metadata area on all disks in disk groups
- Holds backup copy of current controller configuration
- Holds the following appliance controller events and controller data:
 - Management agent events
 - Event configuration list
 - Trap host list
 - Elements view list
 - License list (for all storage systems)

Initializing an HSV Storage System

To start this wizard, click *Initialize* in the Content pane of the *Uninitialized HSV Storage System*. To initialize the storage system:

1. Enter a name — You are required to enter a name for your HSV storage system, up to a maximum of 20 characters. Display field help for more information on the acceptable format.
2. Enter the number of disks — This is optional and refers to the number of disks to include in the default disk group. Specify a number between 8 and the number of disks on your system, up to a maximum of 240 disks.
3. Click **Finish** — The storage system enters the name and specifies the default group.

Alternatively, you can click *Advanced options* to set other parameters.

4. Click **OK** to confirm you actions.

Example

The example is Page 1 of the wizard, which shows **hsv3lab** assigned as the HSV storage system name with a default disk group of eight disks.

Initialize an HSV Storage System

Page 1 Page 2 Page 3 Page 4

Finish Advanced options Cancel ?

Complete this step and click **Finish** to initialize your HSV storage system in the simplest way possible. If you'd like more control over the initialization of your HSV storage system, complete the step and click **Adv Options** instead.

STEP 1: Enter a Name

Enter a name for your HSV storage system.

 ?

STEP 2: Enter the number of disks

Enter a number of disks between 8 and 43. (You can add more disks later, if you wish.)

 ?

If you selected *Advanced options*, Page 2 appears as follows. Some of the fields for the advanced options are in a later topic.

Initialize an HSV Storage System

Page 1 Page 2 Page 3 Page 4

Previous

Next

Cancel

?

Continue with this step to initialize your HSV storage system using advanced options. Click the **Next Step** button to move to the next page.

STEP 3: Set the system date/time

☒ Use management appliance date/time

19 Apr 2003 09:07:12

?

☐ Use local (browser) date/time

19 Apr 2003 09:07:34

☐ Use existing controller date/time setting:

19 Apr 2003 09:07:12

☐ Use a custom date/time setting

01

 —

Jan

 —

2001

00

 :

00

 :

00

STEP 4: Enter the Console LUN ID

0

?

By clicking on *Next*, you can set protection levels on Page 3 as shown.

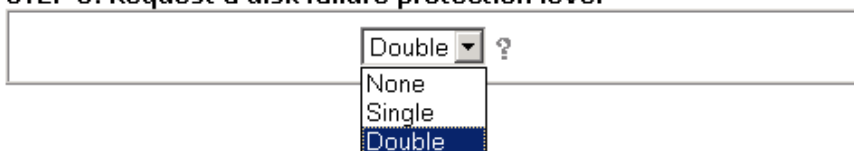
Initialize an HSV Storage System

Page 1 Page 2 Page 3 Page 4

Previous Next Cancel ?

When you initialize your HSV storage system, a default disk group is created. Continue with these steps to specify the attributes of the default disk group. Click the **Next Step** button to move to the next pages.

STEP 5: Request a disk failure protection level



Note

The default protection level is None.

By clicking on *Next*, you can enter comments on Page 4 as shown.

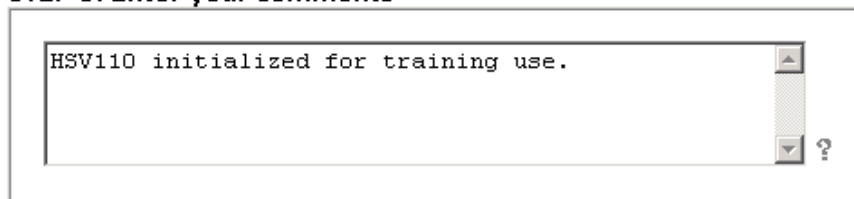
Initialize an HSV Storage System

Page 1 Page 2 Page 3 Page 4

Previous step Finish Cancel ?

Continue with these steps to create the default disk group and initialize your HSV storage system using advanced options.

STEP 6: Enter your comments



STEP 7: Initialize your HSV storage system

Click the **Finish** button to initialize your HSV storage system.

By clicking *Finish*, you receive the warning as follows.

Initialize an HSV Storage System

Page 1 Page 2 Page 3 Page 4


Previous step

Finish

Cancel

?

Continue with these steps to create the default disk group and initialize your HSV storage system using advanced options.



STEP 7: Initialize your HSV storage system

Click the **Finish** button to initialize your HSV storage system.

Click *OK* to complete initialization.

Advanced Options

Advanced options let you specify the characteristics described in this table.

Field	Description
System date and time	Set the system date and time to use the appliance date/time so the Command View EVA event log timestamps are synchronized with the Enterprise Virtual Array.
Console LUN ID	<p>A SCSI-3 virtual object that enables the controller pair to be accessible by the host before any Vdisks are created. The console LUN is what the SCSI-3 specifications call a "communication LUN."</p> <p>Operating systems that require a unique ID for their devices must set this ID by entering a number between 1 and 32767. The default setting is 0 (disabled).</p>
Disk failure protection level	<p>A method by which the controller pair sets aside reserved capacity to take over the functionality of a failed or failing disk drive.</p> <p>Disk groups allocate two spindles for single disk group protection and four spindles for double disk group protection.</p> <p>Select one of the three levels:</p> <p>None</p> <p>Single — Capacity of one physical disk is reserved.</p> <p>Double — Capacity of two physical disks is reserved.</p> <p>The default value is None.</p>

Initialized Storage System Properties Page

After the initialization process completes, verify that the system is initialized by selecting the storage system name in the Navigation pane to display the Properties page.

Example

This is an example of the Initialized Storage System Properties page, slightly different from the Uninitialized HSV Storage System Properties page.

Initialized Storage System Properties

Save changes	Set options	View events	Uninitialize	?
Code load	Shut down			

Identification		Condition/State	
Name:	hsv3lab	Operational state:	<input checked="" type="checkbox"/> Good (Initialized)
Node World Wide Name:			
5000-1FE1-5000-2C00			
UUID:			
6005-08b4-0001-4523-0002-d000-00a6-0000			
Licensed features		System	
Basic:	Yes	Type:	HSV110
Snapshot:	Yes	Version:	3000
Data replication:	Yes	Console LUN ID:	0
		Time:	19 Apr 2003 09:12:46
Policies		Capacity	
Device addition:	Manual	Total:	135.15 GB
Disk replacement delay:	1 mins	Used:	0.00 GB
		Available:	135.15 GB
Comments			
HSV110 initialized for training use.			

Note the additional fields under Policies:

- Disk replacement delay — The amount of time between a drive failure and when the controllers search for empty space to replace it. If you are moving disks around frequently, increase the Disk replacement delay setting.
- Device addition — Governs whether newly added physical disk drives are automatically added to a disk group or require user intervention.

Note the following regarding the total disk capacity:

- The default disk group was chosen as eight drives of 18.2GB. This would be 145.6GB total capacity.
- The 18.2GB is a decimal view of storage capacity. Because the hosts rely on a binary view, you must reduce this amount by about 7%, giving you 135.4GB.
- Accounting for metadata, reduce by another 0.2%, to give you 135.1GB, which matches the sample.

Note

Refer to Module 12, Best Practices, and Appendix C, Space Allocation for more detail about disk capacity calculations and capacity utilization.

Performing a Disk Drive Firmware Code Load

Disk drive firmware is not included in the SuperImage file and must be loaded separately from the controller and storage system firmware. This process loads the files to the appliance and then downloads them to each disk drive in the Enterprise Virtual Array.

You must download the disk drive firmware file (has a **.lod** extension) to one disk drive at a time. The disk drive to receive the code load may be a member of a disk group.

The disk drive is spun down for three minutes during code load. The files are loaded from the browser client machine using a CD provided with the Enterprise Virtual Array storage system. Command View EVA provides a Code Load wizard to guide you through this process.

To launch the wizard for this operation navigate to the Disk Enclosure Bay Properties page by selecting a disk drive bay in the Navigation pane. Click *Code Load* in the Content pane.

Example

Disk Enclosure Bay Properties

Save changes		Code load		Locate		?	
Disk Bay				Disk Drive			
Bay Identification				Condition/State			
Name:	Disk Bay 11			Operational State:	<input checked="" type="checkbox"/> Populated (data disk)		
Bay ID:	11						
Enclosure ID:	10			Location			
Loop Pair:	LoopPair2						
Comments							
<input type="text"/>							


The Code Load wizard steps you through the Disk Drive Firmware Code Load similar to the code load previously discussed.

Creating a Disk Group after Initialization

You can create additional disk groups after you initialize the storage system if the Default Disk Group does not contain all the disks in the array. The Disk Groups folder contains properties pages for all disk groups. When you select the folder, you see Default Disk Group, Ungrouped Disks, and disk groups names (if defined yet).

Example

Disk Group Folder Properties

Create disk group	?
Disk Group Folder Properties	
Name:	Disk Groups
Operational state:	 Good
Total disk groups:	1
Total grouped disks:	8
Total ungrouped disks:	35

This Disk Group Folder Properties page displays the following characteristics:

- **Name** — Disk Group folder
- **Operational state** — Operational state of all disk groups in this folder
- **Total disk groups** — Total number of disk groups in this folder
- **Total grouped disks** — Total number of disks that are assigned to disk groups
- **Total ungrouped disks** — Total number of disks that are not assigned to disk groups

Select the *Create disk group* button in the Content pane to start the Create a Disk Group wizard.

Using the Create a Disk Group Wizard

The only required step to creating a disk group is to assign a name. To create a disk group in the simplest way:

1. Enter a name — You are required to enter a name for your disk group. Display field help for the acceptable format.
2. Click *Finish* to create the disk group and use the default value of eight disks for the disk group.

Alternatively, you can click *Advanced options* to set other parameters.

Example

This example is Page 1 of the wizard that shows **Scratch Disks** assigned as the disk group name.

Create a Disk Group

Page 1 Page 2 Page 3

Finish Advanced options Cancel ?

Complete this step and click **Finish** to create a disk group in the simplest way possible. If you'd like more control over the creation of your disk group, complete the step and click **Adv Options** instead.

STEP 1: Enter a Name

Enter a name for your disk group.

Scratch Disks ?

If you selected *Advanced options*, Page 2 of the wizard displays. Some of the fields for the advanced options are in a later topic.

Create a Disk Group

Page 1 Page 2 Page 3

Previous step

Next step

Cancel

?

Continue with these steps to create a disk group using advanced options. Click the **Next Step** button to move to the next page.

STEP 2: Enter the number of disks

Enter a number of disks between 8 and 10.

 ?

STEP 3: Select a requested disk failure protection level

None ?

None
Single
Double

By clicking on *Next step*, Page 3 of the wizard is displayed.

Create a Disk Group

Page 1 Page 2 Page 3

Previous step

Finish

Cancel

?

Continue with these steps to create your disk group using advanced options.

STEP 4: Enter the occupancy alarm level

 % ?

STEP 5: Enter your comments

Just some scratch disks.

?

STEP 6: Create your disk group

Click the **Finish** button to create your disk group.

Note

If you use more space for your virtual disks than the percent in the occupancy alarm level, then you will see a triangle (attention) in the disk group. The latest recommendation is to set the occupancy alarm level to 90.

Advanced Options

Advanced options let you specify the characteristics described in this table.

Field	Description
Number of disks	As described previously. The default value is 8 .
Disk failure protection level	As described previously. The default value is none .
Occupancy alarm level	This value is the percentage of total disk capacity used. When the amount of data in the disk group reaches this level, an event code is generated on the SMA. For example, if the capacity of a disk group is 5GB, and the occupancy alarm level is 80%, the event code is generated when the amount of data in the disk group reaches 4GB. The default value is 95% .
Comments	Enter text that you want tracked.

Viewing and Changing Disk Group Properties

You can view and change disk group properties from the Disk Group Properties page for the default disk group or another disk group you created.

The disk group properties that you can change are:

- Number of disks in a disk group
- Disk failure protection level
- Occupancy alarm level

Example

This example shows the default disk group properties.

The screenshot shows the 'Disk Group Properties' window for the 'Scratch Disks' group. The left sidebar shows a tree view of the storage network, including 'Virtual Disks', 'Hosts', 'Disk Groups', and 'Scratch Disks'. The main panel has tabs for 'General' and 'Vdisks'. The 'General' tab is active, showing various properties and status indicators.

Disk Group Properties	
Save changes Add disks Locate Delete ?	
General Vdisks	
Identification	
Name:	Scratch Disks
UUID:	6005-08b4-0001-4523-0002-d000-0072-0000
Capacity	
Total:	270.7 GB
Available:	Vraid0 270.49 GB Vraid1 135.28 GB Vraid5 216.42 GB
Total disks:	10
Condition/State	
Operational state:	Good
Leveling state:	Inactive
Leveling progress:	n/a
Disk failure protection	
Requested level:	Single
Actual level:	Single
Occupancy	
Total:	0 GB
Alarm level:	95 %
Comments	
Just some scratch disks.	

The buttons on the Disk Groups Properties page enable you to perform the following:

- **Save changes** — Save changes made on this page.
- **Add disks** — Increase the capacity of the disk group by adding physical disks.
- **Locate** — Flash the LEDs on all physical disks that are members of this disk group.
- **Delete** — Delete the disk group. You cannot delete the Default Disk Group.

Change the disk group name, disk failure protection level, occupancy alarm level, and comments by entering new values and clicking *Save changes*.

Creating a Vertical Disk Group

A vertical disk group configuration has one disk drive per shelf. This configuration is automatically done through VCS, however, these steps are included as a reference in case you need to do some of these manipulations:

1. Create a disk group as discussed previously.
2. Perform a locate on the disk group by clicking the *Locate* button on the Disk Group Properties page. This identifies the disk group members in the Enterprise Virtual Array. Write down the location.
3. Navigate to the hardware disk enclosure where there are presently no members of the disk group.
4. Select an ungrouped disk from that enclosure to add to the disk group.
5. Once you have at least one disk from each enclosure, remove disks from the disk group where more than one exists per enclosure.

Note

If the number of disks in the disk group exceeds the number of shelves, then you have a greater risk of failure.

Deleting Disk Groups

You can delete disk groups from the Disk Group Properties page by clicking *Delete* and then *OK* to confirm your actions.

You cannot delete a disk group if it has virtual disks. You must first delete the virtual disks. You cannot delete the default disk group unless an additional disk group has been created to protect metadata.

Creating a Host

A host is a collection of Fibre Channel adapters (FCAs) that belong to the same virtual server, and use one or more of the virtual disks created and presented by the HSV controller pair. Before a host can use the storage system virtual disks, you must make the host known to the storage system by adding a host entry using Command View EVA.

Command View EVA allows you to perform the following functions:

- Create a host folder.
- Add a host.
- Modify host properties:
 - Add an FCA.
 - Delete an FCA.
- Delete a host.

Before you can add a host, you need to know the following:

- LAN name
- IP address (optional)
- World Wide ID (WWID) of one FCA
- Operating system

Adding a host to a storage system involves the following processes:

1. Collect the host information listed above.
2. Create a host folder (if desired).
3. Add a host (to a folder if desired).
4. Add more FCAs.
5. Verify that the host has been added.

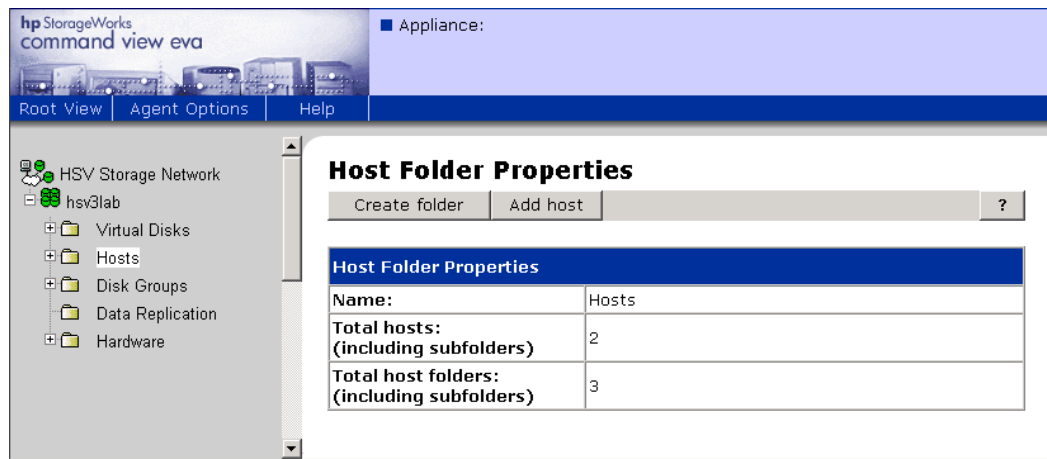
You can add a host any time after you initialize the storage system.

Creating a Host Folder

When the storage system is initialized, it shows only a top-level host folder named **Hosts**. Clicking on this folder displays its Properties page. Use this page to create a new host folder. The new host folder is a subfolder of the folder named **Hosts**.

Example

In this example, **Hosts** is selected in the Navigation pane under the storage system named **hsv3lab**. The Host Folder Properties page displays in the Content pane and shows that two hosts and three host folders have been defined.



The buttons on the Host Folder Properties page launch the following functions:

- **Create folder** — Starts the Create a Folder wizard so you can create a host subfolder.
- **Add host** — Starts the Add Host wizard, through which you can specify a new host in this folder.

To create a host folder, click *Create folder* on the Host Folder Properties page. To create the folder:

1. Enter a name — You are required to enter a name for your folder.
2. Click *Finish*.

Example

This is an example of entering a host folder name and comments.

Create a Folder

Finish	Cancel	?
--------	--------	---

Complete the steps below to create your folder.

STEP 1: Enter a Name

Enter a name for your folder.

HP Servers

STEP 2: Enter comments

Enter any comments you'd like to attach to your folder.

This folder contains all HP host definitions.

STEP 3: Create Your Folder

Click the **Finish** button to create your folder.

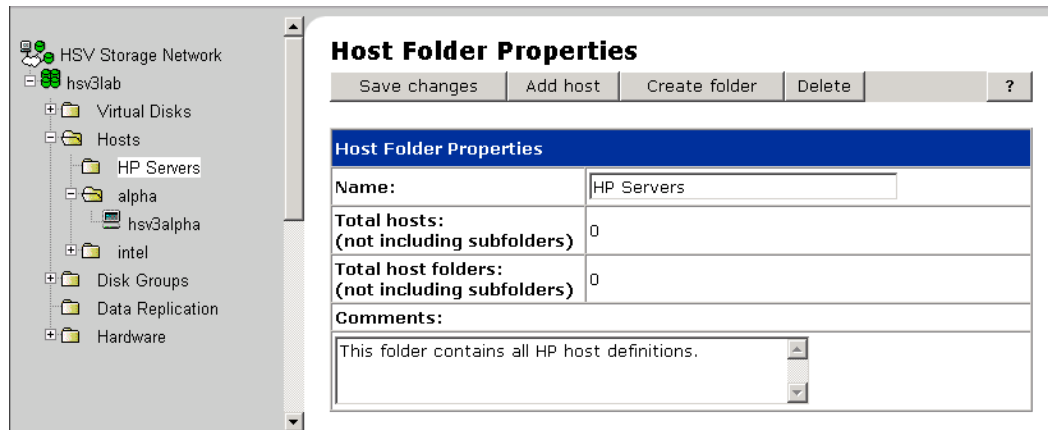
The new host folder you create becomes a subfolder of Hosts.

Adding a Host

Before a host can use the storage system virtual disks, the host must be known to the storage system. Adding a host creates a path from the storage system to one FCA. More FCAs can be specified as a modification to the host properties.

To add a host, navigate to the host folder you created.

Example



Using the Add a Host Wizard

To add a host, click *Add host* in the Content pane of the Properties page. The wizard leads you through the steps of adding a host to the list of hosts that can potentially use the storage system.

Example

This example shows Page 1 of the wizard. The host name **HP1** is assigned.

Add a Host

Page 1 | Page 2 | Page 3

Next step

Cancel

?

Complete these steps and click **Next Step** to add your host.

STEP 1: Enter the Name

Enter your host's LAN node name.

HP1

?

STEP 2: Enter the IP address

If your host uses a static LAN IP address, enter the address. Skip this step if your host uses dynamic IP addresses.

Dynamic IP Assignmen

?

Click *Next step* to proceed to Page 2 of the wizard.

Example

The adapter port WWN and Host OS are selected from drop-down lists.

Add a Host

Page 1 Page 2 Page 3

Previous step	Next step	Cancel	?
---------------	-----------	--------	---

Complete this step and click **Next Step** to continue adding your host.

STEP 3: Enter an adapter port World Wide Name

Click to select from list Click to select from list 1000-0000-c926-99ae 1000-0000-c927-6243 1000-0000-c92c-1d3b	—OR—	Enter a port WWN ?
Host OS Hewlett Packard HP-UX ? Tru64 UNIX VMS Microsoft Windows Sun Solaris Hewlett Packard HP-UX IBM AIX Unknown Custom	System attributes Custom mode number: n/a ?	Direct eventing: n/a ?

Click *Next step* to proceed to Page 3 of the wizard.

Example

An optional comment is added.

Add a Host

Page 1 Page 2 Page 3

Previous step	Finish	Cancel	?
---------------	--------	--------	---

Continue with these steps to add your host.

STEP 5: Enter your comments

Place relevant comments here, such as which other FCAs are in the host.

STEP 6: Add your host

Click the **Finish** button to add your host.

Click *Finish* to add the host.

Field Definitions

You can specify values for fields as described in this table.

Field	Description
Name	Assign a name to the host.
IP address	If your host uses a static LAN IP address, enter the address; otherwise, skip this step.
Host adapter WWID	Select an FCA WWID from the drop-down list or enter it in the text box.
Operating system	Select a host operating system the server is running from the drop-down list.
Custom mode number	Hex string that specifies an OS that is not listed in the host OS drop-down list. HP provides custom mode numbers, as needed.
Direct eventing	Enable or disable inband host event reporting.
Comments	Enter useful comments about the host you are adding.

Viewing and Changing Host Properties

View and change host properties from the *Host Properties* page. The only host properties you can change are:

- Operating system type, custom type, and direct eventing
- Comments

For any other changes, you must delete and add the host again.

Example

This is an example of the Host Properties, General page for host **hsv3alpha**.

Host Properties

Save changes Delete host Move ?

General Presentation Ports

Host Properties

Node name:	hsv3alpha		
IP Address:	Dynamic IP Assignment		
Operating System:	Type:	Tru64 UNIX	
	Custom type:	n/a	
	Direct eventing:	Disabled	
UUID:	6005-08b4-0001-4523-0002-c000-01a3-0000		
Comments:	<input type="text"/>		

The buttons on the Hosts Properties, General page enable you to perform the following:

- **Save changes** — Save changes you make to this page.
- **Delete host** — Delete this host from the list of hosts that can access the storage system.
- **Move** — Moves the host to a desired folder location using a navigation tree.

The second tab is the Presentation tab, which shows virtual disks that are presented to the host. There are no host properties that can be changed on this screen.

Example

This is an example of the Host Properties, Presentation page for host **hsv3alpha**.

Host Properties

		?
General		Presentation
Ports		
Vdisk	LUN	
vdalpha\labvd1alpha\ACTIVE	1	
vdalpha\labvd2alpha\ACTIVE	4	
vdalpha\labvd3alpha\ACTIVE	6	

The third tab is the Ports tab where you can add or delete an FCA.

Example

This is an example of the Host Properties, Ports page for host **hsv3alpha**.

Host Properties

Add port	Delete port	?
General	Presentation	Ports
Total adapter ports: 2		
FC Adapter Port WWN		
1000-0000-c92a-43aa		
1000-0000-c92a-45ff		

Click *Add port* to bring up a screen to add an FCA, or click *Delete port* to delete a specific FCA (prevent it from accessing the storage system).

Adding a Host FCA

You must have at least two paths to the host for multipathing. You can add an additional port to the host using the Add a Host Port function. Navigate to the host's Properties page.

Adding a Host Port

When you click on *Add port* on the Host Properties, Port page, you can add a host port. Then use the following procedure.

1. Select or enter an FCA WWID.
2. Click *Add port*.

Example

Add a Host Port

Add port	Cancel	?
----------	--------	---

Select an unassigned World Wide Name from the list or enter the World Wide Name of a host adapter port into the text box. To skip to a particular list entry, place your cursor in the selection box and repetitively enter the first character of the entry you wish to find.

<div style="border: 1px solid black; padding: 2px;"> <div style="display: flex; justify-content: space-between; align-items: center;"> Click to select from list ▼ </div> <div style="background-color: #e0e0e0; padding: 2px;">Click to select from list</div> <div style="padding: 2px;"> 1000-0000-c926-99ae 1000-0000-c927-6243 1000-0000-c92c-1d3b </div> </div>	—OR—	<div style="border: 1px solid black; padding: 2px;"> Enter a port WWN ? </div>
---	------	---

Note

The host list is read from the HSV controller by way of a SCMI Get Host command: SCMI_F_NSC_GET_UNASSIGNED_HOST_PORTS.

Creating a Virtual Disk

A virtual disk is a simulated disk drive that the storage system HSV controllers create for hosts to use. The virtual disk characteristics defined by you or a storage administrator provide a specific combination of capacity, availability, performance, and accessibility. The controller pair simulates these characteristics by deploying the disk group specified for the virtual disk. The host sees the virtual disk like it would see a physical disk with the same characteristics.

Whenever you create a new virtual disk, a virtual disk family is automatically created. The new virtual disk is the active virtual disk of the family.

A virtual disk family consists of a virtual disk and its snapshot, if a snapshot exists. The original virtual disk is called the active disk. When you first create a virtual disk family, the only member is the active disk. There is only one active disk per virtual disk family. A snapshot is not an active disk.

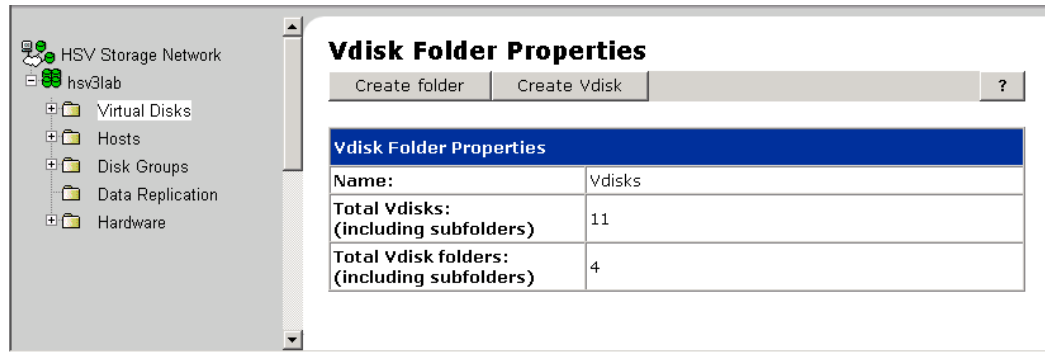
Note

The Enterprise Virtual Array supports dynamic on-line expansion of a virtual disk; that is, you can expand the capacity of a virtual disk without having to unpresent the virtual disk to all hosts. This is available for Windows 2000 and Sun Solaris only.

Creating a Virtual Disk Folder

A virtual disk folder organizes the virtual disk families. Create it from the Vdisk Folder Properties page. Navigate to this page by selecting *Virtual Disks* in the Navigation pane. The folder you create becomes a subfolder of the Virtual Disks folder.

Example



The buttons on the Virtual Disk Folder Properties page launch the following functions:

- **Create folder** — Create a virtual disk subfolder.
- **Create Vdisk** — Create a virtual disk (virtual disk family).

To create a virtual disk folder, click *Create folder* and do the following:

1. Enter a name — select a meaningful name.
2. Enter a comment (optional).
3. Click *Finish*.

Create a Folder

Finish Cancel

?

Complete the steps below to create your folder.

STEP 1: Enter a Name

Enter a name for your folder.

Scratch Disks

STEP 2: Enter comments

Enter any comments you'd like to attach to your folder.

This is a folder of scratch disks.

STEP 3: Create Your Folder

Click the **Finish** button to create your folder.

Creating a Virtual Disk Family

You can create a virtual disk family from the Virtual Disk Folder Properties page. Navigate to the Virtual Disk Folder Properties page by selecting the folder name in the Navigation pane you created.

To create a virtual disk family, click *Create Vdisk* and do the following:

1. Assign a virtual disk name — This is really the virtual disk family name. The active virtual disk created in this folder when this operation is complete always has the name *Active*.
2. Specify the virtual disk characteristics (see definitions in next topic).
3. Click *Finish*.

Example

This example shows virtual disk characteristics for a virtual disk named **Vdisk003**.

Create a Vdisk Family

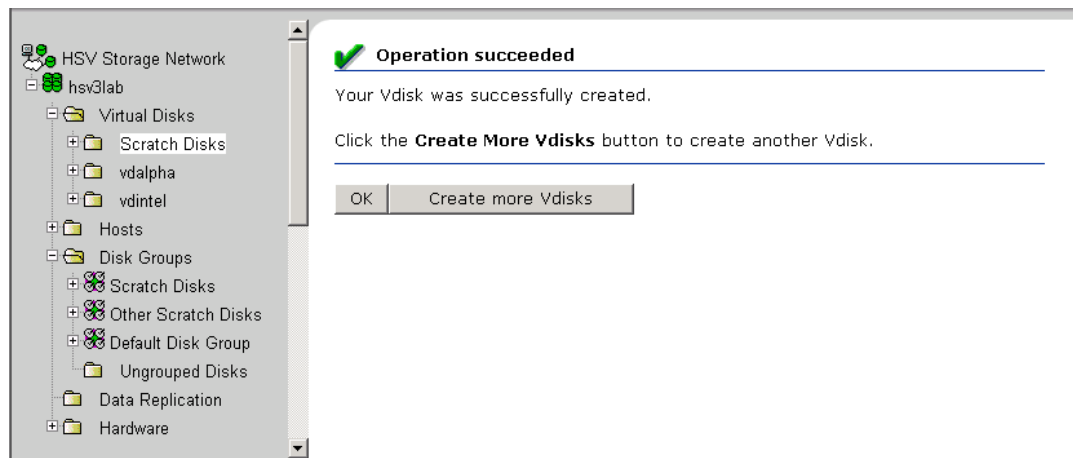
Finish		Cancel		?	
Vdisk name: <input type="text" value="Vdisk003"/> ?					
Disk group name		Available GB: Vraid0/Vraid5/Vraid1			
<input type="text" value="Scratch Disks"/>		<input type="text" value="270.49"/>		<input type="text" value="216.42"/>	
		<input type="text" value="135.28"/>		?	
Redundancy:					
<input type="radio"/> Vraid0 ?		<input type="radio"/> Vraid5 ?		<input checked="" type="radio"/> Vraid1 ?	
Space available		Space available		Space available	
270.49 GB		216.42 GB		135.28 GB	
Size: <input type="text" value="0"/> GB					
World Wide Name: <input type="text" value="Default WWName"/> ? (format: 6xxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx)					
Write Cache policy:			Read Cache policy:		
<input type="text" value="Mirrored write-back"/> ?			<input type="text" value="On"/> ?		
<input checked="" type="radio"/> Read/write <input type="radio"/> Read only?			OS Unit ID: <input type="text" value="0"/> ?		
Present to host: <input type="text" value="None"/> ?			Preferred path/mode: <input type="text" value="No preference"/> ?		



Important

If the size of the virtual disk is greater than 100GB, do not present it to a host at the time of creation, allow the virtual disk to normalize. You can monitor this from the property page of the virtual disk.

The next screen that appears notifies you if your virtual disk was successfully created.



Virtual Disk Characteristics

Specify characteristics as described in this table.

Field Name	Description
Virtual disk name	Assign a name with a maximum of 32 characters.
Disk group name	Select a disk group from the drop-down list.
Data protection level	Select VRAID 0, VRAID 5, or VRAID 1.
Size	Enter a size in GB (whole number). The maximum size depends upon the unused capacity of the disk group in which it is created and the selected redundancy (data protection level) up to 2TB.
Write-cache policy	<p>Select a write cache policy.</p> <ul style="list-style-type: none"> ■ Mirrored write-back ■ Unmirrored write-back <p>Note: A write cache policy cannot be changed while the virtual disk is presented to the host.</p>
Read cache policy	Select on or off. The default is on .
Write protection	Select Read/write or Read only.
OS Unit ID	<p>Some operating systems require a unique ID to assign a virtual disk a device name, therefore each virtual disk may be assigned an operating system-specific unit ID.</p> <p>Caution: OpenVMS requires that each virtual disk be assigned an OS Unit ID. If a virtual disk is not assigned an OS Unit ID, Open VMS cannot see it. In addition, OpenVMS requires that these OS Unit IDs be unique across the entire SAN.</p> <p>Note: The Tru64 UNIX operating system may use an OS Unit ID but it is not mandatory.</p> <p>Note: The OS Unit ID must be unique in the host cluster.</p>
Present to host	Select none (if not presenting a virtual disk to a host now) or select a host WWID.
Preferred path/mode	<p>Allows you to designate one controller, arbitrarily, as path A, and the other as path B. This provides a way of initially distributing the virtual disk load evenly between controllers. During operation, this path preference can be overridden by commands from the host.</p> <p>Select none for no preference, or one of the following modes to determine the behavior of the controllers upon failure:</p> <ul style="list-style-type: none"> ■ Failover only — Controller turns over control of all its virtual disks to the other controller and remains inoperative. ■ Failover and failback — Controller turns over control of all its virtual disks to the other controller. If the condition that caused the failover no longer exists, the controller automatically takes over control of its own virtual disks again.

Viewing and Changing Virtual Disk Properties

At the Vdisk Active Member Properties page, you can view and change virtual disk properties. The following properties can be changed:

- Write cache policy
- Read cache policy
- Capacity
- Write protection
- Comments
- OS Unit ID

This is mandatory for OpenVMS and must be unique across the SAN. For Tru64, it is suggested but not mandatory.

- Preferred path/mode

Note

There is a new Data Replication tab for VCS V3.0.

Example

This example shows the properties of the active virtual disk in the virtual disk family **Vdisk003**. The General tab is chosen.

Vdisk Active Member Properties

Save changes		Create snapshot		Create Snapclone		?	
General		Presentation		Data Replication			
Identification				Condition/State			
Name:		ACTIVE		Operational State:		<input checked="" type="checkbox"/> Good	
Family Name:		Vdisk003					
World Wide LUN Name:				Date/Time			
<input type="text" value="6005-08b4-0001-4523-0002-d000"/>				Created:		18-Apr-2003 10:29:50	
UUID:				Cache Policies			
<input type="text" value="6005-08b4-0001-4523-0002-d000-0080-0000"/>				Write:		<input type="text" value="Mirrored write-back"/>	
				Read:		<input type="text" value="On"/>	
Attributes							
Type:		Original					
Disk Group:		Scratch Disks					
Capacity Req:		<input type="text" value="100"/> GB					
Capacity Used:		100 GB					
Redundancy:		Vraid1					
Write Protect:		<input type="text" value="No"/>					
Comments							
<input type="text"/>							

The buttons on the Vdisk Active Member Properties page enable you to perform the following:

- **Save changes** — Save changes made on the properties page.
- **Create snapshot** — Launch the wizard to create fully-allocated and demand-allocated snapshots.
- **Create Snapclone** — Launch the wizard to create snapclones.

To make changes, enter the new values and click *Save Changes*.

The following shows the display for the Presentation tab.

Vdisk Active Member Properties

Save changes		Present		?	
General		Presentation		Data Replication	
Presentation Properties					
OS Unit ID:		<input type="text" value="0"/>			
Preferred path/mode:		<input type="text" value="No preference"/> ▼			
Presentations					
Host		LUN		FC Reservation	
Not presented					

The buttons on this page enable you to perform the following:

- **Save changes** — Save changes made on the properties page.
- **Present** — Launch the wizard to present a virtual disk to a host.
- **Unpresent (not shown)** — Launch the wizard to unpresent a virtual disk to a host.

The Present and Unpresent options are covered in the next topics.


Presenting a Virtual Disk to a Host

A virtual disk can be presented to a host during or after virtual disk creation. To present it during virtual disk creation, you must select a host, and optionally, a LUN address.

Present the virtual disk after virtual disk creation from the Vdisk Active Member Properties page, Presentation tab:

1. Click *Present*.
2. Select the host or hosts.
3. Select the LUN (optional) at which the virtual disk will be presented to the host.
4. Click *Present Vdisk*.

Present Vdisk



Select one or more hosts to which your Vdisk will be presented. Only those hosts without existing presentations to your Vdisk are shown. Click the **Present Vdisk** button after you have made your selections.



Note the following buttons:

- **Select none** — Removes current host selections for presentation
- **Select all** — Selects all hosts in the display for presentation
- **Assign LUN** — Allows selection of a LUN

Clicking the *Assign LUN* button presents the following screen.

Present Vdisk

Page 1 Page 2

Previous step

Finish

Cancel

?

Complete the steps below to present your Vdisk.

STEP 2: Select a LUN:

Available LUNs for this host: 4 ?

STEP 3: Present your Vdisk

Click the **Finish** button to present your Vdisk.

As long as the virtual disk is enabled, the host can access the virtual disk at the LUN address.

Unpresenting a Virtual Disk to a Host

Unpresenting a virtual disk to a host makes it unavailable to the host to which it is currently presented. This prevents a host or hosts from accessing the virtual disk, or from seeing the virtual disk at a specific LUN.

Unpresent a virtual disk to a host from the Vdisk Active Member Properties page by doing the following:

1. Click *Unpresent*.
2. Select one or more hosts to be unpresented. Only those hosts with existing presentations to your virtual disk are shown.
3. Click *Unpresent host(s)*.

Creating a Snapshot

A snapshot is a virtual disk that reflects the contents of another virtual disk at a particular point in time.

- A snapshot operation can only be done on an active virtual disk.
- Seven snapshots of an active virtual disk can exist at any point.
- The active virtual disk and its snapshot constitute a virtual disk family.
- All snapshots within a virtual disk family must have the same allocation policy.
- You cannot create a snapshot if a snapclone is in a copy state using the same active virtual disk.
- A snapshot is intended to be temporary.

Using the Snapshot Wizard

Start the Create a Vdisk Snapshot wizard at the Vdisk Active Member Properties page by clicking *Create snapshot*.

To create a snapshot in the simplest way:

1. Enter a name — Enter a meaningful name for the snapshot that differentiates it from the active virtual disk from which it is created.
2. Select a host.
3. Click *Finish*.

Example

Create a Vdisk Snapshot

Page 1 Page 2 Page 3 Page 4 Page 5

Finish

Advanced options

Cancel

?

Complete these steps and click **Finish** to snapshot your Vdisk family's Active member in the simplest way. For more control, complete the steps and click **Advanced Options** instead.

STEP 1: Enter a Name

Snap003

?

STEP 2: Select a host

None

?

None

hsv3alpha

hsv3intel

If you click on *Advanced options*, Page 2 of the wizard displays.

Create a Vdisk Snapshot

Page 1 Page 2 Page 3 Page 4 Page 5

Previous step

Next step

Finish

Cancel

?

Continue with these steps to snapshot your Vdisk family's Active member using advanced options.

STEP 3: Select a LUN Address

 ?

STEP 4: Select an allocation policy

 ?
 Fully-allocated snapshot
 Demand-allocated snapshot

STEP 5: Enter the OS Unit ID

 ?

Create Traditional Snapshots or Virtually-Capacity Free Snapshots by specifying one of the following allocation policies:

- **Fully-allocated** (Traditional Snapshot) — The space a virtual disk requires on the physical disks is reserved, even if the virtual disk is not currently using space.
- **Demand-allocated** (Virtually Capacity-Free Snapshot) — The space a capacity-free virtual disk requires on the physical disks is not reserved until needed.

Other characteristics default to the same values as the source virtual disk unless otherwise specified. These characteristics can be specified at creation time or later in the Properties page.



Important

All snapshots within a virtual disk family must have the same allocation policy, that is, if one snapshot is created using allocation on demand, all snapshots must use allocation on demand.

Clicking on *Next step* displays the third page of the wizard.

Create a Vdisk Snapshot

Page 1 Page 2 **Page 3** Page 4 Page 5

Previous step

Next step

Finish

Cancel

?

Continue with these steps to snapshot your Vdisk family's Active member using advanced options.

STEP 6: Select a read-cache policy

On ▼ ?

The read cache is specified as either on or off for each virtual disk. The default state is on.

Clicking on *Next step* displays the fourth page of the wizard.

Create a Vdisk Snapshot

Page 1 Page 2 Page 3 **Page 4** Page 5

Previous step

Next step

Finish

Cancel

?

Continue with these steps to snapshot your Vdisk family's Active member using advanced options.

STEP 7: Set write protection

☒ Read/write

☐ Read only

?

STEP 8: World Wide LUN Name

Default WWName

(format : 6xxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx)

Clicking on *Next step* displays the last page of the wizard.

Create a Vdisk Snapshot

Page 1 Page 2 Page 3 Page 4 Page 5

Previous step

Finish

Cancel

?

Continue with these steps to snapshot your Vdisk family's Active member using advanced options.

STEP 9: Enter your comments

This is a snapshot of Vdisk003.

?

STEP 10: Create your snapshot

Click the **Finish** button to create your Vdisk snapshot.

Clicking on *Finish* completes the snapshot creation.

Virtual Disk Active Properties Page for Snapshot

The following is an example of the virtual disk snapshot properties page.

Vdisk Snapshot Properties

Save changes		Delete	?
--------------	--	--------	---

General		Presentation	
---------	--	--------------	--

Identification		Condition/State	
Name:	Snap003	Operational State:	<input checked="" type="checkbox"/> Good
Family Name:	Vdisk003	Date/Time	
World Wide LUN Name:	6005-08b4-0001-4523-0002-d000-008c-0000	Created:	18-Apr-2003 11:47:34
UUID:	6005-08b4-0001-4523-0002-d000-008c-0000	Cache Policies	
		Write:	Mirrored write-back
		Read:	On <input type="button" value="v"/>

Attributes	
Disk Group:	Scratch Disks
Capacity Used:	0 GB
Initial Allocation:	20 GB
Allocation policy:	Fully allocated
Redundancy:	Vraid1
Write Protect:	No <input type="button" value="v"/>

Comments
This is a snapshot of Vdisk003. <input type="button" value="v"/>

Creating a Snapclone

A snapclone is a copy of an active virtual disk. During the copy process, the following rules apply to the source virtual disk:

- No other copy of the source virtual disk is in process
- You can create multiple snapclones as long as you have disk space
- You can create a snapclone in a different disk group
- The preferred path of the snapclone will be the same as the source virtual disk if you choose a preferred path



Important

During the snapclone creation process, the user may set a preferred path. Because of its sharing relationship with the source virtual disk, the snapclone will retain the same path as the source. HP recommends that you choose *No preference* in the creation wizard (see next page) if you want the source virtual disk and the snapclone to have different paths. Then set the preferred paths for the virtual disk and snapclone to be different after the snapclone operation has completed.

Using the Snapclone Wizard

Start the Create a Snapclone wizard at the Vdisk Active Member Properties page by clicking *Create Snapclone*.

To create a snapclone in the simplest way:

1. Enter a name.
2. Select a host.
3. Click *Finish*.

Example

Create a Snapclone

Page 1 Page 2 Page 3 Page 4 Page 5

Finish Advanced options Cancel ?

Complete these steps and click **Finish** to create a Snapclone of your Vdisk's Active member in the simplest way. For more control, complete the steps and click **Adv Options** instead.

STEP 1: Enter a Name

Clone002 ?

STEP 2: Select a host

None ?
 None
 hsv3alpha
 hsv3intel

The family that is created can reside in the same or in a different disk group as the source virtual disk family. Before creating the snapclone, ensure that the disk group is large enough to accommodate both families.

If you click on *Advanced options*, Page 2 of the wizard displays.

Create a Snapclone

Page 1 Page 2 Page 3 Page 4 Page 5

Previous step Next step Finish Cancel ?

Continue with this step to create a Snapclone of your Vdisk's Active member using advanced options.

STEP 3: Select a LUN Address

 ?

STEP 4: Enter the OS Unit ID

 0 ?

Clicking on *Next step* displays the third page of the wizard.

Create a Snapclone

Page 1 Page 2 Page 3 Page 4 Page 5

Previous step Next step Finish Cancel ?

Continue with these steps to create a Snapclone of your Vdisk's Active member using advanced options.

STEP 5: Select a disk group

Name	Available capacity (GB): Stripe/Parity/Mirror			
Scratch Disks	150.42	120.34	75.22	?
Scratch Disks	150.42	120.34	75.22	
Other Scratch Disks	385.84	308.75	192.96	
Default Disk Group	315.28	252.28	150.23	

STEP 6: Select a read-cache policy

 On ?

Clicking on *Next step* displays the fourth page of the wizard.

Create a Snapclone Page 1 Page 2 Page 3 Page 4 Page 5

Previous step Next step Finish Cancel ?

Continue with these steps to create a Snapclone of your Vdisk's Active member using advanced options.

STEP 7: Set write protection

☒ Read/write ☐ Read only ?

STEP 8: World Wide LUN Name

Default WWName
(format : 6xxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx)

Clicking on *Next step* displays the last page of the wizard.

Create a Snapclone Page 1 Page 2 Page 3 Page 4 Page 5

Previous step Finish Cancel ?

Continue with these steps to create a Snapclone of your Vdisk's Active member using advanced options.

STEP 9: Enter your comments

This is a snapclone of Vdisk002. ?

STEP 10: Create your Snapclone


Click the **Finish** button to create your Snapclone.

Clicking on *Finish* completes the snapclone creation.

Snapclone Properties During Background Process

The following is an example of a snapclone properties page during the snapclone creation. Notice that the capacity required and the capacity used are different during the background process.


Vdisk Active Member Properties

		?
<div> <div>General</div> <div>Presentation</div> <div>Data Replication</div> </div>		
Identification		Condition/State
Name:	Active	Operational State:  Snapclone in progress
Family Name:	Clone002	
World Wide LUN Name:		Date/Time
6005-08b4-0001-4523-0002-d000-0096-0000		Created: 18-Apr-2003 12:28:01
		Cache Policies
		Write: Mirrored write-back
		Read: On
Attributes		
Disk Group:	Scratch Disks	
Capacity Req:	20 GB	
Capacity Used:	2 GB	
Redundancy:	Vraid1	
Write Protect:	No	
Comments		
This is a snapclone of Vdisk002.		

Virtual Disks with Snapshots and Snapclones

The following is an example of the Navigation pane showing snapclone and snapshot creation.

Note difference
between
snapclone and
snapshot



Learning Check

1. What information displays in the Command View EVA Session pane?

.....

.....

.....

2. Describe how to access system help, page help, and field help.

.....

.....

.....

3. How do you locate the Command View EVA version number?

.....

4. What do you load to the Enterprise Virtual Array when you perform a code load operation and a disk drive firmware code load operation?

.....

.....

.....

5. On which page do you launch the wizard to perform a disk drive firmware code load?

.....

.....

.....

.....

6. What is created when you initialize the storage system?

.....

7. What host information must you gather before you can add a host to the storage system?

.....

.....

.....

.....

8. What Command View EVA pages allow you to modify disk group, virtual disk, and host properties?

.....

9. How do you add a host with two FCAs to the storage system using Command View EVA?

.....

.....

.....

.....

10. Which of the following statements is correct? (Choose one).

- a. You can create a snapshot and a snapclone from the same active virtual disk.
- b. You can delete a disk group if it has a virtual disk.
- c. An active virtual disk and its snapclone constitute a virtual disk family.
- d. Up to seven snapshots of an active virtual disk can exist at any point.

Overview

This module discusses Command View EVA resources available for managing and monitoring the Enterprise Virtual Array. It describes the management agent options that you can set in Command View EVA and details the options you can use to set up passwords and add licenses. Finally, hardware property displays are presented to observe configuration and operational states of a storage system's components: rack, controllers, disk enclosures, and disk drives.

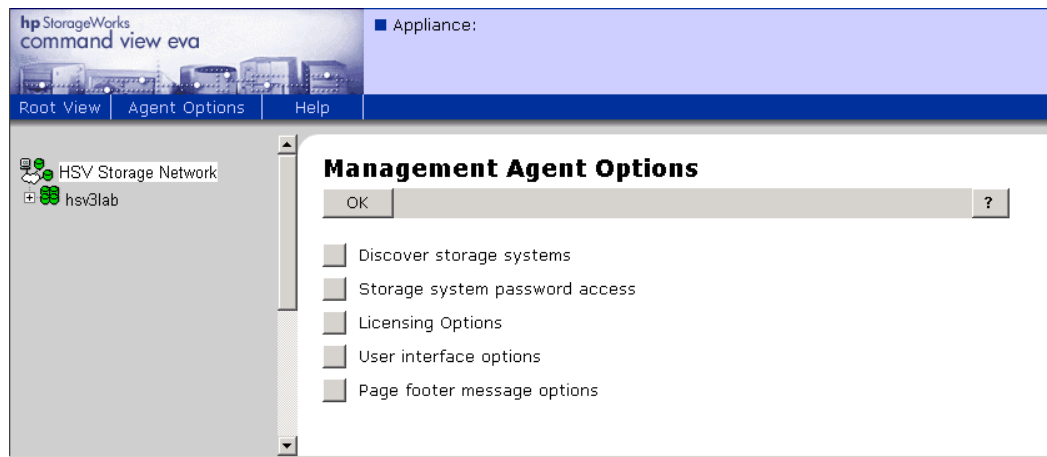
Objectives

After completing this module you should be able to:

- Identify the management agent options that you can set in Command View EVA.
- Describe how to use the management agent options to set up passwords and add licenses.
- Identify the information that displays on a storage system's rack, controller, and disk drive properties pages.

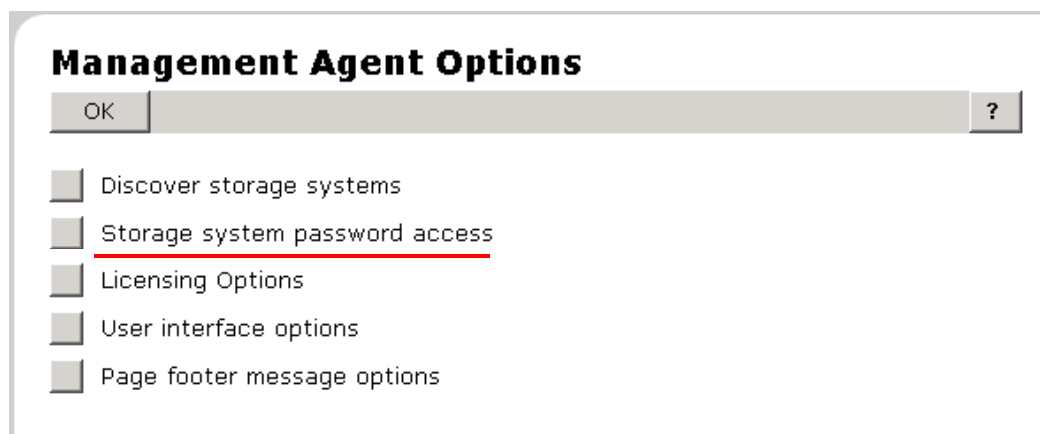
Viewing Management Agent Options

You can access the Management Agent Options page by selecting *Agent Options* in the Session pane. The following screen appears.



Storage System Password Access

When you click on the button next to **Storage system password access**, you will be able to set passwords for your storage systems.



Command View EVA provides password protection to the storage systems. Without password protection, any Command View EVA on the fabric can access any storage system on the fabric if not properly zoned.

The storage system becomes password-protected when a password is entered into the operator control panel (OCP) of one of the controllers. This locks the storage systems with the password. This password must be entered in the Command View EVA you want to access that storage system.

Once a password is entered for a specific storage system, Command View EVA stores that password to continually associate it with that storage system. If the storage system password changes, you can change the stored password.

Enable a Storage System

Once you have entered the password thru the OCP, this page lets you enable a storage system, change a storage system password, or delete one or more storage systems. In Command View EVA, click the *Enable* button to enter a password and bind it to the one you entered through the OCP.

Storage System Password Access

OK ?

Enable Enable password access to a storage system

Change Change a storage system password

Disable Disable password access to storage systems

When you click the *Enable* button, you will see the following screen.

Enable Password Access to a Storage System

OK ?

To manage a storage system that uses password access, your management agent must know its password. Your agent maintains a list of storage system World Wide Names and the passwords associated with them. You can add a system to the list by entering its WWN/password combination below. You may select a WWN from the drop-down list or enter one directly (xxxx-xxxx-xxxx-xxxx). Enter the 8-letter, alphabetic password associated with your WWN.

When you enter the 8th letter of your password, a popup dialog will appear. Click the **OK** button to add your WWN/password combination to your agent's system list. (The Add and Erase buttons below are not used.)

Enter World Wide Name and Password

Storage System List Manual World Wide Name Entry

World Wide Name [Text Field]

Password [Text Field]

Add Erase

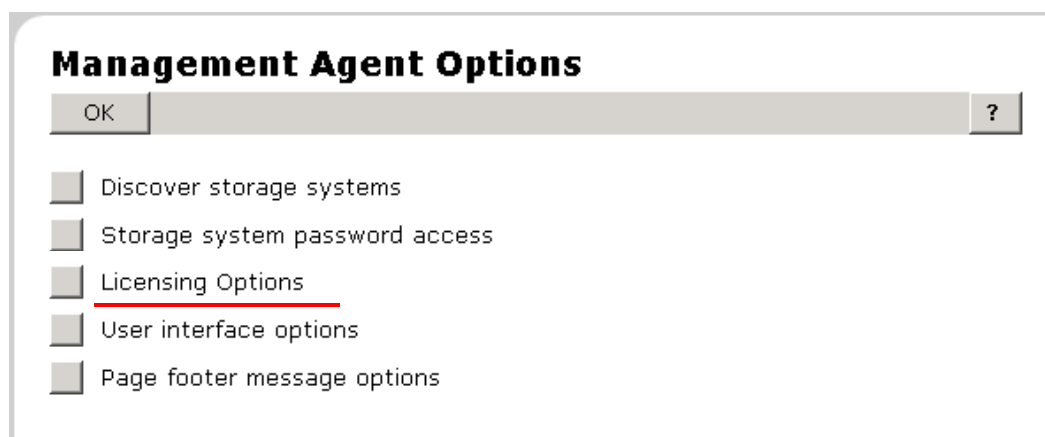
On this page, select a storage system, enter the WWID, and enter an eight alpha character password. The password must match the one entered in the OCP.

Click *Add* to register the storage system and password in Command View EVA. This binds the password to the one you entered through the OCP.

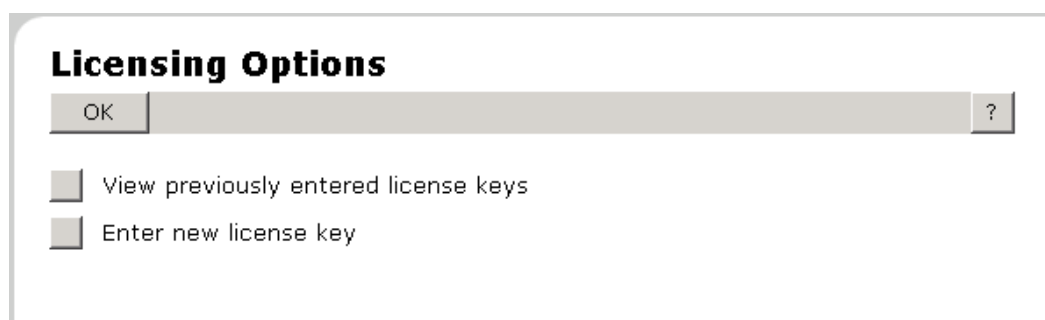
Click *Change* to delete a storage system password and click *Disable* to disable password access to a storage system.

Licensing Options

The licensing options allow you to view current license keys or enter new ones. Click the button next to **Licensing Options**.



You will see a screen with options to view previously entered license keys or to enter a license key.



Click the button next to **View previously entered license keys** to view license keys that have been entered.

View License Keys

Cancel	?
--------	---

The keys entered for the licensed features on this system are shown below.

```
INCREMENT HSV110-BASIC Compaq 2.0 permanent uncounted 270DD7FD7DC5 \
HOSTID=HSVWWN=5000-1FE1-5000-2CD0 NOTICE="Authorization \
=DD020CREAMER17426431, Qty 1, QM-6RNAA-AA2.0, VCS PKG V2.0 \
DUAL HSV CNTLR"
```

```
INCREMENT HSV110-BASIC Compaq 3.0 permanent uncounted 6FAB023B815E \
HOSTID=HSVWWN=5000-1FE1-5000-2CD0 NOTICE="Authorization = \
BM02WHITMORE85493546, 60 Day TEMPORARY Key. Replace these keys \
with permanent keys to continue operation."
```

```
INCREMENT HSV110-SNAPSHOT Compaq 3.0 permanent uncounted FE4C229F4EA2 \
HOSTID=HSVWWN=5000-1FE1-5000-2CD0 NOTICE="Authorization = \
BM02WHITMORE85493546, 60 Day TEMPORARY Key. Replace these keys \
with permanent keys to continue operation."
```

```
INCREMENT HSV110-DRM Compaq 3.0 permanent uncounted 1AE60DF2429D \
HOSTID=HSVWWN=5000-1FE1-5000-2CD0 NOTICE="Authorization = \
BM02WHITMORE85493546, 60 Day TEMPORARY Key. Replace these keys \
with permanent keys to continue operation."
```

Click the button next to **Enter new license key** to add a license.

Add a license

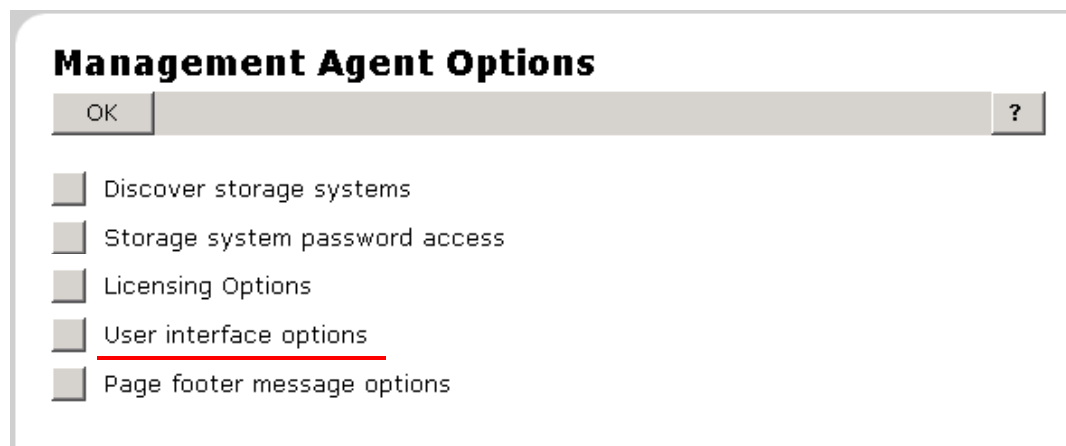
Add license	Cancel	?
-------------	--------	---

Enter a license key and click the **Add License** button to activate special features on your storage system.

```
1F1C377E0B3B
HOSTID=HSVWWN=5000-1FE1-0013-A220 NOTICE="Authorization = \
DID1DADDIEC021364684, Qty 1, QM-ENTRP-RI.SE - Enterprise \
Controller Software - SNAPSHOT, Quickspec 78*90*12 ck=129|
```


User Interface Options

The user interface options allow you to choose how objects are displayed in Command View EVA. Click the button next to **User interface options** to see the options.

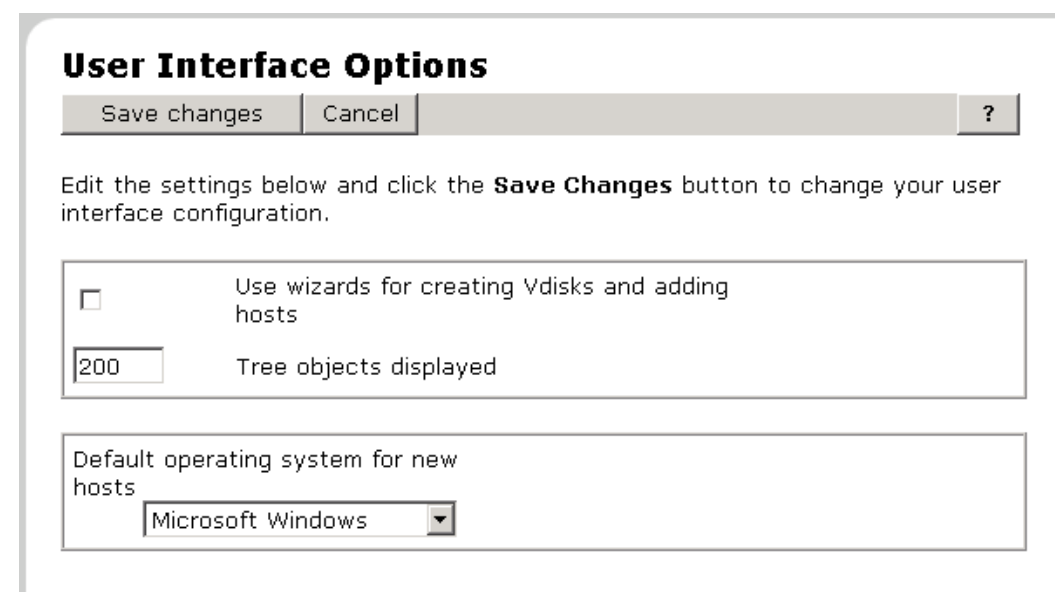


The **Management Agent Options** dialog box features a title bar with an **OK** button on the left and a help icon (?) on the right. Below the title bar is a list of five options, each preceded by a small square icon:

- Discover storage systems
- Storage system password access
- Licensing Options
- User interface options
- Page footer message options

The **User interface options** entry is highlighted with a red underline.

You have the option to use wizards for virtual disk creation and adding hosts, or to see a certain number of objects in the Navigation pane. Click on the options you want and then *Save changes*.



The **User Interface Options** dialog box has a title bar with **Save changes**, **Cancel**, and a help icon (?) button. Below the title bar is a text instruction: "Edit the settings below and click the **Save Changes** button to change your user interface configuration."

The settings are organized into three sections:

- A checkbox labeled "Use wizards for creating Vdisks and adding hosts".
- A text input field containing "200" followed by the label "Tree objects displayed".
- A section titled "Default operating system for new hosts" containing a dropdown menu currently set to "Microsoft Windows".

Example with Virtual Disks Wizards Not Selected

One of the user interface options is whether or not to use wizards for virtual disk creation. The following screen shows the Command View EVA screen for creating virtual disks **without** a wizard.

Finish

Cancel

?

Vdisk name:

Vdisk005

?

Disk group name

Available GB: Vraid0/Vraid5/Vraid1

Default Disk Group 124.51 99.60 62.27

?

Redundancy:

☐ Vraid0 ?

Space available

124.51 GB

☐ Vraid5 ?

Space available

99.60 GB

☒ Vraid1 ?

Space available

62.27 GB

Size:

0

GB

World Wide Name:

Default WWName

(format: 6xxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx)

Write Cache policy:

Mirrored write-back

?

Read Cache policy:

On

?

☒ Read/write

☐ Read only?

OS Unit ID:

0

?

Present to host:

None

?

Preferred path/mode:

No preference

?

Example with Virtual Disks Wizards Selected

The following screen shows the Command View EVA screen for creating virtual disks **with** a wizard.

Create a Vdisk Family

Page 1
Page 2
Page 3
Page 4
Page 5

Finish
Advanced options
Cancel
?

Complete these steps and click **Finish** to create your Vdisk family. For more control, complete the steps and click **Adv Options** instead.

STEP 1: Enter a Name

Vdisk005
?

STEP 2: Select a disk group

Name	Available capacity (GB):	Vraid0/Vraid5/Vraid1
Default Disk Group	124.51_99.60_62.27	?

STEP 3: Select a Redundancy Level and Size

<input type="radio"/> ?	<input type="radio"/> ?	<input checked="" type="radio"/> ?
Vraid0	Vraid5	Vraid1
Space available:	Space available:	Space available:
124.51 GB	99.60 GB	62.27 GB

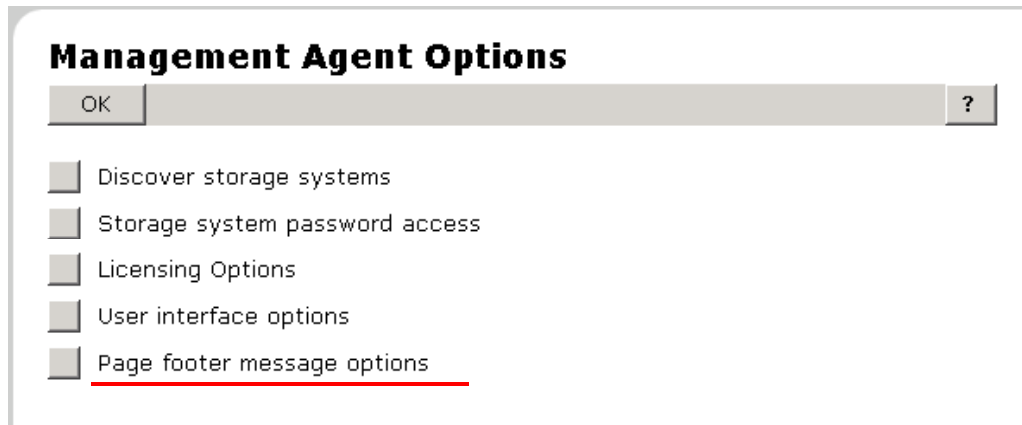
Desired Size: 0 GB

STEP 4: Select a host

None
?

Page Footer Message Options

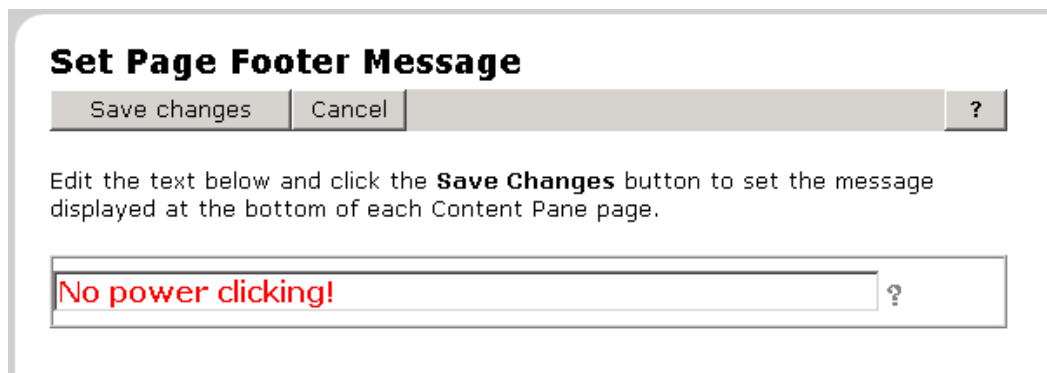
The page footer message options allow you to place a footer on every page of the Command View EVA display. Click the button next to **Page footer message options** to see the options.



The **Management Agent Options** dialog box features a title bar with an **OK** button and a help icon. Below the title bar is a list of five options, each preceded by a checkbox. The option **Page footer message options** is selected and underlined with a red line.

Management Agent Options	
<input type="checkbox"/>	Discover storage systems
<input type="checkbox"/>	Storage system password access
<input type="checkbox"/>	Licensing Options
<input type="checkbox"/>	User interface options
<input checked="" type="checkbox"/>	<u>Page footer message options</u>

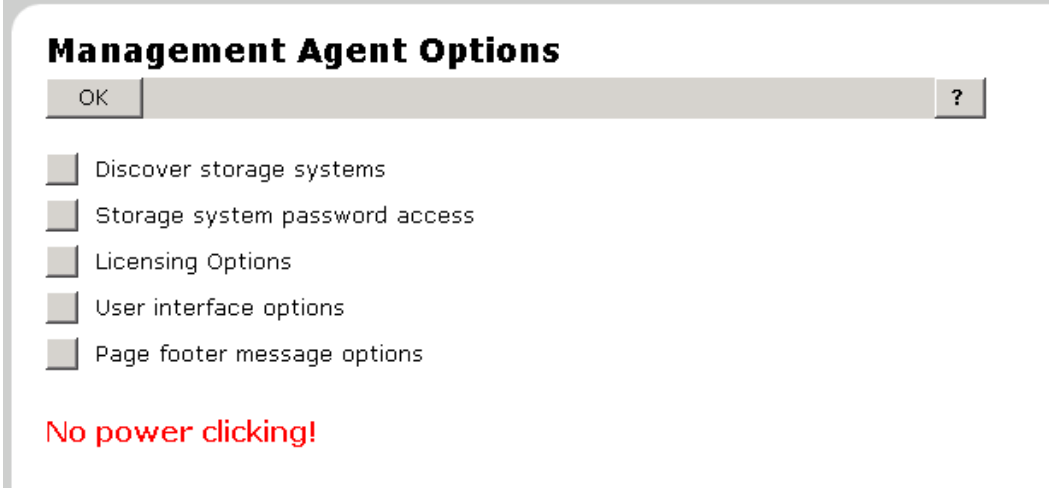
You have the option to enter any text you want. Click *Save Changes* to create the footer.



The **Set Page Footer Message** dialog box has a title bar with **Save changes**, **Cancel**, and a help icon button. Below the title bar is a text area for editing the footer message. The text **No power clicking!** is entered in red. Below the text area is a help icon button.

Set Page Footer Message	
Save changes	Cancel
Edit the text below and click the Save Changes button to set the message displayed at the bottom of each Content Pane page.	
No power clicking!	?

The following shows the footer displayed.



The image shows a dialog box titled "Management Agent Options". It has a title bar with "OK" and "?" buttons. The main area contains five options, each with a checkbox and a label:

- ☐ Discover storage systems
- ☐ Storage system password access
- ☐ Licensing Options
- ☐ User interface options
- ☐ Page footer message options

Below the list, the text "No power clicking!" is displayed in red.

Viewing Hardware Properties

You can view hardware properties using Command View EVA. This allows you to monitor the Enterprise Virtual Array components and assist in diagnosing problems. The following property displays are available:

- Rack properties
- Controller properties
- Disk enclosure properties
- Disk drive properties

To view these properties, you must select a storage system and then a rack within the storage system.

Note

The FC loop switches are not recognized by Command View EVA.

Rack Properties

The Rack Properties page displays:

- Operational state of the rack
- Number of each rack component
 - Controllers
 - Disk enclosures
 - Disks
- Comments

Example

This example shows that Rack 1 contains two enclosures, six disk enclosures, and 21 disk drives.

hp StorageWorks
command view eva

Appliance:

Root View Agent Options Help

HSV Storage Network
hsv3lab
Virtual Disks
Hosts
Disk Groups
Data Replication
Hardware
Rack 1
Unmappable Hardware

Rack Properties

Save changes ?

Name:	Rack 1
Operational state:	<input checked="" type="checkbox"/> Good
Controllers:	2
Disk enclosures:	6
Disks:	21
Comments:	

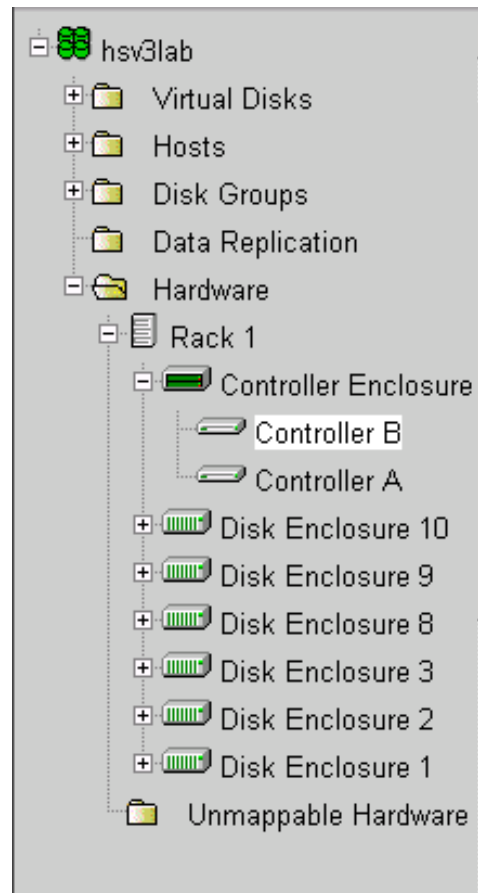
You can change the following rack properties on this page by entering the new information and clicking *Save changes*. The following properties can be changed:

- Rack name
- Comments

Controller Properties

To display the Controller Properties page, select a controller enclosure from a rack folder and select a controller.

Example



The Controller Properties page displays the following properties:

- General
- Controller host ports
- Controller devices
- Controller enclosure

General Controller Properties

General controller properties include the following:

- VCS version
- Controller serial number
- WWID of the storage system
- Operational state of the controller
- Amount of cache memory
- Mirror path and port states
- Enclosure number

Example

Controller Properties

Save changes		Locate		Shut down		Code load		?	
General		Host Ports		Device Ports		Enclosure			
Identification					Condition/State				
Name:		Controller B			Operational State:		<input checked="" type="checkbox"/> Good		
Type:		HSV Storage System			Cache Memory		Operational State: <input checked="" type="checkbox"/> Good		
Manufacturer:		Hewlett-Packard Company							
Model Number:		HSV110			Write capacity:		256 MB		
Software version:		V3_W030414-3000			Read capacity:		512 MB		
Serial Number:		P4889B49ILV00N			Mirror Path:		<input checked="" type="checkbox"/> Good		
World Wide Node Name:					Mirror Port State:		<input checked="" type="checkbox"/> Good		
5000-1FE1-5000-2CD0					Location				
UUID:									
5005-08b4-0001-4529-0000-0000-0000-0000					Enclosure number:		7		
Comments									

Note

You can use this page to view the VCS version (V3-W030414-3000).

Controller Identification

You must issue a *Locate* command to determine which physical controller is the A or B controller. The settings are as follows:

- Controller A FP1, WWN + 9
- Controller A FP2, WWN + 8
- Controller B FP1, WWN + D
- Controller B FP2, WWN + C

The WWN assignments for A and B will not change once set.



Important

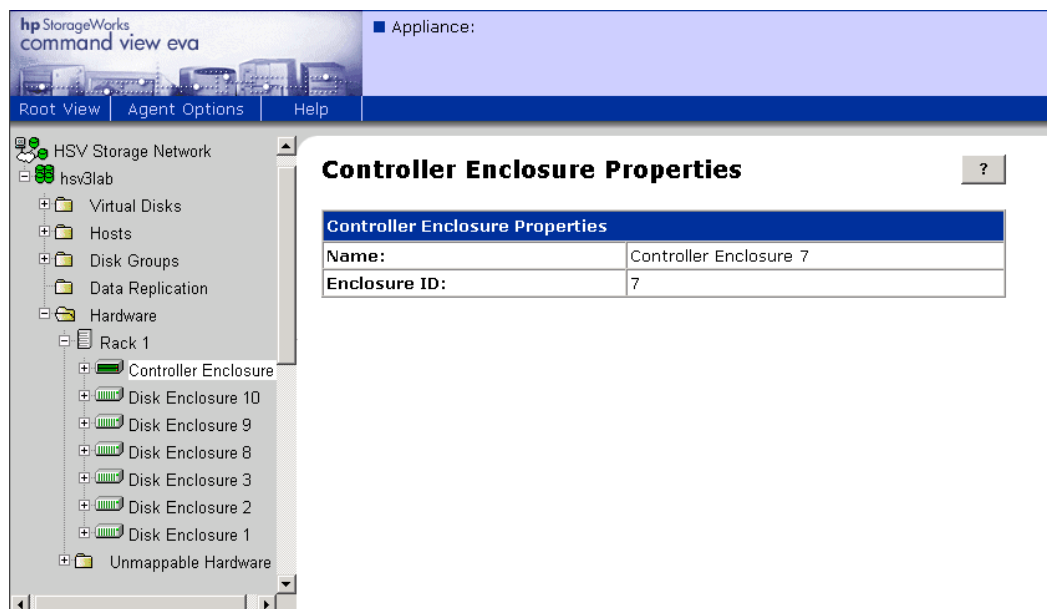
Various events may cause controller priority to reverse. Once the primary controller is determined, do not assume that it retains that designation.

A new feature of VCS V3.0 is that the Controller Termination Event Log will have events corresponding to both controllers.

Controller Enclosure

Selecting *Controller Enclosure* in the Navigation pane brings up the properties of the enclosure. This does not indicate the physical position of the controller.

Example



Controller A Properties

The following screen will show the controller A general, host port, device port, and enclosure properties.

Example

Controller Properties

Save changes		Locate		Shut down		Code load		?	
General		Host Ports		Device Ports		Enclosure			
Identification					Condition/State				
Name:		Controller A			Operational State:		<input checked="" type="checkbox"/> Good		
Type:		HSV Storage System			Cache Memory				
Manufacturer:		Hewlett-Packard Company			Operational State:		<input checked="" type="checkbox"/> Good		
Model Number:		HSV110			Write capacity:		256 MB		
Software version:		V3_W030414-3000			Read capacity:		512 MB		
Serial Number:		P4889B49ILV033			Mirror Path:		<input checked="" type="checkbox"/> Good		
World Wide Node Name:		5000-1FE1-5000-2CDD			Mirror Port State:		<input checked="" type="checkbox"/> Good		
UUID:		5005-08b4-0001-4523-0000-0000-0000-0000			Location				
					Enclosure number:		7		
Comments									
<input type="text"/>									

Both VCS and Command View EVA are object-oriented. Each object, whether hardware element or software structure, has a UUID. The various objects in a storage system use the UUID to identify each other.

Controller A Host Ports Properties

The Controller A Host Ports page shows the WWID, address (fabric-assigned, 24-bit), and operational and connection states for each host port of controller A.

Example

Controller Properties

Locate

Shut down

Code load

?

General

Host Ports

Device Ports

Enclosure

Port 1:


World Wide Name:

50001FE1 50002CD9


Address:

07-11-00

Operational State:

 Good

Connection state:

 Logged into fabric

Speed:

1Gb/s

Port 2:


World Wide Name:

50001FE1 50002CD8


Address:

06-11-00

Operational State:

 Good

Connection state:

 Logged into fabric

Speed:

1Gb/s

On the HSV controllers, note the last character of the port WWID:

- 9, controller A, host port 1
- 8, controller A, host port 2





Controller A Device Ports Properties

The Controller A Device Ports page shows the following:

- WWID of the controller on each loop
- Loop ID (AL_PA)
- Operational state for each port

Example

Controller Properties

Locate	Shut down	Code load	?
General	Host Ports	Device Ports	Enclosure
Loop Pair 1			
Loop A:			
World Wide Name:	00508B40 00145231		
Loop ID:	125		
Operational State:	 Good		
Loop B:			
World Wide Name:	00508B40 00145232		
Loop ID:	125		
Operational State:	 Good		
Loop Pair 2			
Loop A:			
World Wide Name:	00508B40 00145233		
Loop ID:	125		
Operational State:	 Good		
Loop B:			
World Wide Name:	00508B40 00145234		
Loop ID:	125		
Operational State:	 Good		

Determining the Extended Port LUN WWID

To determine the port LUN WWID:

- Remove the last digit (1, 2, 3, or 4).
- Precede the port WWID with a 6.

Example

The WWID of loop pair 1, port B is 00508B40 00100DC2. This means that the WWID of virtual disk Data Disk 1 (Properties Page displays 0001-3000-003A-0000) is:

6 + 00508B4000100DC + 0001-3000-003A-0000, which makes

6005-08B4-0001-00DC-0001-3000-003A-0000

Controller A Enclosure Properties

The Controller Properties page for Controller A shows the following:





- Blower speed
- Power
- Temperature sensors
- Cache battery status

Example

Controller Properties

Locate	Shut down	Code load	?
--------	-----------	-----------	---

General	Host Ports	Device Ports	Enclosure
---------	------------	--------------	-----------

Location		Temperature	
Enclosure number:	7	I2C Sensor 1:	68°F / 20°C
Blowers		I2C Sensor 2:	77°F / 25°C
FRU Type	Blower only	Overtemp Threshold:	131°F / 55°C
Blower 1		Cache Battery System	
Installed?	Yes	Operational State:	 Good
Actual speed:	4050 RPM	Battery voltage:	2.32 VDC
Blower 2		Battery Module 1	
Installed?	Yes	Installed?	Yes
Actual speed:	3930 RPM	Operational state:	 Good
Power		Battery Module 2	
12 VDC Voltage:	11.87 VDC	Installed?	Yes
5 VDC Voltage:	4.99 VDC	Operational state:	 Good
3.3 VDC Voltage:	3.28 VDC	Battery Charger	
2.5 VDC Voltage:	2.50 VDC	Operational state:	 Good
2.0 VDC Voltage:	2.01 VDC		

Controller Locate

You use the *Locate* button to locate a controller by turning its indicator on or off.

Example

Locate Hardware Device

OK	Locate ON	Locate OFF	?
----	-----------	------------	---

Click the appropriate button to turn the location indicator on the hardware device you have selected on or off. Click the **OK** button to exit. Be sure to turn the location indicator off before you exit.

The locate indicator on this hardware device is now **ON**.

Controller Shutdown

You use the *Shut down* button to shut down the controller.

Example

Shut Down Controllers

Cancel	?
--------	---

Controller Shutdown	
Restart	Shut down and restart this controller only
Power down	Shut down this controller and power off. Disable cache battery backup power.

System Shutdown	
Power down	Shut down both controllers and all disk drives and power them off. Disable all cache battery backup power.
0	Shutdown delay (0—60 minutes)

You use the *Power down* button to power down the controllers immediately. The disk enclosure will power down 60-90 seconds later. The default setting is 0. The user can also set the time up to one hour.

Disk Enclosure Properties

To display the Disk Enclosure Properties page, select a disk enclosure from the controller enclosure folder. Select the tabs on the page to display the following properties:

- General
- Power
- Cooling
- I/O ports and communication buses

General Disk Enclosure Properties

The properties that display on this page include:

- Loop pair where the disk enclosure is located
- Operational state
- EMU firmware version

Example

The example that follows shows that Disk Enclosure 10 is located on **LoopPair2**. The EMU firmware version is **02020070**.

Disk Enclosure Properties

Save changes		Locate		?	
General		Power		Cooling	
I/O-Comm					
Identification			LED Display		
Name:	Disk Enclosure 10		Operational state:	<input checked="" type="checkbox"/> Good	
World Wide ID:	5000-1FE1-0016-75E0		Language:	English	
Location			EMU		
Loop Pair:	LoopPair2		Operational state:	<input checked="" type="checkbox"/> Good	
Audible Alarm			Firmware version:	02020070	
Operational State:	<input checked="" type="checkbox"/> Good				
Alarm:	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled				
Last Enclosure Event					
No events logged					
Comments					

Disk Enclosure Power Properties

In the Disk Enclosure Power Properties page you can view properties for:

- Power Supply 1
- Power Supply 2

Click *Power Supply 1* or *Power Supply 2* to view properties for each.

Example for Power Supply 1

Disk Enclosure Properties

Locate		?	
General		Power	
Cooling		I/O-Comm	
Power Supply 1		Power Supply 2	
Power Supply 1 Properties			
Condition/State:		Alarm Thresholds:	
Operational state:	<input checked="" type="checkbox"/> Good	5 VDC Voltage Alarms	
Output:		OverVolt Warn:	5.7 VDC
5 VDC		OverVolt Critical:	5.7 VDC
Voltage:	5.5 Volts	UnderVolt Warn:	5.1 VDC
Current:	11.7 Amps	UnderVolt Critical:	5.1 VDC
12 VDC		5 VDC Current Alarms	
Voltage:	12.5 Volts	OverCurrent Warn:	20.0 Amps
Current:	4.5 Amps	OverCurrent Critical:	24.0 Amps
Alarms:		12 VDC Voltage Alarms	
AC Failure:	No	OverVolt Warn:	13.1 VDC
5 VDC		OverVolt Critical:	13.3 VDC
Voltage:	None	UnderVolt Warn:	11.9 VDC
Current:	None	UnderVolt Critical:	11.8 VDC
12 VDC		12 VDC Current Alarms	
Voltage:	None	OverCurrent Warn:	22.5 Amps
Current:	None	OverCurrent Critical:	26.5 Amps

Disk Enclosure Cooling Properties

On the Disk Enclosure Cooling Properties page, you can view properties for:

- **Sensor-General** — Status and temperature of disk drive.
- **Sensor-Thresholds**
- **Blowers**
 - **Blower 1** — Shows the status of the blower on the left (looking at the back) of the disk enclosure.
 - **Blower 2** — Shows the status of the blower on the right-hand corner in front of the enclosure.
 - Characteristics include operational state, requested speed (as requested by the EMU), and actual speed.

Click each name at the top to view its properties.

Sensor — General Properties

This page displays the state and temperature of the disk drives noted by the temperature sensors in the disk enclosure. There are 17 temperature sensors located as follows:

- One in each enclosure power supply
 - Power Supply 1 (PS1TEMP)
 - Power Supply 2 (PS2TEMP)
- One in the EMU (EMUTEMP)
- One in each disk drive (DiskxTemp)

The statuses are:

- **OK** — Temperature is within specification
- **Critical** — Temperature is outside the specification, making continued normal operation impossible.
- **Noncritical** — Temperature is outside the specification, but normal operation is still possible.
- **NotInstalled** — The module containing the sensor is not installed.
- **NotAvailable** — The power supply is installed, but the AC is unplugged or failed. This applies only to the power supply sensors.

Example

This example shows the Sensor-General properties, which include status and temperature detected by each sensor.

Disk Enclosure Properties

Locate

?

General

Power

Cooling

I/O-Comm

Sensor-General

Sensor-Thresholds

Blower 1

Blower 2

Temperature Sensor General Properties

Sensor	Status	Temperature	Alarms
PS1TEMP	OK	96.8°F / 36.0°C	None
PS2TEMP	OK	91.4°F / 33.0°C	None
EMUTEMP	OK	77.0°F / 25.0°C	None
DISK1TEMP	OK	87.8°F / 31.0°C	None
DISK2TEMP	OK	82.4°F / 28.0°C	None
DISK3TEMP	OK	82.4°F / 28.0°C	None
DISK4TEMP	OK	75.2°F / 24.0°C	None
DISK5TEMP	OK	77.0°F / 25.0°C	None
DISK6TEMP	OK	73.4°F / 23.0°C	None
DISK7TEMP	OK	75.2°F / 24.0°C	None
DISK8TEMP	OK	80.6°F / 27.0°C	None
DISK9TEMP	OK	80.6°F / 27.0°C	None
DISK10TEMP	OK	84.2°F / 29.0°C	None
DISK11TEMP	OK	80.6°F / 27.0°C	None
DISK12TEMP	OK	75.2°F / 24.0°C	None
DISK13TEMP	OK	87.8°F / 31.0°C	None
DISK14TEMP	OK	93.2°F / 34.0°C	None

Sensor — Threshold Properties

The following shows Sensor-Threshold properties. These are not settable.

Example

Disk Enclosure Properties

Locate					?		
General		Power		Cooling		I/O-Comm	
Sensor-General		Sensor-Thresholds		Blower 1		Blower 2	
Temperature Sensor Thresholds							
	Undertemperature			Overtemperature			
Sensor	Warning Level	Critical Level	Warning Level	Critical Level			
PS1	48.2°F / 9.0°C	41.0°F / 5.0°C	122.0°F / 50.0°C	140.0°F / 60.0°C			
PS2	48.2°F / 9.0°C	41.0°F / 5.0°C	122.0°F / 50.0°C	140.0°F / 60.0°C			
EMU	48.2°F / 9.0°C	41.0°F / 5.0°C	111.2°F / 44.0°C	118.4°F / 48.0°C			
DISKSLOT1	48.2°F / 9.0°C	41.0°F / 5.0°C	114.8°F / 46.0°C	134.6°F / 57.0°C			
DISKSLOT2	48.2°F / 9.0°C	41.0°F / 5.0°C	114.8°F / 46.0°C	134.6°F / 57.0°C			
DISKSLOT3	48.2°F / 9.0°C	41.0°F / 5.0°C	114.8°F / 46.0°C	134.6°F / 57.0°C			
DISKSLOT4	48.2°F / 9.0°C	41.0°F / 5.0°C	114.8°F / 46.0°C	134.6°F / 57.0°C			
DISKSLOT5	48.2°F / 9.0°C	41.0°F / 5.0°C	114.8°F / 46.0°C	134.6°F / 57.0°C			
DISKSLOT6	48.2°F / 9.0°C	41.0°F / 5.0°C	114.8°F / 46.0°C	134.6°F / 57.0°C			
DISKSLOT7	48.2°F / 9.0°C	41.0°F / 5.0°C	114.8°F / 46.0°C	134.6°F / 57.0°C			

Blower Properties

The blower properties allow you to see the operational state and speed of blowers 1 and 2.

Example showing Blower 1

Disk Enclosure Properties

Locate				?	
General		Power		Cooling	I/O-Comm
Sensor-General		Sensor-Thresholds		Blower 1	Blower 2
Blower 1					
Condition/State:					
Operational state:		<input checked="" type="checkbox"/> Good			
Speed:					
Actual Speed:		Second Lowest			

Disk Enclosure I/O Communication Properties

On the Disk Enclosure I/O communications page, you can view properties for:

- I/O ports
 - Module and link operational state
 - Transceiver installation state
- Communication buses
 - Operational state of internal and external buses
 - Enabled or disabled status

Click *I/O Ports* or *Communication Buses* at the top to view properties of each.

Example — I/O Communication

This example shows the I/O Ports properties:

- Module A and B operational statuses are OK.
- Link operational status for each port is good.
- Transceiver installation status for each port is OK.

Disk Enclosure Properties

Locate

?

General

Power

Cooling

I/O-Comm

I/O Ports

Communication Buses

I/O Modules

	Module A:	Module B:
Operational state:	Good	Good

Fibre Channel Ports


	Port A1:	Port A2:	Port B1:	Port B2:
Link Operational state:	Good	Good	Good	Good
SFP Installation state:	OK	OK	OK	OK

Note

Point out that in switched hardware configurations (loop switches) for VCS version 2.002 and before, if the unused I/O module transceiver slot actually has a transceiver, then a link state of Failed will be reported. If the unused I/O module does **not** contain a transceiver, the link state will be Not Installed or Not present.

I/O Module Port Numbers

Command View EVA refers to the physical ports oppositely to their physical configuration, that is, the top and bottom hardware ports are port 1 and port 2, respectively; the top and bottom ports viewed through Command View EVA are port 2 and port1, respectively.



Disk Enclosure Communication Buses

The **Locate** button at the top of this page enables you to flash the locate LED in the lower right-hand corner on the front of the enclosure.

Example

Disk Enclosure Properties

Locate

?

General

Power

Cooling

I/O-Comm

I/O Ports

Communication Buses

EMU Buses

	Internal Bus:	Rack Bus:
Operational state:	Good	Good
Enabled/Disabled:	Enabled	Enabled

Disk Enclosure Bay Properties

To display the Disk Enclosure Bay Properties, select a bay from the disk enclosure folder and then select a disk bay. Select the tabs on the page to display the following properties:

- Disk bay properties
- Disk drive properties

Disk Bay Properties

On the Disk Enclosure Bay Properties, Disk Bay page, you can view the location of each disk bay in a disk enclosure and determine if it is currently populated.

Example

This example shows where Disk Bay 11 is located in the disk enclosure.

Disk Enclosure Bay Properties

Save changes		Code load		Locate		?															
Disk Bay				Disk Drive																	
Bay Identification				Condition/State																	
Name:		Disk Bay 11		Operational State:		<input checked="" type="checkbox"/> Populated (data disk)															
Bay ID:		11																			
Enclosure ID:		10		Location																	
Loop Pair:		LoopPair2		<table border="1"> <tr> <td>1</td><td>2</td><td>3</td><td>4</td><td>5</td><td>6</td><td>7</td><td>8</td><td>9</td><td>10</td><td>11</td><td>12</td><td>13</td><td>14</td> </tr> </table>				1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	2	3	4	5	6	7	8	9	10	11	12	13	14								
Comments																					
<input type="text"/>																					

The *Code load* button enables you to download firmware to the EMU.

Disk Enclosure Bay Disk Drive Properties

On the Disk Enclosure Bay Properties, Disk Drive page, you can view properties for a specific disk drive located in the bay such as:

- Identification
- System information
 - Disk drive usage
 - Disk group name
 - Disk Group
 - Occupancy
 - RSS ID and RSS Index (new for VCS V3.0)
- Condition and state of disk drive and loops
- Physical information
 - Physical drive type
 - Firmware revision
 - Formatted capacity
- Bay location

Example

This example shows that Disk 053 is located in Bay 11. The physical characteristics are:

- Type — Fibre Channel Disk
- Firmware revision — 3BE6
- Formatted capacity — 33.91GB

Disk Enclosure Bay Properties

Ungroup	Locate	Code load	Remove	?
---------	--------	-----------	--------	---

Disk Bay	Disk Drive
----------	------------


Identification		Condition/State	
Name:	Disk 053	Operational state:	<input checked="" type="checkbox"/> Good
Loop Pair:	LoopPair2	Migration state:	<input checked="" type="checkbox"/> Not migrating
Node World Wide Name:	2000-0004-CFA7-DE74	Failure prediction:	<input checked="" type="checkbox"/> No
UID:	2000-0004-cfa7-de74-0000-0000-0000-0000	Media accessible:	<input checked="" type="checkbox"/> Yes
Physical		Ports	
Type:	Fibre Channel Disk	Loop A:	
Manufacturer:	COMPAQ	Port World Wide Name:	2000-0004-CFA7-DE74
Model number:	BF03654564	Assigned LUN:	0
Firmware version:	3BE6	Loop ID:	80
Formatted capacity:	33.91 GB	Operational state:	<input checked="" type="checkbox"/> Good
		Loop B:	
		Port World Wide Name:	2000-0004-CFA7-DE74
		Assigned LUN:	0
		Loop ID:	80
		Operational state:	<input checked="" type="checkbox"/> Good
System			
Requested usage:	Grouped		
Actual usage:	Grouped		
Disk group:	dick		
Occupancy:	1.67 GB		
RSS ID:	4		
RSS index:	5		
Location			
Enclosure ID:	10		
Bay ID:	11		

Disk Drive Removal

To remove a disk, use the *Remove* button.

Example

Remove Disk



The location indicator on the disk drive you have selected is flashing.

Remove the disk drive from its enclosure and click the **OK** button to complete the removal operation.



Important

Always use the *Remove* button before removing drives. If the drive is a member of a disk group, you will get a message to remove it from the disk group before removing it.

Learning Check

1. Name the four management agent options that you can set from Command View EVA.
.....
.....
2. Which of the following elements is required to set a storage system password in Command View EVA?
 - a. Storage System World Wide Name
 - b. Storage system name assigned when the system is initialized
 - c. Ten digit password
 - d. Host system World Wide Name
3. How do you navigate to the Rack Properties page?
.....
.....
4. From the Controller Properties page, you can view which of the following?
 - a. Blower properties
 - b. VCS revision number
 - c. Power supply 1 and power supply 2 properties
 - d. Temperature sensors
5. For the HSV controllers what is the last character of the port WWID for each of the following?
 - a. controller A, host port 1
 - b. controller A, host port 2
 - c. controller B, host port 1
 - d. controller B, host port 2

Overview

This module discusses the resources available for detecting, interpreting, and analyzing Enterprise Virtual Array storage system errors. It includes a look at the common event logs, that is, the Management Agent Event Log, Controller Event Log, Controller Termination Event Log, and NT Event Log, how to decode and process the information using various tools, and how to interpret them. Event filtering, logging, and downloading are described. This module also explains how monitoring tools, such as Proactive Remote Services (PRS), are used with the Enterprise Virtual Array. Common diagnostics are described.

Objectives

After completing this module you should be able to:

- List activities that can cause events to be written to a log.
- Identify the types of monitoring tools used for event reporting and fault analysis.
- Describe the four types of event logs and how they are viewed.
- Describe the process used to view events from a downloaded controller event file.
- Name the fields in the controller event logs.
- Describe the tools and processes used to translate and read controller events.
- Describe the steps involved to set up notification events for SNMP traps.
- List the diagnostics used for the HSV controller.

Event Handling

An event is any change that is significant to the storage system. Events include:

- State change in hardware
- State change in a logical element, such as a virtual disk
- Procedure completion
- Environmental change
- Operational failure
- Management agent software events

The storage management appliance and Command View EVA log significant events that occur within the Command View EVA software, and those that occur within the initialized storage system.

The basic flow of controller events and disk enclosure events occurs as follows:

1. A software component (such as the Fault Manager, Fibre Channel Services, and so forth) reports an event by passing an event-specific data packet to the Fault Manager component in the VCS software.
2. The Fault Manager component adds some information that is common to all events, such as the UUID of the reporting controller and the report time, to the packet. The Fault Manager then places the packet in the Controller Event Log in the storage system metadata. This metadata is stored on a virtual disk called the Cell State Logical Disk (CSLD) on the Enterprise storage system.
3. Command View EVA polls the storage system for events at 45-second intervals and when the Controller Event Log is opened by a user (either for viewing on the browsing station monitor or downloading of the log file). Therefore, the latest events are retrieved when a user requests the Controller Event Log file.

Note

Refer to the Service Manual for the process used to report controller termination events.

Event Logs

Command View EVA logs display events reported by the Enterprise Virtual Array controllers and disk enclosure, and those that occur within the Command View EVA software.

You can display four types of event logs, three of which you can display using Command View EVA:

- Management Agent Event Log
- Controller Event Log
- Controller Termination Event Log

You can view the fourth, the NT Event Log, through terminal services; by using the Remote Insight Board (RIB); by attaching a keyboard, monitor, and mouse to the appliance; or by using a processing tool.

Note

The Command View EVA-to-controller communication is done through StorageCell Management Interface (SCMI) commands. You will need to know how to interpret SCMI events.

All of the log files, except the NT Event Log, are downloaded using the web browser **Get log file** or **Get event file** buttons. The Management Agent Event Log is saved as an ASCII file, while the Controller and Controller Termination Event Logs are saved as binary files. These binary files are decoded using tools such as Event View EVA, which are tools for HP Services only.

Monitoring and Notification Tools

There are several tools that allow monitoring and reporting of events in the Enterprise Virtual Array environment, including:

- Proactive Remote Services (PRS)
- Simple Network Management Protocol (SNMP) traps
- System Event Analyzer (SEA)
- Open Service Event Manager (OSEM)

Proactive Remote Services

PRS is an end-to-end remote service offering that allows “phone home” call generation and service. How PRS is used depends on whether the latest service offering, SEA, is used in the environment.

If SEA is installed on the storage management appliance controlling the storage system, then it directly monitors the NT Event Log. When it determines that a serviceable event has occurred, it sends a problem report to PRS.

If SEA is not installed on the storage management appliance controlling the storage system, PRS receives SNMP traps directed by Command View EVA to the server where the PRS software is running. PRS looks up the trap in the Management and Service Action Database (MSADB) filter set and Management Information Base (MIB) to determine if there is a serviceable event. The MSADB contains the filtering rules that determine which events get forwarded to the service provider. It also contains additional information to be added to SNMP traps. This additional information is determined by the serviceability engineer when he determines which traps are serviceable issues, and, if so, what the correct service response (rule) is. Each rule is extended with supporting information such as whether the problem is hardware or software, whether it is a warranty or contract issue, FRU identifiers, and so on.

Note

There is more information on PRS, and how it is implemented with the Enterprise Virtual Array, later in this module. There is also PRS training required available on the HP Services Learning Utility.

Event Forwarding Using SNMP Traps

All events, except controller termination events, can be forwarded to an arbitrary receiver such as PRS using a SNMP MIB. Termination events are not forwarded because they are too large. SNMP events contain summary information.

This feature is optional. If you want to enable SNMP forwarding, you must use Command View EVA to specify the traps that will be enabled through the Set Event Notification Options page and you must specify the SNMP receivers using the fully qualified host name and optional port number on the Modify Host Notification List page. The definition of the traps is in the **cpqhsv110v3.mib** file.

System Event Analyzer

System Event Analyzer (SEA) is a fault analysis utility that provides basic analysis for single and multiple error/fault events as well as complex analysis and corrective action recommendation using analysis rules.

The Enterprise Virtual Array with VCS V3.0 is supported by and **requires** the V4.2 Web-Based Enterprise Services (WEBES)/SEA suite to perform PRS operations. HP Services strongly recommends the use of WEBES/SEA in VCS V2.x environments as well.

Open Service Event Manager

Open Service Event Manager (OSEM) is a service offering that allows customer systems to automatically send service event notifications as email messages within their own intranet.

OSEM V1.2 and higher supports the EVA attached to storage management appliances. Setup of OSEM in the EVA environment is very similar to PRS setup, covered in a later topic. Apply the same principles for PRS to the OSEM architecture.

Viewing Event Logs

There are two ways to view event logs:

1. Directly on the Command View EVA console
 - The controller and controller termination events show an abbreviated form of the event.
 - Clicking the *More details* icon downloads the full event to the console.
2. Download to a file and view with HP Service applications such as Event View EVA
 - Contains the same information
 - Uses a slightly different format

A limited subset of HTML tags may appear in some descriptions to aid in formatting. These tags also appear in the parse file itself. The following descriptions may contain these tags:

- Event code descriptions
- Termination code descriptions
- Corrective action code descriptions

HTML tags that may appear in some event descriptions are listed in the following table.

HTML Tag	Significance
	Beginning of an unnumbered list
	Beginning of an unnumbered list item
	End of an unnumbered list
<P>	Paragraph marker

Viewing Logs from the Command View EVA Console

When you view an event log from the console, note the following:

- Use the View Events screen and select your selected event log type.
- A flashing Loading Events indicator appears.
- The display may take some time to load.
- You can view ranges of controller events.
- The controller event log file may contain up to 20,000 events.
- HP does not recommend using the **More details** icon, so use the **Get log file** or **Get event file** buttons instead.



Important

Selecting *More details* on an event can take up to seven minutes and you may encounter a significant delay. Therefore, be patient when downloading and viewing events and do not power click.

Viewing Logs from a Downloaded File

When you view an event log from a downloaded file, note the following:

- You get a detailed view of the event file.
- You use the **Get log file** or **Get event file** buttons and save the file to disk.
- Select the appropriate application to view the log:
 - For controller and controller termination events (raw hex data), you can view the logs with an ASCII text editor such as Notepad or WordPad, or with MS Excel (if the option is selected).
 - For Management Agent events (ASCII text), you can view the log with an ASCII text editor such as Notepad or WordPad.

View Events Page

To launch the View Events page, select the storage system and click *View events* at the top of the content pane.

Initialized Storage System Properties

Save changes		Set options		View events		Uninitialize		?	
Code load		Shut down							

Identification		Condition/State	
Name:	hsv3lab	Operational state:	<input checked="" type="checkbox"/> Good (Initialized)
Node World Wide Name:			
5000-1FE1-5000-2CD0			
UUID:			
6005-08b4-0001-4523-0002-c000-017e-0000			
Licensed features		System	
Basic:	Yes	Type:	HSV110
Snapshot:	Yes	Version:	3000
Data replication:	Yes	Console LUN ID:	0
		Time:	18 Apr 2003 08:22:58
Policies		Capacity	
Device addition:	Manual	Total:	978.30 GB
Disk replacement delay:	1 mins	Used:	271.64 GB
		Available:	706.66 GB
Comments			
<input type="text"/> <input type="button" value="OK"/> <input type="button" value="Cancel"/>			

The View Events page displays as shown here.

View Events	
OK	?
<input type="checkbox"/> Management Agent Event Log <input type="checkbox"/> Controller Event Log <input type="checkbox"/> Controller Termination Event Log	

Click the appropriate button to display the desired log.

The following buttons may display at the top of the log page, depending on which log you request:

- OK — Exits the log.
- Get event (log) file — Displays the raw data log file. Not all log files are in readable format.
- Filter events — Launches the Filter Events page.
- Get parse file — Downloads the parse file.
- Send parse file — Copies the parse from the local server to the storage management appliance.
- Clear log — Clears the log file.

You can only view one log at a time. Once you have finished viewing the log, click *OK* at the top of each log page.

Management Agent Event Log

A Management Agent event is a significant occurrence within Command View EVA, such as the following:

- Change in state of management agent software
- Change in state of an initialized storage system

Management Agent Event Log Features

The Management Agent Event Log has the following features:

- Reports configuration changes.
- Supports 1000 entries per storage cell (initialized storage system).
- Filterable based on category for viewing
- Supports SNMP forwarding and posting to the NT Event Log
- Saved as an ASCII text file when **Get event file** button is used.

Management Agent Event Log Characteristics

The Management Agent Event Log has the following characteristics:

- Log file name is:
 - **UUID.bridgelog.txt** for Command View EVA versions before V3.0.
 - **handle.bridgelog.txt** for Command View EVA V3.0.

Note

Handle represents a string of numbers uniquely identifying the storage system. The file **wwn2handlemap.txt** on **c:\hsvmafiles** contains a list of WWNs of the storage systems on the storage management appliance and their corresponding handles.




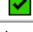



- Log file is located on the storage management appliance in:
 - **c:\hsvmafiles** for Command View EVA versions before V3.0
 - **c:\hsvmafiles\WWN** for Command View EVA V3.0
- Log file is initialized or cleared when the:
 - Storage system is initialized.
 - Storage system UUID associated with the WWN has changed.
 - *Clear Log* button is clicked.

- When the log file reaches the maximum number of events, contents are saved to:
 - ***UUID.lastbridgelog.txt*** for Command View EVA versions before V3.0
 - ***handle.lastbridgelog.txt*** for Command View EVA V3.0

Management Agent Event Log Example

The following is a sample Management Agent Event Log:

Management Agent Events

OK	Get event file	Filter events	Clear log	?
Date/Time	Severity	Alarm	Code	Appliance/Storage System
13:26:58 4-Apr-2003		Yes	24157	SCell:hsv3lab
SC Event Code: 9 cc 51 5 - The state of the internal Logical Disk associated with the Virtual Disk identified in the handle field has transitioned to the DEVICE DATA LOST state.				
13:26:59 4-Apr-2003		No	9012	HSV3SANAPP
Operation on a fast virtual disk deletion completed				
13:26:39 4-Apr-2003		No	10024	HSV3SANAPP
Disk Group Data Lost Error Resolution Sequence 1 - Marked				
13:26:48 4-Apr-2003		No	10025	HSV3SANAPP
Disk Group Data Lost Error Resolution Sequence 2 - Affected Units Deleted				
13:23:04 4-Apr-2003		Yes	24077	SCell:hsv3lab
SC Event Code: 9 6 bf 5 - The state of the Volume identified in the handle field has transitioned to the MISSING state.				
13:23:04 4-Apr-2003		Yes	24130	SCell:hsv3lab
SC Event Code: 9 3e 42 e - The Physical Disk Drive identified in the handle field has disappeared.				
13:23:04 4-Apr-2003		Yes	24154	SCell:hsv3lab
SC Event Code: 9 c9 51 5 - The state of the Disk Group identified in the handle field has transitioned to an INOPERATIVE state.				

When you click the *Get event file* button, you can save the file to a specific location.

Management Agent Events

OK Get event file Filter events Clear log ?				
Date/Time	Severity	Alarm	Code	Appliance/Storage System
13:26:58 4-Apr-2003		SC Eve		Virtual
13:26:59 4-Apr-2003		Operat		
13:26:39 4-Apr-2003		Disk Gr		
13:26:48 4-Apr-2003		Disk Gr		
13:23:04 4-Apr-2003		SC Eve		transitioned to the MISSING state.
13:23:04 4-Apr-2003		Yes	24130	SCell:hsv3lab
13:23:04 4-Apr-2003		Yes	24154	SCell:hsv3lab

Here is a file opened in WordPad.

```

GetEventFile - WordPad
File Edit View Insert Format Help
[Icons]
4 23 2003 14 17 15 2055 0 0 0 0 1 1 2Kool NSASubsystem *
8.7.16.1610942644.65828.524288.36175872E~n~d000000710b408056024010100000008000
0006a02E~n~d
4 23 2003 14 17 15 2048 0 0 0 0 0 1 2Kool NsaControl *Nsa object: - deleted
successfully
4 23 2003 13 57 40 2056 0 0 0 0 1 1 2Kool NsaControl *
8.7.16.1610942644.65828.524288.36175872E~n~d000000710b408056024010100000008000
0006a02E~n~d8E~n~d
4 23 2003 13 57 40 2051 0 0 0 0 0 0 2Kool NsaControl *Create Folder -
Completed
4 10 2003 22 26 19 17007 0 4 0 0 0 0 2Kool DriveCodeLoadMgr *Drive Code Load
Drive Firmware Update Failed
4 10 2003 22 26 19 0 0 0 0 0 1 0 2Kool NscManager *Successful Status
4 9 2003 18 7 4 28225 0 4 1 1 0 0 SCell:Local Subsystem *SC Event Code: d 83
89 11 - A Drive Enclosure Environmental Monitoring Unit has experienced a
hardware failure.
4 4 2003 16 55 45 0 0 0 0 0 1 0 2Kool NSASubsystem *Successful Status
4 4 2003 16 55 45 0 0 0 0 0 1 0 2Kool NSASubsystem *Successful Status
4 4 2003 16 55 44 0 0 0 0 0 1 0 2Kool NSASubsystem *Successful Status
4 4 2003 16 55 44 0 0 0 0 0 1 0 2Kool NSASubsystem *Successful Status
4 4 2003 16 55 44 0 0 0 0 0 1 0 2Kool NSASubsystem *Successful Status
4 4 2003 16 55 44 0 0 0 0 0 1 0 2Kool NSASubsystem *Successful Status
For Help, press F1
NUM

```








Event Log Filtering

Command View EVA provides a filtering function that filters logs based on your specifications. This does not change the log file organization, only the display presentation. You can filter events by:

- Date
- Time
- Severity
- Event code
- Status
- Appliance name

You access the Filter Events page by clicking the *Filter events* button.

Management Agent Events

OK	Get event file	Filter events	Clear log	?
Date/Time	Severity	Alarm	Code	Appliance/Storage System
	Description			
13:26:58 4-Apr-2003		Yes	24157	SCell:hsv3lab
	SC Event Code: 9 cc 51 5 - The state of the internal Logical Disk associated with the Virtual Disk identified in the handle field has transitioned to the DEVICE DATA LOST state.			
13:26:59 4-Apr-2003		No	9012	HSV3SANAPP
	Operation on a fast virtual disk deletion completed			
13:26:39 4-Apr-2003		No	10024	HSV3SANAPP
	Disk Group Data Lost Error Resolution Sequence 1 - Marked			
13:26:48 4-Apr-2003		No	10025	HSV3SANAPP
	Disk Group Data Lost Error Resolution Sequence 2 - Affected Units Deleted			
13:23:04 4-Apr-2003		Yes	24077	SCell:hsv3lab
	SC Event Code: 9 6 bf 5 - The state of the Volume identified in the handle field has transitioned to the MISSING state.			
13:23:04 4-Apr-2003		Yes	24130	SCell:hsv3lab
	SC Event Code: 9 3e 42 e - The Physical Disk Drive identified in the handle field has disappeared.			
13:23:04 4-Apr-2003		Yes	24154	SCell:hsv3lab
	SC Event Code: 9 c9 51 5 - The state of the Disk Group identified in the handle field has transitioned to an INOPERATIVE state.			

The Filter Events page looks like the following.

Filter Events

View filtered events

Cancel

?

You can change the way your event log is displayed using the filter parameters below. To filter your event log, change any parameter's default value to the value of your choice. To refine your event display, you may use as many filter parameters as you wish.

To display the filtered log, click the **View Filtered Events** button.

Filter by date

☒ Show events occurring on all dates

☐ Show only events occurring:

from / /

to / /
(mm/dd/yyyy)

Filter by time

☒ Show events occurring at all times

☐ Show only events occurring:

from :

to :
(hh:mm, 24-hr format)

Filter by severity

Show only events with severity code:

All

Filter by event code

Show only events with event code:

All (Enter 4/5-digit number)

Filter by status

Show only events with an alarm status

of: All

Filter by appliance name

Show only events occurring at

appliance: All
(Enter appliance name)

Controller Event Log

Controller events are significant occurrences to any hardware or software component of the storage system. The controller pair reports these events. The management agent polls the controller pair for event log information on a regular basis.

The Controller Event Log reports storage system state changes such as:

- EMU-detected events
- Disk enclosure and drive events
- Configuration changes such as creation of a disk group
- Controller rebooting

Controller Event Log Features

The Controller Event Log has the following features:

- Varies in the number of allowable events (has a buffer of 2MB)
- Stored on a virtual disk in the storage system called the Cell State Logical Disk (CSLD)
- Filterable on importance level and category
- Supports forwarding to SNMP and posting in the NT Event Log
- Saved as a binary file when the **Get log file** button is used

The **Get log file** button in this log retrieves the raw data file but it is not in readable format. You would need an HP Services tool such as Event View EVA to translate the raw event data.

Controller Event Log Characteristics









The Controller Event Log has the following characteristics:

- Log file name is:
 - **WWN.sceventfile.ascii** for Command View EVA versions before V3.0
 - **WWN.sceventfile.bin** for Command View EVA V3.0
- Log file is located on the storage management appliance in:
 - **c:\hsvmafiles** for Command View EVA versions before V3.0
 - **c:\hsvmafiles\WWN** for Command View EVA V3.0
- Log file is initialized or cleared when the:
 - Storage system is initialized.
 - Storage system UUID associated with the WWN has changed.
- When the log file reaches the maximum number of events, contents are saved to:
 - **WWN.last.sceventfile.ascii** for Command View EVA versions before V3.0
 - **WWN.last.sceventfile.bin** for Command View EVA V3.0

Controller Event Log Example

The following is a sample Controller Event Log.

Controller Events (Initialized system)

OK	Get log file	Get parse file	Send parse file	?
Display Range:	1 - 400	Previous group	Next group	
Controllers:	Controller B: 5005-08b4-0001-4529-0000-0000-0000-0000 Controller A: 5005-08b4-0001-4523-0000-0000-0000-0000			
Date/Time Controller	Severity	Event Code	Sequence #	
Description				
17:43:23:368 15-Apr-2003 Controller A		06014a08	#1965	
A Fibre Channel port on the HSV110 controller has failed to respond. ✖ Corrective action code: 4a				
17:42:53:806 15-Apr-2003 Controller A		09070005	#1964	
The state of the Fibre Channel port identified in the attribute.value.str field and located on the rear panel of the HSV110 controller identified in the handle field has transitioned to the NORMAL state. ✖ Corrective action code: 00				
17:42:11:21 15-Apr-2003 Controller A		0df00011	#1963	
The status has changed on one or more of the drive enclosures. ✖ Corrective action code: 00				
17:41:48:742 15-Apr-2003 Controller A		060f4013	#1962	
The Drive Enclosure Environmental Monitoring Unit is able to communicate with a physical disk drive but this HSV110 controller is unable to communicate with that physical disk drive on the Fibre Channel bus. ✖ Corrective action code: 40				

The fields on this log are described in the next table.

Field	Description
Date/Time Controller	Date and time the event occurred on the specific controller.
Event Code	The event code consisting of the SWC ID, event number, CAC, and the EIP type.
Sequence #	Number representing the sequence of the event.
Description	The problem that generated the event. Most of the useful information is in the first few lines of the description, but you can click the icon next to More details in the description column to display more details.

When you click the *More details* icon, you see the results on your console.

```

GetSCEventAddData[1] - WordPad
File Edit View Insert Format Help

The state of the Logical Disk identified in the handle field has changed. The value.ul1
field contains the new state. The value.ul2 field contains the old state. The state
values that may be found in the value.ul1 and value.ul2 fields are as follows:
<UL>
<LI>1 = Normal
<LI>2 = Replacement delay in progress
<LI>3 = Redundancy lost restore in progress
<LI>4 = Redundancy lost restore stalled
<LI>5 = Failed
<LI>6 = Creation in progress
<LI>7 = Snapshot is inoperative due to overcommit
<LI>8 = Deletion in progress
<LI>9 = Capacity change in progress
</UL>

EIP05 - Storage System Management Interface Entity State Change
The state of a Storage System Management Interface entity has changed.

Event Log Packet Event Specific Information

flags (Flags)
Flag field bits: 00000011

Time has been set on this HSV110: TRUE
Time has been synchronized with all HSV110s in the Storage System: TRUE
Event sequence number reset occurred: FALSE
Event reported out of sequence due to Final Event Block reconciliation: FALSE
For Help, press F1

```



Important

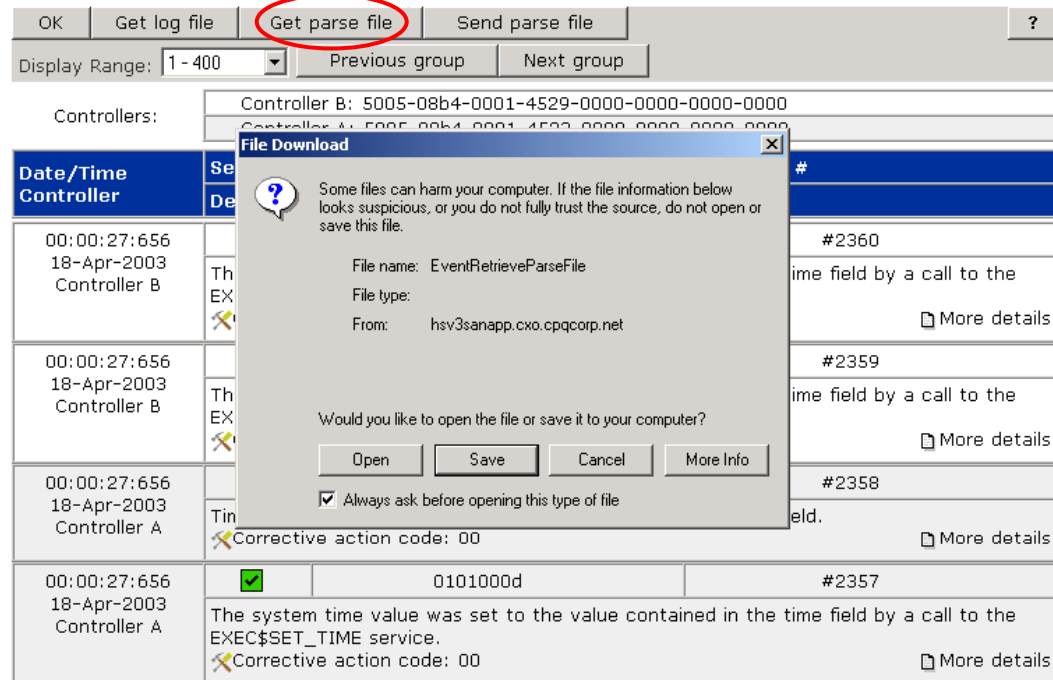
Using *More details* is not the recommended way to view the log. You may get 600 or more pages of information. Instead, download the log file and translate it using an HP Services tool such as Event View EVA.

Get Parse File Button Example

The parse file defines the format and content of the raw controller and termination events, including their codes, corrective actions, and structure. Details about parse files are given when describing translation and decoding tools in a later topic.

Download the parse file for use with an HP Services tool from the controller by using the **Get parse file** button. The file is also saved in **c:/hsvmafiles/** on the storage management appliance.

Controller Events (Initialized system)



The screenshot shows the 'Controller Events (Initialized system)' window. At the top, there are buttons: 'OK', 'Get log file', 'Get parse file' (circled in red), and 'Send parse file'. Below these are 'Display Range: 1 - 400', 'Previous group', and 'Next group'. The main area displays a table of controller events. A 'File Download' dialog box is open in the foreground, displaying the following information:

- File name: EventRetrieveParseFile
- File type:
- From: hsv3sanapp.cxo.cpqcorp.net

The dialog box asks: 'Would you like to open the file or save it to your computer?' and provides buttons for 'Open', 'Save', 'Cancel', and 'More Info'. There is also a checkbox labeled 'Always ask before opening this type of file' which is checked.

Date/Time	Controller	Se	De	#
00:00:27:656 18-Apr-2003 Controller B		Th	EX	#2360
00:00:27:656 18-Apr-2003 Controller B		Th	EX	#2359
00:00:27:656 18-Apr-2003 Controller A		Tin		#2358
00:00:27:656 18-Apr-2003 Controller A				#2357

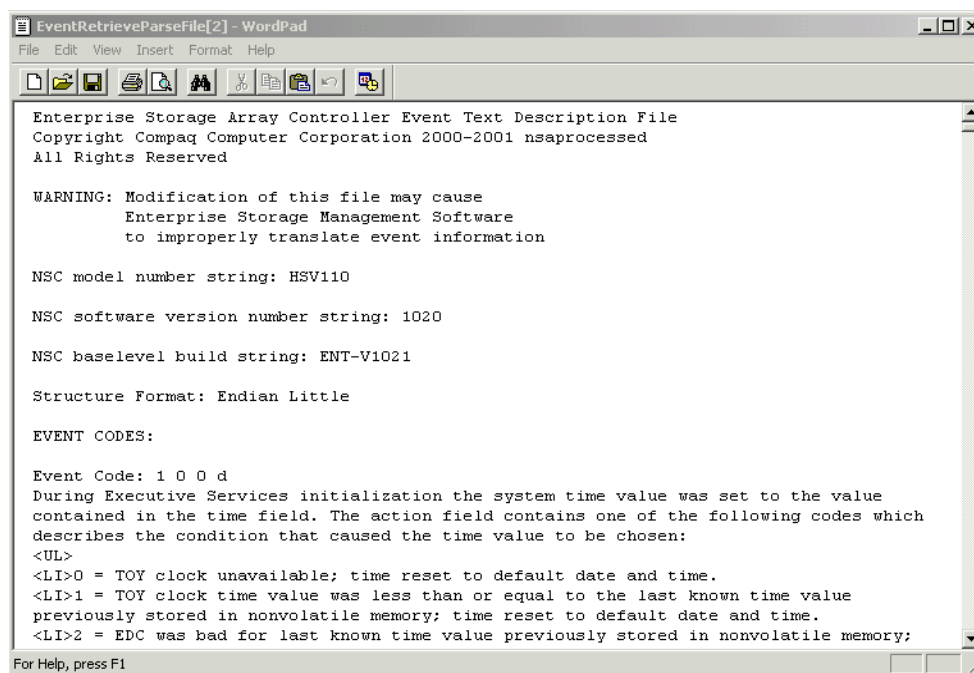
Corrective action code: 00

The system time value was set to the value contained in the time field by a call to the EXEC\$SET_TIME service.

Corrective action code: 00

Parse File Example

The following is a sample parse file.



```
Enterprise Storage Array Controller Event Text Description File
Copyright Compaq Computer Corporation 2000-2001 nsaprocesed
All Rights Reserved

WARNING: Modification of this file may cause
         Enterprise Storage Management Software
         to improperly translate event information

NSC model number string: HSV110

NSC software version number string: 1020

NSC baselevel build string: ENT-V1021

Structure Format: Endian Little

EVENT CODES:

Event Code: 1 0 0 d
During Executive Services initialization the system time value was set to the value
contained in the time field. The action field contains one of the following codes which
describes the condition that caused the time value to be chosen:
<UL>
<LI>0 = TOY clock unavailable; time reset to default date and time.
<LI>1 = TOY clock time value was less than or equal to the last known time value
previously stored in nonvolatile memory; time reset to default date and time.
<LI>2 = EDC was bad for last known time value previously stored in nonvolatile memory;
```

Note

The parse file for use is located in **c:/hsvmafiles/** on the storage management appliance, for example, **HSV110_event_ENT-2001_2001.txt**.

Send Parse File Button Example

By using the **Send parse file** button, you can select and copy the parse file from the local server to the **c:/hsvmafiles/** directory on the storage management appliance.

Send Event Parse File

Send parse file	Cancel	?
-----------------	--------	---

Enter the complete path or browse to the file containing your event parse information. The file may be a .TXT parse file, or a .SSS system code distribution file that contains a parse file component.

<input type="text"/>	Browse...	?
----------------------	-----------	---

Controller Termination Event Log

Controller termination events are occurrences that cause the storage system to stop operation. The controller reports these events. The management agent polls the controller pair for event log information on a regular basis.

Because there is a time lapse between polls, you usually cannot detect the problem until the system starts back up.

Controller Termination Event Log Features

The Controller Termination Event Log has the following features:

- Reports HSV shutdown messages:
 - Fatal errors (firmware problems)
 - Graceful shutdowns
- Events are logged and stored on both controllers in the pair
 - Events are retrieved from both primary and secondary controllers
 - Access to merged log for both controllers in VCS V3.0. Previous versions have separate logs per controller.
- Does not enable a crash dump — This is enabled from the VCS Field Services page
- No forwarding capability to SNMP or posting to NT Event Log
- Allows 32 event entries for each controller (64 maximum)

The **Get log file** button in this log retrieves the raw data file, but it is not in readable format. Use an HP Services tool such as Event View EVA to translate the raw event data.

Controller Termination Event Log Characteristics









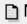





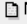



The Controller Termination Event Log has the following characteristics:

- For Command View EVA versions before V3.0:
 - Log file is maintained for an initialized storage system only
 - File name is ***UUID.termeventfile.txt***.
- For Command View EVA V3.0:
 - Log file is maintained for an initialized or uninitialized storage system only
 - File name is ***handle.termeventfile.txt***
- Log file is located on the storage management appliance in:
 - **c:\hsvmafiles** for Command View EVA versions before V3.0
 - **c:\hsvmafiles\WWN** for Command View EVA V3.0
 - When Command View is launched, the **termeventfile** is updated with the current events from the controller's termination event log. Thereafter, the file is updated with the current events every 45 minutes.
- When the log file reaches maximum number of events:
 - For Command View EVA versions before V3.0, the last 32 events are stored in the log and any previous events are dropped.
 - For Command View EVA V3.0, the last 32 events (64 maximum) for each controller are stored in the log and any previous events are dropped.

Controller Termination Event Log Example

The following is a sample Controller Termination Event Log.

Controller Termination Events (Initialized system)

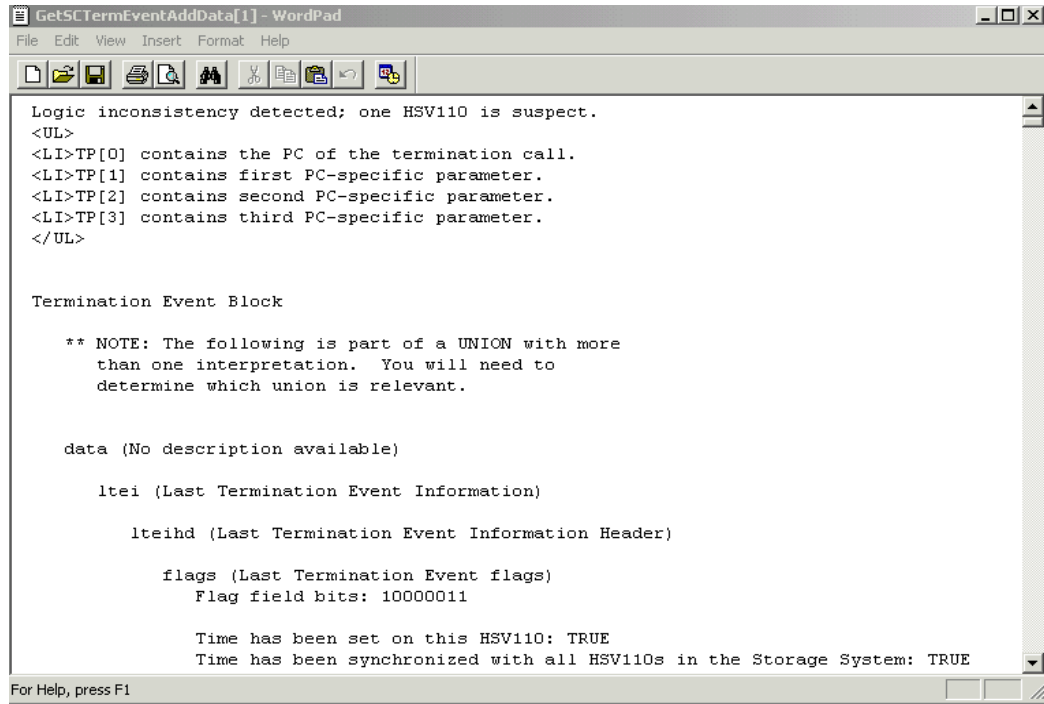
OK	Get log file	Get parse file	Send parse file	?
Controllers:		Controller B: 5005-08b4-0001-4529-0000-0000-0000-0000		
		Controller A: 5005-08b4-0001-4523-0000-0000-0000-0000		
Date/Time Controller	Severity	Event Code	Sequence #	
Description				
17:38:04:09 15-Apr-2003 Controller A		04240960	#100	
Power failed.  Corrective actions code: 09  More details				
17:04:30:762 15-Apr-2003 Controller A		03640020	#99	
This HSV110 controller was requested to terminate operation and then restart.  Corrective actions code: 00  More details				
13:56:16:505 9-Apr-2003 Controller A		03640020	#98	
This HSV110 controller was requested to terminate operation and then restart.  Corrective actions code: 00  More details				
11:36:57:516 4-Apr-2003 Controller A		03640020	#97	
This HSV110 controller was requested to terminate operation and then restart.  Corrective actions code: 00  More details				
14:44:28:332 3-Apr-2003 Controller A		011e0120	#96	
Console requested restart without dump (not coupled).  Corrective actions code: 01  More details				
14:38:47:351 3-Apr-2003 Controller A		011e0120	#95	
Console requested restart without dump (not coupled).  Corrective actions code: 01  More details				

Note

The Controller Event Log uses an EIP Type in the Event Code field, while the Controller Termination Event Log uses a Code Flag field, consisting of control codes and parameter counts. These differences are explained in the Event Message Formats section.

The typical size of a binary file is 600KB to 800KB.

When you click the *More details* icon, you see the results on your console.



```
GetSCTermEventAddData[1] - WordPad
File Edit View Insert Format Help

Logic inconsistency detected; one HSV110 is suspect.
<UL>
<LI>TP[0] contains the PC of the termination call.
<LI>TP[1] contains first PC-specific parameter.
<LI>TP[2] contains second PC-specific parameter.
<LI>TP[3] contains third PC-specific parameter.
</UL>

Termination Event Block

** NOTE: The following is part of a UNION with more
than one interpretation. You will need to
determine which union is relevant.

data (No description available)

ltei (Last Termination Event Information)

lteihd (Last Termination Event Information Header)

flags (Last Termination Event flags)
Flag field bits: 10000011

Time has been set on this HSV110: TRUE
Time has been synchronized with all HSV110s in the Storage System: TRUE

For Help, press F1
```



Important

As with the Controller Event Log, using *More details* is not the recommended way to view the log. Instead, use the **Get log file** button and translate the file by running an HP Services tool such as Event View EVA.

NT Event Log

The appliance polls the storage system for events, writes them to logs accessible through Command View EVA, and writes them into the NT Event Log on the appliance. All events, except termination events, can be posted to the NT Event Log. This log is used to get an overall aggregate view of logs that Command View EVA has seen. Full logs are posted here, unlike SNMP.

You can view the NT Event Log through terminal services, by using the RIB board, or by attaching a keyboard, monitor, and mouse to the appliance.

You can also view the controller events from the NT Event Log by using Event View EVA in continuous mode to process the file.

Example — NT Event Log

Date	Time	Source	Category	Event	User	Computer
6/19/2001	9:55:40 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:55:40 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:55:40 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:55:40 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:55:40 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:55:39 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:55:39 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:55:39 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:55:39 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:39:09 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:39:06 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:36:53 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:36:53 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:35:24 AM	HSV	HSV EMU	60000	N/A	PROWLER2
6/19/2001	9:35:24 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:33:50 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:33:50 AM	HSV	HSV Controller	50000	N/A	PROWLER2
6/19/2001	9:33:50 AM	HSV	HSV Controller	50000	N/A	PROWLER2
6/19/2001	9:33:49 AM	HSV	HSV Controller	50000	N/A	PROWLER2
6/19/2001	9:33:49 AM	HSV	HSV Controller	50000	N/A	PROWLER2
6/19/2001	9:33:49 AM	HSV	HSV Controller	50000	N/A	PROWLER2
6/19/2001	9:33:49 AM	HSV	HSV Controller	50000	N/A	PROWLER2
6/19/2001	9:33:10 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:29:52 AM	HSV	HSV Management	40000	N/A	PROWLER2
6/19/2001	9:29:52 AM	HSV	HSV Management	40000	N/A	PROWLER2

From this screen, you can select an event and get more details.

Tools for Translating Controller Events

Whereas the Management Agent Event Log is a human-readable file and can be saved as a plain text file, the Controller Event Log and Controller Termination Event Log contain binary data. Even though a **Get log file** saves the raw data, it is not translated.

To analyze controller and controller termination event codes, translate the code to human readable form by using the following tools:

- Not a Post Processor (NAPP)
- Event View EVA
- HSV Event Translation (HSVET) — Prior to VCS V3.0

An overview of each of these HP Services tools are described later in this section.

Parse Files

The parse file defines the format and content of the raw controller and termination events, including their codes, corrective actions, and structure. Note the following regarding parse files:

- The parse file is generated when VCS is built. Each release of VCS controller operating software includes its own unique parse file in the VCS superfile.
- The parse file is extracted by the Command View EVA for its use in formatting and displaying event logs for these events.
- Parse files are downloaded for use with bit-to-text translation tools. See the topics on NAPP and HSVET for specific information about obtaining the parse files for use with these tools.
- The parse file is in ASCII text format. You can view the version-specific structures (Event Information Packets, Corrective Action Codes, Termination Codes, and so forth) for the VCS version of interest by downloading and reading the appropriate parse file. This file is useful when decoding OCP termination events.

NAPP

NAPP is a standalone tool that translates raw event log data into several formats, including a spreadsheet and several formats of text file. It has numerous options for delimiting the data.

NAPP uses the VCS parse file to perform the bit-to-text translation and is available at the following locations:

- <http://penguin.cxo.cpqcorp.net/napp>.
- The TECHBB website: <http://techport.tay.cpqcorp.net/>. Navigate to the HSV forum and enter “NAPP” into the search function to find the threads containing the NAPP executable, parse files, and documentation.



Important

NAPP cannot use any parse file downloaded from Command View EVA. Specific NAPP-usable parse files are required for correct NAPP operation. NAPP-specific parse files for all released VCS versions are available on the NAPP website.

NAPP has the following features:

- Fast, efficient translation engine
- Interactive mode and ability to be run from a script or batch file
- Options for limiting the ranges of data to be translated (including single events)
- Options for limiting the amount of data output to salient data only
- Options for creating output formats similar to YAPP and Chrono directly from the raw data (eliminating the need for full translation)
- Option for creating several output formats, including one like HSVET (for use with existing post-processing tools)
- Command to pass a command to the operating system as a system call
- Embedded help function with list of commands and their usage.

Event View EVA

Event View EVA is a standalone tool that translates the raw event log data into several formats. It also provides monitoring and reporting capabilities.

Event View EVA has the following features:

- Runs on the storage management appliance
- Uses rules (thresholds, filters, timers, event collation) to intelligently deduce events meriting notification
- Can be run as a standalone utility to analyze an event log offline without notification
- Can be run as a continuous monitoring mechanism to deduce and report events
- Can monitor multiple Enterprise Virtual Array storage systems
- Extracts events from the NT Application Event Log on the storage management appliance
- Provides the following notification methods:
 - Email (via distribution list)
 - Links to Instant Support Enterprise Edition (ISEE) to provide direct notification to HP Services
 - HTML (updates a remotely accessible web page with new events)
 - SNMP traps
 - DTCS (provides an interface to the Disaster Recovery Manager)

- Provides output in the following formats:
 - Text
 - HTML
 - Microsoft Excel spreadsheet
- Provides three modes of operation:
 - Single run: operates on a controller event log without notification to analyze the log offline. Event View EVA reads the input file (event log), performs bit-to-text translation, and presents the results in the output formats.
 - Simulation: operates on a controller event log and generates notification as if it were running in real time. This mode is used primarily for demos and for rules development.
 - Continuous: operates on the NT Application Event Log and generates notification in real time.
- Accepts the controller event log in three different formats:
 - HSVET output file
 - Raw controller event log for both VCS V2.x (ASCII format) and VCS V3.x (binary format)
 - NT Application Event Log

Event View EVA is still in development. The latest version and information on this tool can be downloaded from **<http://storage.jgo.cpqcorp.net/EVA-tools>**. Follow the links to Event View EVA and YEAH. The download file contains the latest *Guide of Operations* as well as the executable files. The *Guide of Operations* describes installation as well as operation of this tool.

HSVET

This standalone tool translates the raw event log data into text format. Post-processing tools operate on the HSVET output file.

HSVET has the following features:

- HSVET does not perform event analysis.
- HSVET is compatible with browsing stations running Windows operating systems only (**hsvetwin.exe**).
- Each HSVET release is backward compatible; that is, it can read log files for all versions of VCS.
- The HSVET executable is located on the TECHBB website:
<http://techport.tay.cpqcorp.net/>. Navigate to the HSV forum and enter “decoder” into the search function to find the threads containing the HSVET download file. Then look for the entry containing the latest version of HSVET.

Note

HSVET is being replaced by NAPP. Support for HSVET ends on August 1, 2003.

Event Message Formats

The format and content of the controller and controller termination events (that is, their codes, corrective actions, and structure) are defined by parse files. The parse file is generated when VCS is built. Each release of VCS controller operating software includes its own unique parse file in the VCS superfile. The parse file is extracted by Command View EVA for its use in formatting and displaying event logs for these events. You can download parse files for use with NAPP and HSVET by accessing the appropriate websites.

Controller Termination Event Format

A controller termination event is:

- A 32-bit integer.
- Displayed as 4 bytes separated by spaces, leading zeroes replaced by spaces, for example: **6 f 40 13**

A termination event has the following format:

31.....24	23.....16	15.....87.....	6.....5	4.....0
Software Component ID Number	Event Number	Corrective Action Code	Couple Crash Control Code	Dump/ Restart Control Code	Parameter Count

Controller Event Format

A controller event is:

- A 32-bit integer.
- Displayed the same way as a termination event.
- Defined with a different interpretation of bits 0 through 7.

A controller event has the following format:

31.....24	23.....16	15.....8	7.....0
Software Component ID Number	Event Number	Corrective Action Code	EIP Type

Definition of Fields

The following defines the fields used in both controller event logs:

- Software Component ID (SWCID) — Number identifying the software component that generated the particular event
- Event Number — Unique event number for a given SWCID
- Corrective Action Code (CAC) — Indicates corrective action for a specific component problem
- EIP (Event Information Packet) Type — Defines the event information format
- Coupled crash control code (CCCC) — Specifies whether both controllers in the pair are terminating operation
- Dump/restart control code (DRCC) — Specifies whether a crash dump is occurring
- Termination parameter count — Specifies how many entries in Termination Parameters array are valid

Note

The EIP Type is used only for controller events. The CCCC, DRCC, and termination parameter count are used for controller termination events only.

Sample Controller Event Code

For the controller event code: **d 83 89 11**

- SWCID = **d**. Parse file defines software component as Disk Enclosure Environmental Monitoring Service.
- Event number = **83**. Parse file defines the event number with the SWCID to provide a unique identifier. An EMU hardware failure occurred.
- CAC = **89**. Parse file defines the corrective action, that is, to replace the EMU.
- EIP type = **11**. Parse file defines this as a Disk Enclosure Services Environment Monitoring Unit Services Status Change — the status of a disk enclosure element has changed.

Sample Controller Termination Event Code

For the controller termination event code: **4 2d 20 13**

- SWCID = **4**. Parse file defines software component as Fault Manager.
- Event number = **2d**. Parse file defines the event number with the SWCID to provide a unique identifier. The PowerPC timed out.
- CAC = **20**. Parse file defines the corrective action, that is, to replace the HSV controller.
- CCCC = **0**. Parse file defines this as: do not perform coupled crash.
- DRCC = **0**. Parse file defines this as: perform a crash dump, then restart.
- Parameter count = **13**. Nineteen (decimal) parameters are valid.

Interpreting Messages with Physical Addresses

The event log usually provides you with the enclosure and bay number of the disk on which the event occurs. However, sometimes the data for the rack, drive enclosure, and bay are invalid, leaving you with only the disk AL_PA to identify it.

Command View EVA does not show the AL_PA of the disk; it gives the Loop ID instead. To determine the enclosure and bay number of the disk, you need to know the Loop ID corresponding to the disk physical address. The **Disk Drive Properties** page shows the Loop ID for each disk by enclosure and bay.

The following table correlates the AL_PA (in hexadecimal) to the Loop ID (in decimal as shown in the Command View EVA display and in hexadecimal for reference).

AL_PA	Loop ID (dec)	Loop ID (hex)	AL_PA	Loop ID (dec)	Loop ID (hex)	AL_PA	Loop ID (dec)	Loop ID (hex)
EF	0	00	A3	43	2B	4D	86	56
E8	1	01	9F	44	2C	4C	87	57
E4	2	02	9E	45	2D	4B	88	58
E2	3	03	9D	46	2E	4A	89	59
E1	4	04	9B	47	2F	49	90	5A
E0	5	05	98	48	30	47	91	5B
DC	6	06	97	49	31	46	92	5C
DA	7	07	90	50	32	45	93	5D
D9	8	08	8F	51	33	43	94	5E
D6	9	09	88	52	34	3C	95	5F
D5	10	0A	84	53	35	3A	96	60
D4	11	0B	82	54	36	39	97	61
D3	12	0C	81	55	37	36	98	62
D2	13	0D	80	56	38	35	99	63
D1	14	0E	7C	57	39	34	100	64
CE	15	0F	7A	58	3A	33	101	65
CD	16	10	79	59	3B	32	102	66
CC	17	11	76	60	3C	31	103	67
CB	18	12	75	61	3D	2E	104	68
CA	19	13	74	62	3E	2D	105	69
C9	20	14	73	63	3F	2C	106	6A
C7	21	15	72	64	40	2B	107	6B
C6	22	16	71	65	41	2A	108	6C
C5	23	17	6E	66	42	29	109	6D
C3	24	18	6D	67	43	27	110	6E
BC	25	19	6C	68	44	26	111	6F
BA	26	1A	6B	69	45	25	112	70
B9	27	1B	6A	70	46	23	113	71
B6	28	1C	69	71	47	1F	114	72
B5	29	1D	67	72	48	1E	115	73
B4	30	1E	66	73	49	1D	116	74
B3	31	1F	65	74	4A	1B	117	75
B2	32	20	63	75	4B	18	118	76
B1	33	21	5C	76	4C	17	119	77
AE	34	22	5A	77	4D	10	120	78
AD	35	23	59	78	4E	0F	121	79
AC	36	24	56	79	4F	08	122	7A
AB	37	25	55	80	50	04	123	7B
AA	38	26	54	81	51	02	124	7C
A9	39	27	53	82	52	01	125	7D
A7	40	28	52	83	53			
A6	41	29	51	84	54			
A5	42	2A	4E	85	55			

Interpreting SCMI Event Codes

The Command View EVA-to-controller communication occurs through the Storage System Management Interface (also called the StorageCell Management Interface or SCMI). Controller Event Logs occasionally contain events indicating that SCMI communication failed, as in the following example:

```
Severity: Normal -- informational in nature. The HSV110
controller identified in the handle field processed a
StorageCell Management Interface command with the result of
non-success return code. The value.ul1 field contains the
command type. The value.ul2 field contains the return code.
```

Two methods are used to decode returned values to determine the rejected command and reason for rejection (the return code):

- Use NAPP and examine the Additional Information column.

SW Ver/ Flags	Seq Nr	DateTime	EventCode	Ctrl	Severity	Brief Description	Additional Information
2002 time_set time_syn ched	17	2003-02-13 10:38:05.433	0x092f0005	'4528'	Normal	The HSV110 controller identified in the handle field processed a Storage System Management Interface command with the result of non-success return code.	Controller: 5005-08b4-0211-4528-0000-0000-0 000-0000; Storage System Management Interface command: "Controller Get Storage System"; Return code: "No Storage System"; Internal target and command version: 00000002

- Use the SCMI Code-to-Mnemonic Translation Table from the Service Manual to correlate command types and return codes.

Analyzing Event Logs

The following are some important considerations when analyzing event logs:

- The controller event logs are formatted in different structures depending on the type of error.
- All error entries contain a couple of basic structures:
 - The title bar with the date, time, controller ID, and event code.
 - Below the title bar is the following:
 - ◆ Description of the event
 - ◆ Event log packet event specific data, including the state of the time, and event sequence numbers
 - ◆ Controller state including event number, controller software version, controller model, reporting controller and event time.
 - The next parts depend on the type of error and what data is being included.
 - ◆ For events that are controller detected, the event log will contain the reporting NSC (network storage controller).
 - ◆ For events that are disk drive specific, the data will contain disk drive specific UUID.
 - For each piece of data, there can be a couple of different methods to format the data.
 - ◆ These are identified as **unions**, where the data is formatted in at least two different structures.
 - ◆ The data presented is identical, and in many cases the different unions provide the same definition but with different format.
 - The end of the controller event packet normally contains the error codes or error type count.

These packets contain several important parts of data including the controller that logs the event, the controller that reports the event (some entries) and the device that is associated with the event. It also contains the event that initiated the entry.

When analyzing event logs, the following are some fields of particular interest:

- Unique Universal Identifier (UUID) of the controller:

This ID is usually reported under the following heading:

```
reporting_nsc (Storage System Management Interface  
Handle of HSV110 that reported the event)
```

The field called **id_value - Entity UUID** contains the UUID of the controller reporting the event.

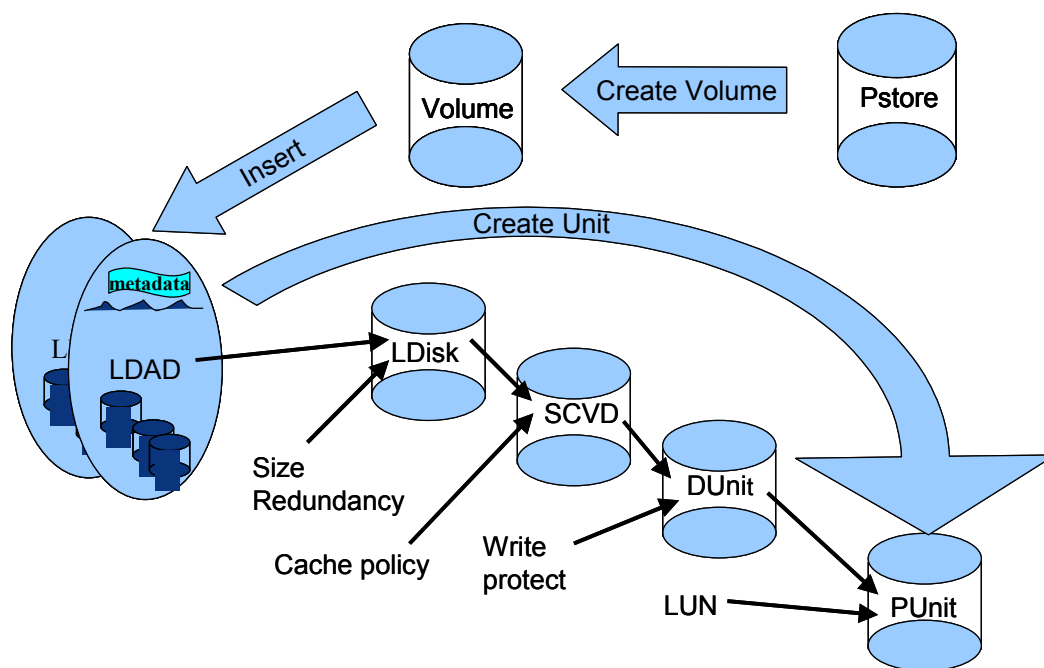
Note

Events are reported to Command View EVA by the primary controller only. Controller A and B UUIDs are displayed in Command View EVA on the Controller Properties page.

- Event log fields
 - report_time — Time event was reported
 - report_location — Location of event report call
 - eiptype — Event Information Packet Type Code
 - cac — Corrective Action Code
 - evnum — Event number
 - scid — Software Component ID
 - revision — Packet revision number
 - type — Packet type
 - count — Number of bytes in packet
- The following is some useful information when interpreting union data:
 - Use the EIP type to determine which set of union data is valid.
 - For example, when interpreting an event of EIP type 4 (Fibre Channel Services Physical Disk Drive/ Mirror Port Error), view the data as an `hdr_cdb` rather than an `evt_peq`.

Enterprise Virtual Array Objects

The following diagram summarizes the Enterprise Virtual Array objects and their relationships. You will see these object names in the event logs.



Objects include the following:

- Pstores (physical stores) — Raw FC-AL disks
- Volumes — Pstores + metadata
- LDAD — Logical Disk Allocation Domain (disk group)
- Ldisks — Logical disks (containers)
- SCVDs — Storage cell virtual disk
- Dunits — Derived units
- Punits — Presented units

Sample Event Listing Showing Objects

The following is a sample portion of an event log:

Event 138. of 1775.

The **Volume** identified in the handle field was created.

Handle: UUID: 6005-08b4-0001-486c-0000-7000-0105-0000

Add_Handle: UUID: 2000-0004-cf96-4b5f-0000-0000-0000-0000

Event 295. of 1775.

The internal **Logical Disk** associated with the **Virtual Disk** identified in the handle field was created

Handle: UUID: 6005-08b4-0001-486c-0000-7000-019d-0000

Event 299. of 1775.

The **Derived Unit** identified in the handle field was created.

Handle: UUID: 6005-08b4-0001-486c-0000-7000-01a0-0000

Add_Handle: UUID: 6005-08b4-0001-486c-0000-7000-019f-0000

Inband Host Event Reporting

VCS V2.x and higher implements a feature called inband host event reporting. This feature allows events reported by the storage system to be correlated with host system events.

Controller events are sent to the host system for storage in the host system's event log. Event information is sent to the host system through the Fibre Channel FP1 and FP2 ports in the controller in the form of vendor-unique SCSI sense data. Event transmission uses the SCSI-3 auto sense data delivery mechanism.

The vendor-unique SCSI sense data contains:

- Sense key of 6x (Unit Attention)
- Additional Sense Code (ASC) of C3x
- Additional Sense Code Qualifier (ASCQ) consisting of 50x plus the EIP number of the event.

Command View EVA contains a Bridge Event Translator Engine (BETE) that uses the parse file to perform bit-to-text translation of the event in an EIP. BETE is written in C++, and its source files can be made available for incorporation into host-based Services event translation and analysis tools.

Note

The format and definitions of the vendor-unique SCSI sense data are defined in the Service Manual.

Feature Status by VCS Version

The following describes the capabilities for inband host event reporting for various VCS and element manager versions:

- VCS V2.002 and lower using HSV Element Manager V2.0A
 - Feature is on
 - Cannot be changed
- VCS V2.003 using HSV Element Manager V2.0A.
 - Feature is off
 - Can enable using the the Command View EVA Field Service page (see overview in the following topics)
- VCS V3.0 and higher using Command View EVA V3 and higher
 - Feature is off
 - Can enable by navigating to the Host Properties page and selecting *Enabled* in the **Direct eventing** box.

Enabling and Disabling Inband Event Host Reporting for VCS V2.003

In VCS V2.003, inband host event reporting is disabled by default. You can enable or disable this feature for either specific hosts or all hosts running an eligible operating system. Only hosts running Tru64 UNIX or Windows-based operating systems (running SecurePath) are eligible to have this feature enabled.

Procedures for enabling inband host event reporting for Tru64 UNIX or Windows are detailed in the Service Manual.

Event Notification

A storage system can be configured to report event information to specified computers using SNMP. The computers must be SNMP enabled—designated to be a collecting point for SNMP traps. Any computer on the same fabric or network as the storage system can be SNMP enabled.

Each storage system contains a configurable event notification file in which you can specify which events or types of events to be reported to the SNMP-enabled host. The event notification file controls event reporting only for the storage system on which it resides. Unless altered, the event notification file contains default settings.

Set Event Notification Options

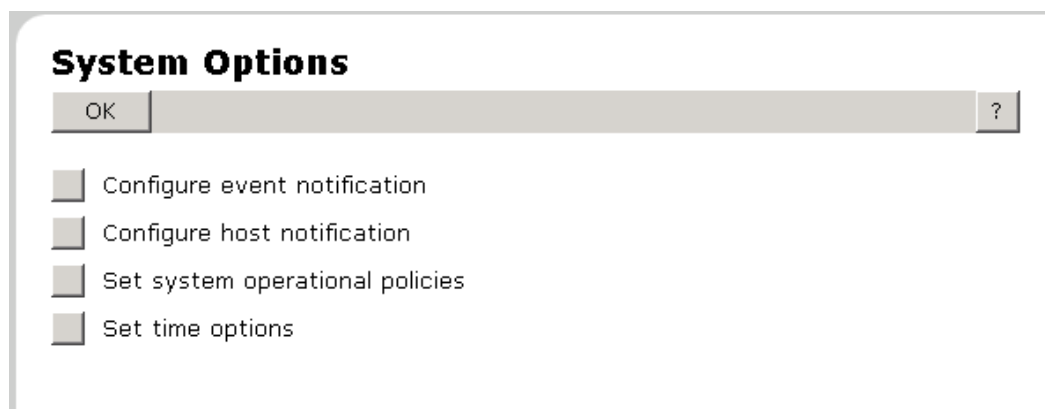
Follow these steps to set event notification options:

1. Select the storage system.
2. Click *Set options* on the Initialized Storage System Properties page.
3. Click the button next to **Configure event notification**.
4. Select the events you want to use for traps.
5. Click the button next to **Configure host notification**.
6. Add the host to receive the SNMP trap of events.

A MIB is required to decode the events.

System Options Page

The System Options page looks like this.



You can configure the events individually or use a configuration file.

Example

All events is selected by default. This means that all events are included, but it is not set until you select the Configure button.

Configure Event Notification

OK Get configuration file Download MIB ?

Restore defaults

You can configure the way your Management Agent notifies you when events occur. You can set each event individually or you can use a configuration file to set them all automatically.

Configure events individually:
Select the event level or levels you wish to view and click the **Configure** button to set each event's notification option individually.

☐ All events
☒ Critical events
☐ Warning events
☐ Normal events

Configure

Configure events using a configuration file:
Browse to a properly formatted configuration file or enter a complete file path and click the **Configure** button to configure your event options.

Browse... **Configure**

The Download MIB button allows you to display the MIB file or save it to the system where the browser is running. This also allows the use of Insight Manager.

Select which events you want sent as traps.

- **Critical Events** — May prevent further operation of the management agent and the storage system.
- **Warning** — If not attended to, could lead to the management agent and storage system becoming nonoperational and could cause the management agent and storage system to run in a degraded fashion.
- **Normal** — Is within the normal operational parameters of the management agent and the storage system.

Some traps are default for any host selected, such as critical issues on Command View EVA, HSV controller, or the initialized storage system.

Notification Options Page

When you click the *Configure* button, as in this case, for critical events, you get a notification list as shown below.

Set Event Notification Options

Save changes	Notify all	Notify none	Restore defaults	?
Cancel				

Check the Notify box for any event you'd like your hosts to be notified of.

Event Notification List		
Notify	Code	Desc
<input type="checkbox"/>	4	Insufficient Available Data Storage
<input checked="" type="checkbox"/>	77	Volume Failure Predicted
<input type="checkbox"/>	88	Storage system inoperative
<input type="checkbox"/>	96	PLDMC failed
<input type="checkbox"/>	97	Storage system could not be locked. System busy. Try command again.
<input checked="" type="checkbox"/>	1001	Storage system is FULL **critical**
<input type="checkbox"/>	2001	Startup Failed
<input checked="" type="checkbox"/>	2038	Control Memory Allocation Failure
<input type="checkbox"/>	3003	Event Manager - Startup Failed
<input checked="" type="checkbox"/>	3036	EM Memory Allocation Error
<input type="checkbox"/>	3050	EM: EMU event - this description will be overwritten
<input checked="" type="checkbox"/>	3060	Storage system being shutdown due to a command failure.
<input checked="" type="checkbox"/>	3064	EM: Storage system pointer is NULL!

Click each event type or select *Notify all* to select all event types.

The Code field is the same as the field used in the Management Agent Event Log. If you generate an event, look in the event log for the code and then enable it here.

The SWC ID, Event Code, CAC, or EIP type number can be found in the Desc (Description) field of the event. Search the code and then select that event to generate SNMP traps.

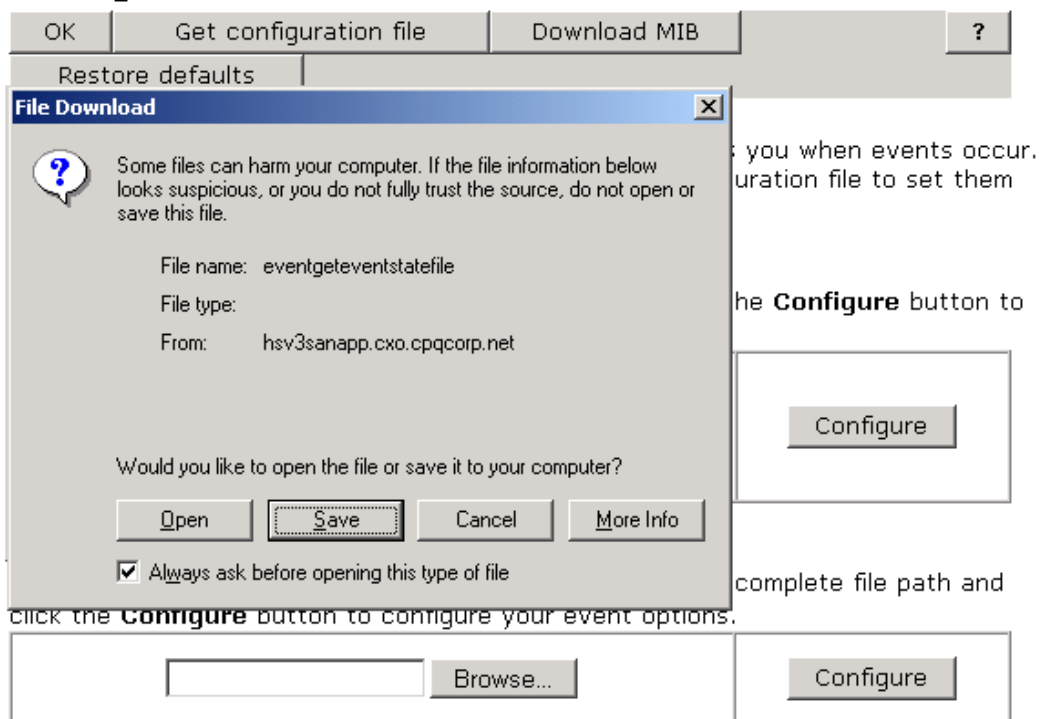
The generation of an SNMP trap logs an event in the Management Agent Event Log. Therefore, you know the trap was sent.

Use the SNMP utility to detect the trap on the trap system. Then set up Insight Manager to process the trap. A utility called **Snmputil** watches for SNMP traps and is readily available.

Saving the Configuration to a File

When you select the *Get configuration file* button from the Configure Event Notification page, you get this screen. This allows you to save the SNMP configuration to a file.

Configure Event Notification



Configuring Events Using a Configuration File

Reinsert the saved setting from the **Get configuration file** operation by selecting the *Configure* button from the **Configure events using a configuration file** option.

Configure Event Notification

OK

Get configuration file

Download MIB

?

Restore defaults

You can configure the way your Management Agent notifies you when events occur. You can set each event individually or you can use a configuration file to set them all automatically.

Configure events individually:
Select the event level or levels you wish to view and click the **Configure** button to set each event's notification option individually.

☒ All events
☐ Critical events
☐ Warning events
☐ Normal events

Configure

Configure events using a configuration file:
Browse to a properly formatted configuration file or enter a complete file path and click the **Configure** button to configure your event options.

Browse...

Configure

In this example, the file **eventgeteventstatefile** will be loaded to configure SNMP traps. The file that comes with the VCS kit is **hsv110_events_for_services_Vx**. This file is applicable to VCS V2.000 only and if used with VCS V2.002 or V2.003, may cause inconsistent event viewing and reporting. Details are in a Customer Advisory.

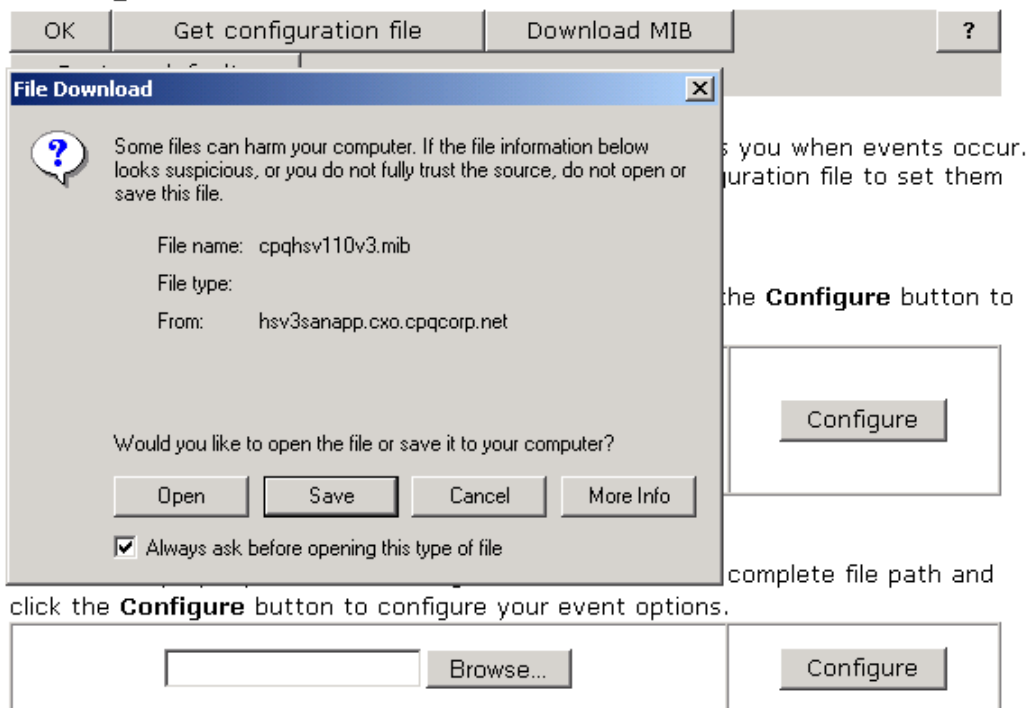
Note

There is not an **hsv110_events_for_services_V3000** file.

Using the Download MIB Button

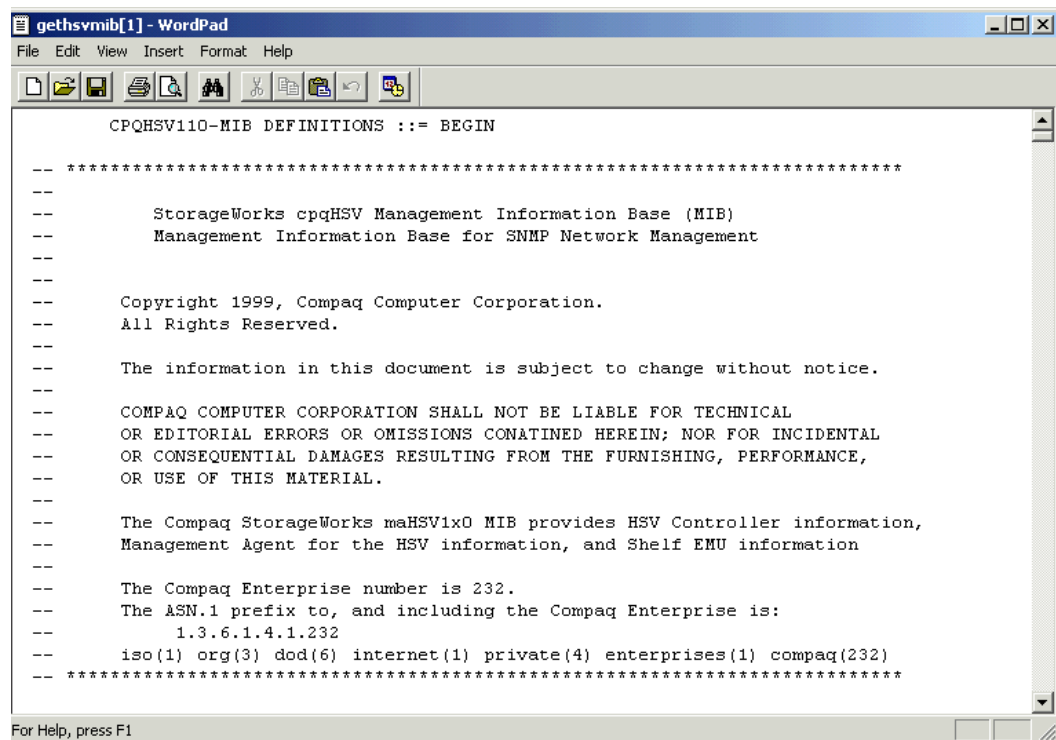
To save the HSV MIB file to a local server, use the Download MIB button.

Configure Event Notification



Once you have the file, you can copy it to the Insight Manager server and install it.

The following is a sample MIB file.



```

gethsvmib[1] - WordPad
File Edit View Insert Format Help

CPQHSV110-MIB DEFINITIONS ::= BEGIN

-- *****
--
--      StorageWorks cpqHSV Management Information Base (MIB)
--      Management Information Base for SNMP Network Management
--
--
--      Copyright 1999, Compaq Computer Corporation.
--      All Rights Reserved.
--
--      The information in this document is subject to change without notice.
--
--      COMPAQ COMPUTER CORPORATION SHALL NOT BE LIABLE FOR TECHNICAL
--      OR EDITORIAL ERRORS OR OMISSIONS CONATINED HEREIN; NOR FOR INCIDENTAL
--      OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE,
--      OR USE OF THIS MATERIAL.
--
--      The Compaq StorageWorks maHSV1x0 MIB provides HSV Controller information,
--      Management Agent for the HSV information, and Shelf EMU information
--
--      The Compaq Enterprise number is 232.
--      The ASN.1 prefix to, and including the Compaq Enterprise is:
--          1.3.6.1.4.1.232
--          iso(1) org(3) dod(6) internet(1) private(4) enterprises(1) compaq(232)
-- *****

For Help, press F1

```

Configure Host Notification Options

Each storage system contains a host notification file that specifies SNMP-enabled hosts that can receive event notification directly over the network. You can configure in one of two ways:

- Directly by clicking *Modify host list*.
- Using an existing host notification file on another storage system by browsing to the file and clicking *Send list file*.

You must specify the network name for the host computer and notify port.

Configure Host Notification

The Host Notification List contains the hosts you have configured to receive event notifications directly via the network. Click the **Modify List** button to change the list.

Click the **Get List File** button to download the list file. After editing the file in a text editor, you can replace the list by sending it back to the server, as described below.

Host Notification List

Host Name	Notify Port
eniac.americas.cpqcorp.net	65535
eniac.americas.cpqcorp.net	SNMP

To replace the host notification list:

You can replace the host notification list above by sending a properly formatted host notification file to the server. Enter the complete file path or browse to the host notification file of your choice and click the **Send List File** button to send the file.

<input type="text"/>	<input type="button" value="Browse..."/>	<input type="button" value="Send list file"/>
----------------------	--	---

Modify Host Notification List

This example shows that adding the host named **hsv1.cxo.cpqcorp.net** to the list will enable it to receive SNMP notifications from the storage system.

Modify Host Notification List

The Host Notification List displays the hosts you have configured to receive event notifications directly via the network. You can modify the list by editing any of its entries. You can also add an entry using the dialogue box below the list. Click the **Save Changes** button to place your changes in effect.

Note: All host names must be fully qualified, network names.

Host Notification List

Host Name	Notify Port
eniad.americas.cpqcorp.net	65535
eniad.americas.cpqcorp.net	SNMP

Add a host to the list:

Enter a host name and notify port and click the **Save Changes** button to add a new entry to the list.

Host Name	Notify Port
hsv1.cxo.cpqcorp.net	SNMP

Get and Send Host List Options

Get a host list by selecting the **Get list file** button. After editing the list file, replace it by sending it back to the server using the **Send list file** button.

Configure Host Notification

[OK](#)[Modify host list](#)[Get list file](#)[?](#)

The Host Notification List contains the hosts you have configured to receive event notifications directly via the network. Click the **Modify List** button to change the list.

Click the **Get List File** button to download the list file. After editing the file in a text editor, you can replace the list by sending it back to the server, as described below.

Host Notification List

Host Name	Notify Port
eniac.americas.cpqcorp.net	65535
eniac.americas.cpqcorp.net	SNMP

To replace the host notification list:

You can replace the host notification list above by sending a properly formatted host notification file to the server. Enter the complete file path or browse to the host notification file of your choice and click the **Send List File** button to send the file.

 [Browse...](#)[Send list file](#)

Notification Using Proactive Remote Services

Proactive Remote Services (PRS) is the secure and automated end-to-end remote service offering by HP. Customer systems can self-monitor and securely report problems and events to a service provider. Service representatives can securely connect back to a remote customer system for non-disruptive repair and maintenance.

Note

The current version that is compatible with VCS V3.0 is PRS V5.0.

HP provides and installs the PRS software components required for the remote support of ProLiant, AlphaServer, and selected StorageWorks products, including the Enterprise Virtual Array. This service is available at no extra charge with HP service contracts. This includes the installation and configuration of the Intelligent Services Link (ISL) as required for the automated reporting of serviceable events and remote problem management of HP networked servers and storage.

PRS allows the automatic reporting of serviceable events to HP and provides remote access to the problem system for diagnosis and repair. This capability speeds up problem resolution times and contributes to higher levels of overall system availability. Additionally, as part of this service, HP provides orientation on the functionality provided by the user interface to PRS. This service covers PRS installation and configuration activities for the system designated as the Customer Service Gateway (CSG) and the PRS agents on up to five managed systems.

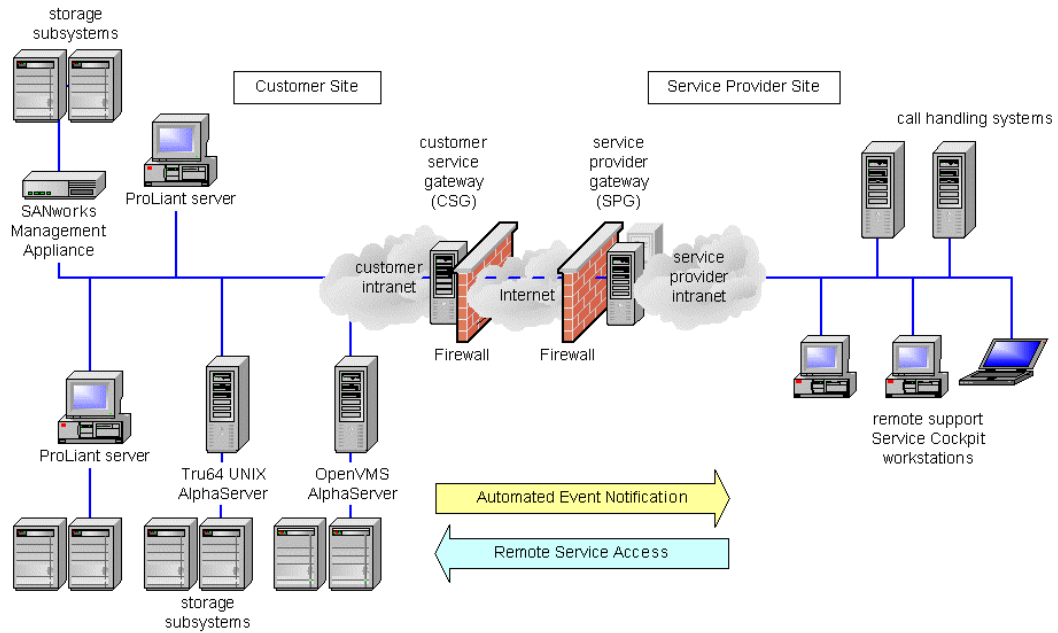
INTRA|NET

Supplementary information and documentation is available through Customer Support R&D Technical Publications at <http://techpubs.cxo.cpqcorp.net/>, and in particular, in the *Proactive Remote Service Customer Setup Guide*.

PRS courses are available through the HP Service Learning Utility.

PRS Functional Overview

The following block diagram details the key components and interactions of PRS.



The diagram shows the basic model of the PRS architecture:

- The left side depicts the customer site, and the right is the service provider site, either HP or a service partner.
- Internet connectivity exists between both sites, offering bandwidth and cost advantages over phone lines.
- The connectivity between the customer site and the service provider site is through WorldWire, a secure Virtual Private Network. This offers the advantage of full-time connect and high bandwidth.

Primary Functions

PRS provides two main functions as denoted in the diagram:

- Automated event notification
 - Allows customer systems to automatically send service event notifications (detected system problems) to the service provider call-handling center.
 - Customer “Managed Systems” self-report by running automated software agents that detect and forward events to a CSG.
 - The CSG sends events to a Service Provider Gateway (SPG), which routes the events to call handling systems.
- Remote access services
 - PRS allows service representatives with customer permission to remotely connect to affected systems for diagnosis and maintenance.
 - Service Cockpit gives them access to a wide variety of diagnosis and maintenance tools, on both the customer system and service provider system.

PRS Operation in the Enterprise Virtual Array Environment

The following describes how PRS components complete the call-handling process:

- The HSV controllers pass an event to Command View EVA through the SAN.
- Command View EVA puts the Event Information Packet (in binary format) into the NT Event Log on the storage management appliance.
- If event notification in Command View EVA is enabled, Command View EVA generates an SNMP trap according to the enabled events and sends the trap to the defined host or hosts.
- PRS receives Enterprise Virtual Array event notification by one of two methods:
 - If SEA is installed on the Management Appliance controlling the storage system, SEA monitors the NT Event Log on the storage management appliance. When SEA determines that a serviceable event has occurred, it sends a problem report to PRS.
 - If SEA is **not** installed on the storage management appliance controlling the storage system, PRS receives SNMP traps directed by the Command View EVA to the server where the PRS/CSG software is running. PRS looks up the trap in the Management and Service Access Database (MSADB) filter set (current version is V1.504) to determine if a serviceable event has occurred.

Note

If the CSG is a ProLiant-type server, the SNMP traps may pass through the optional Insight Manager (IM) before being received by PRS.

- If the events are diagnosed as serviceable, then they are forwarded to the CSG for transmission to the Support Center.
- A call is logged into the Support Center queue.
- The log number is passed back to the Common Remote Support Module (CRSM) viewer for the customer to review.

Prerequisites for PRS on the Enterprise Virtual Array

The following are the prerequisites for PRS on the Enterprise Virtual Array:

- Enterprise Virtual Array
 - HSV controller running VCS firmware
 - On-site product warranty period or a current on-site storage hardware support contract in place
- CSG
 - HP 32-bit Intel based system such as a ProLiant or Prosignia running Windows 2000 SP2, attached to the SAN storage system
 - CRSM installed on both CSG and service provider gateway (SPG), with configurations in each location
 - Insight Manager 7 installed on the CSG Intelligent Services Link
- Managed systems
 - Storage management appliance
 - Management agents installed on affected managed servers

Note

It is assumed that the CSG is a ProLiant-type server and the managed system is the storage management appliance. Other configurations are detailed in the *Proactive Remote Service Customer Setup Guide*.

- An Internet connection approved for use with the CSG and SPG software

Note

The Common Remote Support Module (CRSM) is the PRS event-delivery application. CRSM is installed on both customer service gateway and service provider gateway, with different configurations that define its functions in each location. Consider “CRSM” to be synonymous with “PRS” whenever you encounter the term.

Configuring PRS with the Enterprise Virtual Array

You can find detailed directions for configuring PRS in the *Proactive Remote Service Customer Setup Guide*. The following topics condense the portions of the guide that apply to configuring the Enterprise Virtual Array for use with PRS.

Setting Up and Configuring the Storage Management Appliance

The following describes the general procedure for configuring the managed system, in this case the storage management appliance, for use with PRS on the Enterprise Virtual Array:

- Enable SNMP support on the appliance:
 - Add trap destinations (IP address of CSG).
 - Define communities as required.
 - Set the SNMP trap service startup type to Automatic.
- Set up management event reporting:
 - With SEA, do not enable event reporting. Ensure that Command View EVA event reporting is disabled.
 - Without SEA, enable event reporting
 - ◆ Under the Set Event Host Options page, select *Modify List* and add a host to the list with the IP address of the CSG.
 - ◆ Extract and download the HSV MIB.



Important

If the Enterprise Virtual Array is running VCS V2.000, you must load the `hsv110_events_for_services_2000` file through the Set Event Notification Options page. Refer to the *Proactive Remote Service Customer Setup Guide*, the section called *Managed System: SWMA*, for detailed instructions on setting up the storage management appliance as the managed system.

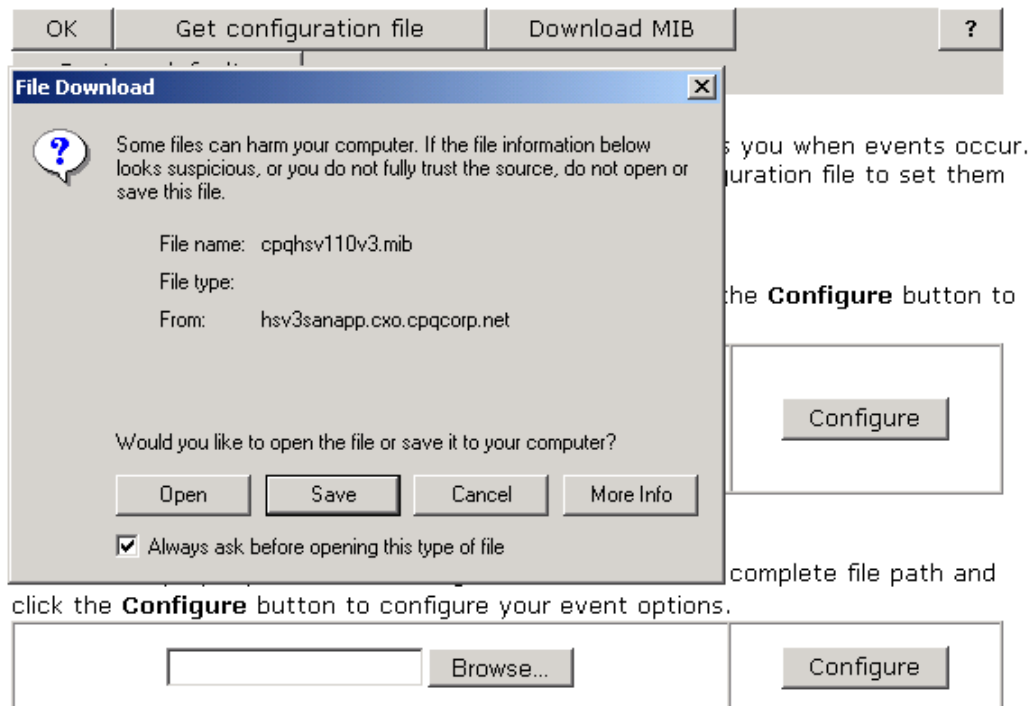
Extracting and Downloading the MIB (Non SEA)

Assuming a PRS environment without SEA, the following describes the general procedure for extracting and downloading the MIB, required for later upload to the CSG:

- Log in to the appliance and the Command View EVA.
- Select the subsystem.
- Select the Set Event Notification Options page.
- Select the *Download MIB* button:
 - Wait for the file download dialog to open.
 - Save the file to disk and name it **CPQHSVG.L.MIB**.

The following is a sample Set Event Notification Options page.

Configure Event Notification



Setting Up and Configuring the CSG

The following describes the general procedure for configuring the CSG, in this case a ProLiant server, for use with PRS on the Enterprise Virtual Array:

- Perform WorldWire configuration.
- Start the CRSM service.
- Enable the CRSM viewer.
- Install Insight Manager (IM) on the server.
- Enable or install SNMP trap support on the server:
 - Define communities.
 - Set the SNMP trap service to Auto Startup.
 - Add the IP address to IM system discovery.
 - For environments with SEA:
 - ◆ SEA sends problem reports directly to PRS.
 - For environments without SEA:
 - ◆ Upload and register the HSV MIB (**cpqhsvgl.mib**).
 - ◆ Initialize IM (connects IM to CRSM and allows events to be passed from IM to PRS).

Note

Students can refer to the *Proactive Remote Service Customer Setup Guide*, the section called *Customer Service Gateway: Windows Intel*, for instructions.

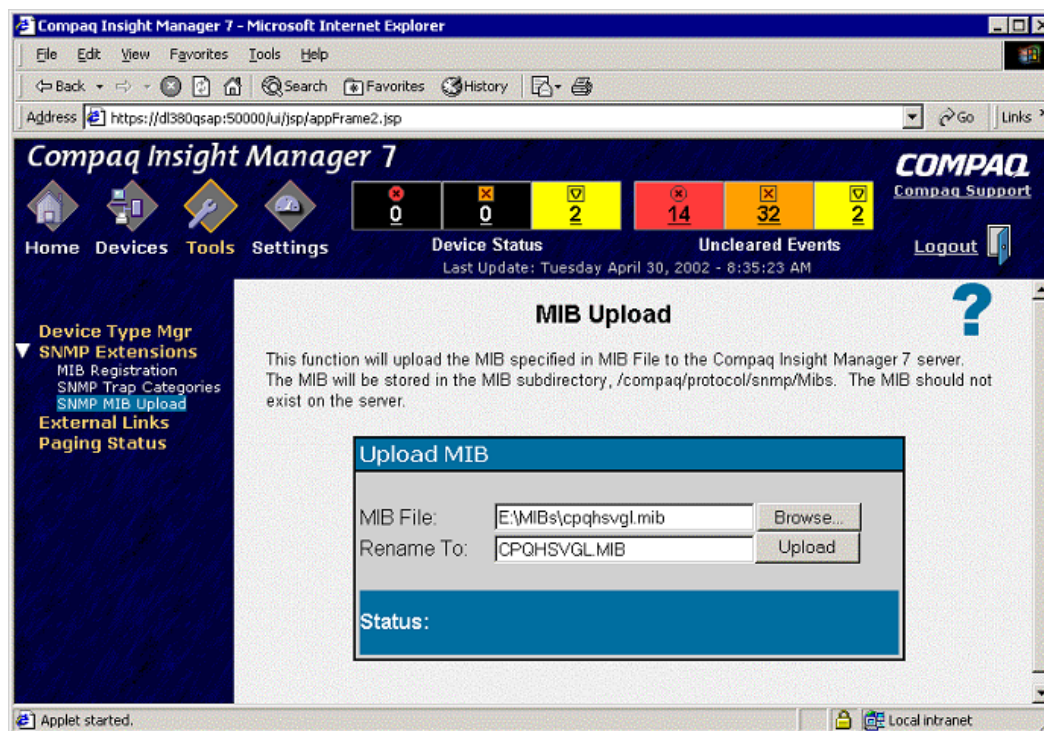
Uploading and Registering the MIB for Use with IM (Non SEA)

Assuming a PRS environment without SEA, the following describes the general procedure for uploading and registering the MIB that was downloaded from the storage management appliance:

- Uploading the MIB:
 - Launch Insight Manager and log in.
 - Select *Tools, SNMP Extensions, MIB Upload*.
 - Browse to the downloaded MIB and select it.
 - Name the file **CPQHSVGL.MIB**.
- Registering the MIB:
 - Launch Insight Manager and log in.
 - Select *Tools, SNMP Extensions, MIB Registration*.
 - Select the file **CPQHSVGL.MIB** in the unregistered box and click *Register*.

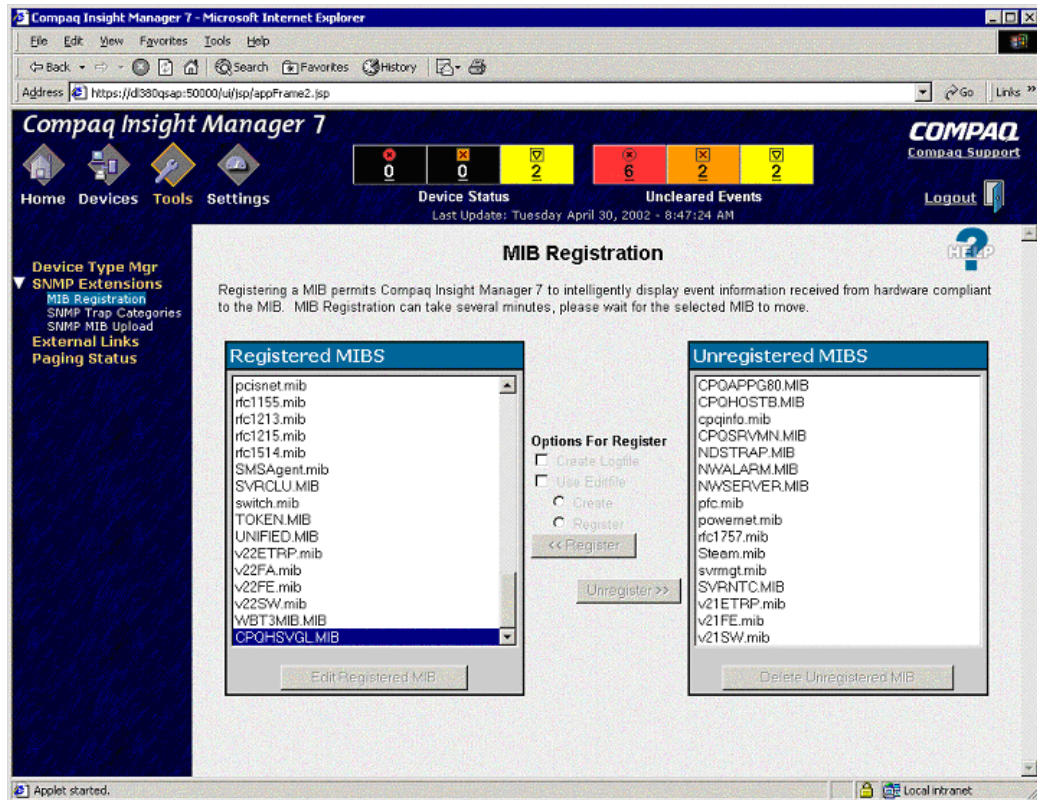
Example

The following is a sample MIB upload page.



Example

The following is a sample MIB registration page.



Notification Using System Event Analyzer

System Event Analyzer (SEA) is a fault analysis utility that has the following features:

- Provides basic analysis for single and multiple error/fault events as well as complex analysis and corrective action recommendation using analysis rules
- Provides background automatic analysis by monitoring the active binary error log and processing events as they occur
 - The events in the binary error log file are checked against the analysis rules
 - If one or more of the events in the binary error log file meets the conditions specified in the rules, the analysis engine collects the error data and creates a problem report containing a description of the problem and any corrective actions required.
 - The problem report is distributed in accordance with the customer's notification scheme (for example, PRS or OSEM).
- Uses a web-based user interface that connects to the SAN switch and can perform a variety of tasks from a remotely connected web browser. In addition, a set of command-line tools enable local viewing and diagnosis of binary event logs.

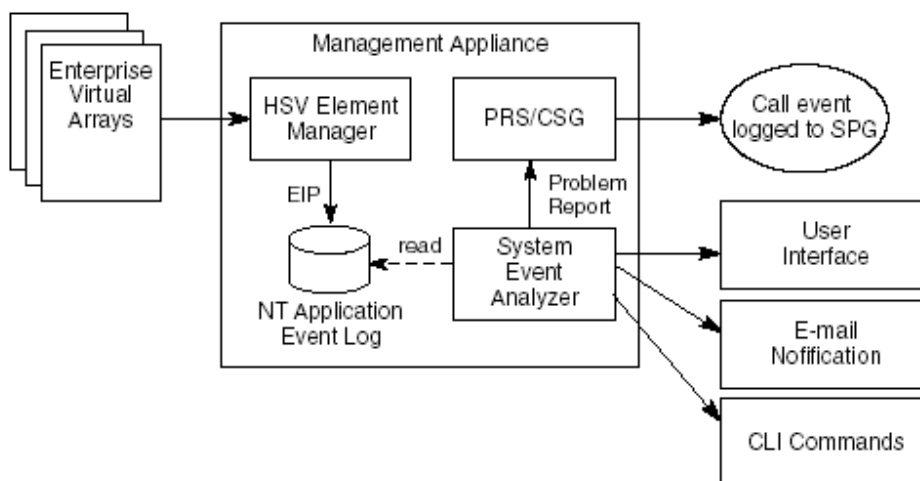


Important

System Event Analyzer reads event information directly from the NT Application Event Log on the storage management appliance; PRS environments without SEA receive event notification by means of SNMP traps.

System Event Analyzer Architecture

The following diagram depicts SEA in the Enterprise Virtual Array environment.



Configuring SEA with the Enterprise Virtual Array

You can find detailed directions for configuring SEA at the following website: <http://techpubs.cxo.cpqcorp.net/>. Reference the material for WEBES.

Disabling SNMP Trap Notification

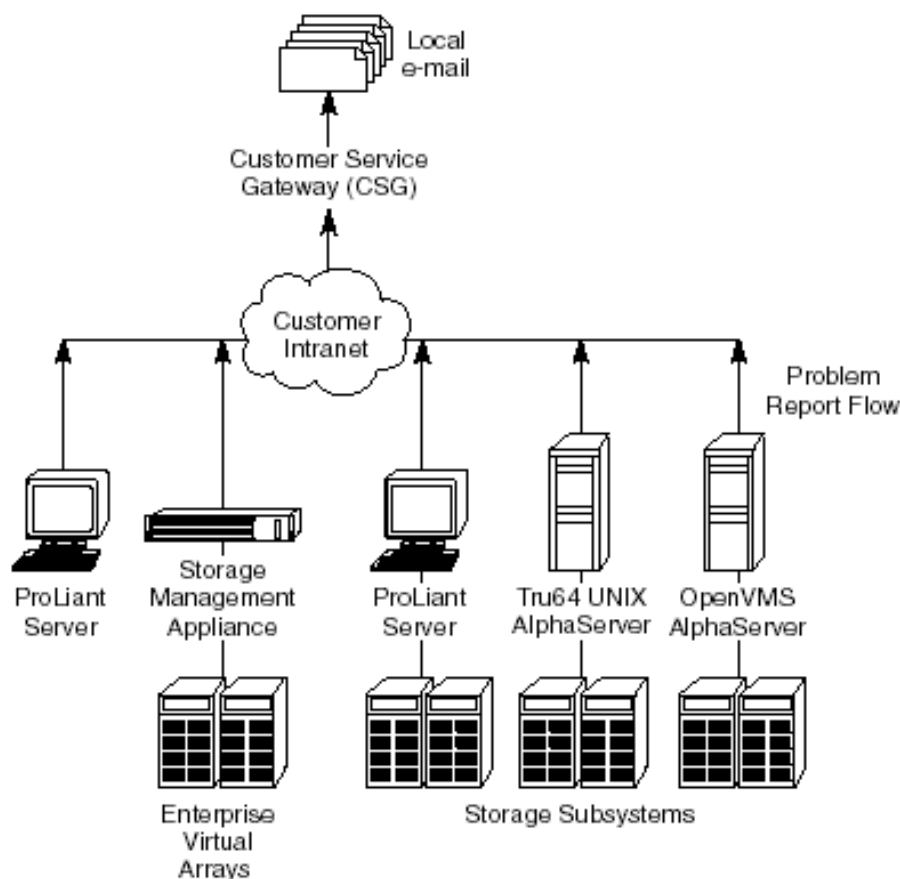
System Event Analyzer reads events directly from the NT Application Event Log rather than receiving them as SNMP traps from Command View EVA. Therefore, in the SEA environment, event notification to the server where SEA is installed (the storage management appliance or a separate ProLiant-type server) must be disabled as described in the following steps:

1. In Command View EVA, navigate to the Configure Host Notification page.
2. Observe the Host Notification List.
3. If the entity running SEA is enabled for notification, remove it as follows:
 - a. Navigate to the Modify Host Notification List page.
 - b. Select the host name in the Host Notification List box and delete it.
 - c. Click the Save changes button.
 - d. When the screen returns to the Configure Host Notification page, click the OK button.

Notification Using Open Service Event Manager

Open Service Event Manager (OSEM) allows customer systems to automatically send service event notifications as email messages within their own intranet.

OSEM uses the Common Remote Support Module (CRSM) and a designated Customer Service Gateway (CSG). Customer systems self-report events by running automated software agents that detect and forward events to CRSM. CRSM formats the problem reports and sends them to specified customer email addresses. The email notification list could include the HP Customer Support Center. This is depicted below.



The *Open Service Event Manager Setup Guide* has complete documentation for installing and configuring OSEM. This manual and other useful documentation is available at the following website: <http://techpubs.cxo.cpqc corp.net>. Navigate to links for PRS, OSEM, WEBES, and other pertinent products.

Diagnostics Overview

Several internal and external diagnostic tools are available for the Enterprise Virtual Array. The diagnostics described here refer to those that run internal software to test critical Enterprise Virtual Array operations.

The following controller diagnostics are described:

- Built-In Self Test (BIST) diagnostics
- Module Integrated Self Test (MIST) diagnostics
- Test entities and test numbers in EIP type 14
- Periodic diagnostics
 - Cache memory scrubber
 - Controller temperature sensing

BIST and MIST

Two primary types of controller diagnostics are used: Built-In Self-Test Diagnostics (BIST) and Module Integrated Self-Test Diagnostics (MIST).

BIST diagnostics are the first software to run when the hardware initializes. The goal of the BIST is to verify that the controller supports the execution of the functional software. To do this, BIST tests all components attached directly to the PowerPC and those accessible through the Quasar chip local bus.

The diagnostics that run after initialization of the functional image are referred to as the Module Integrated Self-Test (MIST). The MIST diagnostics verify all the devices and interconnecting buses that were not tested during BIST.

A controller block diagram is presented and BIST and MIST diagnostics are described in detail in Module 2, HSV Controller.

Test Entities and Test Numbers in EIP Type 14

Diagnostic failures are reported as controller events using EIP type 14. These events include fields called TE_num (test entity number) and Test_num (test number).

- **Test entities** (TE_num) are components undergoing diagnostic testing.
- **Test numbers** (Test_num) are the specific tests being applied to the component.

The TE_num and the Test_num are displayed in both decimal and hexadecimal formats in the on-screen Command View EVA display and the HSVET translation of EIP type 14.

Test Entities (TE_num) for EIP Type 14

The following is a list of test entities used for EIP type 14. Hex values are placed to the right of decimal values.

- TE 1.(0x01) — Temperature Sensor test
- TE 2.(0x02) — HW code check and LCD setup
- TE 3.(0x03) — Serial and WWID Number test
- TE 10.(0x0a) — Near PCI config
- TE 12.(0x0c) — Cache Memory test
- TE 13.(0x0d) — Cache Battery test
- TE 14.(0x0e) — Far PCI config
- TE 17.(0x11) — Port 4 (MP) test
- TE 18.(0x12) — Port 5 (FC1) test
- TE 19.(0x13) — Port 6 (FC2) test
- TE 20.(0x14) — Port 0 (DP2A) test
- TE 21.(0x15) — Port 1 (DP1A) test
- TE 22.(0x16) — Port 2 (DP2B) test
- TE 23.(0x17) — Port 3 (DP1B) test
- TE 24.(0x18) — All Ports test
- TE 25.(0x19) — Config & Init Port regs
- TE 28.(0x1c) — CBIC test
- TE 31.(0x1f) — Hardware Revision test

Sample Test Numbers (Test_num)

The following is a list of sample test numbers (Test_num) performed for TE_num 17 to 23 (ports).

Test_num (hex)	Test Name
01	Register read write test
02	Tachyon data lines test
04	Tachyon DMA read write test
05	Tachyon interrupt test
06	Internal loopback n-port initialization test
07	Internal loopback AL (arbitrated loop) initialization test
08	External pad loopback AL initialization test

Note

The Service Manual contains the complete lists.

Cache Memory Scrubber

Periodic diagnostics run while the controller is running the functional code. The functional code contains two periodic diagnostics, the Cache Memory Scrubber and Controller Temperature Sensing.

Cache scrubbing refers to reading cache memory independent of usage by the functional code to detect and correct ECC errors. All VCS versions contain a periodic cache scrubbing diagnostic. This diagnostic runs at the lowest priority and uses the Surge chip DMA engine to scrub any correctable ECC errors from cache. On an idle system, all cache memory is scrubbed every hour. When this periodic diagnostic detects a failure, a termination event is logged.

The following termination events may be reported by this diagnostic:

- The scrubbing software has had to scrub an excessive number of correctable ECC errors:
`84012003 Excessive correctable errors have been seen in cache memory.`
- The scrubbing software encountered one or more correctable errors that could not be repaired through scrubbing:
`8400200c Cache scrubbing encountered one or more hard correctable memory errors.`
- An Uncorrectable Cache Memory ECC error caused an interrupt:
`84022063 Uncorrectable ECC error in cache memory.`

Controller Temperature Sensing

All VCS versions contain a periodic diagnostic to check the temperature of the HSV controller every 10 seconds. This diagnostic has two event states:

- Temperature close to becoming out of range:

Logs a 090F2E05 event in the Controller Event Log:

A temperature sensor located within the HSV110 controller identified in the handle field is approaching its trip point. The value .ull field contains the current temperature reading from that sensor in Celsius. The secondary_id field contains the temperature trip point in Celsius.

If the current cache policy is writeback mode, it is changed to writethrough mode to prepare for shutdown, which occurs if the trip point is reached. In this case, the cache policy is returned to writeback mode when the temperature returns to an acceptable value.

- Temperature out of range:

Logs a 090E2E05 event in the Controller Event Log:

The temperature trip point for a temperature sensor located within the HSV110 controller identified in the handle field has been reached.

Logs a termination event with code 03660060:

This HSV110 controller was requested to terminate operation and then power off.

Shuts down the controller.

Controller environmental conditions (battery state, system board voltages, and blower speeds) are monitored by a separate PIC chip on the controller board. If this chip detects an out-of-spec condition, it sends an interrupt to the functional software. The functional software generates appropriate event codes and takes appropriate action, which may include adjusting the cache mode. When the PIC chip detects that conditions have returned to within spec, it interrupts the functional software, which generates appropriate event messages and takes appropriate action.

Learning Check

1. A termination event is generated when there is a _____.
 - a. Fatal error that causes the storage system to stop operation
 - b. Procedure completion
 - c. Configuration change
 - d. No AC input for power supply 1
2. Which of the following logs contains events that are significant occurrences within Command View EVA, such as configuration changes?
 - a. Controller Termination Event Log
 - b. NT Event Log
 - c. Controller Event Log
 - d. Management Agent Event Log
3. What are the three types of event logs you can display using Command View EVA?
.....
.....
.....
4. What is the primary method to view event logs.
.....
5. Which button on the event log pages allows you to download raw data files to your local PC?
 - a. Filter Events
 - b. Get log file
 - c. View
 - d. OK
6. What is the only event log seen from Command View EVA that is saved in ASCII format?
.....

7. Name three fields used in the Controller Event Logs.
.....
8. On what page in Command View EVA can you specify which events to capture as SNMP traps?
 - a. Set System Operational Policies
 - b. Set Host Options
 - c. Configure Host Notification Options
 - d. Filter Events
9. The two primary controller diagnostics are _____ and _____.
.....

Overview

This module provides an overview of the Storage System Scripting Utility (SSSU), also called the Scripting Utility. Included are installation and startup instructions, the command structure, naming conventions, and how to get help. Detailed information for each subcommand provided with the Scripting Utility, including command syntax and examples, is also included.

The version of SSSU described here is compatible with VCS V3.0 but also supports scripts written in previous SSSU versions. It also contains data replication functionality, however, those commands are not included in this module. They are, however, in the scope of the Continuous Access (CA) for EVA course.

Note

Refer to the *HP StorageWorks Storage System Scripting Utility, Enterprise Virtual Array V3.0 Reference Guide* for details on all of the information presented in this module.

Objectives

After completing this module, you should be able to:

- Describe the Storage System Scripting Utility (SSSU).
- Describe how to install and start the SSSU.
- Identify the command structure and naming conventions.
- Describe how to create storage systems, folders, disk groups, hosts, virtual disks, clones, and snapshots within the storage system.
- Describe how to select a storage system and display information about the elements in the selected storage system.
- Describe how to rename and remove storage systems, folders, disk groups, hosts, virtual disks, and monitor.
- Describe how to capture the configuration information of a specified storage system into a script.

SSSU Description


The Enterprise Virtual Array offers a command line application, the Storage System Scripting Utility (SSSU), which allows you to configure and control HSV controllers. You can use either Command View EVA or SSSU to issue configuration requests and operations. The scripting utility provides you with a character cell interface that allows you to enter simple, terse commands, and execute repetitious and complex configuration steps.

You can use these scripts to create and modify a configuration, including configuration commands to add, set, and delete storage systems, folders, disk groups, data replication groups, virtual disks, and hosts. You can also capture a configuration into a script file to later re-create the storage system configuration, and display configuration information about various elements of the storage system.

The SSSU is found in the platform kits for the specific operating system installed on the host server.

SSSU Architecture

The following diagram illustrates the SSSU architecture.



The Simple Object Access Protocol (SOAP) is based on an IBM SOAP4J implementation. It is a lightweight protocol for exchange of information in a decentralized, distributed environment.

The XML-based protocol consists of three parts:

- An envelope that defines a framework for describing what is in a message and how to process it.
- A set of encoding rules for expressing instances of application-defined datatypes.
- A convention for representing remote procedure calls and responses.

Terminology

The following table translates the SSSU terms into Command View EVA terms.

Term	Definition
System	A controller or controller pair that is treated conceptually as one controller. In Command View EVA, this appears as the name you give the initialized storage system.
Disk	A physical storage device (disk drive) connected to the system.
Disk Group	A collection of disks from which storage is created.
Folder	A logical organization of objects.
Host	A system that has data path access to the virtual disk through a LUN defined on the system.
LUN	The host-accessible presentation of storage on the system.
Manager	Command View EVA instance that controls configuration of the system.
Monitor	Controller enclosure.
VDisk	A virtual disk, a collection of blocks created on one or more disks, which may be configured as VRAID0, VRAID1, or VRAID5 for data storage.
Spare Policy	Disk failure protection level: NONE, SINGLE, and DOUBLE. <ul style="list-style-type: none"> ■ NONE — Reserves no space within a disk group to allow for data reconstruction for failure of disk drives ■ SINGLE — Reserves space within a disk group to allow for data reconstruction for failure of a single disk drive ■ DOUBLE — Reserves space within a disk group to allow for data reconstruction for failure of two disk drives



Important

Even though CELL, GROUP, and STORAGE have been replaced by SYSTEM, DISK_GROUP, and VDISK, the old references are still valid so that previously created scripts do not fail after an upgrade to version 3.

Installing SSSU

The SSSU is installed from the platform-specific kits and can run on any of the supported host operating systems. See the platform-specific installation guides for details on how to install SSSU for each operating system.

Depending on the operating system, the executable file is named *sssu* or *SSSU.EXE*. Once installed, SSSU can run as a simple executable. For convenience, you may add the directory containing the Scripting Utility executable to your path, or copy the executable to a directory already in your path.

Note

Verify that the attributes of Scripting Utility executable are set with the correct security and execution flags appropriate for your operating system environment.

Starting the SSSU



Important

To execute the Scripting Utility, ensure that the password access to the controller (if enabled) is already set up from the Command View EVA software. You cannot set passwords from within the Scripting Utility.

Start the Scripting Utility at a command prompt window or equivalent by entering the executable name (SSSU or `sssu`). If the Scripting Utility is started without arguments on the command line, a generic application prompt `NoSystemSelected>` is displayed.

If the Scripting Utility is started with arguments on the command line, those commands are echoed to the input terminal and executed; the utility exits when it encounters the EXIT command.

Syntax

SSSU <additional arguments>



Important

Commands or object names with embedded blanks (spaces) must be enclosed in double quotes.

Examples

SSSU

SSSU FILE "script.txt"

The first example starts the Scripting Utility without additional arguments and prompts you for commands. The second example starts the Scripting Utility and executes the file *script.txt* and then exits.

Note

The file extension is not required with Scripting Utility files; *.txt* is used as an example.

Command Structure

Most of the Scripting Utility commands have subcommands or switches. The generic term *options* is used for both subcommands and switches in the Scripting Utility help. The subcommand options are described under each command entry listed in this module. You can find detailed information on each of the switch options in the *HP StorageWorks Storage System Scripting Utility, Enterprise Virtual Array V3.0 Reference Guide*.

Getting Help

To display a list of options for each command, enter a space and question mark (“?”) after the command or option name.

For example, to get information on options available with the ADD command, enter `ADD ?` at the command prompt.

```
NoSystemSelected> ADD ?
```

The following list of options is presented:

```
COPY
DISK_GROUP
FOLDER
HOST
LUN
SNAPSHOT
SYSTEM
VDISK
```

The prompt returns, and you can continue entering commands.



Important

The old references, CELL, GROUP, and STORAGE, will not appear in the help options.

File Paths and Naming Conventions

You use the following root folders to organize your storage system:

- `\Hosts`
- `"\Virtual Disks"`
- `"\Disk Groups"`
- `"Hardware"` -- You cannot create any objects within this folder.



Important

Commands or object names with spaces must be enclosed in double quotes. Be sure to use double quotes (") around the entire path name when the name includes a space.

You cannot create root-level folders, but you can create additional folders below the `\Hosts` and `"\Virtual Disks"` root levels to organize your storage system.

The Scripting Utility requires that you qualify specified names with full paths. However, if you are using the default placement for hosts and virtual disks, you do not need to include the full path.

Note

Commands and subcommands can be abbreviated as long as you enter enough characters to differentiate from another command or subcommand.

SELECT Command

Use the SELECT command to select the desired storage management appliance or storage system. The prompt changes to reflect the selected storage system.

SELECT MANAGER

The SELECT MANAGER command directs the command prompt to Command View EVA. All configuration commands then affect Command View EVA. If the Command View EVA name contains spaces, the name must be surrounded by double quotes (“”).

Note

The SELECT MANAGER operation does not validate login information until a command to the manager is issued. Therefore, perform a SHOW SYSTEM command to validate a successful login.

Example

```
SELECT MANAGER swma31K023 USERNAME=administrator PASSWORD=administrator  
SHOW SYSTEM
```

SELECT SYSTEM

The SELECT SYSTEM command directs the command prompt to the selected system. All configuration commands then affect the selected system. If the system name contains spaces, the name must be surrounded by double quotes (“”).

Note

After selecting a system, the prompt name changes to reflect the selected system.

Example

```
SELECT SYSTEM "HSV Storage System"
```

ADD Command

Use the ADD commands to create the different entities in the system, including a storage system, folders, disk groups, virtual disks, copies, and snapshots, as well as to add LUNs and hosts within the Enterprise Virtual Array.

ADD SYSTEM

The ADD SYSTEM command creates an initialized storage system and its default disk group. Use this command to initialize an uninitialized storage system. If the storage system is already initialized, the command is rejected.

You must select an uninitialized storage system before entering the ADD SYSTEM command. The prompt changes to `Uninitialized Storage System#>`. After initialization is complete, the Scripting Utility changes its default prompt back to `NoSystemSelected>`. You must select the system again using the new name of the system because the system's name changed to the name given with the ADD SYSTEM command.

Attempting to add a storage system without a Basic license entered using the Command View EVA software will return an error message indicating that you need to enter your licensing information in Command View EVA. Licenses cannot be added using the Scripting Utility.

Example

```
ADD SYSTEM Storage_System_1 DEVICE_COUNT=10 SPARE_POLICY=NONE
```

ADD COPY

The ADD COPY command creates a copy of the specified virtual disk. The ADD COPY command is equivalent to creating a clone within the Command View EVA software.

Example

```
ADD COPY Wednesday_Backup VDISK= "\\Virtual Disks\payroll\ACTIVE"
```


ADD FOLDER

The ADD FOLDER command creates a new folder within the current folder to aid in organizing the storage system. Folders can only be created under the “\Virtual Disks” and \Hosts root folders. You cannot create root folders.

For example, if you have a controller that is serving Human Resources and Engineering, you could create four folders—two to separate the virtual disks and two to separate the hosts:

- “\Virtual Disks\Engineering” and “\Virtual Disks\HR”
- \Hosts\Engineering and \Hosts\HR

Creating these folders allows you to put engineering storage and hosts in the Engineering folders and Human Resources storage/hosts in the HR folders. This makes it easier to keep track of the pieces within your storage system.

ADD DISK_GROUP

The ADD DISK_GROUP command adds a disk group to the system. It is good practice to use multiple disk groups because it allows for some failure isolation and possible performance enhancement.

If all the disks are in one group, any disk failure could cause all the virtual disks to run in degraded mode until the data on the disk that was lost is recreated on other disks. By creating multiple disk groups, a disk failure only affects the virtual disks in that single group.

ADD HOST

The ADD HOST command adds the World Wide Name (WWN) of the host port to the list of hosts that can connect to virtual disks within the current system.

The ADD HOST command adds the first FCA only. Each subsequent FCA is added with the SET HOST command.

ADD LUN

The ADD LUN command makes previously created storage available to a specified host.

ADD SNAPSHOT

The ADD SNAPSHOT command creates a fully-allocated or demand-allocated snapshot of the specified virtual disk.

Attempting to add a snapshot without the license returns an error message. This message indicates you must enter your licensing information in Command View EVA.

ADD VDISK

The ADD VDISK command creates a virtual disk with the specified name and parameters.

Example

```
ADD VDISK "Pictures" SIZE=20
```



Important

This command creates a virtual disk family, not the virtual disk. After you create the family (ADD VDISK), an initial virtual disk named *ACTIVE* is placed as the only virtual disk in the family. When you later refer to this virtual disk, specify \ACTIVE after the family name.

SET Command

The SET commands allow you to rename storage systems, virtual disks, folders, disk groups, hosts, and the monitor. It also allows you to reset existing storage settings and application options.

All of the SET commands require the full path from the root directory to the object you want to rename.

SET SYSTEM

The SET SYSTEM command changes the name of the specified storage system.

SET DISK

The SET DISK command changes the name of the specified disk.

SET FOLDER

The SET FOLDER command changes the name of the specified folder.

SET DISK_GROUP

The SET DISK_GROUP command changes the name of the specified disk group.

SET HOST

Use the SET HOST command to add or delete World Wide Names (WWNs), set the operating systems for a host, and comments to a specified host.

Use the SET HOST command to add additional Fibre Channel adapters (FCAs) to hosts added to the storage system with the ADD HOST command.

SET MONITOR

The SET MONITOR command changes the name of the specified monitor.

SET VDISK

The SET VDISK command allows changes to the properties of specific virtual disks (storage):

- **WRITETHROUGH:** Data is written to the physical devices before the write is reported as complete to the host.
- **WRITEBACK:** Once the data is written to the storage system's cache, the write is reported to the host as complete. If the storage system's cache is not in mirrored mode, writes continue in writeback mode.
- **MIRRORED:** Same as writeback, with the exception that if the storage system's cache is not in mirrored mode, writes resort to writethrough mode for reliability.
- **WRITE_PROTECT:** Disables writes on this storage by all presented LUNs.
- **NOWRITE_PROTECT:** Enables writes to this storage by all presented LUNs.

SET OPTIONS

The SET OPTIONS command allows you to change the characteristics of the SSSU. Because of this, it is important to review the following subcommands that you can use:

- **COMMAND_DELAY=** [0-300 seconds], default = 10 seconds
The number of seconds to wait between issuing commands when running a script using the FILE command. This has no effect when you are typing commands.
- **NOCOMMAND_DELAY**
This specifies no wait time between commands issued from within a FILE command.
- **DISPLAY_WIDTH=** [70-500 character line width], default = 80 characters
This sets the amount of characters displayed on a line for SHOW commands. It is useful to set a high line width to keep lines from wrapping and to allow easier cutting, pasting, and parsing of output.
- **ON_ERROR=** [CONTINUE/EXIT_ON_ERROR/HALT_ON_ERROR], default = CONTINUE
This determines the behavior that causes the scripting utility to halt.
- **RETRIES=** [1-10 minutes], default = 4 minutes
This is the number of minutes to attempt long period retries such as when the Command View EVA service is either busy or restarting.
- **NORETRIES**
This specifies that you do not want the scripting utility to retry commands.
- **SAFE_DELETE**, delete all overlying objects first (default)
This requires the user to delete all overlying objects before the object specified is deleted.

DELETE Command

Use the DELETE commands to remove systems, folders, disk groups, data replication groups, virtual disks, hosts, and LUNs from the Enterprise storage system. The behavior of all DELETE commands can be controlled by the SET OPTIONS command. If the SET OPTIONS command uses the default SAFE_DELETE switch, the user is required to delete all overlying objects before the object specified can be deleted.

All of the DELETE commands require the full path from the root directory to the object you want to delete.

DELETE SYSTEM

The DELETE SYSTEM command permanently removes the storage system from the Enterprise Virtual Array. The system is no longer accessible, and all data on it is lost.



Caution

The DELETE SYSTEM command removes the entire storage configuration. All customer data on the entire storage system is lost upon execution of this command. In addition, all information about LUNs and HOSTs is lost. Be sure you can afford to lose all data on the storage system if you use this command.

DELETE FOLDER

The DELETE FOLDER command deletes the specified folder. If you wish to delete multiple layers of folders, you must do so sequentially.

DELETE DISK_GROUP

The DELETE DISK_GROUP command deletes a disk group from the configuration. The command is rejected if any virtual disks (storage) are present in the disk group.

DELETE HOST

The DELETE HOST command removes the specified host from the list of hosts that connect to the Enterprise storage system. The command is rejected if any virtual disks (storage) are present to the specified host.

DELETE LUN

The DELETE LUN command disables access to a virtual disk through the LUN from the specified host.

DELETE VDISK

The DELETE VDISK command destroys the specified virtual disk (storage). The virtual disk is no longer accessible, and all data on it is lost.



Caution

The DELETE VDISK command removes the entire storage (virtual disk) from the storage system. All customer data on the storage specified is lost upon execution of this command. In addition, all information about the LUNs presented from this storage (virtual disk) is lost. Be sure you can afford to lose all data on the storage specified if you use this command.

CAPTURE CONFIGURATION Command

The CAPTURE CONFIGURATION command interrogates the currently selected storage system and creates five SSSU scripts. The script can then be used to recreate the storage system's configuration—including all of the disk groups, folders, hosts, virtual disks, and LUNs.



Important

This command only restores the configuration and does not recover customer data.

The scripts output to the console unless a file is specified. In this case, the scripts generated outputs to five files. Given a file name, **_StepXX** is appended after the file name and before the extension. *XX* is the restore step name, which is 1A, 1B, 1C, 2 and 3.

For example, specifying CAPTURE CONFIGURATION **newyear.txt** causes SSSU to create file files named the following:

- newyear_Step1A.txt
- newyear_Step1B.txt
- newyear_Step1C.txt
- newyear_Step2.txt
- newyear_Step3.txt

SSSU checks to see if any files by the names to be generated exist. If so, the user is prompted to replace the files by the new files or abort the command completely.

When creating files that contain the scripts, periods (.) are printed on the console to show progress since this command may take a long time to complete, depending on the size of the configuration.

Note

The selected system must not be reconfigured while this command is executing.

Description of Script Files

Details of the contents of each step's script follow.

Step1A Script

The Step1A script creates the storage system itself, disk groups, hosts, VDISKS that are not used for DR (either source or destination) and LUNs for the disks created.

Step1A creates a basic system that exhibits no DR VDISKS or groups.

Step1B Script

The Step1B script creates all source VDISKS used in DR groups on this controller.

Step1C Script

The Step1C script presents all source VDISKS (creates LUNS) used for DR groups to their hosts. This step is broken out to assist the user in recovering from a DR failure in which a source site was lost. They can be presented (LUNs) in their original configuration by simply running the correct Step1C script.

Step2 Script

The Step2 script recreates all DR-specific configuration information only, and only that DR-specific information for which this system is the source. This is the configuration's source DR_GROUPS and their members only. Note that this is only the system's DR configuration. Also note that presentations of remote VDISKS are not restored by this command (see Step3).

This is a separate step to add flexibility when a site is completely lost. Scripts Step1A, Step1B and Step1C must be run on both source and destination systems before Step2 can be run.

Step3 Script

The Step3 script recreates a SSSU script that presents all remote VDISKS to their hosts only.

This is a separate file because if disaster recovery is required, script Step2 must be run on both source and destination systems before Step3 can be run.

Example

```
CAPTURE CONFIGURATION c:\sales.txt
```

This example creates the files named sales_Step1A.txt, sales_Step1B.txt, sales_Step1C.txt, sales_Step2.txt, and sales_Step3.txt.

Restoring Configurations

To restore a non-DR system configuration from captured scripts, run the Step1A script.

To restore a DR system configuration from captured scripts, five steps must be performed in a specific order on each system that makes up the DR configuration. First the Scripts 1A, 1B and 1C must be run against **all** systems in the DR configuration. Once all Step1 scripts have run successfully against all systems, then the Step2 script must be run against **all** systems in the DR configuration.

Once the Step2 scripts have been run successfully against all systems, then the Step3 script must be run against **all** systems in the DR configuration.

Note for each step that all systems in the DR configuration must have successfully run the current step before any script for the next step can be run on any other system. In other words the script steps must be run in lockstep on all systems. All Step1 scripts must be run and complete successfully on all systems before running any Step2 scripts, and all Step1 scripts must be run and complete successfully on all systems before running any Step3 scripts.

Once these scripts have been run against all systems, the full DR configuration is restored to its original state. The division of Step1 into three parts is to facilitate DR recovery.

SHOW Command

Use the SHOW commands to display information about various elements in the currently selected storage system.

All of the SHOW commands require the full path from the root directory to the object you want to display.

SHOW SYSTEM

The SHOW SYSTEM command displays the storage systems currently managed by Command View EVA, or, if *<system_name>* is specified, detailed information about that specified system.

SHOW DISK

The SHOW DISK command displays disk configuration information for physical disks connected to the currently selected storage system.

SHOW DISK_GROUP

The SHOW DISK_GROUP command displays the disk groups currently configured by this storage system, or, if *<disk_group_name>* is specified, detailed information about the specified disk group.

SHOW HOST

The SHOW HOST command displays the hosts currently supported by this storage system, or, if *<host_name>* is specified, information about the specified host.

SHOW LUN

The SHOW LUN command displays the LUNs currently configured on this storage system.

SHOW MANAGER

The SHOW MANAGER command displays the managers currently available, or, if *<manager_name>* is specified, detailed information about that specific Command View EVA instance.

SHOW MONITOR

The SHOW MONITOR command displays the monitors currently installed on this storage system, or, if *<monitor_name>* is specified, detailed information about that specified monitor.

SHOW OPTIONS

The SHOW OPTIONS command displays the current SSSU options.

SHOW POWER

The SHOW POWER command displays the controllers currently configured on this storage system, or, if *<controller_name>* is specified, detailed information about that specific controller.

SHOW VDISK

The SHOW VDISK command displays the storage currently configured on this storage system, or, if *<storage_name>* is specified, detailed information about the specified storage.

SHOW WORLD_WIDE_NAME

The SHOW WORLD_WIDE_NAME command displays the host WWNs that are visible to this storage system and not already assigned to a host.

SHUTDOWN Command

The SHUTDOWN command causes any object displayed by the SHOW POWER command to shutdown. A power object is an object that you can control in terms of powering off and on. An example is the HSV controller. No subcommands are available.

The command disables power from one or more of the following:

- Upper or lower controller enclosure
- Some of the number 1 power supplies on the drive enclosures
- Some of the number 2 power supplies on the drive enclosures

Check to be sure of the following:

- The drive enclosure power supply status LED is ON.
- The associated PDM circuit breaker is ON.
- All AC power cords are connected properly.

There are two switches you can use with the command:

- `ALL_PEERS` — All power objects (both controllers) on this storage system will be restarted together.
- `NOALL_PEERS` — Only one power object (the specified controller) will be restarted. This is the default.

RESTART Command

The RESTART command causes any object shown by the SHOW POWER command to restart. A power object is an object that you can control in terms of powering off and on. An example is the HSV controller. There are no subcommands.

FILE Command

The FILE command causes the current mode of input to suspend, and redirects the scripting utility to accept input from the specified file. Either the end of the file or an EXIT command in the specified file causes SSSU to again accept input from the previous input source.

FILE commands can be nested; that is, a file being executed through a FILE command can have FILE commands within its command set. The only limitation on how deep FILE commands can be nested is based on the host system's resources.

EXIT Command

The EXIT command terminates the SSSU session. If SSSU is accepting input from the terminal or the command line, EXIT causes the input to terminate.

Any commands on the command line after an EXIT command are not processed.

If SSSU is processing a script from a file due to encountering a FILE command, the behavior of the EXIT command depends on how the FILE command was entered:

- If the FILE command was entered from the command line, an EXIT command in the file causes SSSU to terminate.
- If the FILE command was entered from the terminal, an EXIT command in the file causes SSSU to return to accepting input from the terminal.
- If the FILE command was entered from a file that is being executed from a previous FILE command, an EXIT command in the file causes control to return to the previous file being executed. The execution continues with the command following the FILE command that started the execution of the file that the EXIT command was found within.

Note

Commands that are found in a file after the EXIT command are not executed. EXIT causes an immediate return to the calling layer.

HELP Command

The HELP command displays information on using the built-in help.

Learning Check

1. What does the Storage System Scripting Utility allow you to do?
.....
.....
.....
2. What does the SSSU term *Spare Policy* translate to in Command View EVA?
.....
.....
3. What must be set up from Command View EVA to execute the Scripting Utility?
.....
.....
.....
.....
4. List three folders used in organizing your storage system with SSSU.
.....
.....
.....
5. What is the command syntax for displaying a list of options available with the SHOW command?
.....
.....
.....
.....
6. What does the SELECT MANAGER subcommand do?
.....
.....
.....
.....

7. Which of the following subcommands allows you to add FCAs after the first FCA has been set?
 - a. ADD HOST
 - b. ADD VDISK
 - c. SET HOST
 - d. SET VDISK

8. Explain what is happening in the following command syntax.

```
ADD SNAPSHOT "\\Virtual Disks\payroll_backup" VDISK= "\\Virtual  
Disks\payroll\ACTIVE"
```

.....

.....

.....

.....

9. Why is the CAPTURE CONFIGURATION command important?

.....

.....

.....

.....

Overview

This module provides information on known best practices for configuring disk groups, sizing disk groups, and improving their availability and performance. The module begins by describing the current rules and limitations of the Enterprise Virtual Array for any configuration. Next is a description of the factors that go into disk group configuration and sizing, including a formula for determining the number of disks in a disk group. Finally, information is presented about current utilities for determining Enterprise Virtual Array storage capacity requirements.

Objectives

After completing this module, you should be able to:

- List Enterprise Virtual Array configuration rules and limitations.
- Describe the key considerations that go into configuring disk groups.
- List the best practices for configuring disk groups.
- Describe the factors that go into sizing disk groups.
- Use the formula for calculating disk counts needed for the given requirements of a disk group.
- Identify some utilities available for calculating Enterprise Virtual Array storage capacity requirements.

Review of Enterprise Virtual Array Rules and Limitations

Before you attempt to configure and optimize an Enterprise storage system, you should be aware of its rules and limitations. These maximums, in some cases, are not hard limits, but are the maximum supportable limits.

Rules for Hosts, Connections, and LUNs

An Enterprise Virtual Array storage system has support for a maximum of:

- 256 hosts, where a host is defined as one or more HBAs.
- 1,024 HBA connections.
- 512 LUNs (virtual disks). The maximum of 512 virtual disks is not a hard stop, but no more than 512 are supported.
- 256 LUNs (virtual disks) on any one HBA.
- 8,192, the sum of all LUN-to-host presentations.

Example

1 LUN presented to 1 host = 1 presentation

1 LUN presented to 256 hosts = 256 presentations

256 LUNs presented to 1 host each = 256 presentations

Note

Presentations include **all** LUN-to-host presentations including snapshots and snapclones. If the **same** unit is presented to two different hosts, that counts as two presentations.

Refer to the *SAN Design Reference Guide* for all of the SAN configuration rules for the Enterprise Virtual Array, including those listed above.

Other Rules for the Enterprise Virtual Array

The following are some other rules and limitations that you should always be aware of while configuring an Enterprise storage system:

- Maximum number of disk groups in a storage system is 16.
- Minimum of eight disks per disk group. (See note.)
- Minimum of four disk drives per enclosure.
- Virtual disk size can be from 1GB to 2TB in 1GB increments. Note that Sun Solaris has a 1.024TB maximum.
- A **single** protection level for a disk group reserves two times the largest drive capacity of any drive in the group; a **double** protection level for a disk group reserves four times the largest drive capacity of any drive in the group.

For example, if your disk group has ten 36GB disks and one 72GB disk, and you need double protection, four 72GB drives are reserved for spare capacity. If you need single protection, two 72GB drives are reserved.

- Maximum of 16 storage systems can be managed per storage management appliance. Note that performance may become so poor that you may need to add a storage management appliance well before the maximum of 16.

Available Tools and Documentation

The following are two white papers that give details about configuring and sizing disk groups:

- *Configuring Disk Groups and Virtual Disks in a StorageWorks Enterprise Virtual Array*
- *Sizing Disk Groups in a StorageWorks Enterprise Virtual Array*

These documents are available through TechBB.

The following utilities are also available through TechBB:

- Disk Group Sizing Utility
- Disk Groups and Virtual Disks Calculator

Both of these utilities are described later in this module.

Disk Group Configuration Overview

You must decide on the number and organization of disk groups in your Enterprise storage system. While virtualization frees the user from most low-level design considerations encountered with HSG storage systems such as bus location, container size, and spindle size, virtualization also requires a new way of thinking about storage planning. Now you must consider items such as capacity, number of spindles, RAID level, availability, and performance.

The only physical boundary to be considered when structuring storage is the disk group.

- Each disk group in an Enterprise Virtual Array subsystem constitutes a separate storage pool, distinct from and independent of the other disk groups in the subsystem.
- Virtual disks must reside entirely in a single storage pool; therefore, disk groups create boundaries that the virtualization algorithms will not cross.

Disk group boundaries have multiple consequences.

- On one hand, they allow the user some high-level control over data placement, and can thus be used to influence availability and performance.
- On the other hand, the benefits of virtualization are diminished by disk group boundaries.
 - The capacity efficiency and load balancing benefits of virtualization increase as the number of disks that can be targeted for a given operation increases.
 - The benefits are maximized if the disk groups are few and their size is large.

Therefore, configuring Enterprise Virtual Array disk groups is an exercise in determining the minimum amount of control necessary to meet extraordinary availability, performance, and capacity requirements, and then letting the subsystem take care of the rest. Each of these requirements is described in detail.

Disk Group Configuration Considerations

You configure disk groups in a certain way because disk group structure can affect the balance between availability, sequential performance, and random performance. You cannot simultaneously maximize all three.

Either of the following two situations will be present:

- Simple requirements—those configurations that require few but large disk groups, will have very high availability, very good random and sequential performance, and very good total cost of ownership. This is a good cost-effective choice for many common situations.
- Unusual requirements—those configurations that require any of the following:
 - Demanding availability requirements such as protection against multiple, simultaneous failures
 - Demanding performance requirements such as sustained, 50MB/s sequential streams or mixed drive performance requirements

If either of these situations arise, then further attention to disk group configuration is required.

Disk Group Availability

Disk groups bind data to a specific set of disk drives, and therefore disk failures will not cross these disk group boundaries. The inability to access any or all of the disks in a specific disk group does not deny access to the data in any other disk group.

The Enterprise Virtual Array is designed to deliver better than 0.99999 availability for VRAID1 and VRAID5 data due to a single failure. It is also designed such that VRAID1 and VRAID5 data will withstand most compound failures.

Also, while simultaneous failure of the active components is also very improbable, simultaneous failure of multiple disk drives or Fibre Channel connections can also cause data to become inaccessible.

The inherent high availability of the Enterprise Virtual Array subsystem meets the needs of many applications without further consideration. For very demanding situations, however, disk groups can be used to improve data availability through:

- Using virtual disk separation

Placing virtual disks in different disk groups divides them into separate failure domains. When the application is keeping two independent copies of the data (such as a database and its transaction log or the two copies of an application-level mirror), placing each copy in a separate disk group ensures that they fail independently.

- Using vertical disk groups

By ensuring separation of data on multiple enclosures, you have extra protection against problems within a single drive enclosure


- Using smaller disk groups if restricted restore times have been set

While compound failures that affect VRAID1 or VRAID5 data are very improbable, most situations require that such a failure be planned for. If the recovery plan requires restricted restore times, and dictates that those abbreviated times be achieved by reducing the amount of data to be restored, then dividing virtual disks among multiple, smaller disk groups will reduce the amount of data affected by a disk group failure.

Note that achieving reduced restore times by dividing the disk pool in this manner can reduce the other benefits of virtualization, and thus raise the total cost of ownership. Many situations will address demanding recovery requirements by other means that do not require restricted disk group sizes (such as site mirrors or more capable nearline hardware).

Note

The manner in which disk groups are arranged across drive enclosures is not considered because Enterprise Virtual Array subsystems are designed to have better than 0.99999 availability for VRAID1 and VRAID5 data regardless of how drives are arranged across enclosures. Only consider disk group layout for the most demanding situations.



Data protection attributes also ensure data availability and are set per virtual disk. The protection level choices are the following:


- VRAID1 — High
- VRAID5 — Moderate (sometimes better than traditional RAID5)
- VRAID0 — None

Parity organization in the VRAID levels protects large groups from multiple drive failures:

- There are a minimum of six drives in a parity group.
- There is an average of eight drives.
- The worst case is 11 drives.

Only multiple drive failures in the same parity group affect data:

- VRAID5 requires a second failure within the group.
- VRAID1 requires a failure of both mirror partners.



Disk Group Performance

The boundaries created by disk groups not only localize failures, they also create boundaries that localize I/O activity.

While the Enterprise Virtual Array is designed such that I/O activity is balanced across all available disks (thus increasing the average throughput for most situations), installations with multiple I/O streams of markedly different character focused on a small number of physical disks can see suboptimal performance. For example, if a highly random, small-transfer I/O stream were directed at a virtual disk in the same disk group as a virtual disk handling a highly sequential, large-transfer I/O stream, the ability of the Enterprise Virtual Array to optimize the disk actuators would be diminished.


The automatic load balancing of the Enterprise Virtual Array subsystem meets the performance needs of many applications without further consideration. For very demanding situations, however, disk groups can be used to improve subsystem throughput as follows:

- Through separate I/O access domains (virtual disk separation)

Creating virtual disks in different disk groups divides them into separate performance domains. By directing sequential I/O streams and random I/O streams, or I/O streams with dramatically different locality of reference, to separate virtual disks in different work groups, disk actuator efficiency can be improved.

The importance of this separation diminishes as disk group size increases. If a disk group is large, sequential and random I/O streams can both be simultaneously serviced with good performance. As disk group size decreases, however, the benefits of separating different I/O streams into different groups increases. In general, large disk groups give best random access performance.

Note also that the performance consequences of disk group configuration are largely limited to the above. Attempting to finely tune performance by balancing I/O activity between Fibre Channel loops or specific disks does not yield much result; the Enterprise Virtual Array algorithms already do this to the point of diminishing returns.



The following are some other important facts to note when configuring virtual disks for performance:

- The Read attributes of all VRAID levels are similar.
- The Write attributes vary, with VRAID5 performing only slightly slower the VRAID1.

Disk Group Capacity Utilization

Boundaries created by disk groups not only localize both failures and I/O activity, they limit the ability to use free capacity, thus impacting the total cost of ownership.

While Enterprise Virtual Array virtualization minimizes the amount of stranded capacity in a subsystem, the failure and performance isolation properties of disk groups require that all capacity used by virtual disks in the disk group come from physical disks in that group. This impacts capacity efficiency in the following ways:

- Spare capacity efficiency

The spare capacity used to recover redundancy due to disk failure currently cannot be shared across groups. Because each disk group must have its own spare capacity, increasing the number of disk groups increases the total spare capacity requirement for the subsystem. See the later subtopic on spare capacity.

- Snapshot and snapclone space efficiency

The capacity required for snapshot or snapclone operations must come from the same disk group as the target virtual disk. Increasing the number of disk groups in the subsystem that are undergoing snapshot or snapclone activity increases the total free space that must be set aside in the subsystem for this function.

- New virtual disk creation flexibility

In addition to snapshots and snapclones, the capacity required for new virtual disks currently must come entirely from the target disk group. Because the subsystem manages unused pool capacity within the disk group boundary, increasing the number of disk groups divides the unused capacity into more, smaller pools, thus reducing overall flexibility.

- VRAID1 and VRAID5 efficiency

The algorithms that control VRAID1 and VRAID5 data layout work best if all the disks in the disk group have the same physical capacity. Their effectiveness also improves if the group has an even number of disks, and is best if the number of disks in the disk group is either a multiple of the number of drive shelves (preferred), or is a multiple of six or eight. In no cases should a system be configured with a disk group of less than eight drives; while it is acceptable for a disk group to temporarily have fewer than eight drives due to failure, this is undesirable as a permanent condition.

Issues with Spare Capacity

Any unassigned capacity is spare capacity. The system draws on unassigned capacity as needed for:

- Writes after Vsnap
- New virtual disk creation
- Freeing a physical disk for removal or reassignment
- Data reconstruction after disk failure

Unassigned capacity increases when virtual disks are deleted or new physical disks are added to disk group.

Using the guideline to run with 5-20% unassigned capacity:

- Leaves room for Vsnap capacity creep
- Leaves room for operational flexibility
- Provides adequate resources for automatic adjustments
- Provides adequate space for data reconstruction after disk failure

Unassigned capacity is monitored with the Occupancy Alarm attribute for the virtual disk. When a threshold is exceeded, you can delete stale virtual disks or add more physical disks to increase unassigned capacity.

Note

Disk failure protection levels do not assign specific capacity or change the amount of unassigned capacity. They ensure that new virtual disk creations and Vsnap writes leave sufficient space for reconstruction. Protection levels are not structurally important if you already have sufficient free capacity.

General Steps for Configuring Disk Groups

The Enterprise Virtual Array performs admirably left entirely on its own. Successful configuration of an Enterprise Virtual Array, therefore, consists in understanding what minimal placement control is necessary to meet any extraordinary requirements, and limiting the number of disk groups to that which is strictly necessary. Leaving the subsystem as many degrees of freedom as possible maximizes the benefits of its virtualization capabilities.

The recommended steps to configuring disk groups are the following:

1. Understand any extraordinary availability requirements.
Consider whether the application requires virtual disk separation or smaller disk groups due to restricted restore time requirements (as discussed above).
2. Understand the performance requirements and the I/O pattern.
 - a. If the application I/O stream is dominated by I/Os of substantially equivalent transfer size and locality (for example, mostly sequential big transfers or mostly random small transfers), or does not require very high sustained sequential throughput, move on to Step 3.
 - b. If the application I/O stream consists of sustained large sequential transfers simultaneous with random transfers, understand how these streams need to be directed to specific virtual disks. Must they all be directed at the same unit, or can they be separated to different units?

Note

Transfer profiles to a given disk that differ over time are not a major consideration. A virtual disk that receives sequential transfers for part of the day and random accesses for the rest of the day works well in both cases. The situation of interest occurs where simultaneous sequential and random streams must be accommodated.

3. Determine the minimum number of disk groups needed.
 - a. Remember that too many disk groups may strand spare capacity.
 - b. Create the fewest number of disk groups consistent with the failure and performance isolation requirements.
 - c. The number of groups should be the greater of the number of disk groups required to get availability separation or the number of groups required for performance locality. For example, if the installation needs four independent failure domains, but only requires that one sequential stream be separated from all the others, then configure four groups. Conversely, if the installation requires only two failure domains but has six sequential I/O streams, configure for six groups.

4. Determine the number of physical disks required for each disk group.

After identifying the number of groups, determine which virtual disks are to reside in each disk group, then size each group for the appropriate capacity.

See the topic called Disk Group Sizing for more detail on how to define the number of disks for a given required user capacity, VRAID level, disk size, and protection level. These disk group size calculations will yield an arbitrary number of disks up to meet other criteria for best practices:

- a. Use disk group sizes that are a multiple of the number of shelves, as this yields the best available failure immunity.

If a disk group is not to be a multiple of the number of shelves, use disk groups that have multiples of 6 or 8 drives.

- b. Always configure disk groups with an even number of drives.
- c. Always attempt to configure groups with drives of identical or similar capacity and performance.

When sizing groups, also consider that the efficiency of the Enterprise Virtual Array virtualization algorithms improves with the amount of free space available. Given the trade-off between the cost of disk capacity and the cost of human management time, extra capacity will often pay for itself in reduced management overheads.



Important

Avoid the temptation to use disk groups in an attempt to tune the system beyond the point of diminishing returns. Do not attempt to calculate capacity requirements too finely; while virtualization improves capacity efficiency considerably, the virtualization algorithms still need free space to work with.

Disk Group Best Practices Summary

The following summarizes the best practices for configuring disk groups:

- For normal situations, use few, but large disk groups.
- Use drives in the disk group of the same capacity and model.
- Create disk groups that have an even number of member disks.
 - Strive for a multiple of the number of shelves.
 - Strive for a multiple of six or eight drives (eight is best).
- Do not run the system with occupancy greater than 90% (95% is the default). This may cause operations with leveling and protection attainment to **not** complete.
- In general, separate sequential I/O stream data (database logs, rich content) from random I/O streams (database information store, file shares).
- Avoid mixing VRAID0 and other VRAID virtual disks in the same disk group.
- When adding disks to expand the storage system:
 - Add disks equally and vertically across all shelves, that is, try to maintain the same number of disks on all shelves.
 - Add disks in even numbers.
 - Set the occupancy alarm level to at least 90%.
 - If you are planning on needing to add disks often, try to add disks all at one time. This increases RSS efficiency.
- Use off peak time to physically add devices to the storage system.
 - Add each drive carefully and firmly into the drive bay.
 - Wait at least 60 seconds between each drive insertion. This allows VCS recognition and reduces FC-AL loop signal changes.

Disk Group Sizing

When an Enterprise Virtual Array subsystem is deployed, one of the configuration questions that must be resolved is the number of physical disks required to deliver the desired *usable capacity*, that is, the capacity of the virtual disk as seen by the hosts to which it is presented.

The raw physical storage capacity available in an Enterprise Virtual Array subsystem is consumed for a number of purposes. The first and most obvious is the storage of the data written by the operating system and applications. Some of the physical storage, however, is used to store the information that makes the subsystem's fault tolerance and virtualization features possible.

Note

The material in this discussion comes largely from the white paper, *Sizing Disk Groups in a StorageWorks Enterprise Virtual Array* by Compaq.

Disk Group Sizing Factors

The following factors impact the translation of raw capacity into usable capacity:

1. Hardware versus software storage representations
2. System metadata overheads
3. VRAID redundancy overheads
4. Spare capacity
5. Snapshot working space

Therefore, sizing a disk group involves determining how much usable capacity is required, accounting for the fixed-system overheads, then factoring in the variable overheads associated with the VRAID type, spare capacity, and snapshot activity.

Hardware Versus Software Storage Representations

Most hardware storage capacity (including that of the Enterprise Virtual Array) is quoted using a decimal representation of bytes. Many operating systems, including those from Microsoft, use a binary representation (power of 2). Whereas 1 gigabyte is 1,000,000,000 bytes as a decimal representation, the binary representation is 2^{30} , or 1,073,741,824 bytes. Similarly, 1 terabyte is 1,000,000,000,000 bytes decimal, whereas the binary representation is 2^{40} , or 1,099,627,776 bytes.

The difference means that 1 physical gigabyte must be consumed for every 0.93 software gigabyte of usable capacity delivered, and that 1 physical terabyte must be consumed for every 0.91 software terabyte of usable capacity delivered. For Enterprise Virtual Array drives, that can mean using either fourteen or sixteen 72GB drives. The disk sizing formula accounts for this difference.

System Metadata Overheads

The subsystem stores its configuration, the tables that map virtual disks to specific physical disks blocks, and other system metadata on the disk drives. At disk group creation, the Enterprise Virtual Array keeps a minimum of five copies of metadata. Thereafter, metadata is duplicated in up to a maximum of 16 disk groups. It has been determined that approximately 0.2% of the physical capacity is consumed for this purpose. The disk sizing formula accounts for this overhead.

VRAID Redundancy Overheads

The parity used to protect VRAID1 and VRAID5 data from a disk failure also requires storage. VRAID1 data is stored twice, and thus consumes two blocks of physical capacity for every block of usable capacity. VRAID5 data stores one block of parity for every four blocks of data, and thus consumes 1.25 blocks of physical capacity for every block of usable capacity. VRAID0 data has no parity protection, and thus consumes only one block of physical capacity for every block of usable capacity.

Spare Capacity

When a disk drive fails, the subsystem reconstructs the missing VRAID1 and VRAID5 data, drawing from unused physical capacity in the affected disk group. To ensure that sufficient spare capacity is available to reconstruct all of the affected data, the system reserves physical capacity equivalent to twice the largest disk in the disk group for each level of disk failure protection selected.

For example, if your disk group has five 36GB disks and five 72GB disks, and you need double protection, four 72GB drives are set aside for spare capacity. If you need single protection, two 72GB drives are set aside.

Note

Spare capacity is not used unless available capacity in the disk group is depleted. Spare capacity is **true** spare space reserved for times when the disk group is nearly full.

Snapshot Working Space

An installation that creates snapshots or snapclones on a regular basis needs free capacity from which these new disks can draw. Snapclones and standard snapshots consume the same physical capacity as the original virtual disk. This additional capacity is consumed as part of the operation that creates the snapclone or standard snapshot.

Virtually Capacity-Free Snapshots (Vsnaps) consume physical capacity only when new data is written and capacity is allocated from the disk group in 1GB increments. When the Vsnap is first taken, both the Vsnap and the original target disk contain the same data, and thus share the same physical storage. Each time a block on the original target disk is written subsequent to the snapshot operation, the snapshot contents are preserved by copying the original data to a new location, after which the original target virtual disk block is updated with the new write data.

Therefore, a Vsnap consumes physical capacity as a function of the rate at which the original target disk is changed. A Vsnap for a virtual disk that changes at a slow rate consumes little additional storage; a Vsnap for a virtual disk that is completely rewritten subsequent to the snapshot operation consumes the same physical capacity as the original virtual disk.

An estimate of the usable capacity required for a Vsnap is calculated in two ways:

1. Use an estimate of the rate at which new data will be written to the Vsnap's original virtual disk, and the length of time the Vsnap will be in existence:

Virtual Disk Write Rate * Snapshot Duration

For example, if new data is being written to a VRAID5 virtual disk at 10 MB/s, and the Vsnap is to exist for two hours, the additional VRAID5 usable capacity consumed by the Vsnap will be:

$10\text{MB/s} * 7200\text{s} = 72\text{GB}$

The above applies no matter how many snapshots exist for a specific target virtual disk, as long as the Vsnaps themselves are not being actively written. The original data need only be copied once, so in this case the estimating rule is insensitive to the number of snapshots.

2. Use the 80/20 rule, that is, usable capacity for a Vsnap will be 20% of the total capacity of the disk group.

Formula for Determining Disk Count

A formula has been determined to define the approximate relationship between hardware disk capacity, number of disks, and usable capacity for each of the three VRAID types. To solve for the number of disks, supply the usable capacity for each of the three VRAID types, the disk drive capacity, and the protection level desired.

The formula assumes disk groups that follow the guidelines described in the topic called Disk Group Structuring. In particular, the groups are assumed to consist of a multiple of six or eight drives of identical or similar capacity.

The formula to determine disk count is:

$$\text{DiskCount} \cong ((\text{UsableV0} * 538) + (\text{UsableV5} * 673) + (\text{UsableV1} * 1076)) / (\text{DiskCap} * 476) + (\text{ProtLevel} * 2)$$

Where each of the variables is defined as follows:

DiskCap	Disk drive capacity in hardware GB
DiskCount	Integer number of disk drives
ProtLevel	0 for None, 1 for Single, 2 for Double
UsableV0	Desired usable VRAID0 capacity in software GB
UsableV1	Desired usable VRAID1 capacity in software GB
UsableV5	Desired usable VRAID5 capacity in software GB

For example, if you need 1TB of usable VRAID0, 1TB of usable VRAID5, and 1TB of usable VRAID1, all using 72GB disks, and you need double protection level, you can calculate the disk count you need:

$$\begin{aligned} \text{DiskCount} &\cong ((1000\text{GB} * 538) + (1000\text{GB} * 673) + (1000\text{GB} * 1076)) / (72\text{GB} * 476) + (2 * 2) \\ &\cong 66.7 + 4 \\ &\cong 70.7 \text{ disks} \\ &\cong 71 \text{ disks} \end{aligned}$$

Note

Remember to convert to the same units, for example, 1TB = 1000GB.

The above formulas are approximations because of the dynamic nature of physical capacity use in a virtualized environment. The variance inherent in dynamic algorithms results in a small variability in the outcome. Most of that variability has been accounted for in the formulas, and while not exact, the results should be generally applicable.

Disk Group Sizing Procedure Summary

To determine the number of physical disks required for the disk group:

1. Determine the physical capacity of the drives to be used, and the disk failure protection level required for the disk group.
2. Determine the number of virtual disks that are to reside in the disk group, the usable capacity required for each, and the VRAID type required for each.
3. Determine the amount of additional usable capacity required for snapshot working space for each virtual disk that will have snapshots. The VRAID type of the snapshots will match the original.
4. Sum the total usable capacity required for each VRAID type. This is the sum of the usable capacity for each virtual disk of that type and its corresponding snapshot working space (if any).
5. Solve for the number of disks using the formula.
6. Add margins appropriate to the circumstances. Subsystem management flexibility improves significantly as free capacity grows.

Storage Calculation Utilities

A few utilities are available for you to determine storage capacities for your Enterprise storage system configuration. These tools enforce the current maximums and minimums known for any Enterprise storage system while providing appropriate factors for virtualization overhead. The TechBB for the Enterprise Virtual Array contains access to these utilities.

Disk Group Sizing Utility

This Microsoft Excel spreadsheet allows you to determine the net storage capacity of an Enterprise storage system resulting from the selection of up to 16 disk groups, their VRAID mixes, and their protection levels. The impact of virtualization overhead is included through a VCS impact factor, and you are allowed to select the disk size (36GB or 72GB).

Disk Groups and Virtual Disks Calculator

This Microsoft Excel spreadsheet allows you to interactively choose an Enterprise Virtual Array configuration with a number of drives and capacity per drive. For this configuration, you interactively select the number of virtual disks per disk group (up to 16), and for each of those virtual disks (up to 512), select the usable capacity of the virtual disk and the VRAID type to determine the capacity used and the resulting free capacity. Virtualization overhead is accounted for in the calculations.

Learning Check

1. List the four major limitations for numbers of hosts, connections, and LUNs on an Enterprise Virtual Array storage system.
.....
.....
.....
2. Fill in the blanks. The maximum number of disk groups in a storage system is _____, the minimum number of disks in any disk group is _____, and the minimum number of disk drives per enclosure is _____.
.....
.....
3. What are the three factors that must be considered when configuring disk groups?
.....
.....
4. Fill in the blanks. When structuring disk groups for performance, use virtual disk separation to separate _____ and _____ I/O streams. This is especially important if the disk group size is _____.
.....
5. List the recommendations for the number and characteristics of the disks in a disk group.
.....
.....
.....
6. List the five factors that determine the number of disks required to deliver a given usable capacity in an Enterprise storage system.
.....
.....
.....
.....
7. To deliver 2TB of usable VRAID0 and 1TB of usable VRAID5, what is the approximate disk count that you would need?
.....
.....
.....

Overview

This module provides detailed information on the sources used for troubleshooting the Enterprise Virtual Array and all of its major components. This material does not describe troubleshooting situations and resolutions, but intends to prepare you for troubleshooting by supplying references and tools. Included in the module are troubleshooting aids and sources, including event logs and various LEDs. Additional troubleshooting tools and lists are described for each type of aid. Various references, both documentation and website, are summarized. Finally, this module includes a list of field replaceable units (FRUs) and customer replaceable units (CRUs).

Objectives

After completing this module, you should be able to:

- Identify the types of event and error data used in troubleshooting.
- Identify the tools used for decoding error messages.
- List the controller and disk enclosure drive status indicators, tools, and references used for troubleshooting.
- Identify the other LEDs and references used for troubleshooting.
- List the switch commands used for troubleshooting.
- Identify supplemental information sources used for troubleshooting.
- List the replaceable units (FRUs and CRUs).

Troubleshooting Background and Approach

The storage systems used in the past were very simple in data flow and structure. Therefore, the error paths and failure modes are relatively simple. The Enterprise Virtual Array has changed that model, and now service engineers must change their traditional models of service. With the Enterprise Virtual Array, you must collect and understand more event data before an analysis is feasible. Proper analysis is not realistic with only one event or event type. In most cases, it is necessary to review all entries for the timeframe associated with the events, and then view the events in the context of all errors and events.

A key change with the Enterprise Virtual Array is that you must view it as both hardware and software. Because the Enterprise Virtual Array tightly couples hardware and software operations, you must view failures in both contexts. The hardware failures can have a significant impact on the software, and that can result in an impact to user access and user data. Many attempts have been made to repair or service the Enterprise Virtual Array without fully understanding the impact a simple activity can have.

Developing a Troubleshooting Plan

The complexity of the Enterprise Virtual Array, its underlying virtualization software, and its access from various hosts through various paths requires a logical and methodical approach to troubleshooting. This requires being aware of and using all of the data that the system has to offer, and using them at the appropriate time.

A general troubleshooting plan consists of the following steps:

- Collect the facts. Identify and define the reported problem, fault, or symptom by using indicators in any of the troubleshooting aids. Define the nature and scope of the problem from a system wide perspective.

Note

You can acquire this perspective by creating a system diagram or map of all interdependencies and interrelationships of all elements within the system. This helps with diagnosis. For example, if symptoms suggest a problem with data transfers between Host B and a virtual disk, a map can identify the affected hosts, switches, interconnecting fibre cables, the WWIDs, and mapping information from the hosts to the virtual disks and LUNs.

- Analyze the system in more detail by checking additional troubleshooting aids (such as LEDs, displays, and reported errors). Do this to further define the nature of the problem and to begin to rule out possible causes.
- Acquire additional data from your troubleshooting aids to help further isolate the cause of the problem.
- Identify all steps necessary to correct the problem.
- Follow your planned corrective actions, being careful to manipulate only one variable at a time to avoid introducing new symptoms.
- Record actions and results.
- Run system diagnostics and checks to ensure that the problem is fixed.
- Identify and perform any steps to prevent recurrence of the problem.

The following section describes all of the troubleshooting aids that you should be familiar with and able to use when diagnosing and correcting problems.

Troubleshooting Aids and Tools

You have a choice of ways to collect troubleshooting data for the Enterprise Virtual Array and its major components. You can view and save event logs for changes in configuration or error conditions of a particular component. You can also use FC switch commands or port LEDs to verify FC connectivity. The following is a list of the sources for troubleshooting information:

- Event and error logs
 - Command View EVA logs
 - Event decoding and viewing tools
 - Other event and host error logs (inband)
 - SNMP error traps
- Command View EVA GUI pages
- SSSU commands
- HSV controller status indicators
- Disk drive enclosure status indicators
- Other LEDs (FC loop switch, SAN switch, FCA, appliance)
- Emulex Configuration Utility on the SMA
- LightPulse Utility (LPUTILNT) for Windows servers
- FC switch commands

This list only refers to physical aids. Also included are soft aids, such as lists, tables, and diagrams, which help you when troubleshooting.

Access to Troubleshooting Aids

The following table describes the means by which you access a specific aid.

Troubleshooting Aid or Tool	Access
Management Agent Event Log	Command View EVA
Controller Event Log	Note: HP Services only tools are also available to translate errors on the controller logs
Controller Termination Event Log	
Management Appliance NT Event Log	Attach monitor, keyboard, mouse to appliance or use
Emulex Configuration Utility	MS Terminal Services
Host error logs	Specific host operating system
Command View EVA GUI pages	Command View EVA
SSSU commands	Enter SSSU at the command prompt window in the directory where it is located or Run directly from CD
HSV controller status indicators	Front and rear of each controller
Disk drive enclosure status indicators	View on rear of disk enclosure EMU
Other LEDs (loop switch, SAN switch, FCA, and appliance)	On each device
LightPulse Utility/NT (LPUTILNT)	Host server console or MS Terminal Services if enabled
FC switch commands	Telnet session to the switch

Event and Error Logs

All of the available event and error logs are critical when determining the source of problems and possible corrective actions on the Enterprise Virtual Array. In general, be mindful of the following when accessing event and error logs:

- As a rule, any error or event on a port should be reviewed.
- When the controllers are running at load, you should not expect the ports to log multiple events.
- While an event log entry each day would not be cause for alarm, multiple entries each day should be investigated and recorded. Keep in mind that the following events create entries in the log:
 - Disk drive insertion and removal (multiple entries)
 - Configuration changes

Event logs are described in the following topics.

Command View EVA Event Logs

You can download three event logs from Command View EVA:

- Management Agent Event Log
- Controller Event Log
- Controller Termination Event Log

In all cases, selecting the **Get log file** or **Get event file** button saves the file for later viewing. In the cases of the Controller Event Log and the Controller Termination Event Log, the file is saved in binary and can be translated using HP Services only tools such as Event View EVA. The Management Agent Event Log is saved in ASCII and does not need translation. You can view all three logs using a standard ASCII text editor such as WordPad or Notepad, or use MS Excel.

Other Event and Error Logs

Two other types of event and error logs are available for troubleshooting: the Management Appliance NT Event Log and host error logs.

The NT Event log is not available through Command View EVA; therefore, you access it either through Terminal Services or by using the RIB with a monitor, mouse, and keyboard.

The host error logs are used for inband error reporting, and errors are logged to certain host types. The following are some properties of the host error log:

- The log is seen through a particular host.
- The log contains errors as generated through the specific operating system for I/Os that have failed.
- You configure this log on the storage management appliance under Notification Options.

Event Decoding

The Controller Event and Controller Termination Event logs contain raw hex data; a Get Log File saves the raw data, but it is not translated.

- To analyze controller and controller termination event codes, forward the files to HP Services to translate the code to human readable form using Services only tools such as Event View EVA.

Monitoring and Reporting Tools

There are several tools that allow monitoring and reporting of events in the Enterprise Virtual Array environment, including:

- System Event Analyzer (SEA)
- Proactive Remote Services (PRS)
 - SEA installed, disable event reporting through Command View EVA
 - SEA not installed, enable event reporting through SNMP traps
- Simple Network Management Protocol (SNMP) traps
- Open Service Event Manager (OSEM)

This information is covered in Module 11.

Command View EVA GUI

By drilling down into the Command View EVA rack, controller, disk enclosure, and other properties, you can determine the status of Enterprise Virtual Array hardware.

Use the GUI to check controller status by viewing the following:

- Host ports
- Device ports
- Controller status
 - Batteries
 - Blowers
 - Temperature

Use the GUI to check disk drive enclosure status by viewing the following:

- Environmental information
- Device bay and disk drive status and information

SSSU Commands

The following are useful SSSU commands for determining the status of the storage system:

- `SHOW MONITOR [FULL]`
- `SHOW DISK [FULL]`
- `SHOW VDISK [FULL]`
- `SHOW LUN [FULL]`
- `SHOW SYSTEM [FULL]`
- `SHOW HOST [FULL]`
- `SHOW WORLD_WIDE_NAME`
- `SHOW POWER`

HSV Controller Status Indicators

The HSV controller status is observed in two ways: the data connector link status LEDs located on the rear of each controller and the OCP front panel.

The following summarizes the status indicators:

- Operator Control Panel (OCP)
 - Four status LEDs
 - LCD, up to 20 alphanumeric characters
 - Control pushbuttons (four-position pushbutton switch), allowing display of:
 - ◆ Storage system menu tree
 - ◆ Firmware versions, such as VCS and Battery PIC
 - ◆ Termination events
- Device port LEDs (rear)
 - One for each device port

Note

All of these sources have been detailed in Module 2.

Disk Drive Enclosure Status Indicators

The disk drive enclosure status is observed in several ways: the EMU display available on the rear of the drive enclosure, and through LEDs viewable on the front and back of the disk drive enclosure.

The following summarizes the status indicators:

- Alphanumeric display
 - EMU-generated condition reports (Er)
 - Three-part error code (displays two digits at a time)
 - May have multiple errors available
- EMU status LEDs
- EMU pushbutton LEDs
- Disk drive LEDs
- Power supply and blower LEDs
- I/O module LEDs

Note

All of these sources have been detailed in Module 3.

Supplemental Drive Enclosure Information

Besides status indicators, the following are some areas of specific interest for troubleshooting drive enclosure errors:

- EMU firmware application code update

EMU firmware contains application code load that is automatically updated through the current superfile. If problems occur with the loading through the superfile, you must use an EMU serial port, special cable, and PC. The complete procedures are in the Service Manual.

Note

The EMU serial cable part number is 17-04875-04.

- EMU-generated condition reports, for example, 0.4.12.02.

These condition reports and resolutions are supplied in Appendix J in the Service Manual and Appendix A of the Student Guide.

Supplemental Drive Enclosure Lists, Tables, and Diagrams

The following are some other supplemental materials in the Service Manual that may aid you when troubleshooting the disk drive enclosure:

- Hard addressing of disk drives in VCS V2.002 or higher
- AL-PA to Loop ID Correlation
- FC loop communication path diagrams (switched or nonswitched)
- Possible EAB failures (several tables)

FC Loop Switch LEDs

Four **System** LEDs indicate the status of the switch, independent of the port LEDs. The following indicates the status of the System LEDs for a correctly functioning switch:

- Power — Green LED should be **ON**.
- Loop operational — Green LED should be **ON**.
- POST fault — Amber LED should be **OFF**.
- Over temp — Amber LED should be **OFF**.

Two **Port** LEDs indicate the status of the port. The following indicates the status of the Port LEDs for a switch:

- Both LEDs **OFF** — Normal status of operation for ports in which SFPs are not installed.
- SFP status **ON** and port bypassed **OFF** — Normal operation. Port and device are fully operational.
- Both LEDs **ON** — The port is nonoperational due to loss of signal, poor signal integrity, or the attached node is sending LIP(F8,xx).
- SFP status **OFF** and port bypassed **ON** — Tx Fault. The port is nonoperational due to an SFP transmitter fault, improperly seated SFP, or another failed device.

FC Switch Commands and Status Indicators

The FC switches contain important details on the FC connections. Remember to always check the port LEDs on the FC switches to verify that the ports are enabled. If the LEDs on the host ports are off, Command View EVA cannot connect.

The following table identifies some useful switch commands for troubleshooting the Enterprise Virtual Array and its components.

Command	Description
cfgShow	Displays the current configuration.
fabricShow	Displays all FC switches (domain) forming one fabric.
nsShow	Displays name server information of FC switch.
nsAllShow	Displays a list of PIDs of nodes still logged in to the fabric.
licenseShow	Displays which optional products are enabled.
loopPortTest	Tests the L-path on a loop.
supportShow	Displays switch information for debugging purposes.
switchShow	Displays all WWIDs that can be seen on the FC switch.
portShow	Displays port status information.
portStatShow	Displays port status counters.
portPerfShow	Displays performance information per port per specified interval.
portLogDump	Displays the port activity log.
version	Displays the switch firmware version.
zoneShow	Shows zone definition.

SAN Switch Port LEDs

The following table describes the Compaq SAN Switch 8 or 16 (1Gb) LED states and descriptions.

Color	Description
Black	No module or cable, power off
Steady Yellow	Receiving light, not online
Slow Yellow	Disabled
Fast Yellow	Faulty port
Steady Green	Online
Slow Green	Segmented
Fast Green	Internal/external loopback
Flickering Green	Frames being transferred
Yellow/Green	Marginal port

The following table describes the Compaq SAN Switch 2/16 (2Gb) LED states and descriptions.

Color	Description
Black	No module or cable, power off
Steady Yellow	Receiving light, not online
Fast Yellow	Faulty port
Steady Green	Online
Slow Green	Segmented
Fast Green	Internal/external loopback
Flickering Green	Frames being transferred

Fibre Channel Adapter LEDs

The following table describes the KGPSA-CA/CB POST LEDs. Note the following:

- The green LED indicates power functions.
- The yellow LED signifies port activity.
- One LED blinks at all times during normal operation.

Green LED	Yellow LED	State
OFF	OFF	Wake-up failure (dead board)
OFF	ON	POST failure (dead board)
OFF	Slow blink (1Hz)	Wake-up failure
OFF	Fast blink (4Hz)	Failure in POST
OFF	Flashing (irregular)	POST processing in progress
ON	OFF	Failure while functioning
ON	ON	Failure while functioning
ON	Slow blink (1Hz)	Normal — Inactive
ON	Flashing (irregular)	Normal — Active
ON	Fast blink (4Hz)	Normal — Busy
Slow blink	OFF	Normal — Link down or not yet started
Slow blink	Slow blink (1Hz)	Offline for download
Slow blink	Fast blink (4Hz)	Restricted offline mode (waiting for restart)

Emulex Configuration Utility

The storage management appliance includes this tool to assist in troubleshooting communication problems between the Enterprise Virtual Array and the appliance, or Command View EVA. Using this utility, you can verify that the appliance is seeing the HSV controller ports. You can also identify which controller ports are SCSI targets of the appliance's FC adapters.



Important

You should use this utility to view states and settings only. You should not change any states, settings, or controls or you could create problems in the network.

You access the utility from the Windows 2000 desktop on the appliance as follows:

1. Make a service connection to the Windows 2000 desktop by attaching a monitor, keyboard, and mouse to appliance or use MS Terminal Services.
2. Click *Start, Programs*.
3. In the programs menu, select *Emulex Configuration Utility*.

A screen similar to the following is displayed.

Available Adapters

- Emulex LP-8000 Adapter, Bus 3 Slot 5 Rev 3.81A1 (Present)
- Emulex LP-8000 Adapter, Bus 3 Slot 6 Rev 3.81A1 (Present)

SCSI Targets

World Wide Port Name	Mapped SCSI ID
50001FE10013A208	0 (Present)
50001FE10013A20C	1 (Present)
50001FE10013A209	2

Adapter Controls

- ☒ Automatically Map SCSI Devices
- ☒ Query name server for all N-Ports
- ☒ Point to Point
- ☒ Allow Multiple paths to SCSI Targets
- ☒ Register For State Change
- ☒ Use Report LUNs
- ☒ Use Name Server after RSCN
- ☒ Lun Mapping
- ☐ Automatic Lun Mapping
- ☒ Scan in Device ID Order
- ☐ Use SLI-1 Mode
- ☐ Report Unknown SCSI Devices
- ☐ Look for Disappearing Devices
- ☐ Translate Queue Full to Busy
- ☐ Use Bus Reset Status for Retries
- ☐ Retry Unit Attention
- ☐ Retry PLOGI Open Failures

Maximum Number of LUNs: 32
Maximum Queue Depth: 8

Static Poll Destination Address

Address:
Add Address
Delete Address

Link Timer: 30 sec Wait Ready Timer: 45 sec
Retries: 64 Retry Timer: 2000 ms
E_D_TOV: 2000 ms R_A_TOV: 2 sec
AL_TOV: 15 ms ARB_TOV: 1000 ms

Buttons: Add Mapping, Modify, Delete Mapping, Lun Map, Performance, Firmware, Memory, Reset Bus, Network, Apply, Exit

LightPulse Utility/NT

The LightPulse Utility/NT (LPUTILNT) is installed on Windows-based host servers in the **system32\drivers** directory as part of the solution kit. This utility displays information that may be useful in troubleshooting communication problems involving the Enterprise Virtual Array and the host. Much of the same information displayed by this utility is also available from the switch, but it may be more convenient to view it from the host server using this utility.

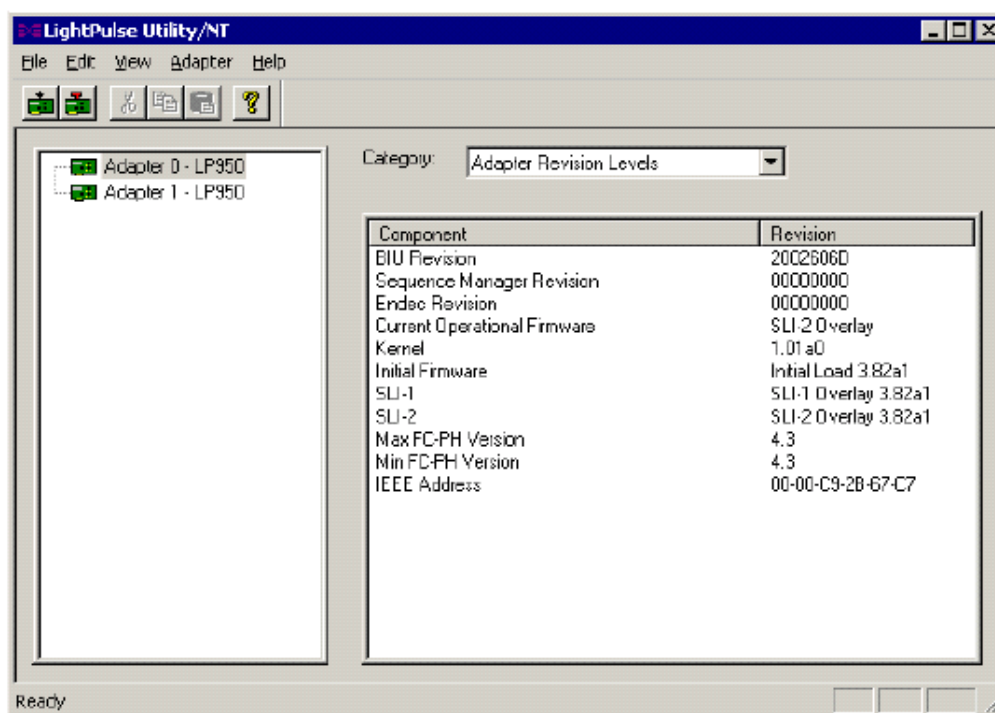


Caution

Use this utility only to view states and settings; do not change any states, settings, or controls or perform any functions using the utility. Doing so could create problems in the network. For example, a known problem of instability in the SAN occurs when LPUTILNT is used to modify the Windows registry.

Execute the LPUTILNT utility from the host server console. Alternatively, if the Server component of Microsoft Terminal Services is enabled on the host, you can connect to the host desktop using this service. You can launch LPUTILNT through a command window or by clicking *Start, Run* and typing `LPUTILNT` in the Open box.

A screen similar to the following is displayed.



To view other information displays, pull down the Category menu and select the desired category of information. The following categories are of interest for troubleshooting:

- Configuration Data
- Link Statistics
- Status and Counters

More details are available in the Service Manual.

Storage Management Appliance LEDs

You should not neglect the OpenView Storage Management Appliance LEDs when trying to diagnose problems that may involve the appliance or Command View EVA.

View the following LEDs:

- G1 front panel
- G2 front panel

Note

For a complete description of these LEDs and supplemental information, including troubleshooting, refer to the Service Manual, Chapter 9, or to the Student Guide, Appendix B.

Troubleshooting References

Troubleshooting references are scattered throughout numerous documents and websites, both HP internal and external. The following topics describe those documents and websites that are most helpful.

Documents

Various documents are scattered throughout Cybrary, HP Technical Bulletin Board (TechBB), and other websites. The most important documents are the following:

- Service Manual
- Release Notes
- Customer and Engineering Advisories
- EVA Repair Techniques document

Websites

Various websites contain important troubleshooting information. The most important websites are the following:

- <http://iq.cca.cpqcorp.net/default.asp> (iQ2000 knowledge base)
- <http://qtip.alf.cpqcorp.net/qtip/default.asp> (internal search engine)
- <http://cybrary.inet.cpqcorp.net/cybrary.html> (Cybrary)
- <http://techport.tay.cpqcorp.net/> (access to TechBB)
- <http://storage.jgo.cpqcorp.net/> (Enterprise Virtual Array Storage Center, The Netherlands)
- <http://milbrn.cxo.cpqcorp.net/> (Storage Performance Engineering home page)
- http://sbuwww.cxo.cpqcorp.net/mvs-hwe/Moongazer/mvs-hwe_emu.asp (M5214 SES)
- <http://hostname:2301/ResEltCpqFusion/fieldservice> (VCS Field Service Utilities)
- http://storage.inet.cpqcorp.net/application/view/menu_products.asp (NSS Product Documentation)
- <http://h18006.www1.hp.com/products/storageworks/enterprise/documentation.html> (Enterprise Virtual Array documentation)
- <http://h18000.www1.hp.com/products/sanworks/managementappliance/> (storage management appliance)

Tips for Analyzing Controller Event Logs

The following are some general suggestions when analyzing Controller Event logs:

- Before starting, try to put the problem within context by determining customer and service actions.
- Ensure that the appliance time is synchronized with the controller time.
- Use timestamps, not sequence numbers, to determine when events were logged. The exception is that if the date Jan XX 1946 is seen in the log (the default base year for storage controllers), then the sequence numbers would be better to identify when the event was logged.
- Do not get lost in any one specific error or event. When reviewing the logs, look at all entries associated with the timeframe or customer issue. Try to understand how the different events occurred to create the big picture. When the big picture is understood, look at the detailed logs to fill gaps.
- Keep in mind that some data is engineering specific (data structure information) and may not contain field usable data.
- The absence of an event or condition will often tell you what's wrong, for example, indications of EMU communication failure resolution on enclosures 8, 10, 11, 12, and 13 might point to a problem with communication on enclosure 9.
- Don't think that just because a specific component is reporting errors, that component is bad. You must look at connections (cables, I/O modules, disk drives, etc.) to and from the suspect component. This is especially true of loop or link errors where the controller or I/O module is suspect.

Replaceable Components

The following table identifies the FRUs and CRUs for the Enterprise Virtual Array. Removal and replacement procedures for some of these components are covered in the Module 13 lab. Complete removal and replacement instructions are contained in the Service Manual.

Component	FRU	CRU
Drive enclosure	X	
▪ Disk drive		X
▪ Drive blank		X
▪ Power supply	X	
▪ Blower	X	
▪ EMU	X	
▪ I/O module	X	
▪ I/O module transceiver	X	
▪ Fiber optic cable	X	
▪ Copper cable	X	
HSV controller	X	
▪ Cache battery assembly	X	
▪ Blower	X	
▪ Fiber optic cable	X	
▪ Copper cable	X	
▪ Transceiver	X	
▪ Power supply/blower*	X	
FC loop switch	X	
▪ Transceiver	X	

* For dual power supply controllers only.

Identifying Spare Part Numbers

All replaceable units have a 6-3 and 2-5-2 part number on the product label beneath the *Compaq Spare* entry (see the following figure). The call handling system identifies the part as a FRU or CRU. The last three characters of the assembly number identify the revision number.

The serial number sticker is at the top rear right of the rack. The number is at the bottom left of the sticker



Learning Check

1. List the three major event logs used to collect information for troubleshooting the Enterprise Virtual Array storage system.

.....

.....

.....

2. Name all of the controller status indicators used for troubleshooting.

.....

.....

3. Name all of the drive enclosure status indicators used for troubleshooting.

.....

.....

4. List five FC switch commands that are useful in troubleshooting the Enterprise storage system.

.....

.....

5. List the CRUs of the Enterprise Virtual Array.

.....

.....

EMU error condition reports

appendix A

Overview

This appendix is to be used in conjunction with Module 3, Disk Drive Enclosure, and Module 13, Troubleshooting. The purpose of this document is to supply a single, condensed source for quickly diagnosing EMU errors. The standard format of an EMU error condition is described, navigation instructions are given, and a description of all error conditions for any element type are detailed, including procedures for resolving the condition.

Note

All removal and replacement procedures are available in the *Enterprise Virtual Array Service Manual*.

Condition Report Format

When the EMU alphanumeric display is **Er**, there are additional displays that identify the element type, the specific element, and the error code, all defining the possible cause of the problem.

- The **first** level display defines the type of element affected with two alphanumeric characters separated by periods, such as 0.1., 0.2., 1.3., F. F., and so forth. A disk drive problem displays 0.1. as its element type number.
- The **second** level display defines the element affected with a two-digit, decimal number, followed by a period. For example, when there is a bay 6 drive error, the element number display is 06.; a display of 14. indicates a bay 14 problem.
- The **third** level display defines a specific problem, the error code, with a two-digit, decimal number. For example, if the problem is either the installation of an incorrectly configured drive or one that cannot operate at the loop link rate, the display is 01.

The full format of an error condition is ***e.t.en.ec***, each level as described above, followed by the severity level. The severity levels, in order by precedence, are:

- UNRECOVERABLE — Most severe condition. Take corrective action immediately.
- CRITICAL — Less severe than UNRECOVERABLE condition. Take prompt corrective action to prevent system degradation.
- NONCRITICAL — Less severe than CRITICAL condition. Take early corrective action to prevent system degradation.
- INFORMATION — A condition exists that does not reduce the capability of an element. Prepare to implement corrective action if necessary.


Navigating the Error Display

An active error changes the alphanumeric display to **Er** and activates the error menu. The pushbutton functions are dedicated to displaying errors. For example, assume the following conditions exist:


- There is a high temperature CRITICAL condition for the disk drive in bay 9.
- There is an EMU NVRAM write failure NONCRITICAL condition.

As soon as one of these conditions is reported, the alphanumeric display is **Er**, the top pushbutton LED is on, and the audible alarm is beeping three times per cycle. Refer to the following procedure to analyze the disk drive condition report.


1. To move from the error display to the element type (temperature sensor) display, press and hold the top pushbutton.



2. To move from the temperature sensor element type display to the element number 12 display (disk drive 9), press and release the top pushbutton.
3. To move from the disk drive 9 element display to the error code (CRITICAL high temperature) display, press and release the top pushbutton.
4. To move from the error code back to the element type, press and release the top pushbutton.
5. Pressing and holding the top pushbutton allows you to move from either the element type display, the element number display, or the error code display to the next most severe element type display, the EMU.



6. Pressing and releasing the bottom pushbutton allows you to move from either the element type display, the element number display, or the error code display directly to the **Er** display.



Element Types

Element type codes assigned to the drive enclosure elements. Elements that do not have an active condition report are shaded.

Element Type	Description/Action
0.1	Disk drives
0.2	Power supplies
0.3	Blowers
0.4	Temperature sensors
0.6	Audible alarm
0.7	EMU
0.C	Controller OCP LCD
0.F	Transceivers
1.0	Language
1.1	Communication port
1.2	Voltage sensors
1.3	Current sensors
8.0	Drive enclosure
8.2	Drive enclosure backplane
F.F	Host


Drive Conditions

The format of a disk drive condition report is **0.1.en.ec**, where:

- **0.1** is the disk drive element type number
- **en** is the two-character disk drive element
- **ec** is the error code

Note

There is a direct correlation between the disk drive element number and the bay number. There is no direct correlation between the disk drive bay number and the device FC-AL physical address. The FC-AL physical address is assigned by negotiation during system initialization.



Disk Drive Bay and Element Numbering

The following table gives a summary of the error conditions:

e.t. en.ec	s	Error Description
0.1.**.01	C	Drive configuration error; unsupported link rate, replace drive.
0.1.**.02	I	Drive missing; install a drive or blank.
0.1.**.03	I	Drive removed when locked; clear software lock before removal.
0.1.**.04	C	Drive link A rate error; spin down/spin up drive.
0.1.**.05	C	Drive link B rate error; spin down/spin up drive.

0.1.**.01 CRITICAL Condition

Description: Drive configuration error; unsupported link rate.

As each drive spins up and comes on-line, the EMU determines if the drive is Fibre Channel compatible and can operate at the link rate (1Gbps or 2Gbps) established by the I/O module. If either of these conditions is not met, the EMU issues a condition report. This condition report remains active until the problem is corrected. The problem affects a unique disk drive; therefore, it is not required to make corrections to prevent the failure of other elements.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the defective drive with:
 - a. A Fibre Channel compatible drive.
 - b. A Fibre Channel drive capable of operating at a link rate supported by I/O modules and transceivers.
3. Observe the EMU to ensure the error is corrected.

0.1.**.02 INFORMATION Condition

Description: Drive missing.

The drive is improperly installed or missing. This could affect the enclosure airflow and cause an over temperature condition for another element.

This error remains active for one minute, or until the problem is corrected, whichever occurs first. Immediate correction is not required. However, correction cannot be delayed indefinitely.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Remove and installing the drive to ensure that it is properly installed.
3. Observe the EMU to ensure the error is corrected.
4. If removing and installing the drive did not correct the problem, install a replacement drive or a drive blank.
5. Observe the EMU to ensure the error is corrected.

0.1.**.03 INFORMATION Condition

Description: Drive software lock active.

Some enclosures have a software-activated lock that prevents physically removing a drive while this feature is active. This feature can be activated even when an enclosure does not have a physical lock. Removing a drive when this feature is active generates a condition report. This error remains active for 15 seconds.

There is no action required to correct this condition.

0.1..04 CRITICAL Condition**

Description: Loop A drive link rate incorrect.

The drive is capable of operating at the loop A link rate, but is running at a different rate. For example, the drive is operating at 1Gbps and the loop is operating at 2Gbps. Only when the drive is operating at the Fibre Channel link rate established by the I/O module can this drive transfer data. This error remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Remove and replace the drive in the enclosure.
3. Observe the drive status LEDs to ensure the drive is operational.
4. Observe the EMU to ensure the error is corrected.
5. If removing and replacing the drive did not correct the problem, replace the drive.
6. Observe the drive status LEDs to ensure the drive is operational.
7. Observe the EMU to ensure the error is corrected.

0.1..05 CRITICAL Condition**


Description: Loop B drive link rate incorrect.

Same as for 0.1.**.04, but loop B instead.

Power Supply Conditions

The format of a power supply condition report is **0.2.en.ec**, where:

- **0.2** is the power supply element type number
- **en** is the two-character power supply element number
- **ec** is the error code



1 — Power Supply 1

2 — Power Supply 2

The following table gives a summary of the error conditions:

e.t. en.ec	s	Error Description
0.2.**.01	N	Power supply no AC input; check AC power to the supply.
0.2.**.02	U	Power supply is missing; install supply! Shutdown imminent!
0.2.**.03	C	Power supply load unbalanced; check blower on power supply.

0.2..01 NONCRITICAL Condition**

Description: Power supply AC input missing.

The loss of the AC input to a power supply makes the remaining supply a single point of failure. This condition report remains active until the AC power is applied to the power supply.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Ensure that there is AC power to the rack PDU, and from the PDU to the PDMs and that the PDU and PDM circuit breakers are not reset. If there is no AC power to the PDU, contact building facilities management. Verify that the power supply AC power cord is properly connected.
3. If AC is present, and the rack power distribution circuitry is functioning properly, the power supply LED should be on.
4. Observe the EMU to ensure the error is corrected.

0.2..02 UNRECOVERABLE Condition**

Description: Power supply missing.

This condition report indicates a power supply is not installed or installed incorrectly. Both of these conditions affect airflow within the enclosure and can cause an over temperature condition. **Enclosure shutdown is imminent!** The operational power supply will automatically shut down after seven minutes, thereby disabling the enclosure. This condition report remains active until either the problem is corrected, or the operational power supply shuts down, whichever occurs first.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Ensure that the power supply is properly installed.
3. If the power supply is properly installed, replace the supply.
4. Verify that the power supply status LED is on.
5. Observe the EMU to ensure the error is corrected.

0.2..03 NONCRITICAL Condition**


Description: Power supply imbalance.

There are no description or repair procedures for this error code at this time. The error indicates that there was a voltage or amperage difference in the current disk enclosure. A predetermined threshold was exceeded.

Blower Conditions

The format of a blower condition report is **0.3.en.ec**, where:

- **0.3** is the blower element type number
- **en** is the two-character blower element number
- **ec** is the error code



1 — Blower 1

2 — Blower 2

The following table gives a summary of the error conditions:

e.t. en.ec	s	Error Description
0.3.**.01	N	Blower speed alert; replace blower soon.
0.3.**.02	C	Blower speed error; replace blower now.
0.3.**.03	U	Blower failure; replace blower now.
0.3.**.04	N	Blower internal status alert
0.3.**.05	N	Blower missing alert; install blower.
0.3.**.06	U	No blowers installed; install blowers! Shutdown imminent!

0.3.**.01 NONCRITICAL Condition

Description: Blower speed.

A blower is operating a speed outside of the EMU-specified range, possibly because of a bearing problem. This can affect enclosure cooling and cause an element to fail. This condition report remains active until the problem is corrected. This error does not normally require immediate correction. However, an error of this type could contribute to an element overheating. HP recommends replacing the blower as soon as possible.

Complete the following procedure to correct this problem.

1. Record all six characters of the condition report.
2. Ensure that the blower is properly installed.
3. If the blower is properly installed, replace the supply.
4. Verify that the power supply status LED is on.
5. Observe the EMU to ensure the error is corrected.

0.3.**.02 CRITICAL Condition

Description: Blower speed.

A blower is operating at a speed that is significantly outside the EMU-specified range, possibly because of a bearing problem. This can cause the loss of cooling and cause an element to fail. The error remains active until the problem is corrected. HP recommends replacing the blower as soon as possible.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the defective blower.
3. Verify that the power supply status LED is on.
4. Observe the EMU to ensure the error is corrected.

0.3.**.03 UNRECOVERABLE Condition

Description: Blower failure.

A blower has stopped. The operational blower now operates at a high speed and is a single point of failure. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the defective blower.

**Caution**

Removing a blower automatically closes flaps over the power supply blower opening. However, the air flow within the enclosure changes and can cause an over temperature condition. Therefore, do not remove a defective blower until a replacement blower is available.

3. Verify that the power supply status LED is on.
4. Observe the EMU to ensure the error is corrected.

0.3..04 NONCRITICAL Condition**

Description: Blower internal.

A power supply reported an internal blower error that could affect enclosure cooling and cause an element to fail. HP recommends correcting the problem before the blower fails. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the defective blower.

**Caution**

Removing a blower automatically closes flaps over the power supply blower opening. However, the air flow within the enclosure changes and can cause an over temperature condition. Therefore, do not remove a defective blower until a replacement blower is available.

3. Verify that the power supply status LED is on.
4. Observe the EMU to ensure the error is corrected.

0.3..05 NONCRITICAL Condition**

Description: Blower missing.

A blower has been removed or is improperly installed. Even though the blower flaps close to maintain the proper airflow, the reduced cooling capability can cause overheating, causing an element to fail. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Ensure that the blower is properly installed.
3. If the blower is properly installed, and not rotating, replace the blower.



Caution

Removing a blower automatically closes flaps over the power supply blower opening. However, the air flow within the enclosure changes and can cause an over temperature condition. HP recommends not removing defective blower until a replacement blower is available.

4. Verify that the power supply status LED is on.
5. Observe the EMU to ensure the error is corrected.

0.3..06 UNRECOVERABLE Condition**

Description: No blowers installed.

When this condition exists, there will be two error messages.

- The first message will be 0.3.en.05 and will identify the first blower.
- The second message will be 0.3.en.06 and will identify the second blower.

The EMU cannot detect any installed blowers. **Shutdown is imminent!** The EMU will shut down the enclosure in seven minutes unless you correct the problem. This condition report remains active until you correct the problem or the EMU shuts down the power supplies, whichever occurs first.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace one blower **now!** Replace the other blower as soon as possible.
3. If no blower is available, proceed to step six.
4. Verify that the power supply status LED is on.
5. Observe the EMU to ensure the error is corrected.
6. If unable to correct the problem it may be necessary to use the controller shutdown procedure to:
 - a. Flush data from the controllers.
 - b. Issue a command to the EMUs to shut down after a 10-30 second delay.
 - c. Shut down the controllers.



Caution

Using the controller shutdown command is a drastic measure that will stop all Enterprise Virtual Array operations. HP recommends using this procedure only when necessary to protect an enclosure from overheating.

Temperature Conditions

The format of a temperature condition report is **0.4.en.ec**, where:

- **0.4** is the temperature sensor element type
- **en** is the two-character temperature sensor element
- **ec** is the error code

Sensor	Sensor Location	Sensor	Sensor Location
01.	Power Supply 1 Exhaust	10.	Drive Bay 10
02.	Power Supply 2 Exhaust	11.	Drive Bay 11
03.	EMU	12.	Drive Bay 12
04.	Drive Bay 1	13.	Drive Bay 13
05.	Drive Bay 2	14.	Drive Bay 14
06.	Drive Bay 3	15.	drive Bay 15
07.	Drive Bay 4	16.	Drive Bay 16
08.	Drive Bay 5	17.	Drive Bay 17
09.	Drive Bay 6		

Temperature Sensor Locations

	Disk Drive	EMU	Power Supply
High Temperature Thresholds			
Limit	60°C (140°F)	43°C (109°F)	60°C (140°F)
CRITICAL			
NONCRITICAL	55°C (131°F)	39°C (102°F)	50°C (122°F)
Low Temperature Thresholds			
Limit	5°C (41°F)		
CRITICAL			
NONCRITICAL	9°C (48°F)		

Temperature Sensor Thresholds

The following table gives a summary of the error conditions:

e.t. en.ec	s	Error Description
0.4.**.01	N	TSensor high temp alert; check for empty bay or air temperature.
0.4.**.02	C	TSensor high temp error; check for empty bay or air temperature.
0.4.**.03	N	TSensor low temp alert; check air temperature.
0.4.**.04	C	TSensor low temp error; check air temperature.
0.4.**.05	U	TSensor high temp error in 2 of 3 groups (EMU, Disk, PS); check air temperature.

0.4..01 NONCRITICAL Condition**

Description: High temperature.

This condition report indicates that an element temperature is approaching, but has not reached, the high temperature CRITICAL threshold. Continued operation under these conditions may result in a CRITICAL condition. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Ensure that all elements are properly installed to maintain proper airflow.
3. Ensure that nothing is obstructing the airflow at either the front of the enclosure or the rear of the blower.
4. Ensure that both blowers are operating properly (the LEDs are on) and neither blower is operating at high speed. Replace a defective blower.
5. Verify that the ambient temperature range is +10°C to+35°C (+50°F to+95°F). Adjust as necessary.
6. Observe the EMU to ensure the error is corrected.

0.4..02 CRITICAL Condition**

Description: High temperature.

This condition report indicates that an element temperature is above the high temperature CRITICAL threshold. Continued operation under these conditions may result in element failure and may reduce an element MTBF. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Ensure that all elements are properly installed to maintain proper airflow.
3. Ensure that nothing is obstructing the airflow at either the front of the enclosure or the rear of the blower.
4. Ensure that both blowers are operating properly (the LEDs are on) and neither blower is operating at high speed. Replace a defective blower.
5. Verify that the ambient temperature range is +10°C to+35°C (+50°F to+95°F). Adjust as necessary.
6. Observe the EMU to ensure the error is corrected.

0.4..03 NONCRITICAL Condition**

Description: Low temperature.

This condition report indicates that an element temperature is approaching, but has not reached, the low temperature CRITICAL threshold. Continued operation under

these conditions may result in a CRITICAL condition. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Verify that the ambient temperature range is +10°C to+35°C (+50°F to+95°F). Adjust as necessary.
3. Observe the EMU to ensure the error is corrected.

0.4..04 CRITICAL Condition**

Description: Low temperature.

This condition report indicates that an element temperature has reached the low temperature CRITICAL threshold. HP recommends correcting this error to prevent affecting other elements. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Verify that the ambient temperature range is +10°C to+35°C (+50°F to+95°F). Adjust as necessary.
3. Observe the EMU to ensure the error is corrected.

0.4..05 UNRECOVERABLE Condition**

Description: High temperature.

This condition reports the EMU has evaluated the temperature of the three temperature groups (EMU, disk drives, power supplies) and determined that the average temperature of two of the three groups exceeds the critical level. Under these conditions, the EMU will start a timer that will automatically shut down the enclosure in seven minutes unless the problem is corrected. **Enclosure shutdown is imminent!**

Complete the following procedure to correct this problem:

1. Ensure that all disk drive, I/O module, and power supply elements are fully seated.
2. Ensure that all blowers are operating properly.
3. Verify that the ambient temperature range is +10°C to+35°C (+50°F to+95°F). Adjust as necessary.
4. If neither steps 1, 2, or 3 revealed a problem, use Command View EVA to request the HSV controller to shut down the drive enclosure. This will halt the drive enclosure data transfers.

EMU Conditions

The format of an EMU condition report is **0.7.01.ec**, where:

- **0.7** is the EMU element type number
- **01** is the two-character EMU element number
- **ec** is the error code



Important

There is only one EMU in a drive enclosure. Therefore, the element number is always 01.

The following table gives a summary of the error conditions:

e.t. en.ec	s	Error Description
0.7.01.01	C	EMU clock error; reset EMU.
0.7.01.02	U	EMU communications interrupted; reset EMU or shelf.
0.7.01.03	U	EMU power supply was shut down.
0.7.01.04	I	EMU internal data error.
0.7.01.05	U	EMU backplane conflict; ESI operations disabled.
0.7.01.10	N	EMU NVRAM read invalid; automatic recovery initiated.
0.7.01.11	N	EMU NVRAM write failed; replace NVRAM.
0.7.01.12	N	EMU NVRAM is unreadable; initialize or replace NVRAM.
0.7.01.13	U	EMU load failure; restart or replace EMU.
0.7.01.14	N	EMU enclosure address incorrect or no address set; check cabinet cables.
0.7.01.15	U	EMU hardware failure; replace EMU.
0.7.01.16	I	EMU internal ESI data corrupt; no action required.
0.7.01.17	I	Power supply shutdown failure.

0.7.01.01 CRITICAL Condition

Description: EMU internal clock.

An internal EMU clock error will remain active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error is corrected.
4. If resetting the EMU did not correct the problem, replace the EMU.

0.7.01.02 UNRECOVERABLE Condition

Description: EMU communication interrupted.

The Inter-IC (I2C) bus is not processing data and the EMU is unable to monitor or report the status of the elements or enclosures. Immediate corrective action is required to ensure proper enclosure operation. This condition report remains active until the problem is corrected.

Complete the following procedure **immediately** to correct this problem:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error is corrected.
4. If resetting the EMU did not correct the problem, replace the EMU.

0.7.01.03 UNRECOVERABLE Condition

Description: Power supply shutdown.

This message only appears on the Command View EVA GUI to report a power supply has already shut down. This message can be the result of the controller shutdown command or an EMU or power supply initiated power shutdown. This message cannot be displayed until after restoration of power.

There is no corrective action required.

0.7.01.04 INFORMATION Condition

Description: EMU internal data.

The EMU is unable to collect data for the SCSI-3 Engineering Services (SES) page. This condition report remains active for 15 seconds. The condition report affects only internal EMU operations. There is no degradation of enclosure operations. The EMU initiates automatic recovery procedures. If the problem is not automatically corrected after one minute, replace the EMU.

0.7.01.05 UNRECOVERABLE Condition

Description: Backplane NVRAM.

Backplane NVRAM errors usually occur during manufacture. At this time, they are identified and corrected. They rarely occur during normal operation. When a backplane NVRAM is not programmed or cannot be read by the EMU, there is no communication with the disk drives. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error is corrected.
4. If resetting the EMU did not correct the problem, replace the EMU.

0.7.01.10 NONCRITICAL Condition

Description: EMU NVRAM invalid read data.

The data read from the EMU NVRAM is invalid. This error initiates an automatic recovery process. This condition report remains active until the problem is corrected. If the automatic recovery process does not correct the problem, complete the following procedure:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error is corrected.
4. If resetting the EMU did not correct the problem, initialize the enclosure by:
 - a. Removing power from the enclosure.
 - b. Applying power to the enclosure.
5. Observe the EMU to ensure the error is corrected.
6. If initializing the enclosure did not correct the problem, replace the EMU.

0.7.01.11 NONCRITICAL Condition

Description: EMU NVRAM write failure.

The EMU cannot write data to the NVRAM. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error is corrected.
4. If resetting the EMU did not correct the problem, initialize the enclosure by:
 - a. Removing power from the enclosure.
 - b. Applying power to the enclosure.
5. Observe the EMU to ensure the error is corrected.
6. If initializing the enclosure did not correct the problem, replace the EMU.
7. Observe the EMU to ensure that the error is corrected.

0.7.01.12 NONCRITICAL Condition

Description: EMU cannot read NVRAM data.

The EMU is unable to read data from the NVRAM. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error is corrected.
4. If resetting the EMU did not correct the problem, initialize the enclosure by:
 - a. Removing power from the enclosure.
 - b. Applying power to the enclosure.
5. Observe the EMU to ensure the error is corrected.
6. If initializing the enclosure did not correct the problem, replace the EMU.

0.7.01.13 UNRECOVERABLE Condition

Description: EMU load failure.

The EMU Field Programmable Gate Array (FPGA) that controls the ESI bus failed to load information required for EMU operation. This condition report remains active until the problem is corrected.

Complete the following procedure **immediately** to correct this problem:

1. Record all six characters of the condition report.
2. Reset the EMU.
3. Observe the EMU to ensure the error is corrected.
4. If resetting the EMU did not correct the problem, initialize the enclosure by:
 - a. Removing power from the enclosure.
 - b. Applying power to the enclosure.
5. Observe the EMU to ensure the error is corrected.
6. If initializing the enclosure did not correct the problem, replace the EMU.

0.7.01.14 NONCRITICAL Condition

Description: EMU enclosure address.

Either the enclosure address is incorrect or the enclosure has no address. Possible causes include a defective enclosure address bus cable, an incorrectly connected cable, or a defective enclosure address bus JB. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Remove and reconnect the cable between the address bus Junction Box (JB) and the EMU.

Note

The EMU display may not display a change in condition for up to 30 seconds.

3. Observe the EMU to ensure the error is corrected.
4. If the problem is not corrected, remove and reinstall the bottom and top terminators, and all the JB-to-JB cables.
5. Observe the EMU to ensure the error is corrected.
6. Reset the EMU.
7. Observe the EMU to ensure the error is corrected.
8. If resetting the EMU did not correct the problem, replace the EMU.

0.7.01.15 UNRECOVERABLE Condition

Description: EMU hardware failure.

The EMU is inoperative and must be replaced **immediately**. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the EMU.
3. Observe the EMU to ensure the error is corrected.

0.7.01.16 INFORMATION Condition

Description: EMU internal ESI data corrupted.

The EMU ESI data is corrupted. This does not affect any other element and there is no action required.

0.7.01.17 UNRECOVERABLE Condition

Description: Power supply shutdown failure.

The power supply did not respond to a controller, EMU, or power supply shut down command. Shutting down the supply is required to prevent overheating.


Complete the following procedure to correct the problem:

1. Record all six characters of the condition report.
2. Move the power cord bail lock to the left.
3. Disconnect the AC power cord 2 from the supply.

Transceiver Conditions

The format of a transceiver condition report is **0.F.en.ec**, where:

- **0.F** is the transceiver element type number
- **en** is the two-character transceiver number
- **ec** is the error code



- 1 — Transceiver 1
- 2 — Transceiver 2
- 3 — Transceiver 3
- 4 — Transceiver 4

The following table gives a summary of the error conditions:

e.t. en.ec	s	Error Description
0.F.**.01	C	I/O transceiver incompatibility; replace transceiver.
0.F.**.02	C	I/O transceiver data signal lost; check cabling.
0.F.**.03	C	I/O transceiver FC-AL bus fault; check transceiver or cabling.

0.F..01 CRITICAL Condition**

Description: Transceiver incompatibility.

The transceivers on this link are not the same type or they are incompatible with the I/O module. This error prevents the controller from establishing a link with the enclosure drives and eliminates the enclosure dual-loop capability. This error remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Check all the transceivers on the loop to ensure they are I/O module compatible.
3. Replace any incompatible transceivers.
4. Observe the EMU to ensure the error is corrected.

0.F..02 CRITICAL Condition**

Description: Transceiver data signal lost.

The transceiver can no longer detect a data signal. This error prevents the controller from transferring data on a loop and eliminates the enclosure dual-loop capability. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the transceiver cable.
3. Observe the EMU to ensure the error is corrected.

0.F..03 CRITICAL Condition**

Description: Transceiver FC-AL bus fault.

The system has detected an FC-AL bus fault involving a transceiver. This error prevents the controller from transferring data on a loop and eliminates the enclosure dual-loop capability.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Check all the transceivers and cables to ensure they are properly connected.
3. Check all the transceivers on the loop to ensure they are I/O module compatible.
4. Replace any incompatible transceivers.
5. Observe the EMU to ensure the error is corrected.
6. Replace the input cable.
7. If replacing the cable did not correct the problem, replace both transceivers attached to the cable.

Voltage Sensor and Current Sensor Conditions

The format of a voltage sensor condition report is **1.2.en.ec** or **1.3.en.ec**, where:

- **1.2** is the voltage sensor element type number
- **1.3** is the current sensor element type number
- **en** is the two-character sensor element number
- **ec** is the error code

Sensor	Sensor Element Location
01.	Power Supply 1 +5 VDC
02.	Power Supply 1 +12 VDC
03.	Power Supply 2 +5 VDC
04.	Power Supply 2 +12 VDC

Voltage and Current Sensor Locations

	+5 VDC	+12 VDC
High Voltage Thresholds		
Nominal Voltage	+5.40 VDC	+12.50 VDC
Limit	+5.73 VDC	+13.25 VDC
CRITICAL		
NONCRITICAL	+5.67 VDC	+13.13 VDC
Low Voltage Thresholds		
Nominal Voltage	+5.40 VDC	+12.50 VDC
Limit	+5.08 VDC	+11.75 VDC
CRITICAL		
NONCRITICAL	+5.13 VDC	+11.88 VDC

Voltage and Current Sensor Thresholds

The following table gives a summary of the error conditions:

e.t. en.ec	s	Error Description
1.2.**.01	N	Voltage sensor high-voltage alert; replace supply soon.
1.2.**.02	C	Voltage sensor high-voltage error; replace supply now.
1.2.**.03	N	Voltage sensor low-voltage alert; replace supply soon.
1.2.**.04	C	Voltage sensor low-voltage error; replace supply now.
1.3.**.01	N	Current sensor high-current alert; replace supply soon.
1.3.**.02	C	Current sensor high-current error; replace supply now.

1.2..01 NONCRITICAL Condition**

Description: High voltage.

This condition report indicates that an element voltage is approaching, but has not reached, the high voltage CRITICAL threshold. Continued operation under these conditions may result in a CRITICAL condition. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the supply.
3. Verify that the power supply status LED is on.
4. Observe the EMU to ensure the error is corrected.

1.2..02 CRITICAL Condition**

Description: High voltage.

This condition report indicates that an element voltage has reached the high voltage CRITICAL threshold. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the supply.
3. Verify that the power supply status LED is on.
4. Observe the EMU to ensure the error is corrected.

1.2..03 NONCRITICAL Condition**

Description: Low voltage.

This condition report indicates that an element voltage is approaching, but has not reached, the low voltage CRITICAL threshold. Continued operation under these conditions may result in a CRITICAL condition. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the supply.
3. Verify that the power supply status LED is on.
4. Observe the EMU to ensure the error is corrected.

1.2.**.04 CRITICAL Condition

Description: Low voltage.

This condition report indicates that an element voltage has reached the high voltage CRITICAL threshold. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the supply.
3. Verify that the power supply status LED is on.
4. Observe the EMU to ensure the error is corrected.

1.3.**.01 NONCRITICAL Condition

Description: High current.

This condition report indicates that an element current is approaching, but has not reached, the high current CRITICAL threshold. Continued operation under these conditions may result in a CRITICAL condition. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the supply.
3. Verify that the power supply status LED is on.
4. Observe the EMU to ensure the error is corrected.

1.3.**.02 CRITICAL Condition

Description: High current.

This condition report indicates that an element current has reached the high current CRITICAL threshold. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the supply.
3. Verify that the power supply status LED is on.
4. Observe the EMU to ensure the error is corrected.

Backplane NVRAM

Backplane NVRAM errors usually occur during manufacturing. At this time, they are identified and corrected. They rarely occur during normal operation. The only corrective action available for these type errors is to replace the drive enclosure.

The format of a backplane NVRAM condition report is **8.2.01.ec**, where:

- **8.2** is the backplane NVRAM element type number
- **01** is the two-character element number
- **ec** is the error code

The following table gives a summary of the error conditions:

e.t. en.ec	s	Error Description
8.2.01.10	N	Backplane NVRAM read invalid; automatic recovery initiated.
8.2.01.11	N	Backplane NVRAM write failure; replace NVRAM.
8.2.01.12	N	Backplane NVRAM read failure; initialize or replace NVRAM.
8.2.01.13	N	Backplane WWN is blank. Initialize NVRAM.

8.2.01.10 NONCRITICAL Condition

Description: Backplane NVRAM invalid read.

An invalid NVRAM read occurred and an automatic recovery process has begun. This condition report is active for 15 seconds. If the automatic recovery process does not correct the problem, replace the enclosure.

8.2.01.11 NONCRITICAL Condition

Description: Backplane NVRAM write failure.

The system is unable to write data to the NVRAM. This problem prevents communication between elements in the enclosure. This condition report is active for 15 seconds.

8.2.01.12 NONCRITICAL Condition

Description: Backplane NVRAM read failure.

The system is unable to read data from the NVRAM. This problem prevents communication between elements in the enclosure. This condition report is active for 15 seconds.

8.2.01.13 NONCRITICAL Condition

Description: Backplane WWN is blank; initialize NVRAM.

The WorldWide ID (WWID) in the drive enclosure backplane is zero. The EMU firmware embedded in the VCS V2.x Superfile flags an event, sounds an alarm, and reports an error via the LCD. The system operates correctly with this condition, but the errors and alarm are undesirable. To correct the error, you must replace the drive enclosure.


Note

This rare error may occur with VCS V2.x only.

I/O Module Conditions

The format of an I/O module condition report is **8.7.en.ec**, where:

- **8.7** is the I/O module element type number
- **en** is the two-character I/O module element number
- **ec** is the error code



1 — I/O Module A (01)

2 — I/O Module B (02)

The following table gives a summary of the error conditions:

e.t. en.ec	s	Error Description
8.7.**.01	C	I/O module unsupported warning; replace I/O module.
8.7.**.02	C	I/O module communication error; reseal or replace I/O module.
8.7.**.10	N	I/O module NVRAM read invalid; automatic recovery initiated.
8.7.**.11	N	I/O module NVRAM write failed; replace NVRAM.
8.7.**.12	N	I/O module NVRAM is unreadable; initialize or replace NVRAM.

8.7.**.01 CRITICAL Condition

Description: I/O module unsupported.

The I/O module Fibre Channel link speed is not supported by the backplane. This error prevents the controller from establishing a link with enclosure drives and eliminates the enclosure dual-loop capability. This condition report remains active until the problem is corrected.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Install a backplane Fibre Channel link speed compatible I/O module.
3. Observe the EMU to ensure the error is corrected.

8.7.**.02 CRITICAL Condition

Description: I/O module communication.

The I/O module is unable to communicate with the EMU.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Remove and re-insert the I/O module.
3. Observe the EMU to ensure the error is corrected.



Important

Multiple erroneous error messages indicating I2C bus errors, such as NVRAM errors, blowers missing, and so forth, could indicate an EMU problem. Replace the EMU when there are multiple erroneous error messages.

8.7.**.10 NONCRITICAL Condition

Description: I/O module NVRAM invalid read.

An invalid NVRAM read occurred and automatic recover was initiated.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Observe the I/O module status LEDs for an operational display. Replace the I/O module when there is an error status.
3. Observe the EMU to ensure the error is corrected.

8.7..11 NONCRITICAL Condition**

Description: I/O module NVRAM write failure.

The system is unable to write data to the I/O module NVRAM.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Observe the I/O module status LEDs for an operational display. Replace the I/O module when there is an error status.
3. Observe the EMU to ensure the error is corrected.

8.7..12 NONCRITICAL Condition**

Description: I/O module NVRAM read failure.

The system is unable to read data from the I/O module NVRAM.

Complete the following procedure to correct this problem:

1. Record all six characters of the condition report.
2. Replace the I/O module.
3. Observe the EMU to ensure the error is corrected.

Host Interconnect

The EMU can display host controller defined condition reports on the EMU alphanumeric display.

The format of a host interconnect condition report is **F.F.en.ec**, where:

- **F.F** is the host interconnect element type number
- **en** is the two-character element number
- **ec** is the error code

The following table gives a summary of the error conditions:

e.t. en.ec	s	Error Description
F.F.**.01	I	Host-generated message

F.F.**.01 INFORMATION Condition

Description: Host generated.

The host controller (HSV) defines the error type and the affected elements. These messages may be displays on the Command View EVA GUI. The proper corrective action is element and error dependent.

troubleshooting the storage management appliance

appendix B

Overview


This appendix is to be used in conjunction with Module 13, Troubleshooting. This appendix assists you in performing the following:

- Identify the initial steps to take when troubleshooting the Storage Management Appliance.
- Use standard troubleshooting techniques to identify a problem with the Storage Management Appliance and determine its solution.
- Identify key troubleshooting areas for applications that reside on the Storage Management Appliance.

Note

Much of the content of this appendix is taken from the *Management Appliance Configuration Guide* and the *Open SAN Manager User Guide*. Refer to both of those documents for more detail.

Operational Status



Rear Panel on the Storage Management Appliance

Number	Connector or LED	Number	Connector or LED
1	External SCSI connector	9	Hot-plug redundant power supply
2	Video connector — Not used	10	Hot-plug power supply
3	RJ-45 Fast Ethernet connector	11	External SCSI connectors
4	Parallel connector	12	Fibre Channel adapter (only one supported)
5	Serial connector B	13	Second Fibre Channel adapter
6	Serial connector A	14	Modem
7	Keyboard connector	15	Service Access Board
8	Mouse connector		


The basic troubleshooting approach consists of three distinct procedures:

1. Devise a troubleshooting plan.
2. Gather and record facts before troubleshooting.
3. Follow an orderly procedure for diagnosis.

The initial step should be to check the status LEDs on the appliance.

System Status LED Indicators


The following graphic and table illustrate the system status LEDs and describe their indications.



Indicator	Status	Meaning	Action
❶ Power ON/ standby status	Green (on)	AC power is connected to the system, and the power supply is functioning properly.	If the green LED is on, no action is required.
	Amber (standby)	System is in standby mode.	If the smber LED remains on, but should not be, it may be a power supply problem. Contact HP Technical Support.
	Off	No AC power is provided to the system.	Determine whether the power should be off. If it should not be: <ul style="list-style-type: none"> Check the power cord to ensure it is plugged into a functional, grounded AC outlet. The power supply may not be operational. Contact HP Technical Support.
❷ Integrated NIC link status	Green (on)	NIC is connected to the network	If the green LED is on, no action is required.
	Off	NIC is not connected.	Determine whether the green LED should be off. If it should not be: <ul style="list-style-type: none"> Ensure the cable is connected properly. Contact your system administrator to determine whether there is a network problem.
❸ Integrated NIC link activity	Green (on)	NIC is communicating with the network	If the green LED is on, no action is required.
	Off	NIC is not communicating with the network at this time.	Determine whether the green LED should be off. If it should not be: <ul style="list-style-type: none"> Contact your system administrator to determine whether there is a network problem. It could indicate a problem in the controller. Contact HP Technical Support.

Hard Drive LED Indicators


The following graphic and table illustrate the hard drive LEDs and describe their indications.



Indicator	Status	Meaning	Action
① Activity	Flashing	Drive being accessed.	None required.
	Off	Drive on standby or off.	Determine whether the green LED should be off: <ul style="list-style-type: none"> ■ If it should not be, reseal the hard drives by unplugging the drives into the drive carriage. ■ If the problem persists, contact HP Technical Support.
② Power/online	On	Indicates one of the following: <ul style="list-style-type: none"> ■ The hard drive is online. ■ Power to the hard drive. ■ The drive is a member of the RAID set. 	None required.
	Flashing	The hard drive is being rebuilt or a RAID set is being created.	If the LED continues to flash for more than three hours, contact HP Technical Support.
	Off	The hard drive is offline.	If there is power to the appliance and the hard drive light is off, do one of the following: <ul style="list-style-type: none"> ■ Reseat the hard drives into the drive carriage. Check for bent pins before plugging in the drive. ■ If the problem persists, contact HP Technical Support.
③ Fault	On	A problem with the hard drive.	Contact HP Technical Support.
	Flashing	If flashing along with the power/online LED, the drive is part of a RAID set being created.	If the light continues to flash for more than three hours, contact HP Technical Support.
	Off	The hard drive is functioning normally.	None required.

RJ-45 Connector Network Status LED Indicators


The following graphic and table illustrate the RJ-45 connector network status LEDs and describe their indications.



Indicator	Status	Meaning	Action
① Activity	Off	No network activity	<p>If in the off status for more than 30 minutes, do one of the following as required:</p> <ul style="list-style-type: none"> Check to ensure no external cables are loose, frayed, or disconnected. Refer to the <i>Management Appliance Configuration Guide</i> to determine whether your connection is correct. Contact your system administrator. If the problem persists, contact HP Technical Support.
	On or irregular	Network activity	None required
② Link	Off	No network link	<p>Determine whether this is an appropriate state. If not, do one of the following as required:</p> <ul style="list-style-type: none"> Refer to the <i>Management Appliance Configuration Guide</i> to determine whether your configuration is correct. Contact your system administrator. If the problem persists, contact HP Technical Support.
	On	Linked to network	None required

Fibre Channel Adapter Indicators

The following graphic and table illustrate the Fibre Channel Adapter LEDs and describe their indications.



Indicator	Yellow LED	Green LED	State
❶ Activity	Off		Normal
	Flashing		POST processing in progress.
❷ Transmit		Off	POST processing in progress.
		On	Normal. Fibre network activity.
❸ Receive		Flashing	Normal. No Fibre activity.
		Off	Processing in progress.
		On	Normal. Fibre network activity.
		Flashing	Normal. No Fibre network activity.

Troubleshooting Open SAN Manager

The following table describes corrective actions for troubleshooting OSM problems.

Problem	Action
Problems browsing.	Verify the following: <ul style="list-style-type: none"> ■ The Internet browser is supported. ■ The browser is configured properly. ■ The URL and port are correct. ■ If all of the above are configured correctly, contact HP Technical Support.
Unable to display Storage Management Appliance splash page.	Browse to the appliance by entering the following URL: http://HOSTNAME:2301. <ul style="list-style-type: none"> ■ Refer to the <i>Management Appliance Configuration Guide</i> to ensure your appliance is configured properly. ■ Reboot the appliance. ■ If the problem persists, contact HP Technical Support.
Open SAN Manager does not launch.	Reboot the appliance. If the problem persists, contact HP Technical Support.
The following appliance folders do not display: <ul style="list-style-type: none"> ■ Configuration Services ■ Installation Services ■ Network Utilities ■ Resource Managers ■ HGS Element Manager ■ Command View EVA ■ CSS2105 Element Manager ■ Volume Management 	<ul style="list-style-type: none"> ■ Browse the appliance ■ Check the status. ■ Log in as the administrator. ■ If the problem persists, contact HP Technical Support.

Troubleshooting OSM Login

If OSM is not operational, check the login procedure by completing the following steps:

1. Review the login procedures in the *Open SAN Manager User Guide*.
2. Check your login user name and password.
3. If the user name and password are correct, but you are still not able to log in, refresh the browser by clicking the browser *Refresh* button, or re-launch your browser.

If these steps do not resolve the problem, refer to the Advanced Help section below. If the Advanced Help procedures do not resolve the problem, contact HP Technical Support.

Advanced Help

This section describes additional troubleshooting procedures to resolve a problem in launching OSM on the Storage Management Appliance. Use these procedures if you continue to enter your user name and password and OSM still does not launch. If these procedures do not resolve the problem, contact HP Technical Support.

Deleting Cookies from Your Browser

Delete the temporary Internet files and the cookie received from the Storage Management Appliance for your browser. If you cannot log in to OSM with user name *administrator* and password *administrator*, refer to your browser documentation for how to delete cookies.

Deleting Files from the Web-Based Management (WBEM) Directory

Use the following procedures to delete files from the WBEM directory:



Important

Deleting files from the WBEM directory returns you to the default login.

1. Attach a terminal to the appliance (keyboard, mouse, monitor).



Important

Ensure you connect the monitor to the Service Access Board (or Remote Insight Board) video connector. Do not connect the monitor to the onboard video connection.

2. From the Microsoft Windows Explorer window, go to the C:\Compaq\WBEM directory.
3. From the WBEM directory:
 - Delete the cache directory.
 - Delete files with the extension .GFG, .DAT, and .ACL.
4. Reboot the appliance. The default login account should now be as follows:

User Name: administrator

Password: administrator
5. Disconnect the keyboard, mouse, and monitor.

Problems Connecting to Storage Management Appliance

This section describes actions to take to resolve problems you may encounter when trying to connect to the Storage Management Appliance.

If your Internet browser does not launch, perform the following steps:

1. Verify the appliance host name by pinging the appliance.
 - Ping host name SWMAxxxxxx.
 - Ping the IP address.
 - Contact your system administrator.
 - Contact HP Technical Services.
2. Verify that your appliance is operational. See the *Management Appliance Configuration Guide*.
3. Ensure that the browser is configured properly. See the *Management Appliance Configuration Guide* to determine how to configure the browser.
4. Verify that you have a supported browser.
5. Ensure that the appliance URL and port are correct. See the *Management Appliance Configuration Guide* to determine the URL and port settings.

OSM Navigation Pane

This section describes the actions to take if the navigation tree does not display in the Navigation pane after you log in to OSM. It also describes the actions to take if the Accounting/Diagnostic Links do not display in the OSM Navigation pane.

Navigation Tree

If the Navigation tree does not display in the Navigation pane after positioning the cursor in the Navigation pane, relaunch your browser. If the problem persists, contact HP Technical Support.

Accounting/Diagnostic Links

If the Accounting/Diagnostic Links do not display in the OSM Navigation pane, perform the following steps:

1. From OSM, launch *Appliance Manager*.
2. In the Appliance Manager Navigation tree, click *Services*.
3. In the Services table, check the application state.
4. If the application state is red, click *Start*.
5. The “Operation Successful!” message displays.
6. Click *OK*.
7. If the problem persists, call HP Technical Support.

Troubleshooting Appliance Manager

This section describes information on troubleshooting Appliance Manager.

Returned to Appliance Manager Home Page

If you click an Appliance Manager selection (for example, Change Password) and the Appliance Manager home page is displayed, proceed with the Appliance Manager selection.

Renaming the Storage Management Appliance

Change the Storage Management Appliance name using only the Appliance Manager name change utility. If you change the name by any other method, the name change compromises the operation of the Storage Management Appliance and applications.



Caution

If you change the appliance name using any other method than the Appliance Manager System Name option, the name change can compromise the operation of the appliance and resident applications.

Troubleshooting Product Installation

Various storage management software applications run on the appliance. The purpose of this section is to provide procedures for troubleshooting the installation of these storage management applications.

Installing or Updating Applications

This section provides general troubleshooting information for installing or updating the storage management applications that run on the appliance. You can install or update storage management applications using a CD-ROM or the network. Both methods are discussed in this section.

After installing an application, the link displays in the Navigation pane. If the link does not display, do one of the following:

1. Close and re-launch your browser.
2. If the link still does not display, launch Appliance Manager. Go to the Services screen and verify the following:
 - If the application box is red, then the application is down. Restart the application.
 - If you do not see the application, reinstall the application.
 - If the application box is green, the application is up.
3. If the link still does not display, reinstall the application. Refer to the online help for how to reinstall an application.
4. Contact HP Technical Support.

CD Install



This section describes errors you may encounter when installing from a CD-ROM and the actions you should take to resolve the error.

The following error message may display when you are trying to install or update storage management applications from a CD-ROM.

`"Error! Check your appliance CD-ROM drive to ensure your CD is seated correctly."`

Follow these steps to resolve the problem:

1. Re-insert the CD into the CD-ROM drive. Wait for the CD-ROM to finish spinning.
2. Click *Install Products* under *Installation Service*.
3. Click one of the options and proceed.

Network Install

The following table describes error messages you may encounter when installing from the network, including actions that you should take to resolve the error. If an application does not display in the Navigation pane, reinstall the application. Refer to the online help for how to reinstall an application. If the problem persists, contact HP Technical Support.

Error Message	Action
"Cannot connect to the FTP site: <ftp site>" where <ftp site> is the name of the FTP site you are trying to connect to.	<ul style="list-style-type: none"> ■ Verify that the FTP server entered is correct and operational. ■ Verify that the username and password entered are correct. ■ Verify that the appliance can access the network where the FTP site resides. <p>If all of the above are correct, and the problem persists, contact your system administrator.</p>
"Error when downloading the file <file path>" where <file path> is the path where the file being downloaded is located.	<ul style="list-style-type: none"> ■ Verify that the <file path> entered is correct. If incorrect, enter the correct file path. ■ Verify that the <file path> entered is the absolute full file path. If incorrect, enter the full file path. ■ Verify that the filename is included in the <file path> entered. If the filename is missing, enter the filename.
"Error in extracting the install kit <file name>" where <file name> is the name of the install kit file.	Verify that the file downloaded is a SANworks install kit.
"Cannot find SANworks product in this file <file path>" where <file path> is the file path to the SANworks product.	Verify that the file path entered in the <file path> is a SANworks install kit.

Application Status

To determine installed applications status, check the *Services* list in Appliance Manager. To access the *Services* List, perform the following steps:

1. Click Services in the Appliance Manager Navigation pane.
2. Check for the State of the service. If an application is listed as “Running,” you can restart it. If an application is listed as “Stopped,” you can start it.



Important

Restarting an application may take from one to five minutes. After you restart an application, open the Services page to confirm that the application is running.

Troubleshooting Resource Managers

This section describes information on troubleshooting Resource Managers including:

- Element Manager for HSG
- Command View EVA for the Enterprise Virtual Array
- SANworks Element Manager for CSS2105

Element Manager for HSG

Element Manager for HSG displays information for the HSG controllers. This section describes troubleshooting information for the HSG controllers.

Note

For more information on troubleshooting Element Manager for HSG, see the online help, the Management Appliance Release Notes, or the *Management Appliance Configuration Guide*.

Subsystem Not Displaying or Communication Status Failed

The subsystem does not display in the HSG Management System options table (under the Options screen) after the initial SAN setup and discovery launch, or the communication state has failed. Perform the following steps to resolve the problem:

1. Refer to the *Management Appliance Configuration Guide* to ensure that your controller is properly configured for the appliance.
2. Ensure that CCL is enabled or that at least one LUN is presented to the appliance. If using a LUN, ensure that it is presented to the appliance through each controller's host port.
3. Ensure that the appliance has at least one online connection. Check the controller connections table by entering the CLI `SHOW CONNECTIONS` command. This step shows online or offline status for each controller's host port connection.
4. If no error is indicated in the set-up or connections, disconnect the Fibre Channel cable from the appliance. This step isolates the appliance from the switch.
5. Reconnect the Fibre Channel cable to the appliance. This step re-establishes connection between the appliance host bus adapters and storage controllers.
6. If the problem persists, contact HP Technical Support.

Note

Rebooting the appliance is not recommended to resolve the subsystem communication state.

Subsystem Data

If any of the subsystem data does not display current data for a specific subsystem, check the Communication State.

If the Communication State is *Failed*, refer to the Subsystem not Displaying or Communication Status Failed section. If the Communication State is *Good*, refresh the subsystem data by doing the following:

1. From the Element Manager for HSG home page, select the subsystem under HSG Network in the Navigation pane.
2. Press the *Refresh* button at the top of the Content pane.

Unable to Communicate with Subsystem

The following states indicate a communication failure with the subsystem:

- CLI commands fail with communication errors.
- Refresh time-outs or unable to refresh subsystem.
- Time-outs when attempting to modify subsystem configuration (for instance, by modifying or creating a virtual disk).

To resolve this problem, check the *Communication State* line. Access the subsystem properties page from the Element Manager for HSG home page. Select the subsystem under HSG Network in the Navigation pane.

If the Communications State is *Failed*, refer to the Communication State Failed section. If the Communication State is *Good*, refresh the subsystem data by doing the following:

1. Attempt the operation again. Occasional time-outs are normal when operations are attempted while Element Manager for HSG operations are in progress.

Note

Ensure that you are running ACS firmware version 8.5 or higher.

2. If CCL is disabled, ensure that a communication LUN is presented to the appliance through each controller's host port (there are two host ports per controller). Refer to the *Management Appliance Configuration Guide* for more information on CCL and LUN.
3. If a Run Command session is open, wait for it to be closed. If a Run Command session was closed incorrectly (without clicking the *Close* button), press the *Unlock* button to recover. The *Unlock* button is located at the top of the Content pane in the CLI window in Element Manager for HSG.



Important

Pressing the *Unlock* button invalidates any Run session in progress. Incorrect use of the *Unlock* button can result in CLI returning undefined data.

4. If a Command Scripter script was running that used the `LOCK` command, wait for the script to issue an `UNLOCK`. If the script terminated without issuing an `UNLOCK`, press the *Unlock* button to recover. The *Unlock* button is located at the top of the Content pane in the CLI window in Element Manager for HSG.



Important

Pressing the *Unlock* button invalidates any lock held by a Command Scripter script. Incorrect usage of the *Unlock* button can result in CLI returning undefined data.

Troubleshooting Open SAN Manager Applications

This section describes information on troubleshooting Open SAN Manager applications including:

- Storage Allocation Reporter
- Network View
- Enterprise Volume Manager

Storage Allocation Reporter

This section describes information on troubleshooting Storage Allocation Reporter.

LUNs Not Found

If LUNs are not found during a scan, check the following:

1. Ensure that CCL is enabled or that at least one LUN is presented to the appliance. If using a LUN, ensure that it is presented to the appliance through each controller's host port.
2. Ensure that the appliance has at least one online connection. Check the controller connections table by entering the CLI `SHOW CONNECTIONS` command. This step shows online or offline status for each controller's host port connection.

Data Is Old, Incorrect, or Missing

If the data in a report is out of date, incorrect, or missing, check the following:

1. If you made changes to your storage resources, ensure that a scan occurred *after* the changes were made. Changes to Storage Allocation Reporter resources do not display until they are scanned.
2. Consult the *Alert Window* for messages related to Storage Allocation Reporter. Take corrective action based on the messages.
3. Ensure that reports have completed loading.

If you select a Single Customer Usage, Single Storage Subsystem, or Single Customer Billing report from a report in which the data has not completed loading, the single-component report might be incomplete. You must wait for the original report to stop loading before selecting a single-component report.

Overview

This appendix presents several key concepts on the subject of capacity allocation on the Enterprise Virtual Array, and supplements materials in Module 7, Concepts and Terminology, and Module 12, Best Practices. You will see how formatted capacity differs from the marketed capacity of hard drives. Sample screens show how formatted capacity and disk group metadata come into play. Next, you will see the affects of various factors when creating virtual disks. The topic of disk failure protection is explored and several screens are used to show what happens to occupancy levels when a disk failure occurs.

Formatted Capacity of Disk Drives

The Enterprise Virtual Array VCS V2 supports 36GB, 72GB, and 146GB disk drives. Remember that these numerical designations are rounded off as follows:

- 36.4GB to 36GB
- 72.8GB to 73GB (but referred to as 72GB)
- 145.6 to 146GB

These formatted drive capacities are decimal (base 10) numbers as prescribed by typical hard drive marketing. Operating systems, however, use binary to represent formatted hard drive capacities. The Enterprise Virtual Array depicts the operating system's view; however, keep in mind that the full capacities are available.

Formatted Capacities Using Binary Equivalents

As shown below, the Enterprise Virtual Array shows the operating system's view of capacity. The decimal equivalent of 36.4GB is 36,400,000,000 bytes. To get the equivalent binary value that the operating system uses, divide by a binary GB, that is, by 1,073,741,824. Therefore, a 36GB disk drive displays a *Formatted capacity* of 33.91GB as shown below.

The screenshot shows the HSV Element Manager web interface in a Microsoft Internet Explorer browser. The address bar shows the URL: <http://smad240fk34d089:2301/ResEtkCpqFusion>. The page title is "SMAD240FK34D089: HSV Storage System - Microsoft Internet Explorer". The main content area displays the "Disk Drive Properties" for "Disk 001".

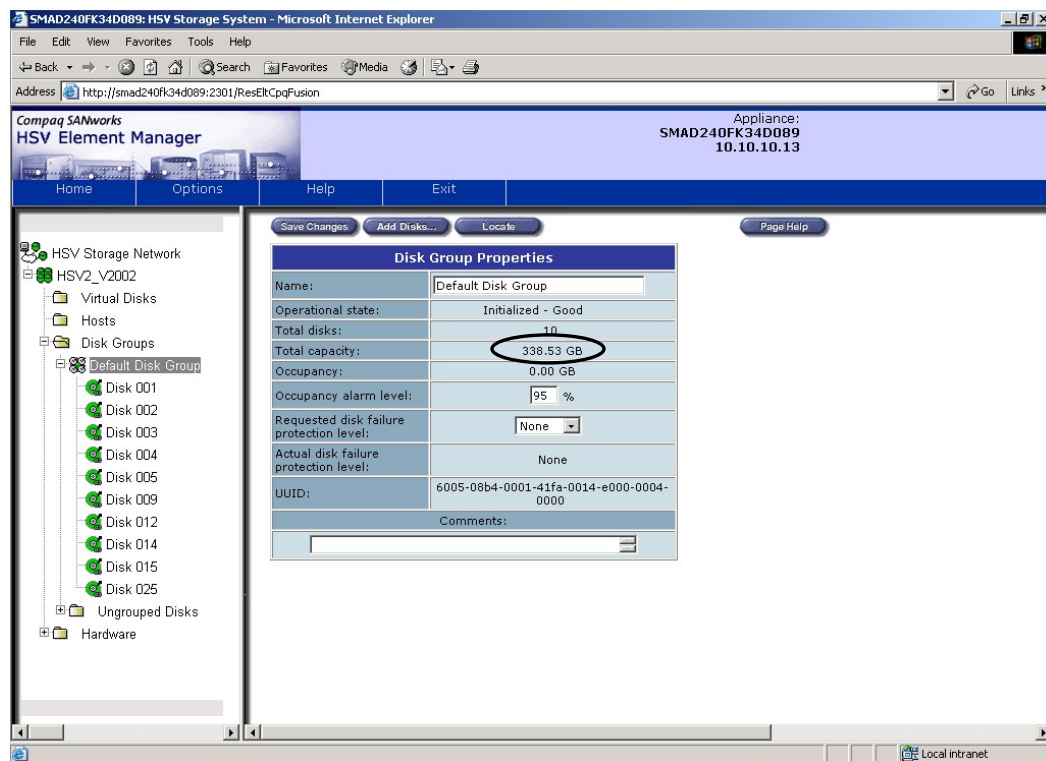
Identification		Condition/State	
Name:	Disk 001	Operational State:	Normal
Loop Pair:	LoopPair2	Migration State:	Not migrating
Node World Wide Name:	2000-0004-CF22-6D05	Failure Prediction:	No
UID:	2000-0004-cf22-6d05-0000-0000-0000-0000	Media Accessible:	Yes
Loop A:		Loop A State:	Normal
Port World Wide Name:	2000-0004-CF22-6D05	Loop B State:	Normal
Assigned LUN:	0	Physical	
Loop ID:	84	Type:	Fibre Channel Disk
Loop B:		Manufacturer:	COMPAQ
Port World Wide Name:	2000-0004-CF22-6D05	Model Number:	BD03654499
Assigned LUN:	0	Firmware Revision:	3BE3
Loop ID:	84	Formatted capacity:	33.91 GB
		System	
		Requested Usage:	Grouped
		Actual Usage:	Grouped
		Disk Group:	Default Disk Group
		Occupancy:	0.00 GB

Formatted Disk Capacity

Even though it appears as though a 36GB disk drive does not give you 36GB of usable disk space, you are using the entire capacity (in decimal equivalent as marketed) of the 36GB disk.

Disk Group Metadata Overhead

When creating a 10-member disk group consisting of 36GB disk drives, you do not get 360GB (36GB x 10) of usable disk space (*Total capacity*). Rather, as the following screen shows, you end up with approximately $(33.91\text{GB} \times 10) \times 0.998 = 338.53\text{GB}$. The 0.2% is overhead used by the controllers to store the disk group metadata. The amount of overhead varies slightly according to the size (number of members) of the disk group.

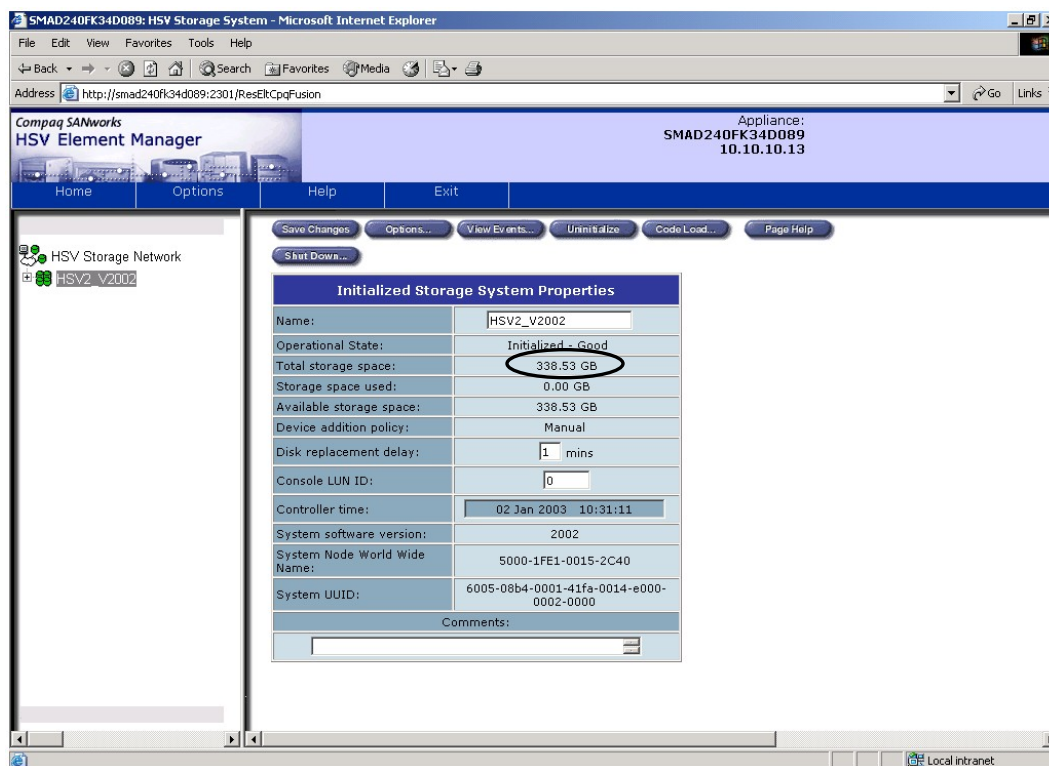


Total Storage Capacity

Total Storage Space and Available Space

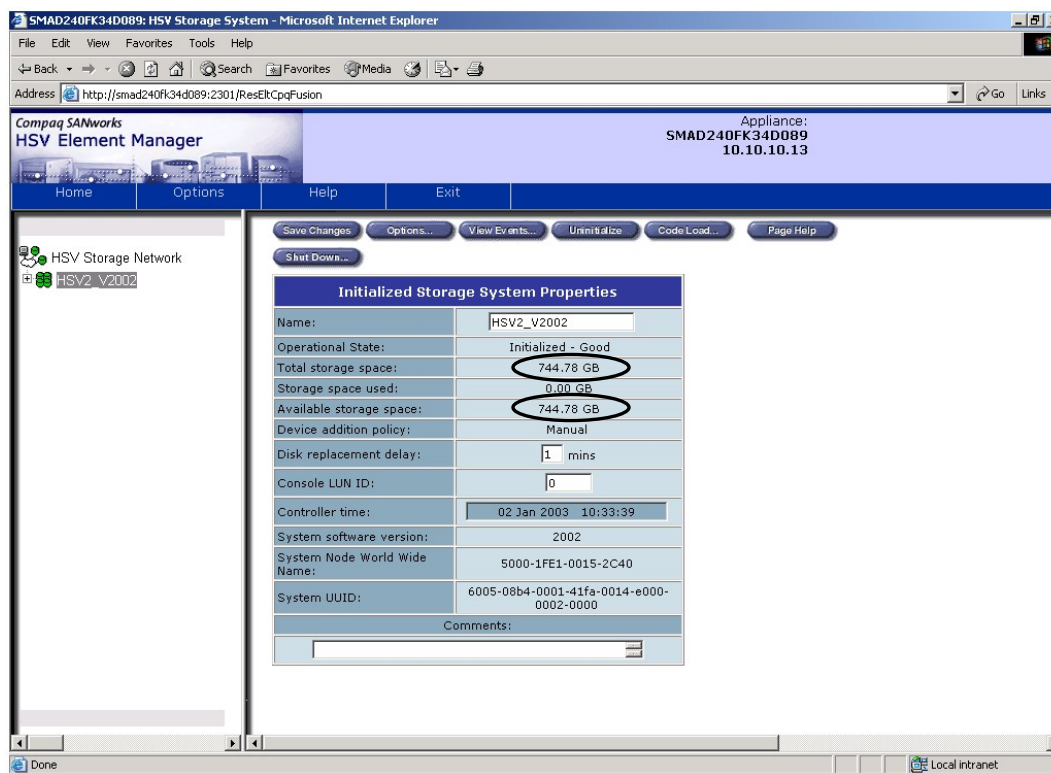
The *Total storage space* shown for any given storage cell shows you how much capacity the storage cell has currently allocated to disk groups.

For example, HSV2_V2002 has a total of 32 installed disks. This yields a total capacity of approximately 1,085GB (32 x 33.91GB). However, as the following screen shows, with only one 338.53GB disk group created, the *Total storage space* for the cell is only 338.53GB.



Total Storage Space with One 338GB Disk Group

When a second 12-member disk group (406GB) is created on the same cell, the *Total storage space* jumps to approximately 745GB ($338.5 + 406.2$).



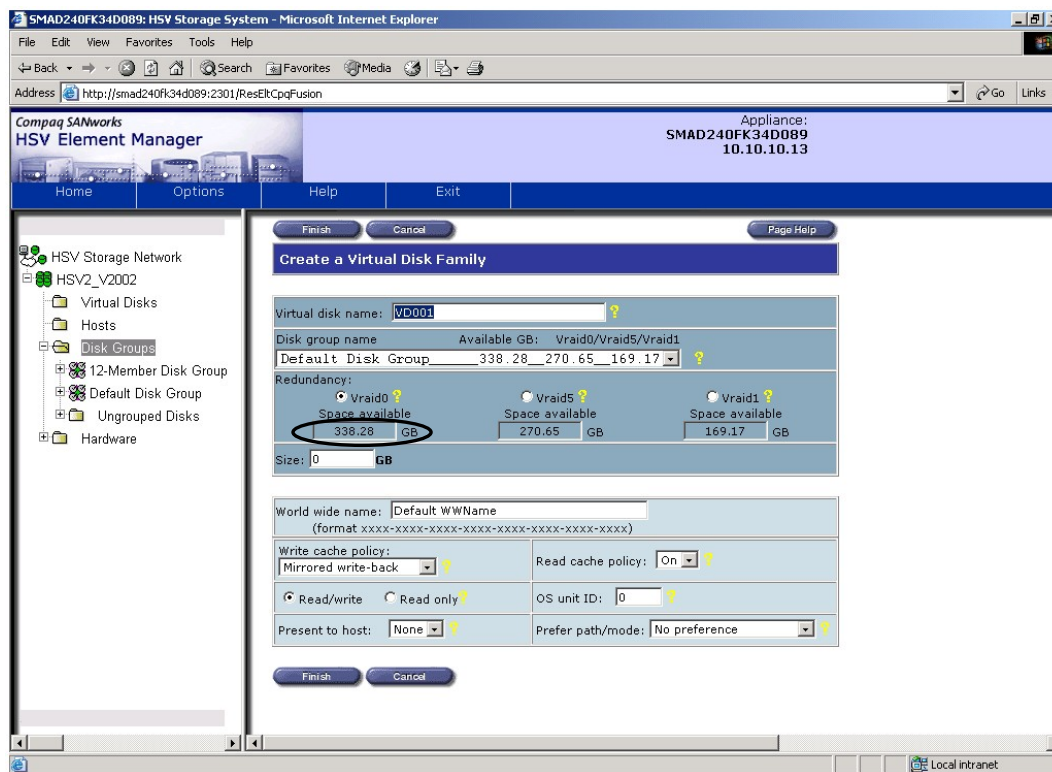
Total Storage Space with Two Disk Groups Present

Notice the field titled *Available storage space*. *Available storage space* is the total amount of currently available cell space from which to create virtual disks.

In this case, no virtual disks have been created so the *Available storage space* matches the *Total storage space*.

Space Available to Create Virtual Disks

Perhaps surprisingly, the *Available Storage space* value doesn't exactly match the total amount of space available from which to create virtual disks. For example, when creating a virtual disk out of the default disk group with a total storage capacity of 338.53GB, you are shown a maximum VRAID0 available capacity of only 338.28GB.



Available Virtual Disk Space

Virtual Disk Metadata Overhead

In the previous example, the physical difference of 0.25GB (338.53GB – 338.28GB) is capacity used by the Enterprise Virtual Array for storing virtual disk metadata. The actual amount of overhead goes up dramatically as the size of the disk group increases.

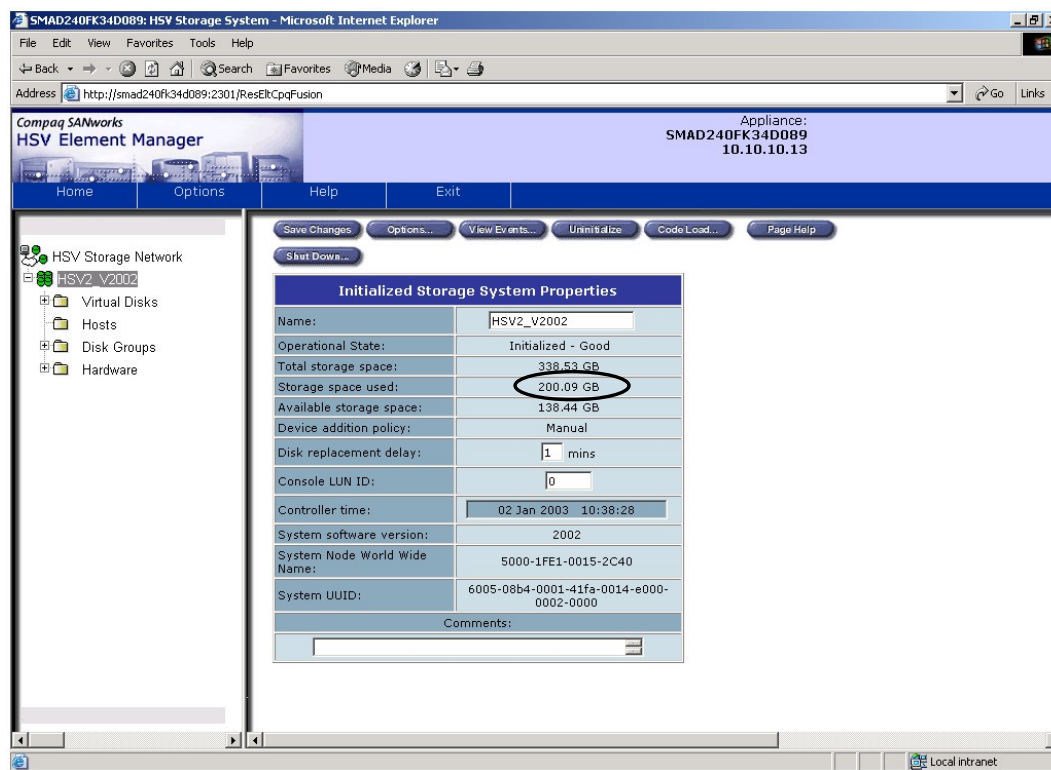
VRAID Redundancy Overhead

Also notice in the previous screen the amount of capacity from which to create the various VRAID levels. Each of the three VRAID levels yields a specific amount of capacity based on its overhead:

- VRAID0
 - 338.28GB is the total available disk group space
 - Striping has no redundancy requirement
 - 0% overhead
- VRAID1
 - 169.17GB is the total available disk group space (total available space divided by 2)
 - Because each disk has to be mirrored, approximately half of the available space will be given up to redundancy
 - If you had an odd number of members in the disk group (for example, 19 or 45), the space of one disk would not be used because each disk **must** be mirrored.
 - 50% overhead
- VRAID5
 - 270.65GB is the total available disk group space (total available space times 0.80)
 - VRAID5 is always based on a 4+1 parity writing scheme
 - 20% overhead

Storage Space Used — Example

Assume a cell called HSV2_V2002 has a 10-member default disk group. If a 100GB VRAID1 virtual disk is created from it (this would require 200GB), the following shows how the *Storage space used* field would look.



Storage Space Used

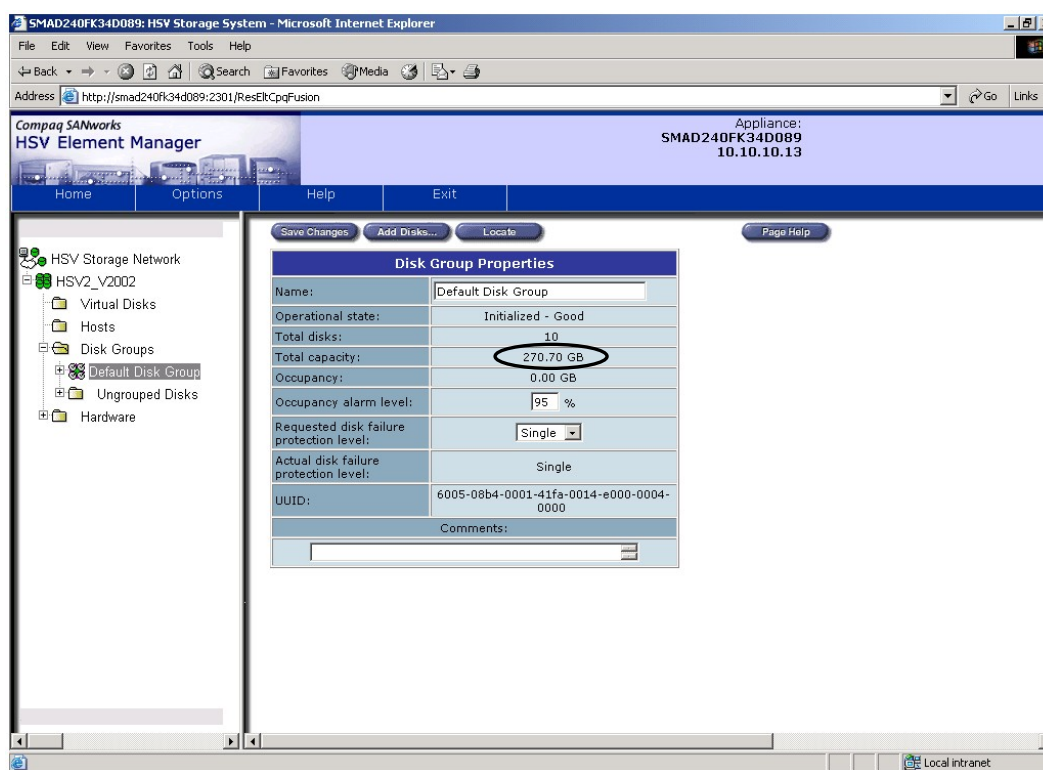
Disk Group Failure Protection Levels

The disk group failure protection level, also known as spare space allocation, will affect the total amount of usable disk group capacity.

In all of the previous examples, the disk groups were created with a failure protection level of none.

Protection Level of Single

When a disk group is created with a failure protection level of single, **two** disks worth of capacity are reserved to deal with disk group member failures. In the example below, the default disk group built with 10 disks of 36GB was created with a failure protection level of single.

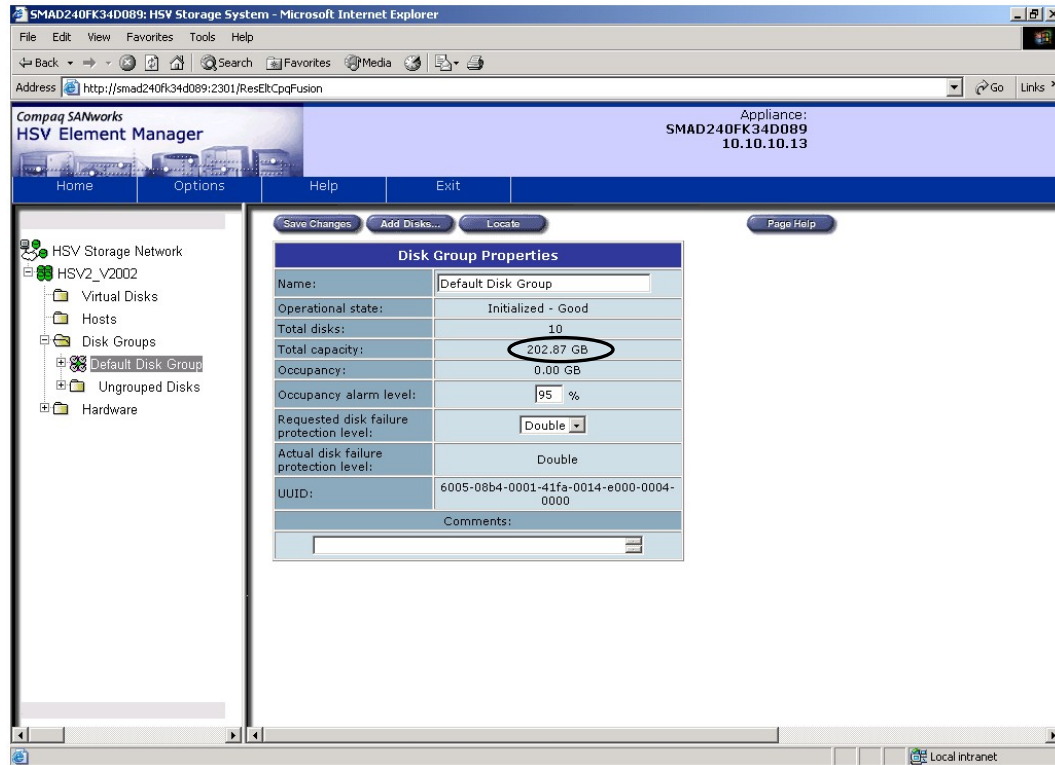


Single Disk Failure Protection

The *Total capacity* of the disk group is not 338.53GB (10 x 33.91GB x 0.998), but is now approximately 270.70GB (8 x 33.91GB x 0.998). Notice that two disks worth of space is subtracted, not one.

Protection Level of Double

When a disk group is created with a failure protection level of double, **four** disks worth of capacity are reserved to deal with disk group member failures. In the example below, the default disk group was created with a failure protection level of double.

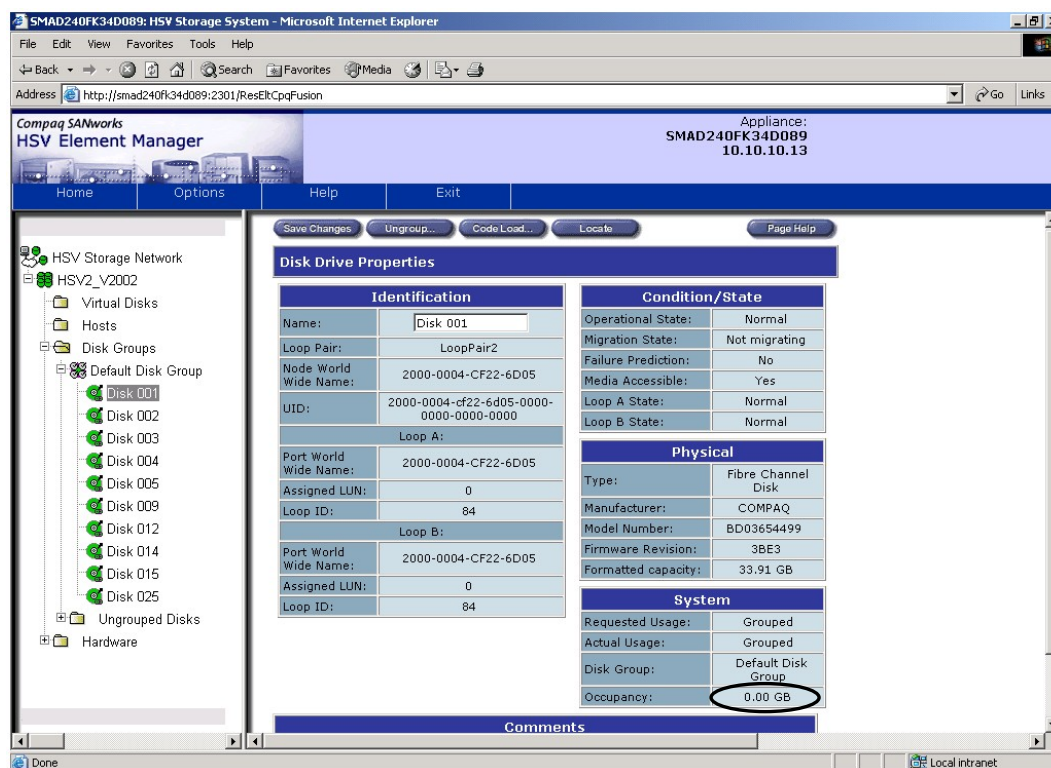


Double Disk Failure Protection

The *Total capacity* of the disk group is now approximately 202.87GB (6 x 33.91GB x 0.998). Notice that four disks worth of space is subtracted, not two.

Disk Occupancy Levels

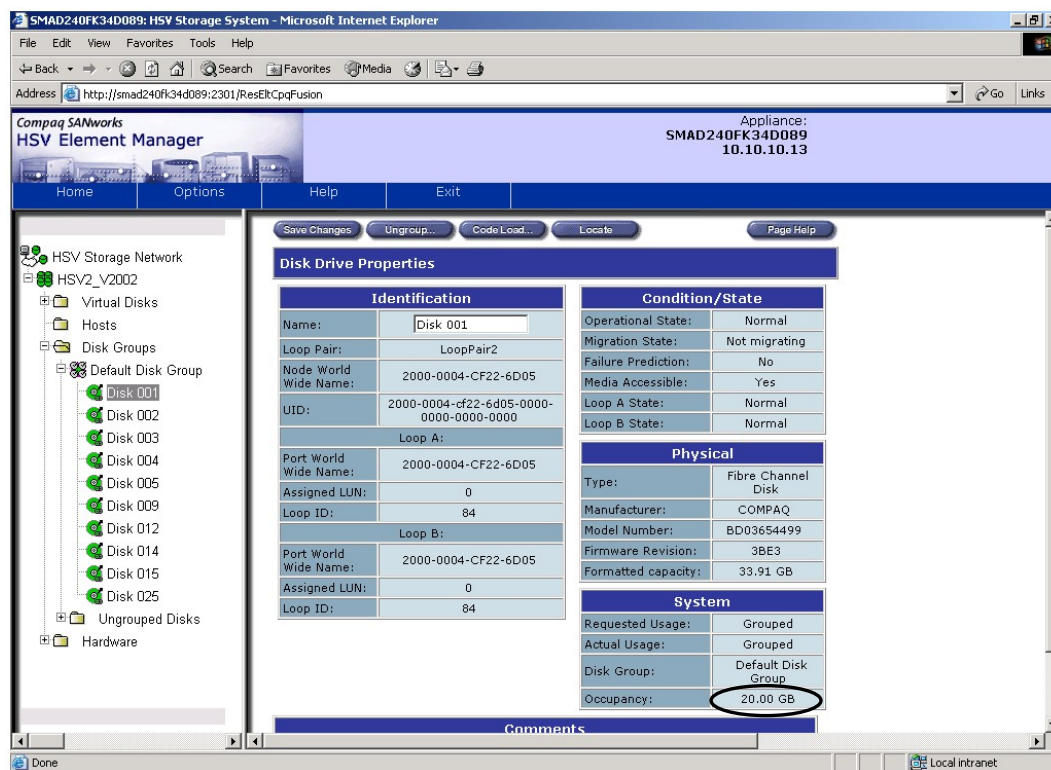
Before you create any virtual disks from a disk group, the *Occupancy* level for each individual member of that disk group will read “0.00 GB”.



Occupancy of a Disk Before Creating Virtual Disks

Once you create virtual disks from a disk group, the *Occupancy* level of each individual disk will reflect the amount of that disk that is being used.

For example, given the ten-member default disk group, if you were to create a 100GB VRAID1 virtual disk from it, the *Occupancy* level for each member of the disk group would be 20GB.



Occupancy with Virtual Disk Created

Note

The yellow warning symbol next to the virtual disk is simply indicating that the virtual disk is in the process of being created.

The 20GB/disk occupancy level was derived as follows:

100GB VRAID1 requires 200GB space.

200GB divided by 10 disk group members is 20GB per disk.



Important

Under normal disk group circumstances (before any of its members fail) the failure protection level of that disk group does **not** impact the occupancy level of its constituent disk drives. That is, regardless of whether the disk group protection level was none, single or double, the occupancy levels of the disk drives above would still only be 20GB.

Adding Up the Numbers

The following is an example designed to show how all of the numbers add up when all factors are considered, including disk failure protection and VRAID levels.

This example uses a ten-member default disk group containing ten 36GB disks with single protection. From that, you create two virtual disks, one a 100GB VRAID1, and the other a 55GB VRAID5.

1. If you create a ten-member (36GB disks) default disk group and use a single disk failure protection level, you lose two disks worth of capacity. Therefore, you would lose:

$$2 * 33.91\text{GB} = 67.82\text{GB}$$

Because you have 10 disks in the disk group, you lose:

$$67.82\text{GB}/10 = 6.782\text{GB per disk}$$

2. Next, create a 100GB VRAID1 virtual disk. This uses a total of 200GB of disk group capacity. You therefore use:

$$200\text{GB}/10 = 20\text{GB per disk}$$

So you have allocated $6.782\text{GB} + 20\text{GB} = 26.772\text{GB}$ per disk

3. Next, create a 55GB VRAID5 virtual disk. This uses a total of 68.75GB ($55\text{GB} \times 1.25 = 68.75\text{GB}$). You therefore use:

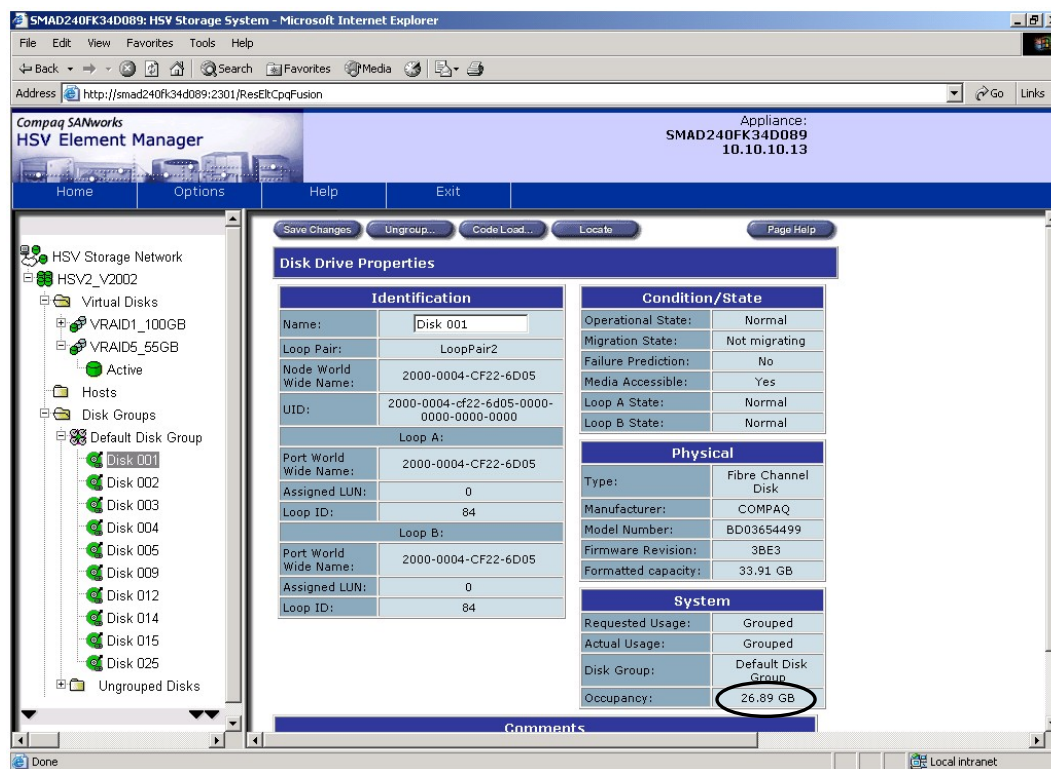
$$68.75\text{GB}/10 = 6.875\text{GB per disk}$$

Note

You use a factor of 1.25 because you need a 25% increase in capacity, that is, for every four disks, you need a parity disk. Another way to compute this is to divide the required capacity by 80%, that is, $55\text{GB}/0.80 = 68.75$. Increasing by 25% is the same as dividing by 80%.

This gives a total of $26.772\text{GB} + 6.875\text{GB} = 33.647\text{GB}$, very nearly the exact 33.91GB formatted capacity of a 36GB disk.

The following screen capture shows the exact numbers that are generated when the example above is implemented.



Example of Real-World Occupancy Level

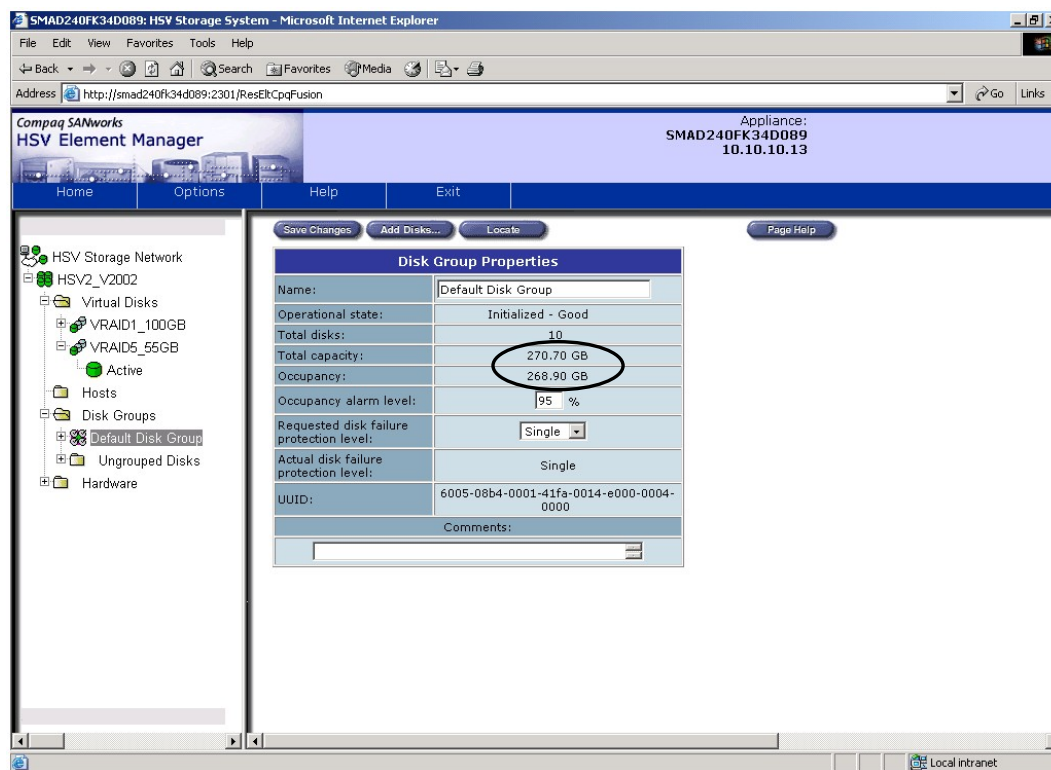
Note

The yellow warning symbol on the disk group is simply indicating the disk group has reached or exceeded its default 95% occupancy alarm level.

Notice in the above output that the disk occupancy level is only 26.89GB.

Remember, the occupancy level for a disk device does **not** take spare space allocation into account until there is a disk group member failure. Therefore, this number is reflecting 20GB per disk for the VRAID1 virtual disk, plus 6.875GB per disk for the VRAID5 virtual disk.

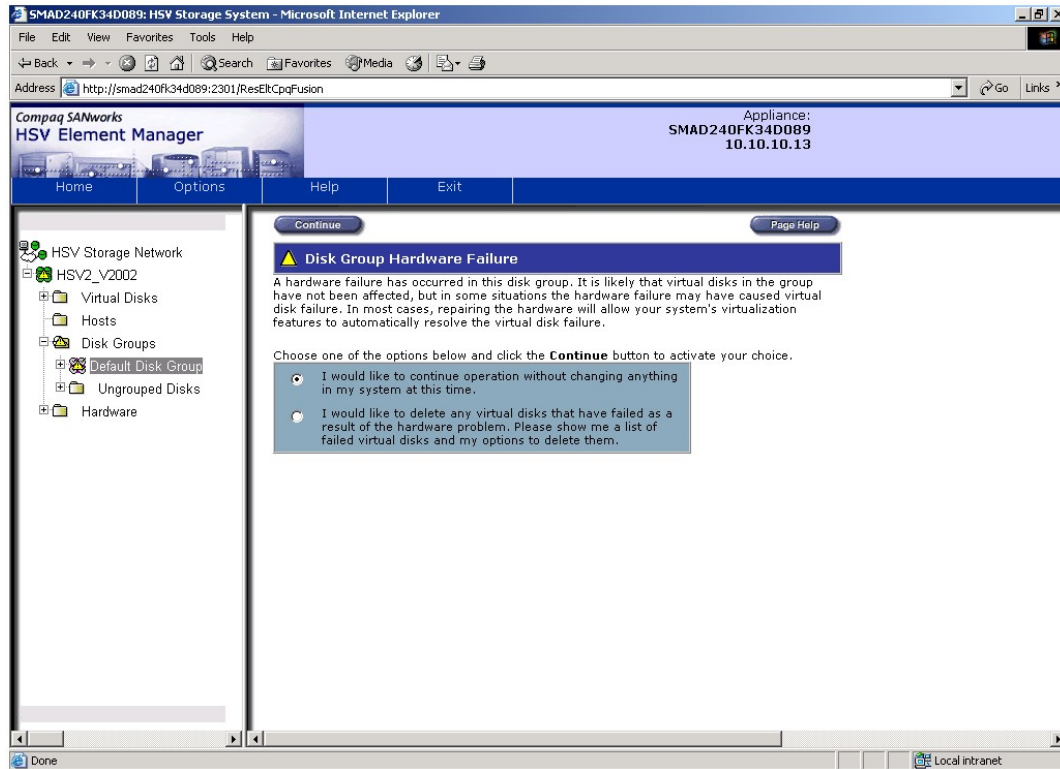
The screen capture below shows that the *Total capacity* of the default disk group is 270.7GB and its *Occupancy* is 268.9GB, indicating only a meager 1.8GB worth of free disk group space left to use.



Total Capacity Accounting for Disk Failure Protection Level

Pulling Out a Disk Group Member Improperly

If you just yank out a disk without first ungrouping the disk and then removing the disk through the GUI, you will get the following disk group error on the screen:



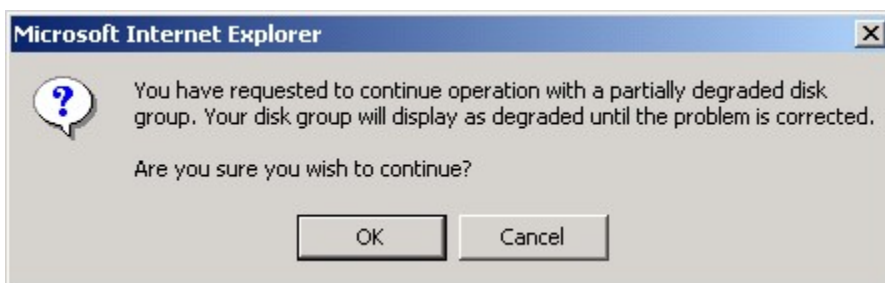
Disk Group Member Failure

There are two different ways to deal with the situation shown above (indicated by the two choices you can select from):

- Manual — The default choice. If you select manual, the Enterprise Virtual Array expects you to attempt to fix the current problem(s) on your own.
- Automatic — The second choice. If you choose this option, the Enterprise Virtual Array will attempt to recover whatever it can from the disk group automatically.

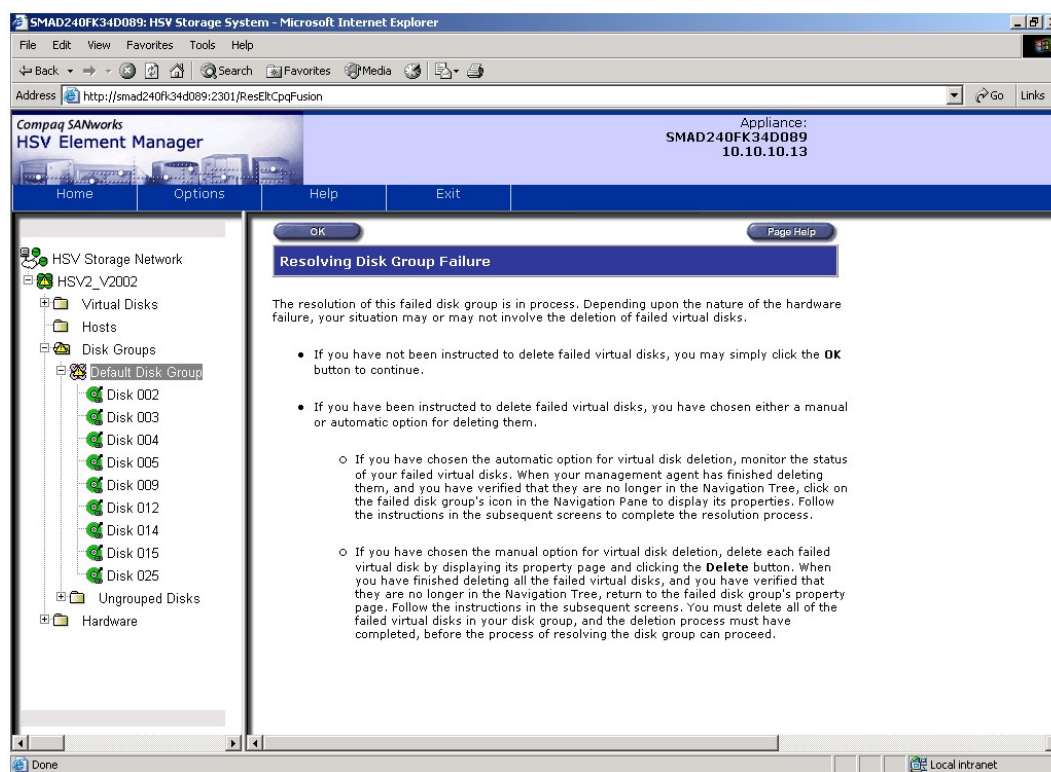
Selecting the Manual Option

If you select the manual option, you will get the following popup window:



Warning for Manual (Default) Option

Selecting *OK* results in the following note regarding resolving disk group failure:



Note on Disk Group Corruption

Click *OK* to continue.

It is likely that the controller will continue to display the disk group error screen until you fix whatever problem caused the error, or until you select the automatic (second) option.

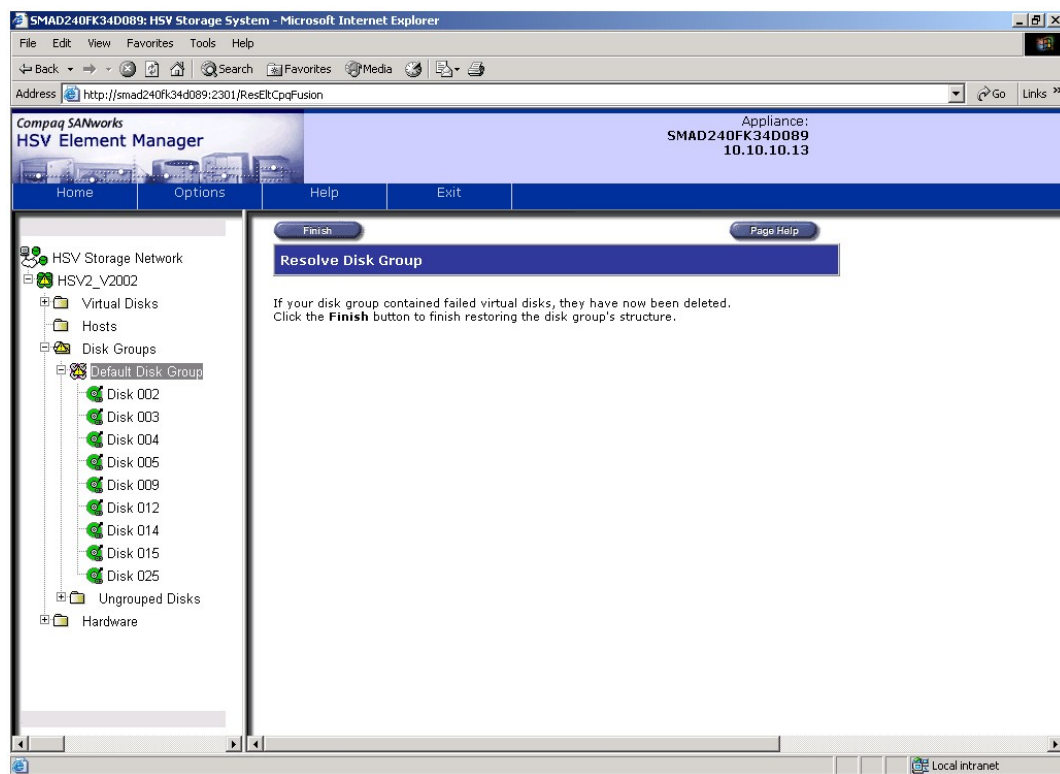
Selecting the Automatic Option

If you select the automatic option, you will get the following popup window:



Automatic Option Selection

Selecting *OK* results in the following note on disk group resolution:



Automatic Disk Group Repair Message



Important

You can avoid these disk group error screens by properly ungrouping (from the disk group) and then removing (by clicking the *Remove* button) the disk group members **before** physically removing them from a disk enclosure.

If you had a real disk group member failure, that is, one that wasn't caused by simply yanking out a disk, the above error screens would be displayed.

Disk Failure Protection Space Usage

The exact usage of disk group disk failure protection space (spare capacity) is an involved subject as testified to by many hours in the lab.

How spare space is used depends on a number of different factors including:

- How many members the disk group has
 - ♦ This will affect Redundant Storage Set (RSS) creation and therefore space allocation.
- Whether or not the disk group contains disks with varying capacities
- What types, sizes, and combinations of virtual disks are created from the disk group
- Whether or not there is adequate free space available in the disk group to use in place of the disk failure protection space

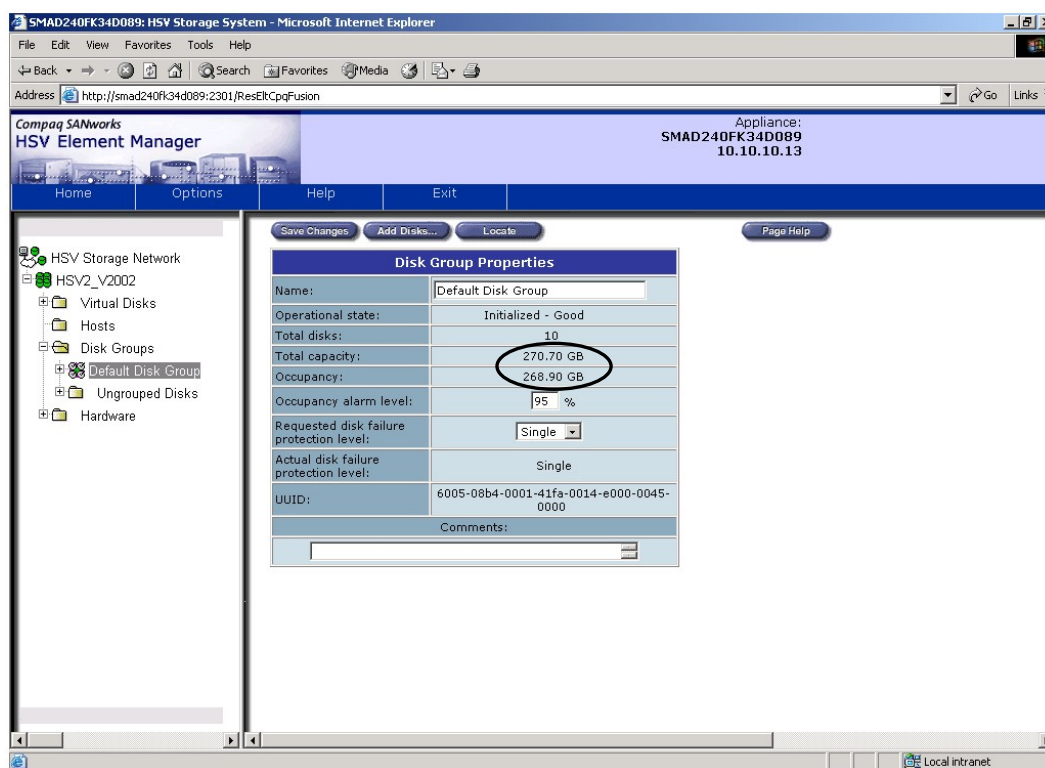
The following is **one** example of the many different ways in which the spare space can and will be used.

Example Using Spare Capacity

For this example, use the same numbers as for the previous example in the topic Adding Up the Numbers, that is, the HSV2_2002 cell with the following:

- 10-member default disk group
- All ten disk group members 36GB in size (33.91GB formatted capacity)
- Single disk failure protection
- A 55GB VRAID5 virtual disk requiring 68.75GB of space
- A 100GB VRAID1 virtual disk requiring 200GB of space

The disk group capacity and occupancy levels for this setup are shown in the following screen capture:

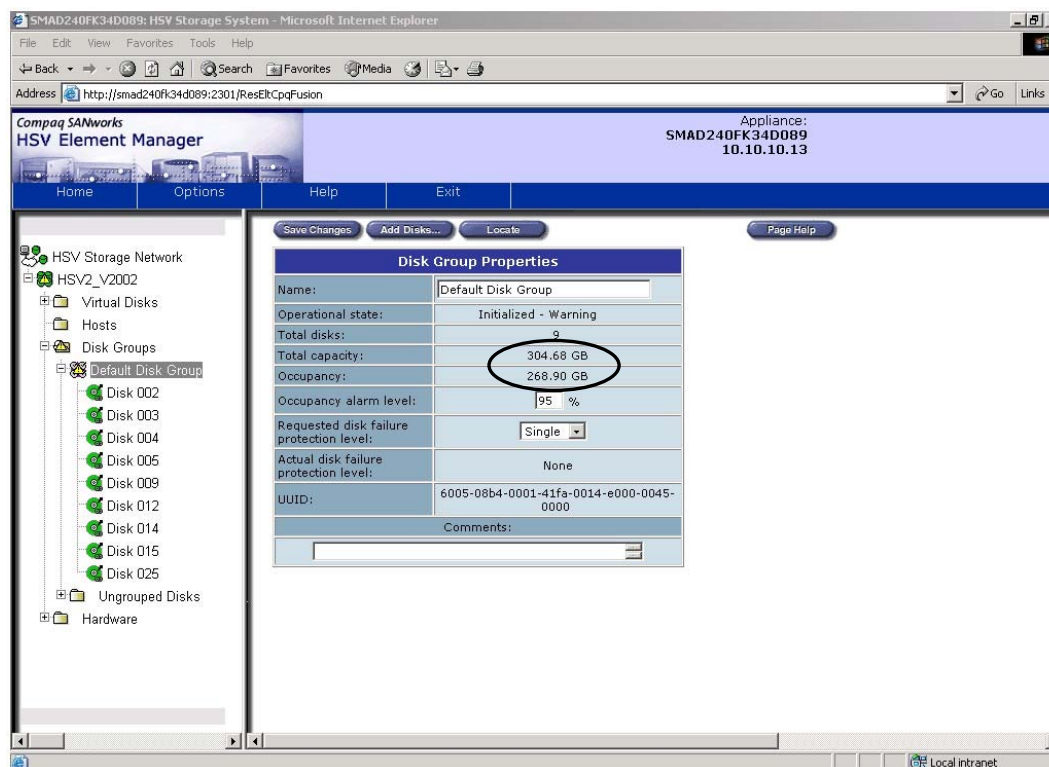


Disk Group Prior to Disk Member Failure

Next a disk drive is pulled from the group.

This will result in the disk group error popup windows discussed earlier. Assume you select the non-default automatic repair action to automatically fix the problem(s) and thus prevent the error popups from reoccurring.

The following screen shows the disk group after the disk member failure:



Disk Group Capacities after Disk Member Failure

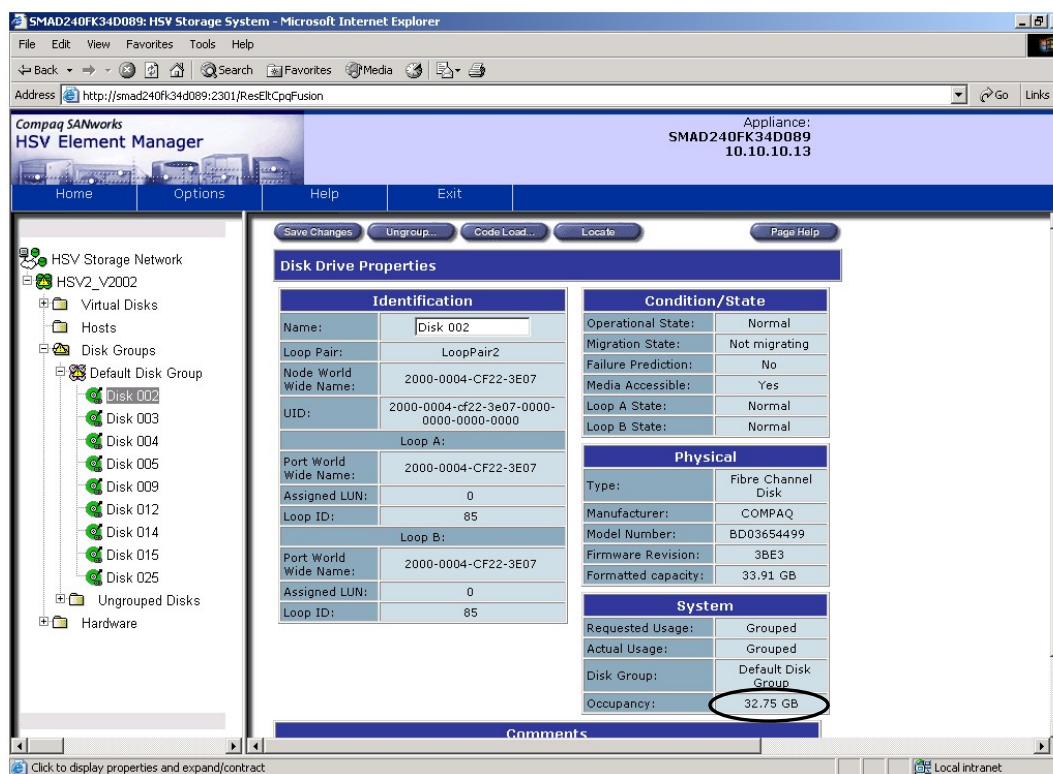
Notice the *Occupancy* level has not changed, but that the *Total capacity* of the disk group is now 304.68GB.

When the disk was pulled out, the *Total capacity* of the disk group went up to 304.68GB. This is reflected by adding the space lost from one disk of the disk group and adding two disks worth of spare space (remember Single protection) for a net gain of one disk's capacity:

$$270.7\text{GB} - 33.91\text{GB (one disk lost)} + 67.82 \text{ (two disks spare space)} = 304.68\text{GB}$$

Therefore, when the disk was pulled, the *Total capacity* of the disk group received a net gain of one disk (33.91GB).

Review the following screen:



Occupancy of a Surviving Disk Group Member (Rising)

This output, taken shortly after the disk group member failure shows that the occupancy level of Disk 002 is rising. Remember, the starting occupancy level of each disk was:

$$20\text{GB} + 6.875\text{GB} = 26.875\text{GB} \text{ (the space required for the VRAID1 and VRAID5 virtual disks combined)}$$

What this indicates is that the data from the failed disk group member (VRAID1 and VRAID5 information) is being reconstructed into the available spare space of the remaining disks.

Note

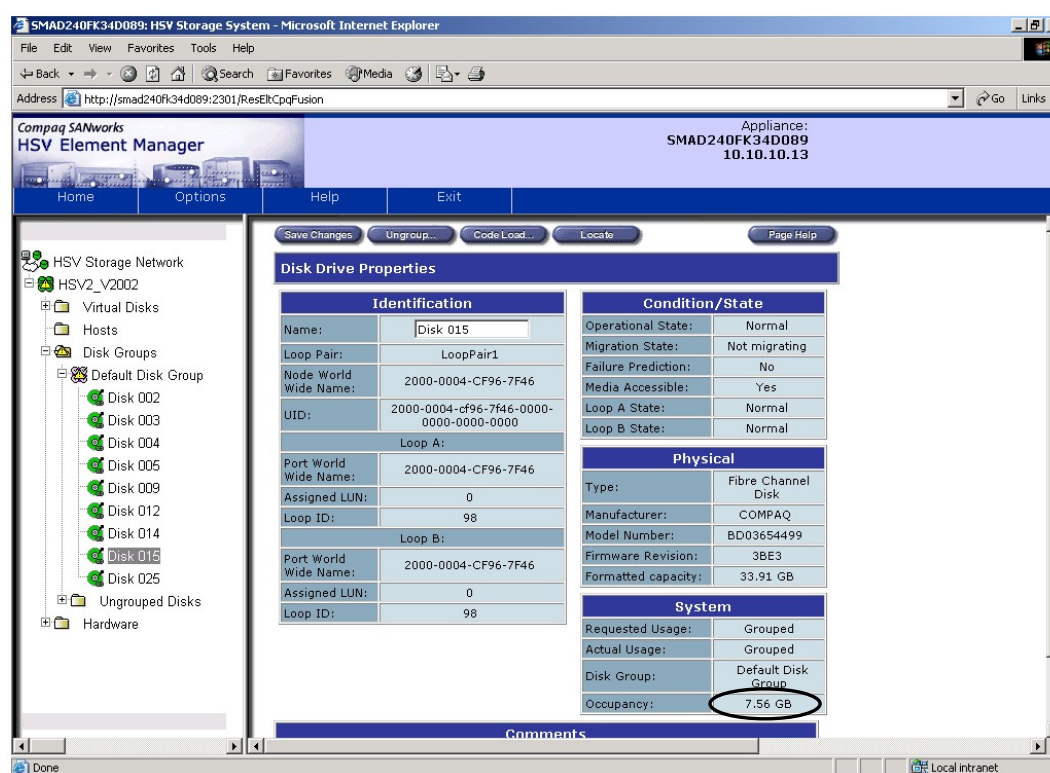
This happens to have been (on purpose) a **very** full disk group. HP recommends attempting to avoid this situation by never having occupancy levels for disk groups exceed approximately 80%-85% of the total disk group capacity.

Note

If this disk group had had sufficient free space on it, the reconstruction process would have taken place within that space. In other words, spare space will only be used if it is the last available space resource.

In reality, not all of the remaining disk group member spare spaces are being used in the same way. The widowed mirror disk will **not** be used for reconstruction of the mirrored data.

Taking a look at the widowed disk (only found by clicking through all of the remaining disks), you would see the following:



Widowed Disk Shortly after the Disk Group Member Failure

Unlike **every** other remaining member of the disk group, this disk group member is going **down** in total occupancy.

What is happening is that as the VRAID1 information is successfully being copied from this disk to the other disks' spare space, that data is **no longer** being stored on the widowed disk.

Note

The data from the failed disk will eventually be fully reconstructed—this is priority one. Once this process has successfully completed, the data is re-leveled.

Reconciling Total Storage Space Used and Occupancy

In this instance, there are eight disks with an *Occupancy* level of 32.75GB and one disk with an *Occupancy* level of 7.56GB.

That yields a total used capacity of: $8 \times 32.75\text{GB} + 7.56\text{GB} = 262\text{GB} + 7.56\text{GB} = 269.56\text{GB}$.

There was originally 200GB capacity for the VRAID1 and 68.75GB capacity for the VRAID5. The total is 268.75GB, very nearly a match.

Disk Sizing Example

When or before an Enterprise Virtual Array subsystem is deployed, the number of physical disks required to deliver the desired *usable* capacity to the hosts must be calculated.

The number of disks required for a projected configuration can be determined by using the following DiskCount formula:

$$\text{DiskCount} \cong ((\text{UsableV0} * 538) + (\text{UsableV5} * 673) + (\text{UsableV1} * 1076)) / (\text{DiskCap} * 476) + (\text{ProtLevel} * 2)$$

where each of the variables is defined as follows:

DiskCap	Disk drive capacity in hardware GB
DiskCount	Integer number of disk drives
ProtLevel	0 for None, 1 for Single, 2 for Double
UsableV0	Desired usable VRAID0 capacity in software GB
UsableV1	Desired usable VRAID1 capacity in software GB
UsableV5	Desired usable VRAID5 capacity in software GB

Another formula with minor differences, as shown in the following scenario, may seem more straightforward.

Scenario: Assume a customer is buying an Enterprise Virtual Array but isn't quite sure how many 36GB disks to purchase. The customer wants three different virtual disks and for arguable reasons, wants each virtual disk in its own disk group. Each disk group has to have a disk failure protection level of double.

The customer wants one virtual disk of each VRAID type. Each virtual disk must be 750GB in size.

Question: What is the minimum number of disks the customer needs to purchase?

Note: assume a total metadata overhead of 2%.

Answer: 110 disks

For each VRAID type, use the following general formula and solve for N:

$$(N - D_{\text{prot}}) \times \text{DriveCap} \times \text{MO} = \text{UsableCap} \times \text{Vfactor}$$

where each of the variables is defined as follows:

N	Number of disks you are trying to calculate
D _{prot}	Number of disks required for disk failure protection
DriveCap	Formatted drive capacity
MO	Metadata overhead percentage
UsableCap	Usable capacity required
Vfactor	VRAID overhead factor

Note

This formula, because it doesn't use such large constants, may seem a little less imposing when calculating by hand. However, disk sizing formulae in MS Excel worksheets use the more precise values of the first DiskCount formula.

For VRAID0, $(N - 4) \times 33.91 \times 0.98 = 750\text{GB} \times 1$

N = 27 disks (rounded up)

For VRAID1, $(N - 4 \times 33.91) \times 0.98 = 750\text{GB} \times 2$

N = 50 disks (rounded up)

For VRAID5, $(N - 4) \times 33.91 \times 0.98 = 750\text{GB} \times 1.25$

N = 33 disks (rounded up)

Total disks required = 27 + 50 + 33 = 110 disks

If using the first DiskCount formula, you would get approximately the same value.

Note

In the real world, you would want extra space available to each disk group. The customer would be much better off purchasing 120 to 130 disks.

learning check answers

appendix D

Introduction

This appendix supplies the answers to the learning checks in each module of this course.

Module 1 — Product Overview

1. Windows NT, Windows 2000, Compaq Tru64 UNIX, Sun Solaris, HP-UX, IBM AIX, OpenVMS, Red Hat and SuSE Linux, Novell NetWare
2. 12.2TB
3. Fibre Channel arbitrated loops
4. There are the PDU and PDMs. There is one dual PDU and eight (four each side) PDMs.
5. Cables and junction boxes
6. Enclosure address bus, drive enclosure, FC loop switches, and the PDMs
7. Two
8. One

Module 2 — HSV Controller

1. All possible answers are: removable bezel, blower compartment (two blowers), operator control panel (four LEDs, LCD, four-position pushbutton switch), cache battery compartment (two lead-acid battery assemblies).
2. All possible answers are: controller pairs (two interconnected HSV series controllers), blowers (two), cache battery assemblies, mirrored write-back cache, operator control panel.
3. Fault LED — Amber light flashing
Controller heartbeat LED — Green LED not flashing
4. Alternates between Storage System Name and Node Worldwide Name
5. Active flashing display
Error condition message
Pressing a pushbutton (user entry)
6. An asterisk (*) in VCS versions before V2, a test entity number in VCS V2.x and greater
7. Four navigation pushbuttons on the OCP
8. Eight alpha characters
9. Blowers, batteries, cables, and transceivers
10. FC loop switches (switched hardware) or disk drive enclosures (nonswitched)
Other controller (inter-controller link — mirror port)
Fabric switches to host

Module 3 — Disk Drive Enclosure

1. All possible answers are: EMU, I/O module A, I/O module B, blower 1, blower 2, power supply 1, power supply 2
2. The drive bay numbers start from 1 at the left of the enclosure and go up to 14. The disk drive numbers correspond to the drive bay numbers.
3. The left end
4. EMU status, enclosure power status, and enclosure fault status
5. Two
6. **Er** represents an error report
7. At the left end of the disk drive enclosure, next to the EMU
8. Three
9. Er (error report) has precedence over all other display groups
10. Enclosure status LEDs

Module 4 — Fibre Channel Loop Switch

1. Expansion panel
2. 2.125Gb/s operating speed, 12 ports, 1U size, UPS, SFP transceivers
3. The controller device ports cable to the FC loop switch ports, which in turn cable to the drive enclosures
4. System LEDs
5. FC loop switch, FC loop switch transceiver

Module 5 — Rack, Cabling, Configuration, and Initial Setup

1. opal 42U, graphite 41U
2. PDUs, PDMs
3. switched, nonswitched
4. 1 = loop A bottom port
2 = loop A top port
3 = loop B bottom port
4 = loop B top port
5. Loop A cables to an I/O module A.
6. Loop pair 2 enclosures are at the top of the rack, below the controller pair
7. The four jumper cables on the expansion panel are removed and eight FC interconnect cables are wired between the two racks
8. All of the drives have spun up
9. 2,1,4,3
10. basic license, WWN

Module 6 — Storage Management Appliance

1. ProLiant DL380 server, 1.26GHz processor, 1GB memory, 2 KGPSA-CB adapters, Remote Insight Board II, 2-18GB Ultra3 drives mirrored
2. Open SAN Manager V1.0C or Storage Management Appliance Software V2.0 SP1A or SP3, HSG Element Manager V1.0E, Command View EVA V2.1 or V3.0
3. All possible answers are:
Initialize the Storage Management Appliance.
Configure the appliance parameters (IP network address, system name, date and time, time zone, and administrator password).
Install and setup SAN applications.
Upgrade system components such as operating system and drivers.
Start and restart the Command View EVA service.
Launch the installed SANworks applications such as Command View EVA.
4. You can only install one Command View EVA instance on one appliance
5. Obtain the licensing agreement, authorization ID, and instruction sheet.
Enter the authorization ID and node WWID of the Enterprise Virtual Array into the HP License Fulfillment Website.
Obtain the license key mailed back from the HP License Fulfillment Website.
Enter the license key into Command View EVA.
6. Establish communication with the appliance.
Configure static IP addresses.
Set time.
Add to a domain if necessary.
Reconnect to the appliance.
Configure the switches (zoning).
Change the appliance name if necessary.
7. Restore the Storage Management Appliance operating system to its original default state

Module 7 — Concepts and Terminology

1. storage system
2. 2MB, 256
3. vertical
4. None — VRAID0, data is striped across all physical disks in the disk group, provides 100% of disk group capacity
Moderate — VRAID5, data is striped with parity across all physical disks in the disk group, always 5 (4 + 1), provides 80% of disk group capacity
High — VRAID1, data is striped mirrored across all physical disks in the disk group, provides 50% of disk group capacity
5.
 - a. Virtual RAID
 - b. Conventional RAID
 - c. Virtual RAID
 - d. Conventional RAID
6. virtual disk leveling
7. None, the default
Single, requires the space of two disk drives per disk group
Double, requires the space of four disk drives per disk group
8. Fully-allocated snapshot — Capacity is reserved up front
Demand-allocated snapshot — Capacity is only used as the original virtual disk changes
Snapclone — Fully-allocated snapshot at creation time, the data can be accessed immediately without having to wait for the copy to complete

Module 8 — Storage System Configuration

1. Menu Options: Home, Options, Help, and Exit
Storage Management Appliance Name and IP Address
2. System help by clicking the Help button on the Session pane.
Page help by clicking the Page Help button on the Content pane.
Field help by clicking the Question Mark next to a field.
3. By clicking on Help, About from the Session pane.
4. Code load — SuperImage file with VCS and EMU firmware
Disk drive firmware code load — Disk drive firmware
5. Disk enclosure bay properties
6. The default disk group
7. LAN name (or alias)
IP address (optional)
WWID of one FCA
Type of operating system
8. Properties page
9. First add a host using the Add a Host wizard. Then, add an FCA using the Add a Host Port page.
10. Up to seven snapshots of an active virtual disk can exist at any time

Module 9 — Storage System Management

1. Storage system access, licensing, user interface, page footer message, display refresh
2. Storage System World Wide Name
3. Select the storage system, then select the rack
4. VCS revision number
5.
 - a. controller A, host port 1, 9
 - b. controller A, host port 2, 8
 - c. controller B, host port 1, D
 - d. controller B, host port 2, C

Module 10 — Error Reporting and Diagnostics

1. Fatal error that causes the storage system to stop operation
2. Management Agent Event Log
3. Management Agent Event Log, Controller Event Log, and Controller Termination Event Log
4. Command View EVA console
5. Get Log File
6. Management Agent Event Log
7. SWCID, Event Number, CAC, EIP Type, CCCC, DRCC, Parameter Count
8. Set Event Notification Options
9. BIST, MIST

Module 11 — Storage System Scripting Utility

1. Configure and control the HSV controllers and storage system
2. Disk failure level protection — None, single, double
3. The password access to the controller (if enabled) must already be set up to execute the Scripting Utility. You cannot set this password from within the Scripting Utility.
4. \Hosts
 "\Virtual Disks"
 "\Disk Groups"
5. SHOW?
6. Directs the command prompt to Command View EVA
7. SET HOST
8. A snapshot of the storage named *payroll* is created to *payroll_backup*
9. It is the only method used to recreate the configuration of a storage system

Module 12 — Best Practices

1. Maximum of 256 hosts
Maximum of 1,024 FCA connections
Maximum of 512 LUNs
Maximum of 256 LUNs per FCA
2. 16, 8, 4
3. Disk group availability, disk group performance, and disk group capacity utilization
4. random, sequential, small
5. Use few, but large, disk groups
Use groups with drives of identical or similar capacity
Use an even number of drives (either multiple of number of shelves or a multiple of 6 or 8 drives)
Use vertical disk groups
Avoid using more than two disk drives per shelf in the same disk group
Avoid mixing VRAID0 and other VRAID virtual disks in the same disk group
6. Hardware versus software data representations
System metadata overhead
VRAID level
Spare capacity
Snapshot working space
7. Approximately 53 disks. Using the formula, you get 51; then add 2 for single protection level.

Module 13 — Troubleshooting

1. Management Agent Event Log, Controller Event Log, Controller Termination Event Log
2. NAPP, Event View EVA, and HSVET
3. operator control panel, four LEDs, LCD, control pushbuttons, device port LEDs
4. alphanumeric display, EMU LEDs, pushbutton LEDs, disk drive LEDs, power and blower LEDs, I/O module LEDs
5. `cfgShow`
`fabricShow`
`nsShow`
`nsAllShow`
`supportShow`
`switchShow`
`portShow`
`portStatShow`
`portPerfShow`
`version`
6. Disk drive
Drive blank