

HP Modular Network Storage Solutions

ESG2069SG10306



HP Modular Network Storage Solutions

ESG2069SG10306

HP Training

Student guide

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

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

This is an HP copyrighted work that may not be reproduced without the written permission of HP. You may not use these materials to deliver training to any person outside of your organization without the written permission of HP.

Printed in USA

HP Modular Network Storage Solutions

Student Guide 1

June 2003

HP Restricted

Table of Contents

Introduction

Product description

Operating Features

Hardware Architecture

Virtual Front Panel

Management Software

Business Copy VA

Secure Manager VA

Auto Path

SAN Integration

Review

Lab Guides

Lab 1 – Virtual Front Panel

Lab 2 – Command View SDM GUI

Lab 3 – Command ViewSDM CLUI / CVUI

Lab 4 – Business Copy

Lab 5 – Secure Manager

Lab 6 – AutoPath VA

Appendix

A00 – Configuration and Order Guides

A01 to A17 – Several Product Briefs and Whitepapers



HP
SurePartner
Training

MNS
Crossupdate

HP Modular Network Storage Solutions

Part 1

Introduction



Student Notes:

Please remember your mobile phone...

- Please switch your mobile to „silent“ or activate vibration signalling
- It is not a problem if you leave the room to take up a call
- Thank you!



Student Notes:



before getting started

- **class times**
 - **breaks**
 - **lunch**
- **restrooms**
- **emergencies**
- **smoking area**
- **phones, fax and general organisational issues**



Student Notes:

HP
SurePartner
Training

MNS
Crossupdate

Welcome to the new HP Introduce yourself

- **name**
- **organisation**
- **title / responsibilities**
- **experience with MNS**
- **expectations**
- **your last movie**



Student Notes:

HP Modular Network Storage Solutions - agenda

Day One :

- Part 1: Introduction
- Part 2: Product Description
- Part 3: Hardware Architecture
- Part 4: Virtual Front Panel
- Part 5: Management Software

Day Two :

- Part 6: Business Copy VA
- Part 7: Secure Manager VA
- Part 8: Auto Path VA
- Part 9: Operating Systems
- Part 10: SAN Integration



Student Notes:

objectives

- Understand HP storage portfolio and strategy
- Understand the HP MNS products
- Understand how to configure the HP MNS products using the configuration and order guides
- Understand the importance of the implementation needs to correctly propose the right solution
- Be able to demonstrate the important features of the HP Storageworks Virtual Array



Student Notes:

hp on the way to #1 in storage



"Today, HP is moving forward with an aggressive storage strategy. It offers a broad range of storage products that it can demonstrate interact effectively. As an HP customer, you should feel comfortable that HP will regain a leadership position such that it may be thought of as a storage company first."

– Bob Zimmerman
Giga Information Group

"We look at three things when selecting data warehouse disks: price per terabyte, throughput capabilities, and availability features. With the HP virtual array, the price per terabyte was extremely competitive and the throughput was phenomenal. The AutoRAID features of the VA exceed our availability requirements."
Mark Dunlap, director of data warehousing amazon.com

**winning over customers
and analysts**



HP is widely recognized as a leader in the storage industry. We can deliver solutions to meet any enterprise storage need.

HP
SurePartner
Training

MNS
Crossupdate

why HP?


- broad heterogeneous – server and storage platform – centralized management capability
- integrated problem management and accounting for real-life datacenters
- integrated service management to provide detailed SLA management
- scalability and performance
- solutions designed to scale to meet any demand without application impact
- heterogeneity and flexibility
- ability to connect all major open systems, windows and mainframe platforms
- virtualization technologies to increase operational flexibility



Student Notes:


External RAID Vendor Revenue Market Shares, 2001

SAN-Attached External RAID Revenue Market Share, 2001



Source: "2001 RAID-based Disk Storage Market Share", Gartner, (May 2002)

External RAID Revenue Market Share, 2001




Source: "2001 RAID-based Disk Storage Market Share", Gartner, (May 2002)



The new HP is the #2 Enterprise Storage vendor with 20% market share in 2001

NAS Vendor Revenue Market Share, 2001




Source "2001 Network-Attached Storage Market Share" Gartner Dataquest – April 2002



While HP trails significantly in the NAS market, the company continues to invest in NAS solutions.

Storage Management Software Vendor Market Shares, 2001




Source: Gartner Dataquest – "2001 RAID Based Disk Storage WW Market Share & Forecast," (May 2002)



With revenue of \$273.5M, the new HP is now the 4th largest storage management software vendor and gaining ground.

Tape Drive Vendor Market Shares, 2001



With revenue of \$273.5M, the new HP is now the 4th largest storage management software vendor and gaining ground.

LTO – Linear Tape Open. Linear Tape-Open (LTO) Ultrium drive support is an industry initiative created by HP, IBM, and Seagate that provides an open tape drive format choice for customers. Ultrium is the LTO format for backup and restore applications. LTO Ultrium tape drives use a best of breed technology approach, taking the best features from other tape technologies and combining them into a single new technology. This inclusion of innovative features such as data rate matching means that the LTO Ultrium tape drive has a proven mean time before failure (MTBF) rating of 250,000 hours at a 100 percent duty cycle, making the LTO Ultrium drive a true enterprise class tape drive. Linear Tape Open (LTO) technology is an industry open-standard format; thus data written on HP StorageWorks Ultrium tape drives can be interchanged directly with other Ultrium tape drives from other vendors.

HP now offers Ultrium tape drives for ProLiant systems and in September 2002 will announce Ultrium offerings to include support of MSL and ESL tape libraries. Ultrium technology leads the “super drive” market in sales, out shipping SDLT 2:1.

HP continues to dominate the LTO drive market, with 46.7% of Ultrium drive shipments. Customers are choosing Ultrium over SDLT for it’s performance and industry standard format.

DLT – Digital Linear Tape. DLT is a high-speed magnetic tape format that uses streaming, serpentine recording instead of helical scan recording.

SDLT – Super Digital Linear Tape

DDS – Digital Data Storage the industry standard for Digital Audio Tape formats, provides a method of recording information in digital form on a small audio tape cassette. 4-millimeter, helical scan drive, which can hold multiple gigabytes of information.

HP
SurePartner
Training

MNS
Crossupdate

HP ENSA Vision



Storage is a utility service to the enterprise ...

**The Enterprise Storage Utility offers
customers storage capacity and access to:**

Anyone, Anytime, Anywhere



What is ENSA and how does it fit into the overall new HP corporate strategy?

ENSA (Enterprise Network Storage Architecture) is the HP storage vision, which is based on the concept of Storage as a Utility. A key element of the HP corporate strategy consists of providing standard building blocks for Internet access and infrastructure. The ENSA storage utility is the foundation of the infrastructure component of the strategy.

Storage Roadmap

pre-merger

the new HP

online storage systems

- Compaq EVA and HP XP high-end storage
- HP VA family, Compaq EMA
- Cluster storage solutions and MSA1000
- Mixed NAS storage offerings

- Solutions spanning EVA and XP
- Consolidate VA and EMA to EVA-based platform over time
- MSA 1000 — DAS-to-SAN (DIS) investment protection
- NAS/SAN fusion

nearline storage

- HP Surestore libraries
- Compaq StorageWorks libraries
- Compaq SDLT
- HP Ultrium technology

- Continue StorageWorks libraries for midrange and enterprise business needs—add LTO support
- Mix of HP & Compaq autoloaders for entry-level
- Best-of-breed (DLT/SDLT) & HP Ultrium tape drives

storage management software

- Mixed storage management software and solutions spanning HP OpenView storage resource management (SRM) and Compaq SRM software, HP Omniback data protection and the HP and Compaq suites for high availability (HA) and replication

- OpenView SAM suite augmented with selected Compaq capabilities for SRM
- Omniback II becomes strategic platform for lifecycle data management in addition to best-of-breed partner solutions for data protection
- Adopt the Compaq HA/Replication suite augmented with selected HP capabilities

virtualization technologies

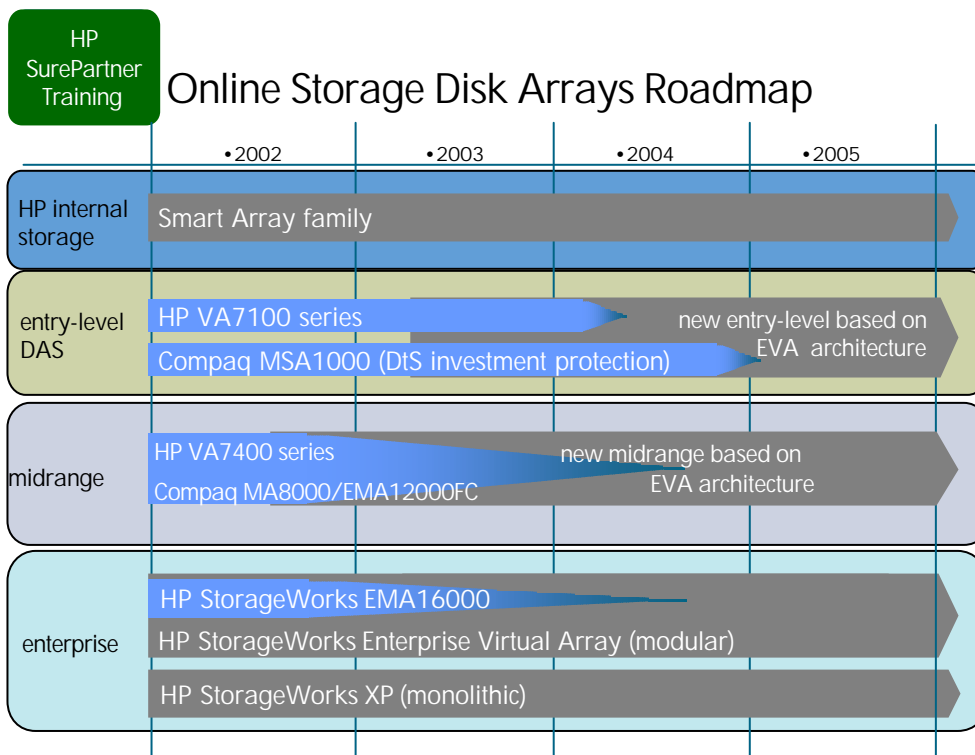
- A variety of virtualization technologies including SANlink, VersaStor, Virtual Replicator, and virtual arrays (EVA/VA)

- VersaStor becomes umbrella technology
- Continue SANlink, converge into single appliance model with SAN-wide VersaStor over time
- Continue with EVA
- Offer host-based for maximum flexibility

storage utility

Both Compaq and HP had very robust storage portfolios so we are in the enviable position of being able to develop a roadmap based on the best of the best. Initially, the new HP will continue to offer virtually the same mix of products. Of course, there is some product overlap -- and we are addressing it immediately and decisively. As you can see, the tough decisions have been made. Our goal now is to mitigate any potential impact our decisions may have on our customers storage environment or our relationship with business partners.

Online: We will provide two categories of high-end storage for maximum choice; the HP XP provides monolithic storage for the enterprise, and the EVA provides high end array provides highly modular, scalable and efficient heterogeneous storage for the enterprise. The EMA family will be available for Compaq centric environments (hp proliant, TRU64R and OVMS) and heterogeneous environments. The VA family will continue to provide best of breed HP-UX solutions and heterogeneous connect, but over time the VA family and MA/EVA will consolidate into a new EVA midrange and an entry level version too. Our entry level MSA1000, will continue the investment protection inherent in our DAS to SAN (DtS) for our Proliant servers. For NAS, both companies have offerings, but neither holds a leading market share position today. We will have a stronger NAS offering that grows into universal networked storage for SAN and NAS fusion, creating a common storage pool that is highly scalable and simpler to manage.



This is a migration slide showing how disk arrays will migrate over time.

Note: Software for online storage disk arrays will track the life cycle of the hardware. The EVA architecture is very flexible, in that, with some changes, we are able to make derivatives products out of the EVA high-end architecture to span the midrange and entry segments as well.

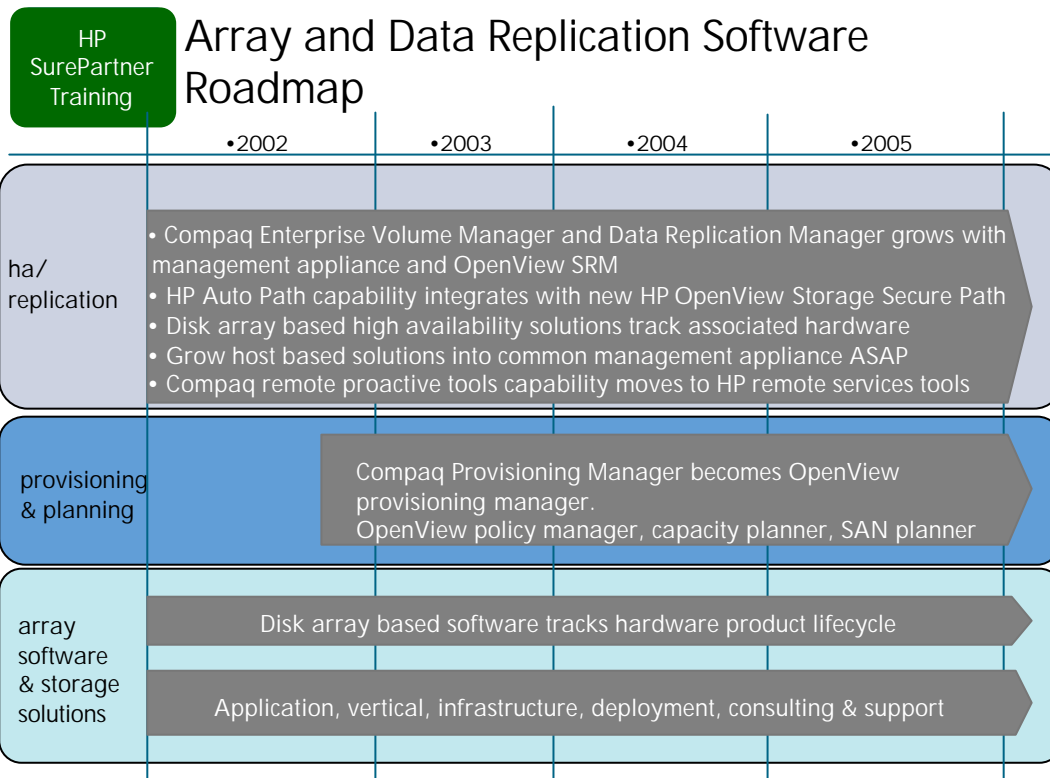
Internal RAID: HP NetRAID products will transition to the Smart Array Family. Smart Array has a much more extensive product mix and a very large installed base.

Entry Level DAS: Both the VA7100 and MSA1000 will coexist. VA7100 is the solutions choice for HP-UX and heterogeneous environments. Eventually the VA7100 will transition to the EVA-entry level disk array over time. The MSA1000 becomes the lead entry level DAS to SAN offering and for hp proliant servers. The Smart Array Cluster Storage Solution for SCSI solutions is available and can easily be upgraded to the MSA1000 for FC and SAN support for maximum investment protection. Easy hardware upgrades and DtS technology for migration and investment protection.

Midrange Storage: The VA family and subsequent releases focuses, primarily, on best of breed HP-UX support and heterogeneous connect. MA8000/EMA12000FC will continue as the preferred midrange solutions for Compaq centric environments (proliant, TRU64R and OVMS) and heterogeneous connect too until its scheduled EO. By mid 2003, we will introduce a midrange disk array based on the EVA architecture that will be a replacement for both offerings.

Enterprise Storage: In the high-end enterprise space, there are really no changes to product plans. We offer maximum choice in high-end storage with the EVA enterprise modular and XP Enterprise monolithic disk array offerings. Both platforms are highly scaleable and offer outstanding solutions for business continuity and storage consolidation.

The XP can be used for solutions requiring mainframe connect. EMA16000 will be replaced by EVA high-end solutions.



This is a migration slide showing how the products will migrate over time.

HA/Replication

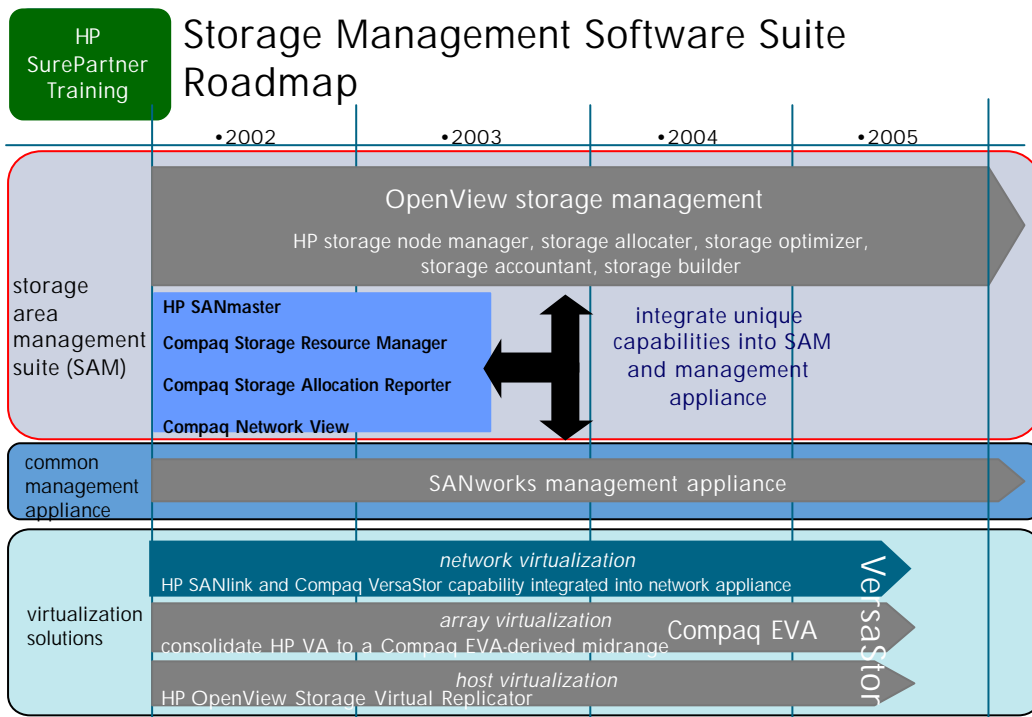
HP Products that are disk array specific will track the product lifecycles of its associated hardware, for example software that resides in firmware for XP arrays, VA, MA/EMA and EVA. Some functionality IP from both companies will be moved to the common management appliance. Example software includes CPQ EVM, DRM, Secure Path and from HP, HP Auto Path, and disk array HA solutions.

Provisioning & Planning

Compaq Provisioning Manager will become OpenView Provisioning Manager and new software services will be offered over time including OpenView Policy Manager, Capacity Planner and SAN Planner.

Array Software and Storage Solutions

All disk array software and solutions will track the natural lifecycle of its associated software. We will continue to offer and develop storage solutions across a variety of areas that are vertically integrated and span partner applications, infrastructure, deployment, servers including consulting and support professional services to meet the needs of the storage utility.



This is a migration slide showing how the products will migrate over time.

Storage Area Management Suite (SAM)

Compaq storage management IP integrates into hp OpenView Storage Area Management.

Common management application platform

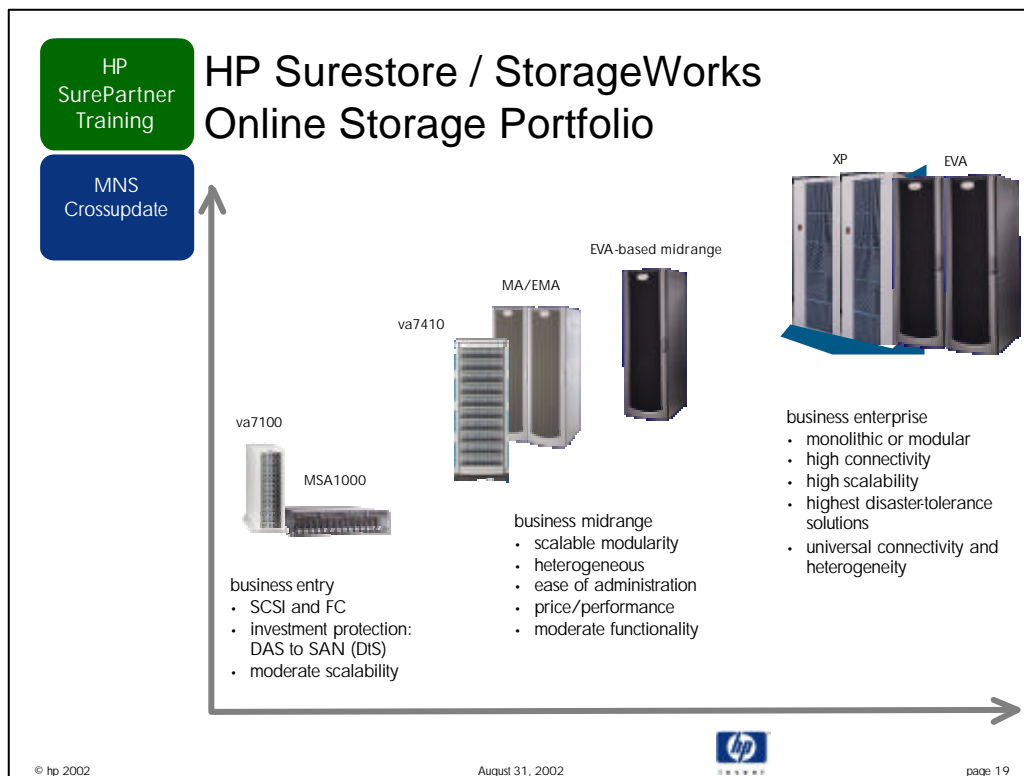
SANworks management appliance IP integrates with OpenView Storage Node Manager to provide a common management application platform sometimes called an appliance for all applicable hp storage (online storage, disk arrays, disk systems, NAS, extended and solutions)

Virtualization

Multi-levels of virtualization solutions will be offered by the new HP in three key areas: network, array and host.

We continue to ship HP virtual array, and target them at the HP-UX installed base. In about 18 months, applicable VA technology will migrate to EVA virtualization technology. We will continue to ship HP SANlink products until applicable functionality is merged with elements of VersaStor in the next 24 months to form a Network Virtualization Appliance which will include EVA technology as a base.

At the host level, we will continue to ship CPQ Virtual Replicator.



enterprise storage:

HP will provide enterprise monolithic and enterprise modular high-end storage solutions for maximum choice. These two enterprise storage solutions provide enterprise customers with two approaches to achieving storage consolidation, business continuity, storage efficiency and reduced TCO. The hp XP series, utilizes a consolidated monolithic approach for achieving high availability and mainframe connectivity. The hp Enterprise Virtual Array uses a highly scalable enterprise modular approach for high availability and heterogeneous connect by exploiting the storage network using virtualization capabilities for maximum storage efficiency and lower TCO.

The EMA16000 will continue to be available to customers. When we have comparable functionality between the EMA16000 and the Enterprise Virtual Array for example remote mirroring and specific OS support, you will be able to completely replace the EMA16000 with Enterprise.

midrange storage:

HP will continue to support the va7400 and will expand the family with new product enhancements. The va7400 platform is an ideal platform for HP-UX environments, but it also supports heterogeneous environments too. By mid 2003, the VA family can start to transition into a new cost reduced midrange platform with firmware derived from the high end Enterprise Virtual Array architecture.

The new hp will continue to offer the MA/EMA for those customers who value its set of specific functionality. By mid 2003, the MA/EMA can begin to transition to the new EVA midrange platform.

By mid 2003, HP will introduce a midrange disk array based on the Enterprise Virtual Array architecture that will be a replacement for both offerings.

entry level DAS:

In the new hp both products are available. HP will continue to offer the MSA 1000 for hp Proliant-based customers. The MSA1000 will be enhanced with additional operating system support, for example new releases of Windows, Netware, Linux and TRU64 and OpenVMS. HP will continue to support the VA7100 and will expand the product with enhancements.

The va7100 platform is an ideal platform for HP-UX, but also supports heterogeneous environments too. By 2004, as part of the natural technology lifecycle, the v a7100 will transition to a new redesigned entry-level array based on Enterprise Virtual Array derived firmware and specific software with a cost reduced hardware platform specific for the entry-level market. Over time, the MSA1000 will also transition to the new redesigned entry-level array based on Enterprise Virtual Array leveraged firmware and specific software.

internal RAID:

HP NetRAID products will transition to the Smart Array Family. Smart Array has a much more extensive product mix and a very large installed base.

HP
SurePartner
Training

MNS
Crossupdate

HP Storageworks Virtual Array 7100

- 1Gb/s native **fibre channel front and back**
- 1 host port per controller
- Single or dual controller
- 256-1024 MB memory per cache controller (upgradeable/replaceable)
- 15 high performance disk drives – 3U package
- 0.93TB usable w/73GB disks – 128 LUNS
- 93MB/s, 12K IOPs cache, 3.4K disk
- Enclosure Level RAID Configuration: RAID 1+0 or AutoRAID
- single redundancy group
- OS support: HP-UX 10.20/11.0/11i, Windows NT, Windows 2000, Red Hat 6.2 & 7.1, Netware 5.0/5.1, Solaris 7 & 8, AIX 4.3.3, MPE/ix, Tru64, OpenVMS

The va 7100 - for small to medium business-critical applications where affordability and simplified management are required and space is a premium.



First product of the new generation of hp virtual storage arrays , released April 2001.

- 15 low profile disks in 3Us - best in the industry
- Up to 1TB capacity with 73GB disks

Comments on heterogeneous os support:

- * Simultaneous access support – multiple servers share box, each with their own set of LUNs - this is hot! One fibre channel - many server types!
- * Support for up to 128 principles (servers) sharing the array


HP
SurePartner
Training

MNS
Crossupdate

HP Storageworks Virtual Array 7400

- **2Gb/s host fibre channel**, 1Gb back-end
- 1 host port per controller
- Dual controller only
- 512 - 1024 MB Memory per Controller (upgradeable/replaceable)
- **105 high performance disk drives**
- 7.3TB usable w/73GB disks – 1024 LUNS
- **180MB/s, 28K IOPs cache, 8.5K disk**
- Enclosure Level RAID Configuration: RAID 1+0 or AutoRAID
- Up to 6 Add-on Enclosures
- two Redundancy Groups
- OS support: HP-UX 10.20/11.0/11i, Windows NT, Windows 2000, Red Hat 6.2 & 7.1, Netware 5.0/5.1, Solaris 7 & 8, AIX 4.3.3, MPE/ix, Tru64, OpenVMS

The va 7400 - for medium to large business-critical applications in a shared heterogeneous data center where affordability and continuous operation are important.



Extends VA7100 architecture and adds:

- * Up to 105 disks
- * 2Gb front end
- * Increased IOPs and MB/s

Note:

The VA7100 cannot be upgraded to VA7400. But easy migration of disks from VA7100 to VA7400 is possible.


Continuous operation – Zero Downtime Backup with Business Copy VA and on-line scalability and mirrored solutions

Plus - optical distance between controller and add-on JBOD (500m) - don't need to reserve floor space for future growth

HP Storageworks Virtual Array 7410

- **2Gb/s host fibre channel, 2 Gb/s back-end**
- **4 host ports** / 2 host ports per controller
- Dual controller only
- 1024 - 2048 MB Memory per Controller (upgradeable/replaceable)
- **105 high performance disk drives**
- 7.3TB usable w/73GB disks – 1024 LUNS
- **330MB/s, 34K IOPs cache, 11K disk**
- Enclosure Level RAID Configuration: RAID 1+0 or AutoRAID
- Up to 6 Add-on Enclosures
- two Redundancy Groups
- OS support: HP-UX, Windows NT, Windows 2000, Linux (Red Hat, SuSE, Debian), Solaris, AIX, HP MPE/iX (secondary release), NetWare (secondary release)

The va 7410 - for medium to large business-critical applications in a shared heterogeneous data center where affordability and continuous operation and higher rate of host connectivity are important.

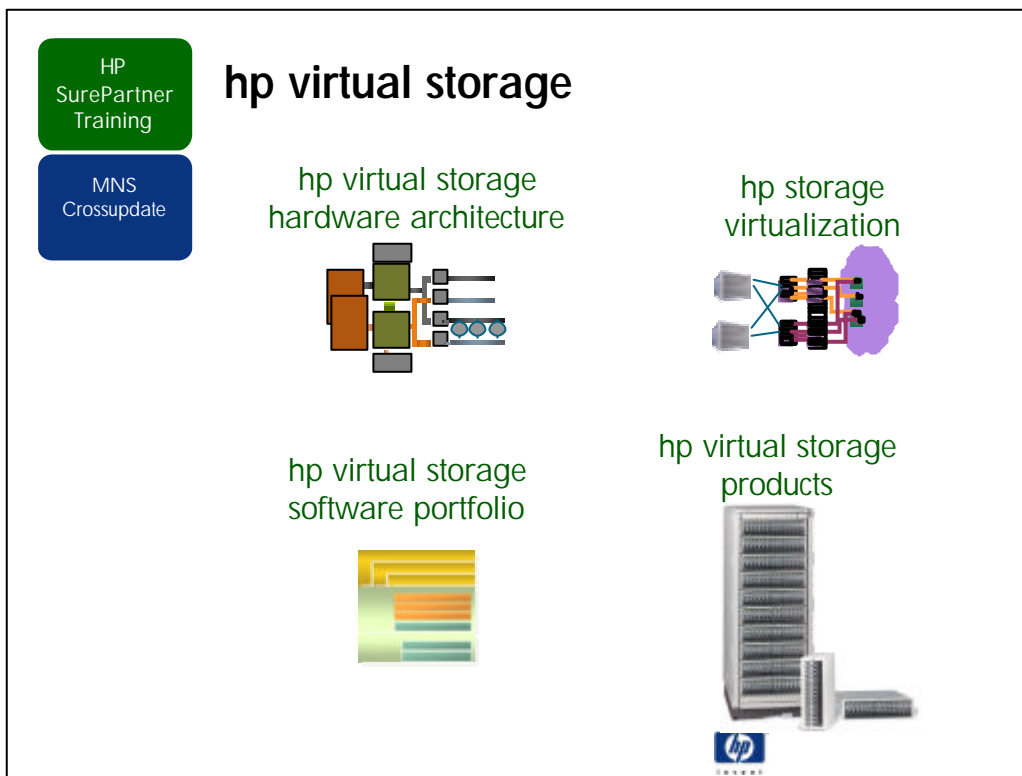


Extends VA7100/VA7400 architecture and adds:

- Additional host ports – a total of 4 ports, 2 per controller
- 2Gb back end
- Larger Cache
- Increased IOPs and MB/s

Note:

The VA7410 is only supported with DS2405 add on JBODs




These are the four major areas of HP's virtualization strategy:

- 1) value of virtualization
- 2) hardware architecture
- 3) Software
- 4) Products

They all participate in customer value - each is a key and unique contributor.

These topics will be discussed in detail in the following modules.

the value of virtualization



Availability-TOC-
Performance

the advantages

- high availability
 - ✍ quickly respond to changing capacity needs
 - ✍ virtual copies for recovery
- increased efficiency and lower cost of ownership
 - ✍ powerful - manage more storage per administrator
- optimized performance
 - ✍ adaptive RAID improves price-performance
 - ✍ traditional configuration mgmt.

the technology

- capacity is treated as a pool of storage blocks, not discrete disks
- powerful mapping techniques present a logical view of storage to host servers

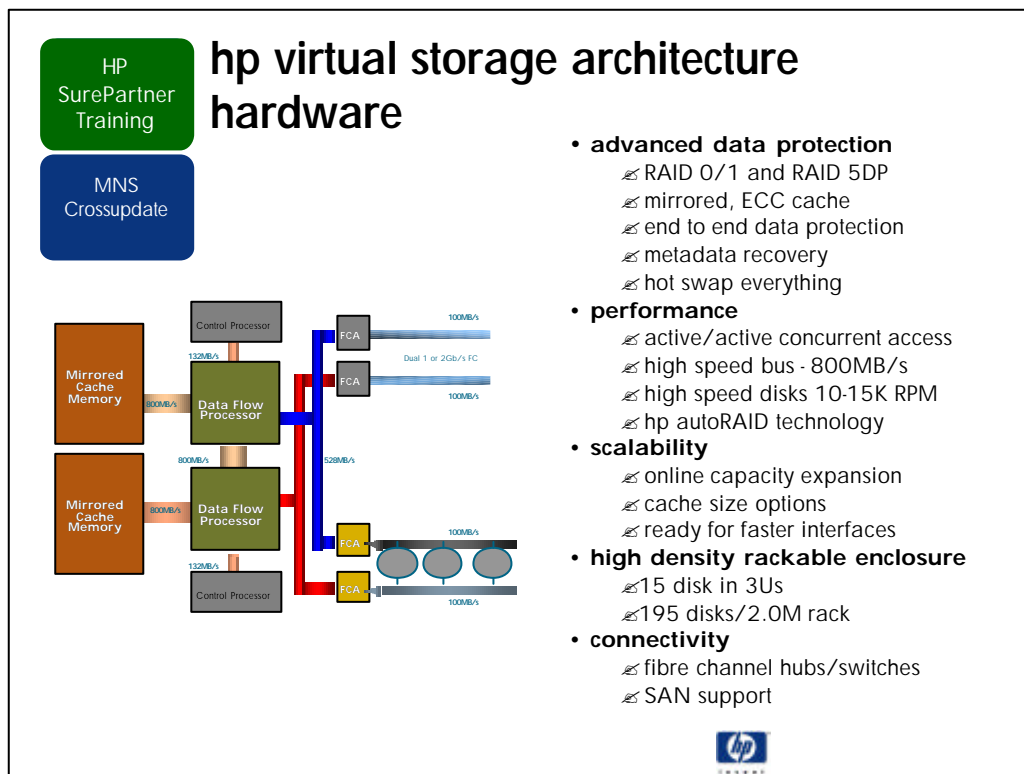


Main advantages for customers:

- High availability – upsizing capacity without array downtime due to reconfiguration, easy data recovery using Business Copies
- High ROI (Return on Investment) and reduced TCO (Total Cost of Ownership) enabled by powerful, but easy-to-manage user interfaces. Enables the customer to manage more storage per administrator
- Optimized performance by automatic configuration and RAID configuration, array adopts configuration automatically.

Main differentiator compared to standard array technologies:

Available capacity is treated as pool of blocks instead of disks. Data distribution mechanism is more efficient.



HP Virtual Array hardware is designed for virtualization and is extending the capabilities of mid-range arrays to state-of-the-art performance and security:

RAID 5DP - unlike standard RAID 5 which uses only one parity disks, RAID 5DP (dual parity) uses two parity disks. This provides far greater data protection than standard RAID 5.

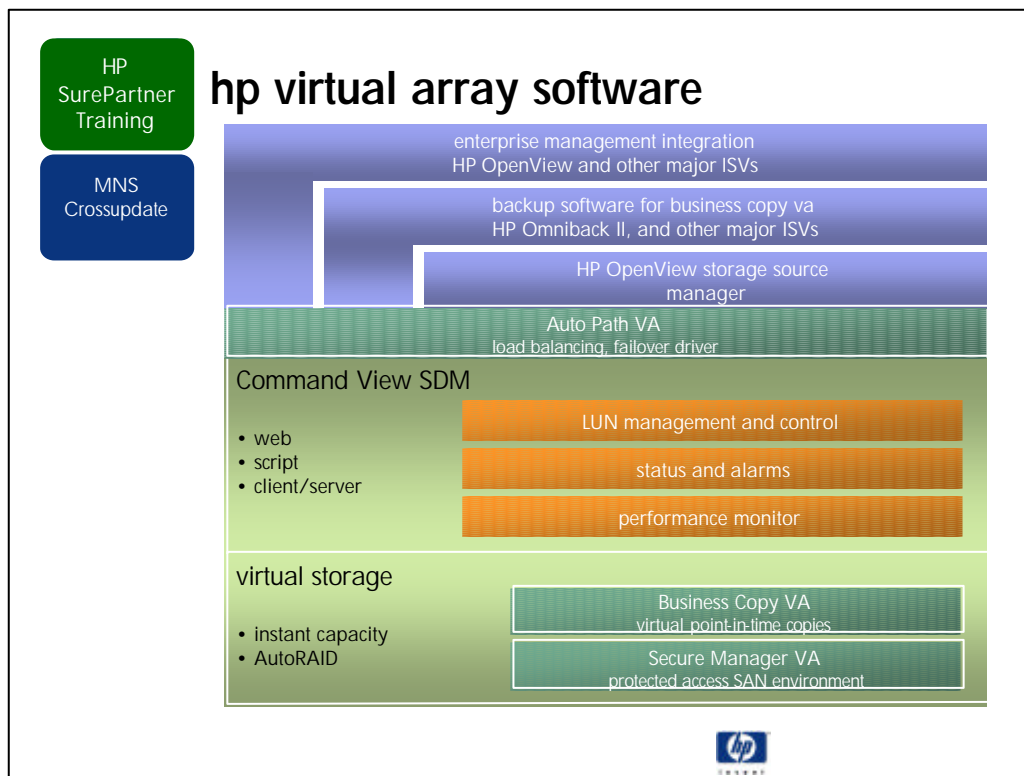
Mirrored memory – the VA architecture implements a tightly coupled mirrored memory design typically found in high-end monolithic arrays - this provides greater availability and higher performance in shared environments.

End-to-end data validation – standard typically only in high-end arrays - protects the data from data corruption.

Metadata recovery - a sophisticated process within the array that protects the virtual mapping information from failures - this technology is very similar to data base transactions and checkpointing - this is a must for true highly available arrays.

Hot Plug - hot swap - everything, anytime - virtual technology even allows disks to be replaced in any location, replaced with different size disks – thus eliminating a common data loss scenario.

Active/Active concurrent - feature typical of high-end monolithic arrays - any controller can access any data – and allows IO load balancing with hp autopath va software.



Virtual Array software portfolio:

command view sdm

manage anywhere/any way – GUI, browser based and scriptable
common look and feel across hp storage

auto path virtual array

HBA failover and load balancing, host based

business copy virtual array

instant local copy & backup
integration with leading 3rd party backup utilities

secure manager virtual array

secure shared storage by masking LUNs from servers by server WWN

hp virtual storage technology

autoRAID or manual control
add capacity instantly
flexible disk configurations

hp SAN integration

enterprise management

command view sdm



- common look and feel
- GUI, command line and APIs
- host or web-browser common look and feel GUI for modular scalable storage
- proactive monitoring and event notification
- active graphic and status at a glance displays
- performance monitoring and logging



Command View SDM

Features three different configuration interfaces:

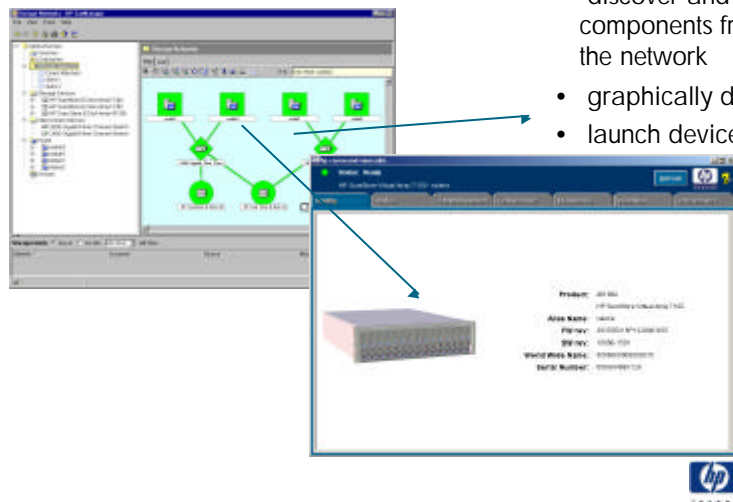
- GUI, installed on SAN host which is attached to VA (HP-UX, Windows or Linux)
- Web based (connects to Command View Server from any Java capable Browser)
- Command Line User Interface (CLUI), which enables powerful scripting capabilities.
Plus: CVUI – Command View User Interface, which is a text based menu system facilitating the use of the CLUI

Features event and performance monitoring and logging via GUI or command line.



seamless integration into hp openview storage node manager

- for larger SAN environments
- discover and manage san components from anywhere in the network
- graphically display topologies
- launch device management



Integration into SAN Management software

The HP Virtual Array products integrate seamlessly with HP Storage Area Manager software for even more powerful storage resource management on SAN wide level, enabling topology mapping and discovery, simple performance and capacity monitoring and trending based on historical data.

Along with powerful allocation and billing tools this product suite further enhances ROI and enables a more effective management of distributed or consolidated storage resources.

Key Customer Need: Solution Uptime

Solution:

- dual, hot swappable redundant fans, power supplies, and virtual array controllers
- no single point of failure
- end-to-end data integrity to provide protection against silent data corruptions
- mirrored, ecc protected battery backed up cache
 - battery is physically located on controller
 - real-time read compare operations and real-time consistency check
- zero-downtime upgrades
 - on-line firmware downloads for functionality enhancements
 - on-line capacity expansion means instant availability



By itself, the VA7x00 provides up to 99.995% uptime contribution from the storage (depending on the model of VA!)

High Availability solutions are available with HP MC/ServiceGuard and Microsoft Cluster Services

7x00's are affordable! So if you mirror them with the appropriate infrastructure (cables, switches, HBAs, OS, etc) you can achieve architectures that provide up to 99.999% uptime contribution from the storage!

Key Customer Need: Array Sharing

Need: To lower infrastructure management costs by sharing the array among heterogeneous operating systems

Solution:

- Simultaneous access for HP-UX 11.0 and 11.i, Windows NT, Windows 2000, Red Hat Linux 6.2 and 7.1, Solaris 7.0 and 8.0, AIX, True64, OpenVMS, MPE/ix and Netware.
- Managed locally, remotely, or integrated with your existing system management environment.



The virtual array product family provides true, *simultaneous access* for popular operating systems.

This is enabled by providing different host behavior modes (using a so called Host Behavior Table) on the same FC port.

In addition array-based LUN security (Secure Manager VA) enables exclusive resenatation of LUNs to a defined set of hosts, therefore restricting LUN access for servers of different operating systems.

Key Customer Need: Effective Management

Need: Cost of Management

Solution:

- surestore command view sdm
 - Tie-ins with industry standard management software
- Self management of RAID levels
- Self management of performance, hot spots
- Zero-downtime management
- Common look-and-feel across hp solutions



The virtual array product family provides a easy to use, easy to understand management software, enabling time effective configuration and managent of LUNs and value added array features.

The virtual architecture reduces management time by more than 80% using AutoRAID, AutoInclude, AutoFormat, AutoRebuild and Active Spare features of the array.

HP
SurePartner
Training

MNS
Crossupdate

module wrap-up





HP
SurePartner
Training

MNS
Crossupdate

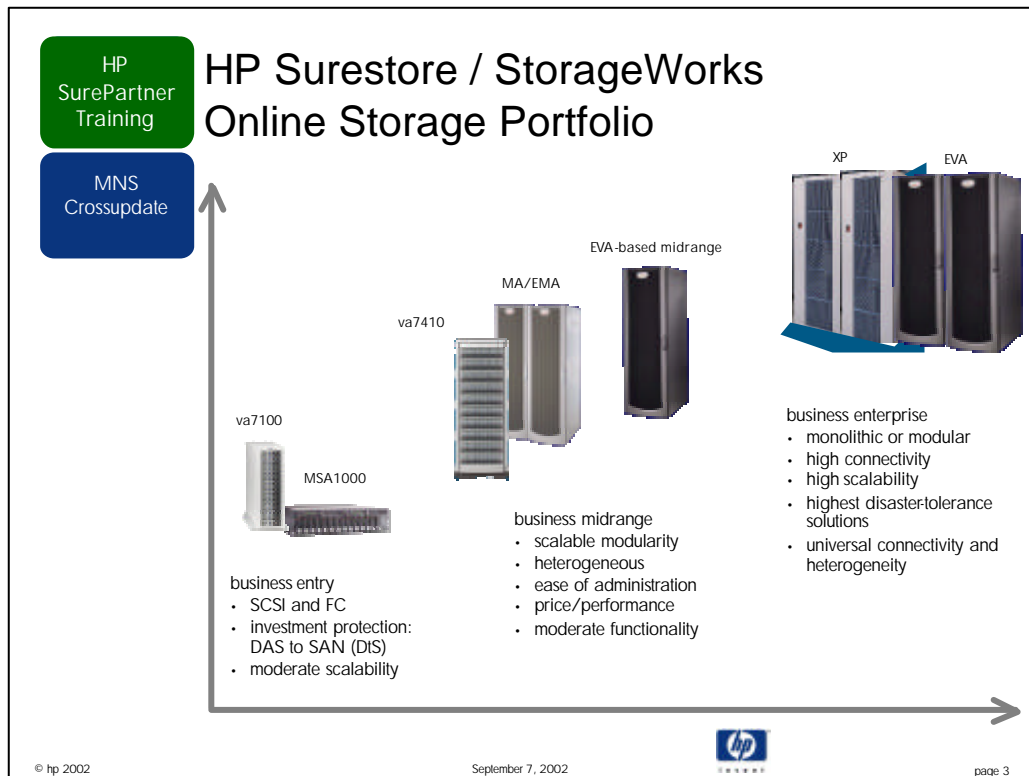
HP Modular Network Storage Solutions

Part 1

Product description



Student Notes:



Student Notes:

VA Family Positioning

	Capacity	Connectivity support	Other factors
VA7100	sub-terabyte	simple SAN or server connectivity	lowest entry cost
VA7400	multi-terabyte	simple or moderate SAN or server connectivity	User has inhibitors like internal qualification that prevent moving to 7410
VA7410	multi-terabyte	moderate or complex SAN or server connectivity	delivers highest performance



Student Notes:

VA Family Positioning

- Sequential Applications (data warehouse)
 - for optimal performance/\$, consider multiple arrays each with limited numbers of disks
 - Low cache, low RPM Drives
- Transactional Applications (OLTP,E-Commerce)
 - for optimal performance/\$, consider multiple arrays each with limited numbers of disks
 - High cache, high RPM Drives
- For lowest \$/MB, use va7410's with the maximum drive count. (up to 105 for 7.7 TBs)



Student Notes:



Virtual Array Family – VA7100

- VA7100 - current release
 - RAID 0+1 or AUTORAID
 - 128 LUNs in 1 Redundancy Group
 - 1Gb Fibre Channel Controller
 - Single or Dual Controller
 - 256 - 1024 MB Memory per Controller (upgradeable/replaceable)
 - 1 or 2 host ports



Student Notes:

Virtual Array Family – VA7400

- VA7400 - current release
 - Enclosure Level RAID configuration
 - RAID 0+1 or AUTORAID
 - 2Gb Fibre Channel Controller
 - 1Gb setting available for legacy HBAs
 - Dual Controller (Only)
 - 512 - 1024 MB Memory per Controller (upgradeable/replaceable)
 - 1024 LUNs
 - 2 host ports
- Up to 6 Add-on Enclosures



One port 2Gb or 1Gb FC connect from the host to the controller. The default setting for host connection will be 2Gb. 2 Gbit FC-HBA for HP-UX is available.

One port 1Gb FC connect to the add-on enclosures when using DS2400, 2 Gbit to the add on enclosures when using DS2405.

Also known as Cronus.



MNS
Crossupdate

Virtual Array Family – VA7410

- VA7410 - current release
 - Enclosure Level RAID configuration
 - RAID 0+1 or AUTORAID
 - 2Gb Fibre Channel Controller
 - 1Gb setting available for legacy HBAs
 - Dual Controller (Only)
 - 1024 - 2048 MB Memory per Controller (upgradeable/replaceable)
 - 1024 LUNs
 - 4 host ports
- Up to 6 Add-on Enclosures



One port 2Gb or 1Gb FC connect from the host to the controller. The default setting for host connection will be 2Gb. 2 Gbit FC-HBA for HP-UX is available.

One port 2Gb FC connect to the add-on enclosures, must use DS2405.

Disk firmware must be updated to the latest revisions to ensure correct operation at 2Gbps. Latest revisions are:

A6191A--HP05

A6192A--HP08

A6193A--HP05

A6194A--HP08

Virtual Array Capacities

VA7100

- 4 to 15 Disk Drives
- 18GB 15K RPM
 - Scalable Capacities from 30GB to 196GB
- 36GB 10K RPM
 - Scalable Capacities from 66GB to 398GB
- 73GB 10K RPM
 - Scalable Capacities from 130GB to 804GB

VA7400 / VA7410

- 10 to 105 Disk Drives
- 18GB 15K RPM
 - Scalable Capacities from 80GB to 1374GB
- 36GB 10K RPM
 - Scalable Capacities from 163GB to 2789GB
- 73GB 10K RPM
 - Scalable Capacities from 130 GB to 5625GB

Technical gigabytes computed as 1024^3 – usable capacity

VA 7410 is not orderable with 18GB drives, however it does support 18GB drives.



Marketing gigabytes computed as 1000^3 – Raw:

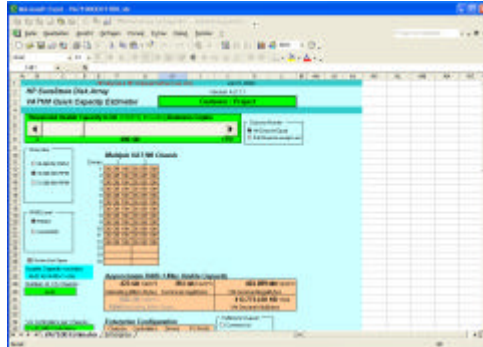
VA7100:

- 18GB 15K RPM scalable capacities from 72GB to 269GB
- 36GB 15K RPM scalable capacities from 143GB to 538GB
- 73GB 10K RPM scalable capacities from 287GB to 1075GB

VA7400/7410:

- 18GB 15K RPM scalable capacities from 86GB to 1882GB
- 36GB 15K RPM scalable capacities from 358GB to 3763GB
- 73GB 10K RPM scalable capacities from 130 GB to 7526GB

Virtual Array Capacity Calculation



Use Virtual Array Quick capacity estimator, a excel spreadsheet for each the VA7100 and VA7400/7410 being available on the ESP website

NOTE:

Technical gigabytes computed as 1024^3 – usable capacity - that's what the customer is asking for

Marketing gigabytes computed as 1000^3 – usable capacity – that's what drive capacities are alike!



A fast method of calculating the appropriate number of drives and enclosures depending on enclosure RAID level and drive sizes.

Available for Enterprise Partners via ESP website.

NOTE:


Technical gigabytes computed as 1024^3 – usable capacity - that's what the customer is asking for

Marketing gigabytes computed as 1000^3 – usable capacity – that's what drive capacities are alike!

Student Exercise 1: Virtual Array Capacity Calculation

Using the Virtual Array Quick capacity estimator, identify the correct number of disk drives for a customer demanding:

- a) 700 GB capacity, customer demanding box level RAID 0/1 and small and fast hard disks
- b) 1,3TB capacity for a file server in the most cost efficient configuration
- c) 670 GB capacity shared by 20 hosts distributed across about 250 LUNs with an average IO rate of more than 51,000 IO/s



Student Solutions:

a)

b)

c)

Virtual Array Availability

- 99.995% system uptime for va7410 single enclosure configurations
- 99.95% system uptime in configurations with add-on JBODs (DS2405)
- 99.95% for a VA7400
- End to end data integrity
- Mirrored write cache & maps
- Three and one-half days Battery
- HP High Availability Observatory
- On-line controller firmware download
- Active Hot Spare
 - Spare is a pool of space located on all drives
- Auto Configuration
 - On line capacity expansion
 - Capacities can be mixed
- Redundant Power Supply/Fan Module



Advertising 99.995% up time

Data integrity checks/alerts the problem; it doesn't fix the problem.

Battery will put out a warning when low, and end of life (ready to fail)

The battery life estimate is a conservative 3.5 day based on worst case scenario.

Active hot spare will save the space of up to the two largest drives depending on user settings.

Multi Operating System Support

- HP-UX 10.20. (only VA7100 and 7400), 11.0 and 11.i
 - Root, Boot, MC ServiceGuard, LVM, AutoPath VA
- NT 4.0 - Enterprise Edition Servers
 - Cluster Certification (2 node), AutoPath va, and Boot
- Windows 2000 - Server, Advanced Server
 - AutoPath VA, Cluster Certification (2 node), and Boot
- Linux Red Hat 6.2 & 7.1
 - Boot, AutoPath VA
- Solaris 7 & 8
 - Veritas volume manager / DMP, veritas VCS
- AIX 4.3.3 and 5L
 - AutoPath for AIX
- Netware 5.0/5.1 and 6.0 (only VA7100 and 7400)
 - No fail over driver, Non Boot
- Tru64 (only VA7100 and 7400)
- MPE/ix (only VA7100 and 7400)



A management station is required with Solaris, Netware, Tru64, MPE and AIX connections to run command view sdm.

Cluster Certification testing has completed NT 4.0 and 2000. The VA 7x00 has received Microsoft Windows 2000 Server Logo and Microsoft Windows 2000 Server Logo Designation on the Hardware Compatibility List (HCL) located at <http://www.microsoft.com/hwtest/hcl>.

Multiple cluster support for MSCS added starting with firmware HP14.

HP
SurePartner
Training


MNS
Crossupdate

Switch Support

- HP 16B
- HP 8B
- Brocade 2800
- Brocade 2400
- HP 6164
- FC 64



These are high level support possibilities, always look to the documentation for the complete details on supported configurations.



VA Family Product Structure

Enterprise	Commercial	Description
• A6261A/AZ	A6261P	VA 7100 w/ Dual Controller, 256 MB Cache
• A6262A/AZ	A6262P	VA 7100 w/ Dual Controller, 512 MB Cache
• A6263A/AZ	A6263P	VA 7100 w/ Dual Controller, 1024 MB Cache
• A6267A/AZ		VA 7410 w/ Dual Controller, 1024 MB Cache
• A6268A/AZ		VA 7410 w/ Dual Controller 2048 MB Cache
...		

Let's have a first look at the configuration and order guides....



Use the Configuration and Order guides to determine the proper and supported configurations and product numbers.

These are chapters of the HP9000 Configuration and Order Guide available on the ESP website.

HP
SurePartner
Training

MNS
Crossupdate

VA Configuration and Order Guide



The HP StorageWorks Virtual
Array 7000 Family
Configuration Guide is
available via the ESP and is
your central point of reference
configuring VA solutions



Your instructor will give you a guided tour through the VA Configuration and order guide.



Student Exercise 2: VA Family Disk Modules



Student Exercise:

**Using the configuration and order guide,
find out the available disk modules (size, speed,
part-no.) for the HP storageworks virtual array
family**



Student Notes:

HP
SurePartner
Training

MNS
Crossupdate

hp virtual array value-added software

- high availability, data security, and on-line backup for mscs-ux, windows nt, windows 2000 and linux!



- with virtual array software you can have:

1. high availability
2. data security, and
3. online backup

for:

- mscs-ux 11.0/11i
- windows nt 4.0
- windows 2000, and
- linux red hat 6.2/7.1

plug into the
enterprise

- openview nnm
- ca-Unicenter-tng
- topTools

common
management

- gui
- clui



we have lots more information, including:

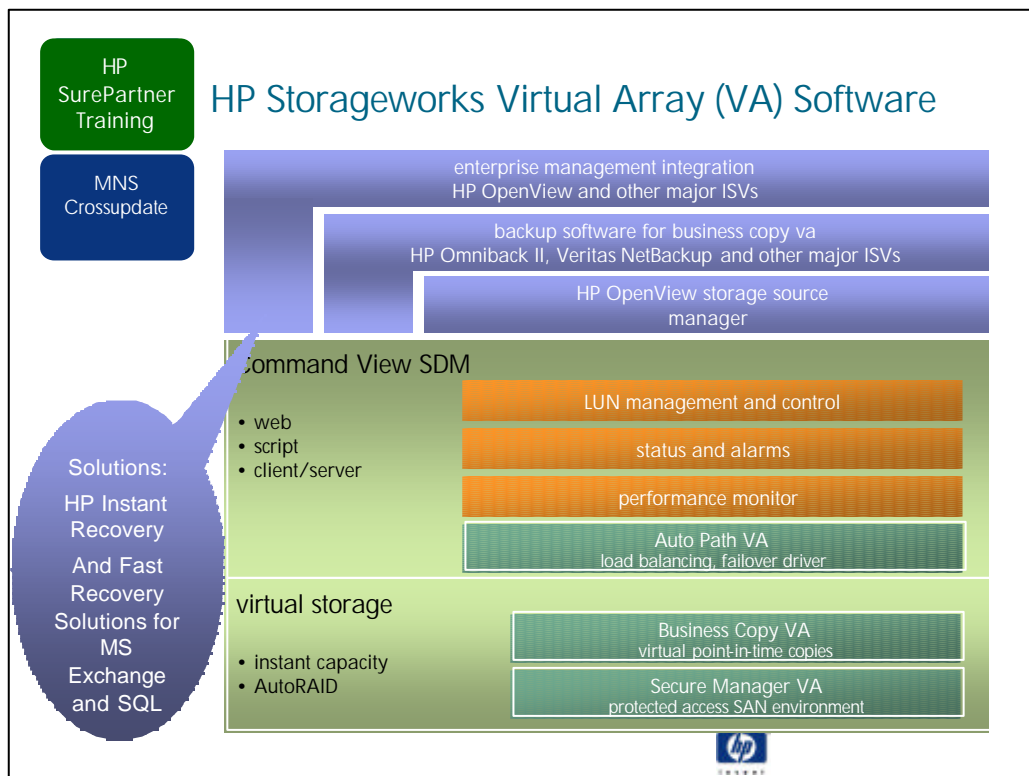
pricing

competitive analysis versus emc and others

feature details

to get more information go to the *ESP website*

Feature details can also be found in the appendix of this student material.



Virtual Array software portfolio:

command view sdm

manage anywhere/any way – GUI, browser based and scriptable
common look and feel across hp storage

auto path virtual array

HBA failover and load balancing, host based

business copy virtual array

instant local copy & backup
integration with leading 3rd party backup utilities

secure manager virtual array

secure shared storage by masking LUNs from servers by server WWN

hp virtual storage technology

autoRAID or manual control
add capacity instantly
flexible disk configurations

hp SAN integration

enterprise management

hp virtual array Value-Added Software

- command view sdm
 - Enables array Device Configuration and Management, and Value-Added Software Products
 - a single host license-to-use and software media kit ships free with each virtual array
- business copy virtual array
 - Enables online data replication or LUN copying within the array for testing and backup, and requires the same physical space to be available in the array as the LUN(s) being copied
- secure manager virtual array
 - Enables LUNs to be locked into a secure shared heterogeneous environment



Remember, always sell Virtual Array software products when selling the Virtual Array.

- Ask customers what their Data Security strategy is and sell Secure Manager VA
- Ask customers what their backup strategy is sell Business Copy VA

hp virtual array Value-Added Software

- auto path virtual array
 - Enables I/O path fail-over in MSCS Windows 2000 and Windows NT, clustered environments with the benefit of I/O load balancing in both failed and non-failed states
- enterprise management smart plug-ins for command view sdm
 - Enables command view sdm in HP Openview NNM for HP-UX, Windows 2000 and Windows NT 4.0, also enables command view sdm for HP TopTools 4 and 5 (5.5 already has the SPI included), CA-Unicenter-TNG, Tivoli Netview and BMC Patrol



Remember, always sell Virtual Array software products when selling the Virtual Array.

- Ask customers what their High Availability strategy is and sell Auto Path VA
- Ask customers if they have a Network Management Application and sell the Smart Plug-Ins

hp surestore command view sdm

- Host or web-based device management application that provides a common user interface for modular networked storage
- Monitors and manages modular scaleable storage resources from a single management console with centralized control of their information resources
- Integrated device management platform for all modular scaleable storage starting with the virtual array product family
- Common look and feel graphical user interface (GUI)



Command view sdm will be the standard device management software platform to configure and diagnose and launch value-added software for MNS arrays and disk systems.

One (1) host license is free with each virtual array, each additional 1 Host license is only \$530 (enterprise) and \$400 (commercial).

Command view sdm supports only modular networked storage. This product does not support xp arrays. Likewise, command view xp does not support modular networked storage.

Command view sdm supports the virtual array 7100, virtual array 7400, DS 2300 and DS2405. New modular networked storage products will be supported as they are released.

Command view sdm will not support the DS2100 disk system because it does not support ems or safe-t. This is required to monitor power and cooling.

hp surestore command view sdm

- Comprehensive user interfaces (web browser, GUI, CLUI, CVUI)
- Native OS support on
 - HP-UX 10.20, 11.0 and 11.i
 - Microsoft Windows NT 4.0,
 - Microsoft Windows 2000
 - Linux Red Hat 6.2 and 7.1
- Integration with MSCS systems, SANs, network and 3rd party management applications
- Historical performance logging and graphing



Command view sdm is not supported on Solaris, Netware, AIX, MPE and True64 hosts, a management station is required for array management functionality in these environments.

hp surestore business copy virtual array



- Online data replication and backup software.
- Creates an online copy of LUNs that can be dedicated for backup or testing environments.
- Integrated with command view sdm storage device management software.
- Integrates into backup software applications that support a custom scripting implementation.
- Comprehensive user interfaces (web browser, GUI, CLUI, CVUI).
- Functionality of the firmware means Native OS support everywhere the array is supported.
- 50GB Buy & Try Version Available.



Business copy is functionality provided by the array firmware.

All the features to manage business copies are provided through the command view sdm interfaces.

The maximum number of business copy LUNs is one less than the maximum number of LUNs for the array (127 for the VA7100 and 1023 for the VA7400).

The media kit includes the user's guide and a CD ROM.

There are utilities that come on the CD to make mounting and unmounting easier on Windows platforms.

Business copy va is the only array-based data replication solution for modular arrays.

Full support with Zero Downtime Backup and Fast Recovery with Omniback II 4.1



hp surestore secure manager virtual array

- LUN security for shared homogeneous and heterogeneous environments.
- Allows LUNs to be locked in a secure shared environment.
- Integrated with command view sdm storage device management software.
- Comprehensive user interfaces (web browser, GUI, CLUI, command view user interface).
- User definable security levels by LUN (read/write, no access, or write/configure).
- Supports 128 secure principals.
- Supports 1024 secure LUNs for each array
- Functionality of the firmware means Native OS support everywhere the array is supported.
- Optional 50GB demo version is available.



Secure manager va functionality is provided by the array firmware.
All the features to manage secure LUNs are provided through the command view sdm interfaces.

128 principles (hosts) are securable on the VA7x00.

hp surestore auto path virtual array



- Fail-over and load-balancing drivers.
- Provides multi-path fail-over and load-balancing of HBAs in a single server and clustered environment.
- GUI user interface for quick set up and management.
- Enables fail-over of all HBAs within a server.
- Enables data load-balancing of all HBAs within a server.
- Native OS support on Microsoft Windows 2000, Microsoft Windows NT, HP-UX and Linux
- MSCS certification allows only for "no loadbalancing".



Auto Path VA is a host based software being available as single host or 5 host licenses.

Supported on Windows, where it is mandatory for several configurations.

Provides add on functionalities for HP-UX to enable dynamic load balancing.

hp surestore enterprise management smart plug-ins for command view sdm



- Allows command view sdm to integrate with:
 - hp openview nnm for HP-UX.
 - hp openview nnm for Microsoft Windows NT 4.0.
 - Hp openview nnm for Microsoft Windows 2000.
 - hp topTools 4.5 and 5.0.
 - CA Unicenter-TNG for HP-UX, Windows NT4.0 and Windows 2000.
 - Tivoli Netview
 - BMC Patrol
- Includes/supports HP SEMI 1.0 agent for HP modular storage and tape devices.



Enterprise management smart plug-ins are only supported on command view sdm for modular networked storage. They are not support on xp arrays.

Enterprise management smart plug-ins support hp openview network node manager. Full integration includes also performance monitoring with vantage point.

Does not support linux red hat 6.2 or Solaris.

Customers can integrate the virtual array into network management applications using standard SNMP scripting without purchasing our smart plug-ins.

hp surestore enterprise management smart plug-ins for command view sdm



- Integration means:
 - Automatic discovery of virtual arrays by the network management application.
 - command view sdm launching from the network management application.
 - Device event communication via SNMP to network management applications.
 - Changes in the status of the device are represented via the device icon which can launch the device software for further diagnostics or configuration status.



Student Notes:

HP
SurePartner
Training

MNS
Crossupdate

module wrap-up






HP
SurePartner
Training

MNS
Crossupdate

HP Modular Network Storage Solutions

Part 2

Operating Features



Student Notes:

Operating Features Topics

- Virtualization Concepts and Features
- Resiliency Settings
- Topology Settings
- Host Port Behavior
- Redundancy Groups
- Understanding Virtualization and Policies



Student Notes:

Virtual Array Objectives

- Lower Management Costs
 - Storage capacities are increasing at a faster rate than skilled storage administrators.
 - Products that allow significantly fewer storage administrators will have a competitive advantage.
- Provide Superior Price / Performance
 - Scale the quantity of arrays to meet any absolute performance requirements
 - Network modular arrays will eventually be the dominant architecture.
- Match or Beat the Competition
 - Provide superior availability and software features for mid-range arrays




Student Notes:



Fundamental Principles for AutoRAID

The cost, performance and availability characteristic vary by RAID level. A dynamic system, like a cache hierarchy, can provide performance characteristics similar to the high performance component of the hierarchy at a cost similar to the lowest cost component of the hierarchy.

Locality of Access - recently accessed data is most likely to be access in the future



- RAID 0+1 : *best for OLTP/general file systems*
 - superior performance for small (<64K) random write workloads
 - most redundancy overhead
- RAID 5DP : *best for OLAP and rich media*
 - highest availability
 - least redundancy overhead (>5 disks)
 - superior (theoretical) sequential read and write performance
 - lowest random write performance
- Either RAID level
 - equivalent performance for random read workloads



The technology of AutoRAID is about 10 years old. First introduced in a product call AutoRAID 12H.


The base problem for customers is the triangle of cost, security and performance and finding a proper tradeoff between these aspects. RAID0 uses every disk without any overhead, you can stripe to achieve performance, but you will have no redundancy. Mirroring is costly because of 50% overhead. With RAID5 or RAID 5DP, you have a fairly good cost. With RAID 5DP you really need at least 6 drives for cost effectiveness. In regards to performance RAID 5DP has a read modify write penalty for two parity blocks per stripe.

The architects solved this by allowing you to have cache. In addition, AutoRAID takes uses available disk storage as a pool of storage. The more active data is stored in RAID 0+1, but less active data is moved to RAID 5DP, dynamically.

RAID 5DP data is managed in an AutoRAID (log structured) manner. This does not update data in-place, but rewrites data to free space. The parity is cached until the end of the stripe is reached.

Dynamic Data Migration (1 of 2)

- Migration between RAID levels is automatic and based on workload
- Controllers are RAID experts, constantly monitoring array performance
- New data typically written to high performance RAID Level 0+1
- Maximum amounts of data are stored in RAID 0+1 to optimize performance
- As data levels increase, less active data migrated to RAID 5DP
- Data is stored redundantly in each case



Dynamic Data Migration is the process of moving data from RAID 0/1 to RAID 5DP and vice versa depending on the array's available capacity and workload.

Basically, data is always written to cache and as the RAID controllers are monitoring performance and array capacity, data will then be written to the disks, typically using RAID 0+1. As long as used capacity is less than 50% of available raw capacity (after deduction of spare capacities) all data will be placed in RAID 0+1. As the array is filling up with data, more and more (inactive) data will be moved into RAID 5DP therefore providing a higher usable capacity.

Dynamic Data Migration (2 of 2)

- Never less than 10% RAID Level 0+1 storage
- Data migration is a background event
- Write Working Set is maintained in RAID 0+1 whenever possible
- Migration limits may also cause data to write directly to RAID 5DP




When an array is formatted, (formatting map, or pointers to the storage) any writes are all written in RAID 0+1. As already mentioned, when the enclosure starts filling up, the controllers will start migrating data to RAID 5DP, giving more capacity.

When the capacity is full, 90% of the data is residing in RAID 5DP. Performance is maintained because the controllers write everything in RAID 0+1 format.

Auto Configuration – Online Capacity Expansion

- Hot plug new disks on-line
- No need to reconfigure the system
- No reloading of data
- Current data automatically balances across all disks, including new ones
- Use new capacity to create new LUN, or to increase RAID 0+1 storage



Student Notes:

Auto Rebuild


- Auto Rebuild allows the array to automatically begin rebuilding a failed disk drive.
- A rebuild priority setting allows you to balance array performance with rebuild speed.
- If a disk fails with Auto Rebuild enabled, the array immediately begins rebuilding the contents of the failed disk on the Active Spare (or other available unallocated capacity).
- To perform a rebuild (automatic or manual), the array must have adequate available capacity. The best way to make this capacity available is by enabling the Active Spare. This ensures that the array always has enough capacity to rebuild even the largest disk in the array.



If you want to have more control over the rebuild process, you can disable Auto Rebuild and start a rebuild manually. This allows you to control when the rebuild is performed, but it can leave your data vulnerable to a second disk failure until the rebuild is complete.

Active Hot Spare (1 of 2)

- Use in conjunction with Auto Rebuild for maximum data protection
- When enabled, space is reserved for rebuilding the largest disk drive
- No disks remain idle, all disks are used to increase performance




Instead of setting aside a spindle, the space of a disk is set aside across all available spindles. The spare space reserved is then built in if a drive fails. This provides a better performance picture since all disks are used to increase performance.

The active hot spare default setting is on, it may be with 4 drives that active hot spare is turned off by default.

In traditional arrays the hot spare disk itself may be prone to undetected failure since it's not in use.

Active Hot Spare (2 of 2)

- Reserved space is spread across the disks
- Reserved space is utilized as additional RAID 0+1 space for further performance increases
- If a rebuild is needed, remaining reserved space is used and RAID 5DP storage is increased



Student Notes:

Resiliency Settings

- Single Controller
 - Used to suppress single controller warning messages, when the configuration is intended.
 - Write cache is turned off, all writes are posted to RMAPs.
- Secure
 - Continually updates disk at a regular interval.
 - Write cache is turned off, all writes are posted to RMAPs.
 - Single controller warnings are not suppressed.
- Normal
 - Higher performance than Secure.
 - RMAPs are updated at 4 second intervals.
 - Does not honor FUAs.



RMAPs – Resiliency Maps - are stored on the image disks in the array.

Forced Unit Access (FUA) - used to indicate that the device shall access the physical medium. For a write operation, setting FUA causes the device to complete the data write to the physical medium before completing the command. For a read operation, setting FUA causes the logical blocks to be retrieved from the physical medium.

Topology

- Private
 - 8 bit addressing private loop topology.
 - AL-Port, Quick-Loop.
- Public
 - 24 bit addressing public loop topology.
 - NL-Port.
- Fabric
 - Public Direct Fabric Attach, Direct Fabric Attach.
 - N-Port.
- Point To Point - not currently supported.
 - Future testing may allow official support in existing virtual array installations.



Topology options can be configured in three different ways:

CLUI

`armmgr -y {Topology} <controller> <array-id>`

Configuration Tab in command view sdm

Virtual Front Panel

`vfpmgr -t <Topology> -c <Controller ID>`

Host Port Behavior

- Configures the controller to communicate with the correct FC “dialect”
- Set globally per controller
 - The default setting is HPUX
- Set specifically per Host node WWN
 - Needed for heterogeneous configurations
 - Demands Host Behavior Table



Default behavior is set on FC-port level.

For heterogeneous environments the behavior must be set for every server running an OS different from default behavior.

This is done using the “armhost” command. This command sets the host port behavior by HBA Node WWN.

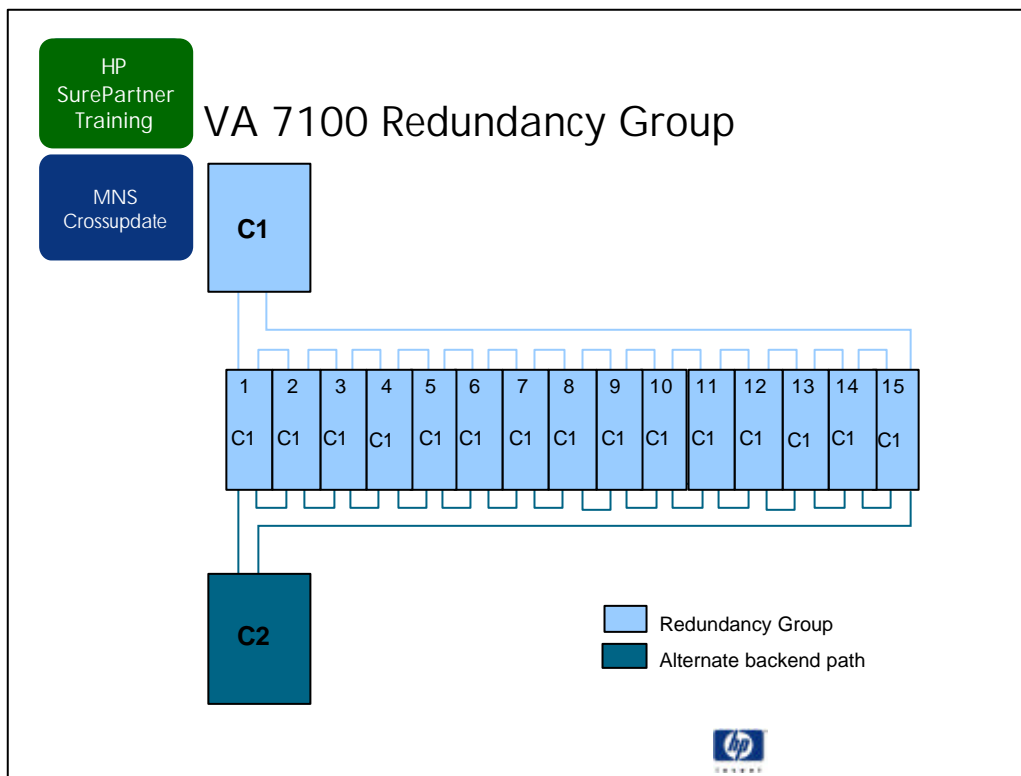
This demands the creation of a so called host behavior table, which has then to be uploaded to the array.

On-line Firmware Download

- Data still running
- Pause in I/O is 10-15 seconds
- Disk drive download is done off-line
- Software and Disk drive firmware downloads are independent of each other

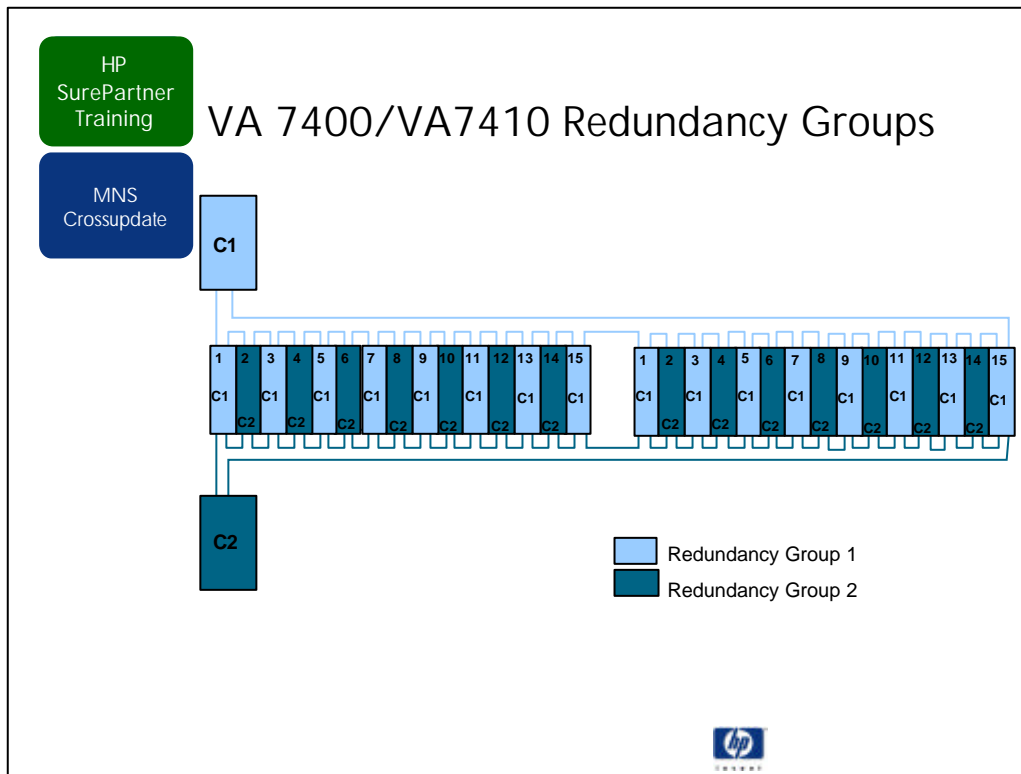


Student Notes:



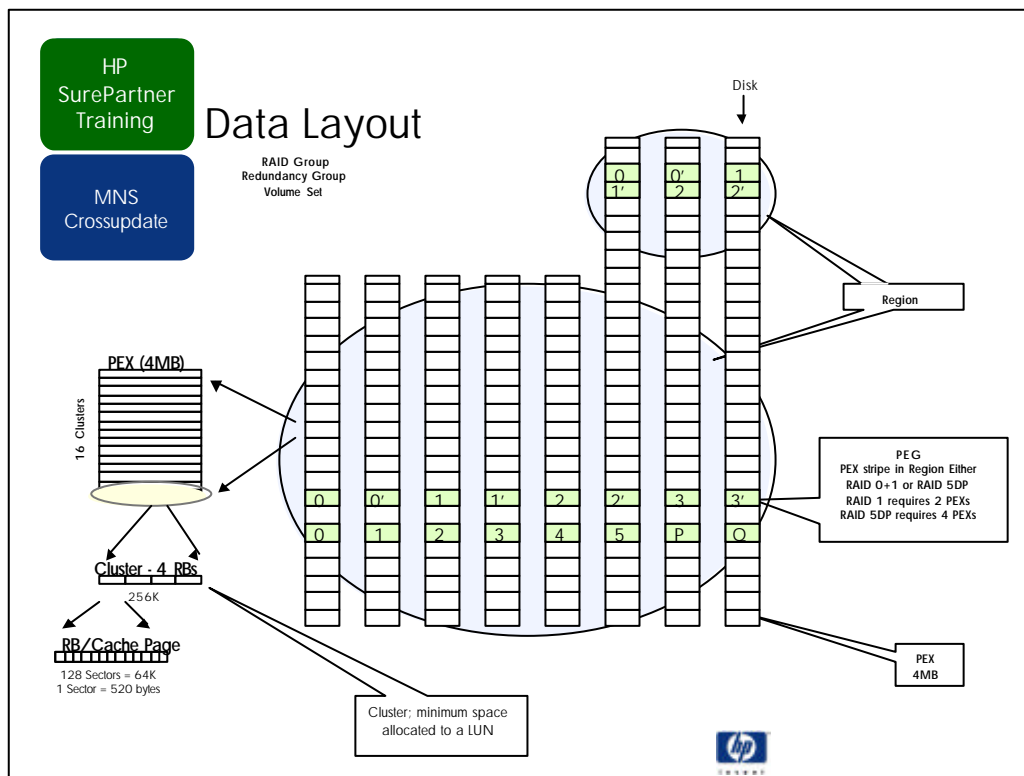
This drawing shows the logical backend I/O path of a VA 7100. If there is any kind of failure with communication to a drive, the drive is bypassed. When you create a LUN you'll assign a LUN number and capacity. The C2 path to the disks is only used when controller C1 fails or is removed.

Minimum 4 drives, maximum 15 drives.



This drawing shows the main and one add on enclosure. The even drives are bypassed on C1's backend loop. They do not delay the frame getting to the odd drives. Same for the other "loop". If there is any kind of communication failure to a drive, the drive is bypassed. To create a LUN assign a LUN number, capacity and RG.

Minimum of 10 drives, maximum of 105.



Physical Space

Each disk is broken up into 4MB spaces called Physical Extents (PEX).

PEXs are broken into 256K segments called a cluster.

PEXs are grouped into Physical Extent Groups (PEGs).

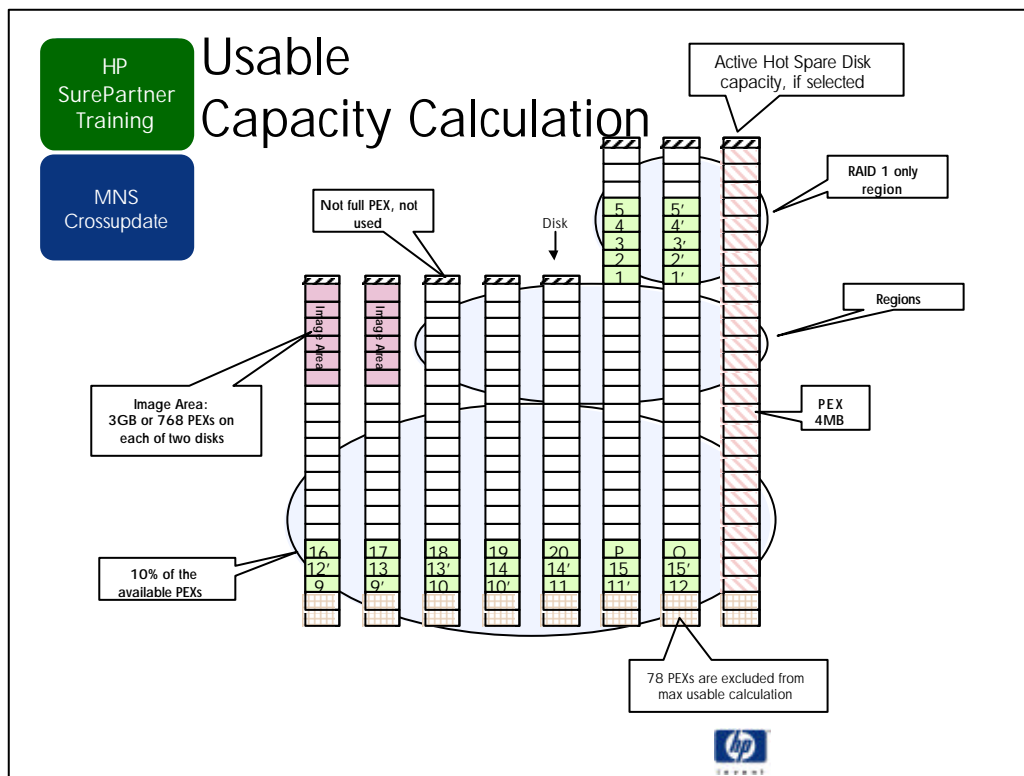
PEGs contain no more than one PEX from each disk.

PEGs are either RAID 0+1 or RAID 5DP

Logical Space

Unlike the 12H which used 64K relocatable blocks (RB), the VA 7x00 arrays use 128K relocatable blocks.

The controllers manage all data mapping, moving RBs to and from RAID levels.



This is only a representation of the capacity, the hot spare space is distributed across all the disks in the array.

Only full PEXs are used (4MB).

2 Disks have 3GB (768 PEXs) reserved for maps/metadata

78 PEXs are reserved as free space = 327.2MB.

A minimum on 10% of total available PEXs are computed as RAID 0+1 capacity. This is rounded up to the next full PEG.

RAID5 DP PEGs require a minimum of 4 PEXs for the stripes.

RAID 0 + 1 PEGs require 2 PEXs (single PEXs are not used).

Active Hot Spare settings removes the disk space of 0,1 or 2 disks from PEGs for each region.

Approximate Disk Capacities:

ST318451FC 18GByte/15Krpm 35000001 (decimal 520 byte sectors)

ST336704FC 36GByte/10Krpm 70000001

ST373405FC 73GByte/10Krpm 140000001

The 'Real' Data Layout

- It's Dynamic - based on workload and rules (policies) of the game
- VA7100 has 1 Redundancy Group (RG),
VA7400/VA7410 has 2 RGs
- Disks are assigned to RG by initial enclosure slot location - odd/even
 - once assigned to RG, disk can be moved to new slot, but will not change RG
- LUNs will be assigned to a RG at LUN creation time
 - each RG is like a separate array



Student Notes:

The Basic Game

- LUN creation reserves space from available clusters
- 'new writes' (space not previously allocated) cause allocation of physical space (cluster)
- array fills with new writes
- reads do not cause data movement – read in place
- re-writes always occur in place (to the same cluster / RAID level)
- policies move clusters between RAID levels or disks

performance - availability - cost



Only when a write is done, the PEG is actually created.

Policies

- **Data Placement** – *the rules on where data is written*
- **Makespace** – *creates more free-space to allow a new-write to complete*
 - foreground conversion of data from RAID 1+0 to RAID 5DP
- **Optimize** – *creates more free-space for future new-writes or promotions by moving inactive write data to R5*
 - background conversion of data from RAID 1+0 to RAID 5DP (not implemented in VA7100)



Student Notes:

Policies Continued


- **Promote** – *get clusters that are recently written randomly to R1*
 - background conversion of data from R5DP to R1
- **Balance** – *re-distributes data after new disks are added*
 - Extend PEGs
- **Scrub** – *tests and corrects latent un-recoverable disk errors*
 - Read every sector
- **Rebuild** – *removes failed disk from RG*
 - Shrink PEGs

Sector – 512 bytes
RB/Cache Page – 128 sectors, 64K
Cluster – 4 RBs, 256K
PEX – 16 Clusters, 4MB
PEG – stripe of PEX's



Student Notes:

Array Performance Utilization Calculation



- The current array performance utilization is used as triggers for *optimize*, *promotion*, *rebuild*, *balance* and *scrub* - only under specific utilization will these policies run
- The controller collects backend workload statistics for the last minute of operation; reads, writes and data transferred
- Utilization is computed (actual/max theoretical)
 - Throughput (MB/s) utilization is the ratio of actual data transferred vs. 90MB/s (per RG)
 - Transaction (IOPs) utilization is the ratio of actual IOs vs. computed capability of disk subsystem based on transfer size



Student Notes:

Cache

- Separate read and write cache
 - Write cache is kept in NVRAM, read is not
 - Reads can be from write cache
- 64K cache page size
- If resiliency timer set
 - Test every 4s - adjustable with SCSI command
 - Cache pages with timer older than 4s is queued for write - time set with each write to cache page
 - Data is kept in write cache after write, but the spaced is marked as available
- Sequential read ahead on 2nd miss in cache page (to end of page)
- Write ahead – queue many cache pages for write when above cache 'high water' mark



The VA 7x00 only uses a portion of the available cache.

There is some overhead in the cache that is used for the maps.

The overhead for the smaller 7100 array is less than that for the 7400/7410 array.

End-to-end Checksum

- 520 byte sectors – extra 8 bytes
- 4 bytes – RBID (Relocate Block ID)
- 4 bytes – checksum (w/RBID as seed)
- Trifid does it all (from host and to host)
- Will detect disk, firmware or data path errors
- Checks only on host read, not internal data movement
- The error is logged then array is reset
 - Host will see power-on, retransmit pending commands
 - No direct reporting of error – must read logs
- Permanent errors must be re-written to clear



Student Notes:

Disk Auto-fail Criteria

- Any smart event
 - Log the event, recommend auto-fail
- Errors vs. successes trigger
 - 32 errors before 32 success
 - 32 success reset the counters
 - Log the event, recommend auto-fail
- Recommend auto-failure
 - Will fail only if no data loss
- Power-on failures
 - Won't respond

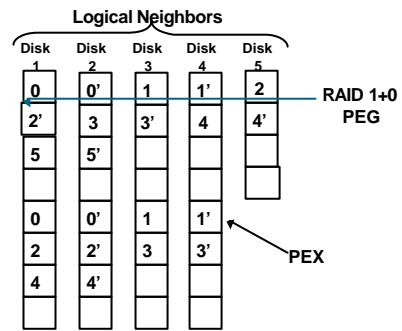


Smart Events are generated by the disk drives. These notify the “host”, in this case the controller that there may be a problem with the disk.

The recommendation for auto-failing a disk will only be accepted if it does not send the array into a data loss situation.

Power-on failures of the disk are not auto-failed since the disks are unable to respond and have not powered on.

RAID 1+0



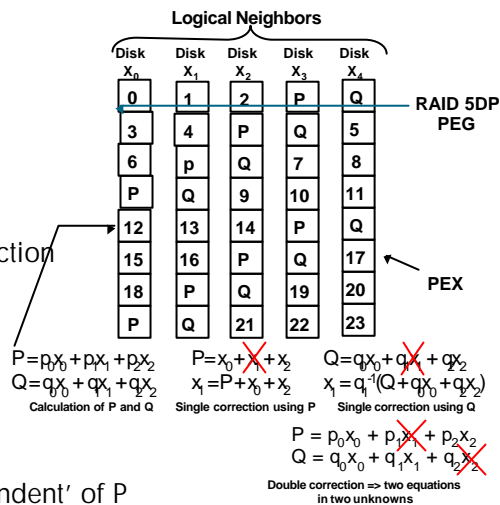
- Striped and mirrored
- Tolerant to some dual failures
 - Even number of disks in PEG
 - Left or right (just one) disk critical
 - Odd number of disks in PEG
 - Both left and right disk critical
- Disk mapping not visible to outside array



The implementation is best described as rotating RAID 1+0.

RAID 5DP

- N+2 rotating error correction
 - similar to RAID5
- Q is Reed-Solomon ECC coding
 - Similar to RAM and Disk error correction
 - Q = 'linearly independent' of P
- P = XOR (Parity)
- Tolerant to Double Failures – two equations in two unknowns
- All in hardware – as fast as XOR



This error correction code is formally categorized as a Reed-Solomon "erasure" code. Erasure means the location of the error is known before correction is attempted.

Galois (finite field) arithmetic is used for the calculations.

Addition is exclusive or multiplication is not normal integer multiplication uses special hardware implemented in Trifid and standard rules of algebra apply commutivity, distributivity, etc.

Coefficients $p_0, p_1, p_2, \dots = 1, 1, 1, \dots$ which reduces to normal exclusive or parity for P.

Coefficients q_0, q_1, q_2, \dots must be linearly independent from the p's so that double correction equations can always be solved. In other words any combination of two q's (qx, qy) cannot be a multiple of the corresponding p's (px, py).

During correction any non-failed segments become constants in the equation. The failed segments (x's not p's or q's) are the variables being solved.

Double failure cases involving one failed parity segment (P or Q) and one failed data segment are handled as a single correction case using the surviving parity segment.

Standard linear algebra techniques are applied to solve the double correction case.

VA7x00 Disk Performance

- Transfer Rates
 - 10K = 26MBs to 40MBs
 - 15K = 38MBs to 47MBs
- Average Seek Time
 - 10K = 5.2ms Read
 - 10K = 6.0ms Write
 - 15K = 2.9ms Read
 - 15K = 4.5ms Write
- Average Latency
 - 10K = 2.99ms
 - 15K = 2.00ms



Student Notes:

VA7410 Performance Enhancements

- Hardware
 - Faster processor
 - 2X fibre channel ports
 - Enabled 2Gb/s operation
- Firmware
 - New Normal Mode cache management algorithm - big reduction in time-outs
 - Improved read-ahead algorithm - better sequential performance for low concurrency workloads (must use vfp to enable)
 - Increased RAID 5DP concurrency - improved IOPs in typical workloads



Student Notes:

The Numbers

VA7410	VA7400	VA7100
34K cache IOPs	28K cache IOPs	12K cache IOPs
11,000 IOPs	8,500 IOPs	3,400 IOPs
330MB/s	180MB/s	93MB/s



Student Notes:

HP15/HP14 Performance Enhancements

- Increased Robustness
- Additional Tuning Mechanisms
 - Queue Full Thresholds
 - Promotion enable/disable
 - New read ahead mode (set w/vpf)
- Performance Increases
 - Improved low demand response time
 - Improved RAID 5DP sequential performance
 - Improved Fill rates



Student Notes:

HP
SurePartner
Training

MNS
Crossupdate

Module wrap-up





HP
SurePartner
Training

MNS
Crossupdate

HP Modular Network Storage Solutions

Part 3

Hardware Architecture



Student Notes:

Hardware Architecture Topics

- Overview of Array and Disk Enclosure Components
- Array and Disk Enclosure Views
- Architecture Block Diagram
- Array and Disk Enclosure Components
- Airflow
- Racking
- Power
- VA7400 Cabling
- VA7410 Architecture
- VA7410 Cabling



Student Notes:

Array Enclosure Chassis Components

- Field Replaceable Units (FRUs):
 - Array Enclosure Chassis
 - Array Enclosure Mid Plane
 - Array Controller Module (1 or 2 per chassis)
 - DIMMs (2 max. per chassis)
 - Controller Battery (1 per controller)
- Customer Replaceable Units (CRUs):
 - Drive Modules (15 max per chassis)
 - Power/Fan Module (2 per chassis- all configurations)
 - GBIC (1 per va7100 controller)
- Filler Modules (CRUs):
 - Controller Filler Modules
 - Drive Filler Modules



Student Notes:

Disk Enclosure Chassis Components

- VA 7400 / VA7410 Only
- Field Replaceable Units (FRUs):
 - Disk Enclosure Chassis
 - Disk Enclosure Mid Plane
 - Disk Enclosure Link Control Card (2 per chassis)
- Customer Replaceable Units (CRUs):
 - Drive Modules (15 max per chassis)
 - Power/Fan Module (2 per chassis- all configurations)
- Filler Modules (CRUs):
 - Link Control Card Filler Modules
 - Drive Filler Modules



Exchangeable Parts

A6214-69001 Link Controller Card (LCC)

A6250-69001 Add on Power Supply / Fan Module The add on power supply fan module is the one that goes into the JBOD.

Replaceable Parts

A6250-67005 Add On Enclosure Mid Plane Kit

A6214-67003 Add On Enclosure Bezel

A6209-60001 System/E Rail Kit (RBII)

A3231-60030 HP Computer Cabinet Rail Kit (RBI)


A5635-60001 Rittal Rack Rail Kit

HP
SurePartner
Training

MNS
Crossupdate


Desk Side

- A6183AD Desk side for Virtual Array Series
- Product label in back underneath the shelf
- Serial number and product number shown in back




Student Notes:

Array Enclosure: Front Panel



Student Notes:

Array Enclosure va7100: Rear Panel




5 minutes replacement time – on everything but the power supply

2 minutes replacement time – on power supply

If you need to replace the fan, you must replace the entire power supply

Disk Array Controller Module - (shown inverted)


- ESD sensitive
- Hot spots
- Five minutes to replace.
- There is foam in the packaging that can be used as a temporary filler for the slot while repairing
- Auto Clone



Disk Array Controller: Responsibilities

- Implementing Hewlett-Packard AutoRAID to ensure optimum performance and cost-efficient data storage
- Managing all data transfers between the host and the disks
- Maintaining data integrity
- Rebuilding the array in the event of a disk failure
- Monitoring the operation of all hardware components, including itself
- Alerting the host in the event of a problem with the disk array

Array Enclosure va7400: Rear Panel




Same rules as on VA7100

5 minutes replacement time – on everything but the power supply

2 minutes replacement time – on power supply

If you need to replace the fan, you must replace the entire power supply

Array Enclosure va7410: Rear Panel




5 minutes replacement time – on everything but the power supply

2 minutes replacement time – on power supply

If you need to replace the fan, you must replace the entire power supply


VA 7410 Controller



Student Notes:

GBICs/SFFP


- Products support optical only
 - multimode shortwave (770 - 860nm)
 - 5v only
- VA7100
 - SC duplex connector
 - Use only supported GBIC modules, i.e., Part Numbers:
 - Commercial – A6241A
 - Enterprise – A6203A
- VA7400/VA7410
 - SFFP connector (LC)
 - Not a field replaceable part



Student Notes:

Controller Memory

- Two DIMM slots on controller (**very ESD sensitive**)
- DIMM sizes must be the same in each slot of the controller
- Two sizes: 256MB, & 512MB
- Three controller memory sizes: 256, 512, & 1024 MB
- Hold AutoRAID maps & write cache
- Memory size in mate controller needs to be the same



Future: 1 GB DIMMS for the VA family of products. This would effectively increase the available cache size from 2 GB physical (1 GB mirrored) to 4 GB physical (2 GB mirrored) per array.

Controller DIMM Support

- Two DIMM slots supporting pc100 SDRAM operating at 100MHz (maximum of 1gb of memory per controller)
- Support 256mb, 512mb, and (future) 1gb DIMMs
- Both slots must be loaded with same size DIMMs
- Paired controllers must have same amount of memory
- DIMMs must have valid serial number and consistent timing specs when two per controller are used
- vfp will warn if an unsupported DIMM configuration is being used
- If valid memory is in DIMM following a hot plug event, must remove battery to invalidate memory



Student Notes:

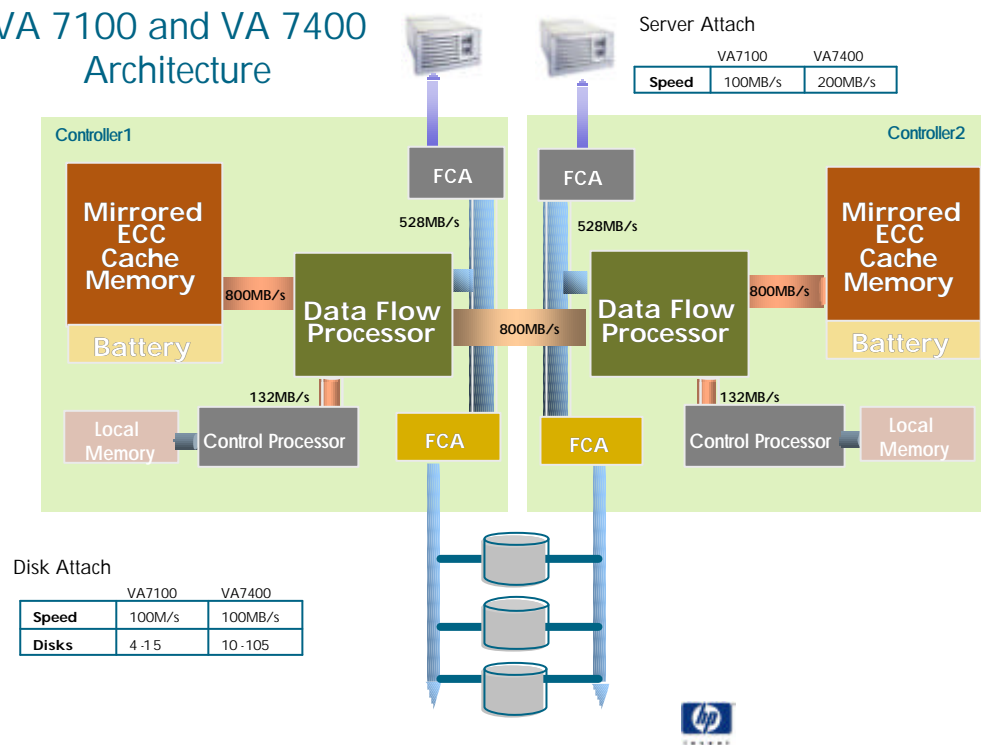
N-way Bus Description

- Internal duplex bus between Trifid asics on paired controllers
 - no single points of failure on memory bus
 - mirrored data memory reads & writes
 - inter-processor communication (ipc)
- 400MB/s data rates in both directions
- Packetized protocol with CRC protection



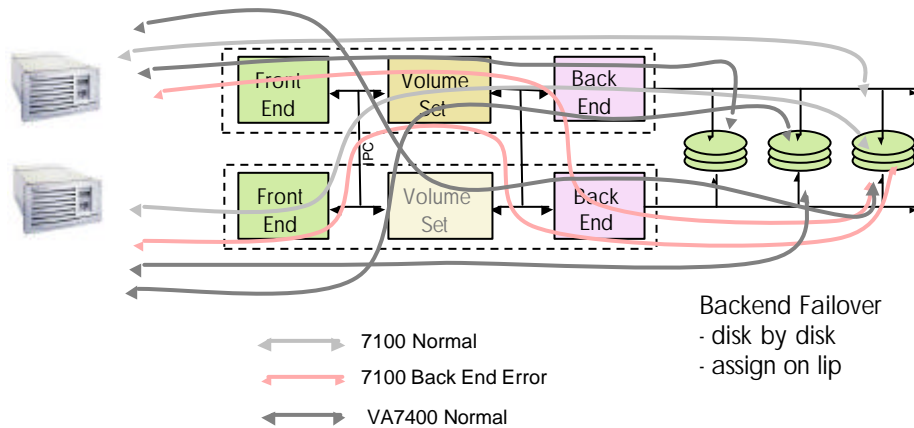
Student Notes:

VA 7100 and VA 7400 Architecture

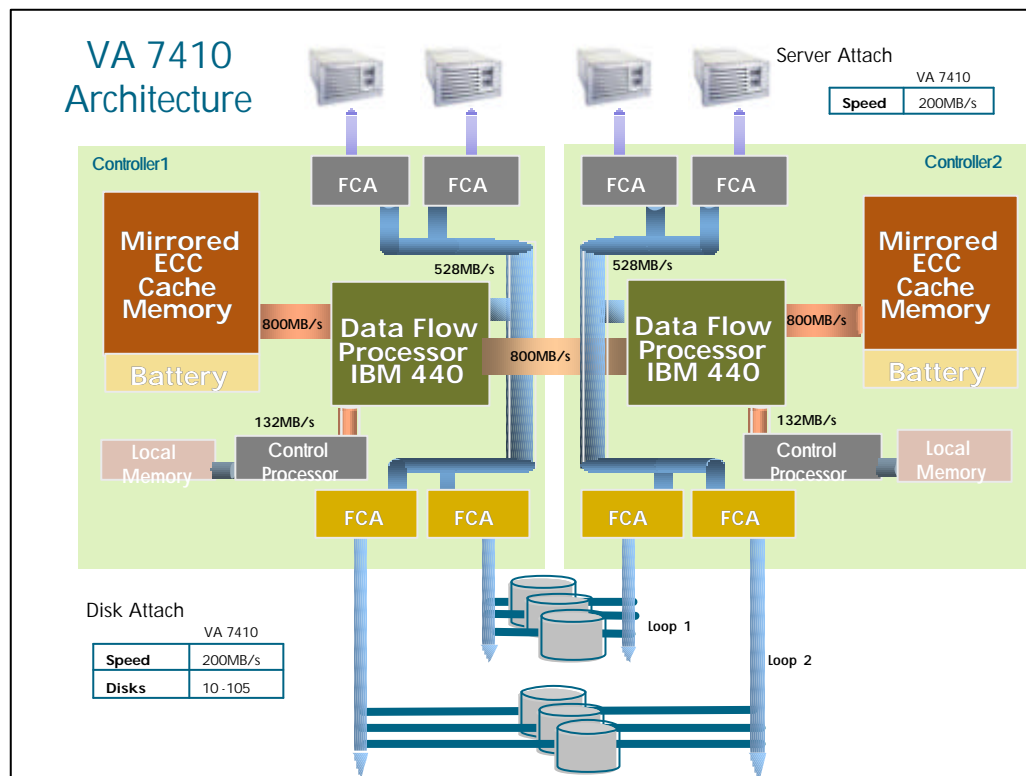


Student Notes:

VA 7100/7400 I/O Flow



Student Notes:



The intent of this slide is to communicate the contribution of the hp hardware. RAID 5DP - unlike standard RAID 5 which uses only one parity disks, RAID 5DP (dual parity) uses two parity disks. This provides far greater data protected than standard RAID 5. Even in our worst case configuration RAID 5DP provides over 60 times the data protect power.

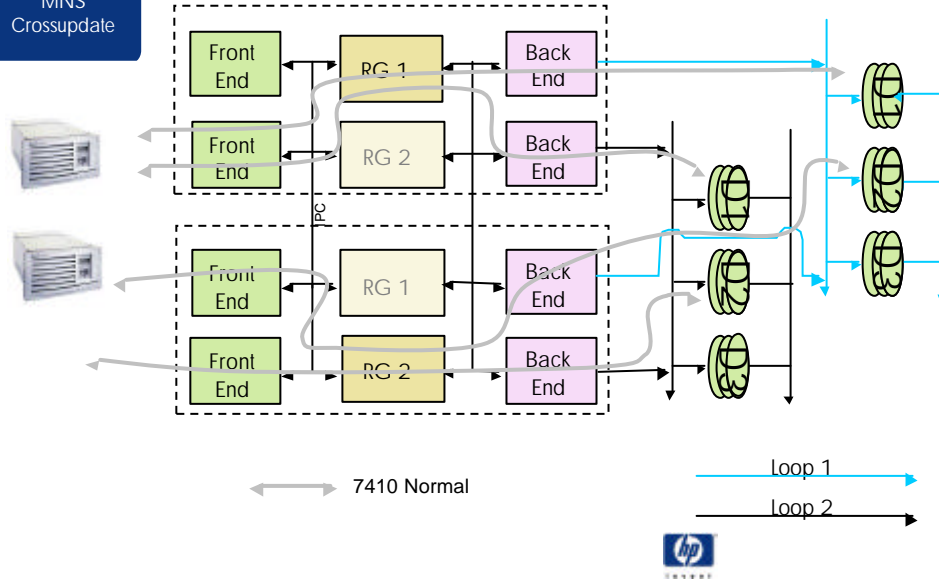
Mirrored memory - unlike competitive mid-range arrays, hp's virtual storage architecture implements a tightly coupled mirrored memory design typically found in high-end monolithic arrays - this provides greater availability and higher performance in shared environments.

End-to-end - seldom found in mid-range modular arrays but is standard in high-end arrays; this feature protects the data from disks that lie! Yes, disks lie - they return the wrong data or bad data that they say is good! Arrays without this feature just pass the bad data on to the server as good!

Metadata recovery - a sophisticated process within the array that protects the virtual mapping information from failures - this technology is very similar to data base transactions and checkpointing - this is a must for true highly available virtual arrays.

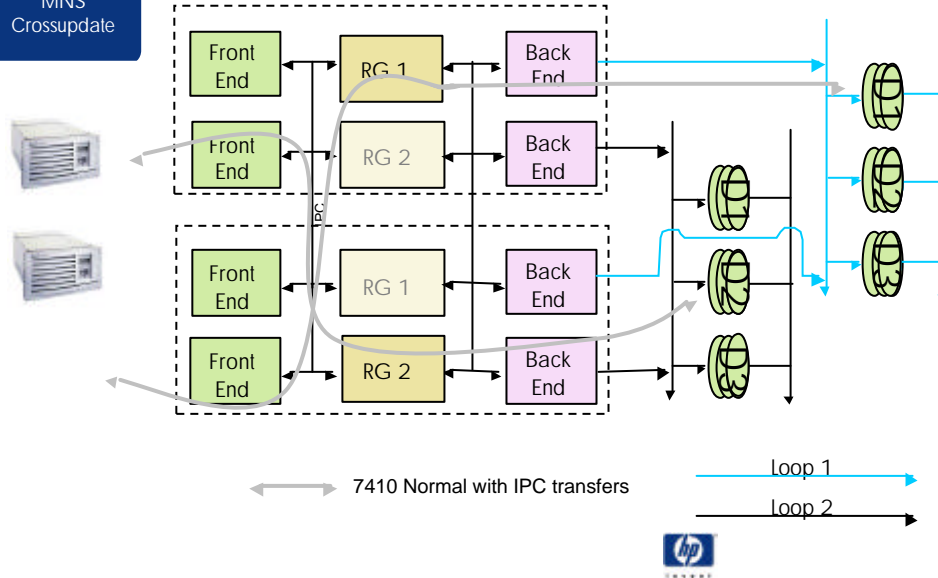
Hot Plug - hot swap - everything, anytime - virtual technology even allows disks to be replaced in any location, replaced with different size disks – thus eliminating a common data loss scenario from service

VA 7410 I/O Flow



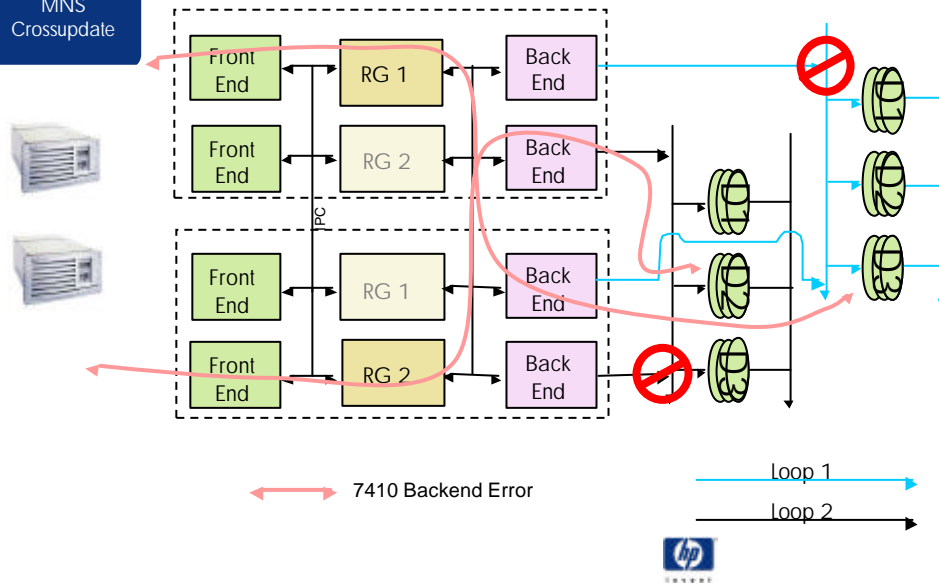
Student Notes:

VA 7410 I/O Flow



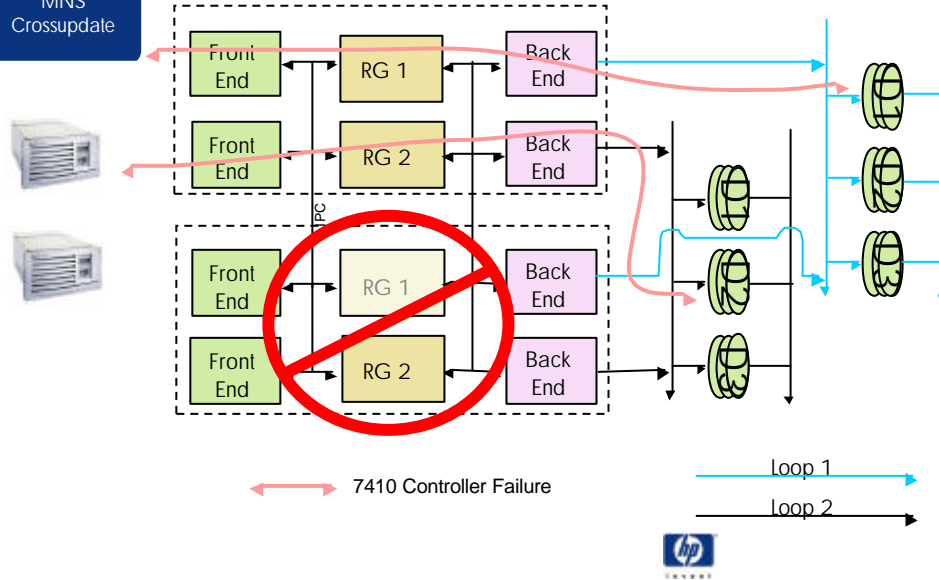
Student Notes:

VA 7410 I/O Flow with Backend Errors



Student Notes:

VA 7410 I/O Flow with Controller Failure



Student Notes:


Battery

- The battery charge must be sufficient to hold NVRAM contents for 3 days. If there is insufficient charge, a battery failure occurs
- If Green LED is flashing at 5% rate, battery is powering DIMM memory
- The battery is shipped with 25% to 50% of rated capacity to maximize its 6-month shelf life
 - If the battery temperature exceeds 25°C during storage, its shelf life will be significantly reduced
- 10.8V Lithium Ion Battery
- Smart Battery. Has own processor
- Do not leave batteries in a place subject to direct sunlight. Exposure to heat will severely reduce battery life



Student Notes:


Disk Array Midplane



There are redundant EPROMs on the midplane. These EPROMS hold:

Licensing information, array alias, enclosure node WWN, serial number, saved and default mode pages for the array, and loop id addresses for the two controllers.(108 and 110 default addresses).

Disk Array Power Supply/Fan Module



Active current sharing (N+1)

Auto-ranging 100-240 VAC

Each module powered by own power cord

CAUTION ! Do not operate without a PS/F for more than 2 minutes or whole unit will probably shut down.

DS2405 Disk Enclosure Chassis Components

- Field Replaceable Units (FRUs):
 - Disk Enclosure Chassis
 - Disk Enclosure Mid Plane
 - Disk Enclosure Link Control Card (2 per chassis)
- Customer Replaceable Units (CRUs):
 - Drive Modules (15 max per chassis)
 - Power/Fan Module (2 per chassis- all configurations)
- Filler Modules (CRUs):
 - Link Control Card Filler Modules
 - Drive Filler Modules



Exchangeable Parts

AXXXX-XXXXX Link Controller Card (LCC)

AXXXX-XXXXX Add on Power Supply / Fan Module The add on power supply fan module is the one that goes into the JBOD.

Replaceable Parts

AXXXX-XXXXX Add On Enclosure Mid Plane Kit


AXXXX-XXXXX Add On Enclosure Bezel

A6209-60001 System/E Rail Kit (RBII)

A3231-60030 HP Computer Cabinet Rail Kit (RBI)


A5635-60001 Rittal Rack Rail Kit

Disk Enclosure: Front Panel



Student Notes:

Disk Enclosure: Rear Panel




5 minutes replacement time – on everything but the power supply

2 minutes replacement time – on power supply

If you need to replace the fan, you must replace the entire power supply

Universal Drive Module



Student Notes:

Disk Drives

- A minimum of four disk modules (VA7100) or ten disk modules (VA7400) are required to operate the array
- Shipped with 520 bytes per sector
- When received as a replacement part there are 512 bytes per sector
 - Controllers will reformat drives from 512 to 520 bytes per sector
- Two drives in enclosure hold NVRAM image
 - Called "Image Disks"



Do not operate the array for more than 5 minutes with a disk drive removed

Disk drive filler panels must be installed in empty disk drive slots to maintain proper cooling in the array enclosure

Supported Fibre Channel Cables

- Optical (2 - 500m)
 - Multimode Shortwave (770 - 860nm)
 - 50.0um Core (2 - 500m)
- VA7100 (SC Duplex Connector)
 - A3583A 2 meter FC fibre optic cable
 - A3531A 16 meter FC fibre optic cable
 - A3735A 50 meter FC fibre optic cable
 - A3736A 100 meter FC fibre optic cable
- VA7400 (LC Duplex Connector)
 - C7524A 2 meter FC fibre optic cable
 - C7525A 16 meter FC fibre optic cable
 - C7526A 50 meter FC fibre optic cable
 - C7527A 200 meter FC fibre optic cable



Student Notes:

Other Supported Fibre Channel Connectors


- C7529A 2M LCm > SCm
- C7530A 16M LCm > SCm
- C7534A SCF > SCF adapter
- C7540A 2M LCm > SCm adapter kit
- A5750A – 16m, 50m, 100m Fibre Channel Host Cables
 - Options 001 – 003 must be ordered and define length
 - Options 004 – 006 LC/LC
 - Options 007-008 LC/SC connections
 - Option 009 – SC Female Adapter for 2M and 16M
 - Option 010 – LC Male Adapter kit



Student Notes:

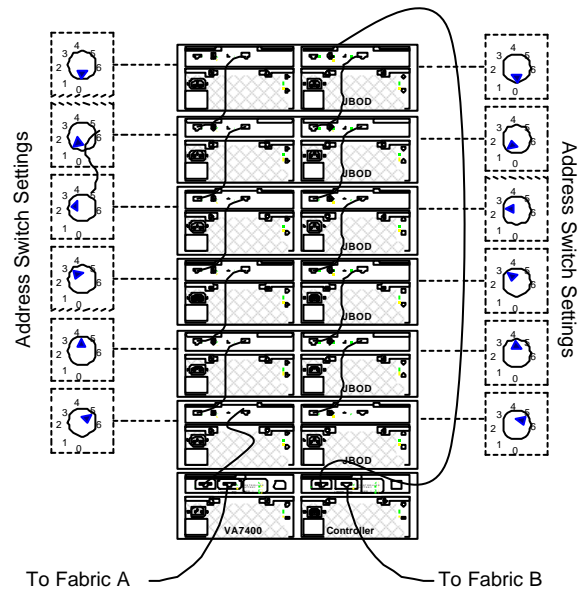
Connecting Power

- Connect each power cord to separate Power Distribution Units (PDUs)
- Connect each PDU to separate AC power sources** to maintain redundant power



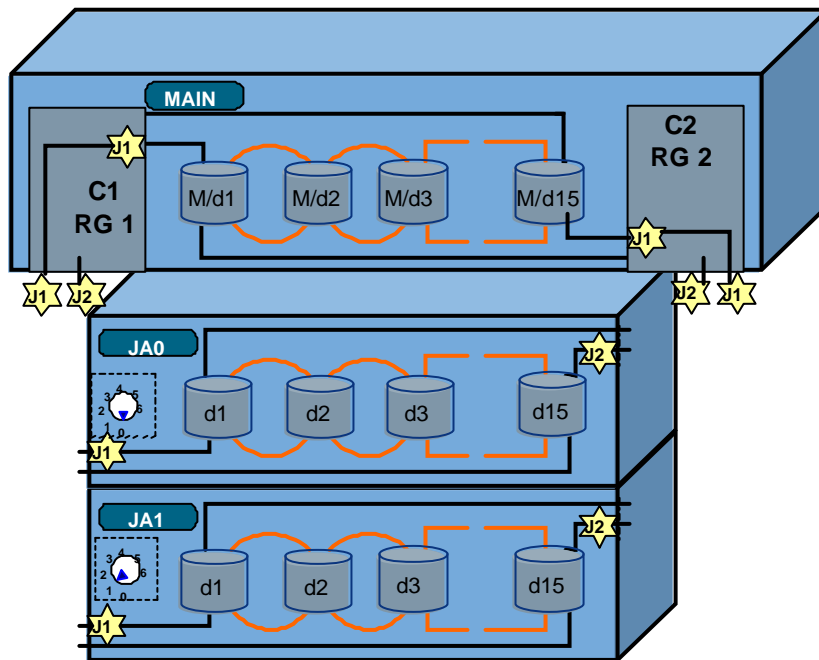
Student Notes:

VA7400 Cabling FC fibre optic cable




Student Notes:

Configuring Backend VA 7410's



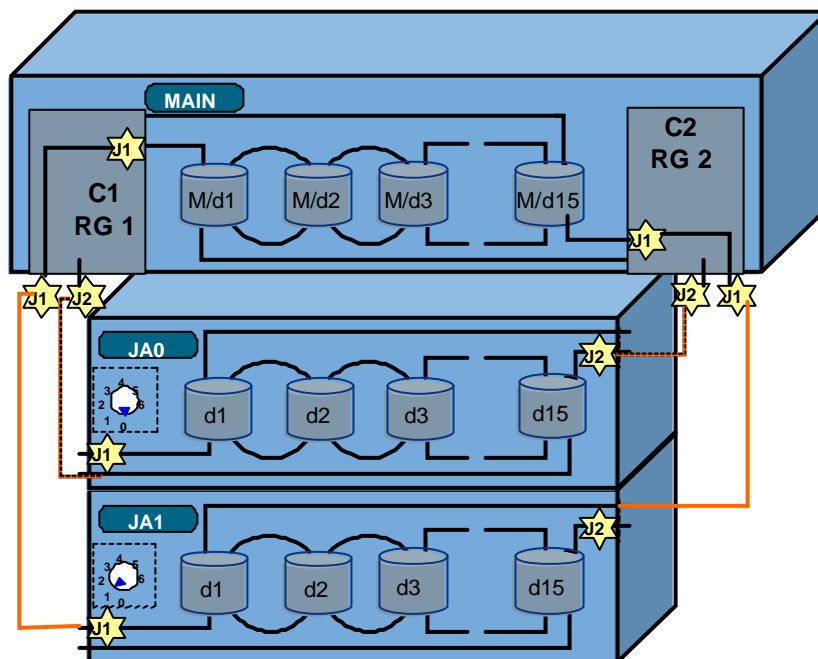
Student Notes:

Add First JBOD (JA0)




Student Notes:

Add Second JBOD (JA1)



Student Notes:


Add Third JBOD (JA2)



Student Notes:

Student Notes:

VA7xxx Cable Options



Student Notes:

HP
SurePartner
Training

MNS
Crossupdate

Module wrap-up





HP
SurePartner
Training

MNS
Crossupdate

HP Modular Network Storage Solutions

Part 4

Virtual Front Panel



Topics

- Virtual Front Panel: Overview
- Virtual Front Panel: Interface
- Connecting to the Virtual Front Panel
- Initialization Steps
- Virtual Front Panel: Commands



Student Notes:

Virtual Front Panel: Overview

- A virtual front panel (VFP) is a dumb terminal or terminal emulator used to issue commands to and display status messages from the array
 - Each controller has a RS232 port to access the VFP
- VFP firmware has three main functions
 - To receive and decode commands
 - Send resulting data and status messages to the RS232
 - Transmit unsolicited, periodic progress messages for certain processes e.g. Initialization, shutdown, download, and recovery information



Array ships with a Serial 9-pin Female to 9-pin Female Cable.
Part number: 24540-80014.

Virtual Front Panel: Interface

- Connect to the RS232 using
 - Dumb terminal
 - Serial interface of any PC using a terminal emulator e.g. VT100 or reflections etc
- Carriage return signals the VFP firmware of a command
- VFP firmware will display a prompt on startup and at the completion of every command. The prompt serves two purposes:
 - To provide a reference for the user
 - To communicate important system information. If all is well, the prompt will be displayed as "Ready >"
- There will be an instance of VFP firmware running on each controller




Student Notes:

Connecting to the Virtual Front Panel

- Connect to the array RS-232 port with a NULL Modem cable. (24540-80014)
- Use default setting for the VT140 or VT100
- Can use reflections

Bits per second	9600
Data Bits	8
Parity	None
Stop bits	1
Flow control	None



Student Notes:

Initialization Steps (1 of 2)

- Initialization steps can be viewed through the Virtual Front Panel (VFP):

Alpha Firmware

Revision

Date firmware was created

36582EX14I1024000903: Initializing:

02 04 06 08 0A 0C 0E 12 16 18 1A 1C 1D 1E 20 22
24 26

36582EX14I1024000903: Ready

Ready >



Student Notes:

Initialization Steps (2 of 2)

- Power on/boot sequence* :

Step	Description of Array Operation
02	All hardware self test is complete. Memory has been tested. In a dual controller system memory is mirrored and the controllers are synced
04 - 08	System initialization
0A	Backend device discovery
12	Maps and Cache uploaded from disk
16-26	Completion of system initialization

* This refers to a correctly operating dual controller system (i.e., the system can find no errors that inhibit dual controller operation; e.g., different FW revs).

Also, all controller testing is complete while the "Initializing" phrase on the VFP has no number associated with it.



Step No. (Hex) - Array Operation

- 02 Power-on self-test complete.
- 04 Check array serial number. Configure NVRAM for maps.
Initialize all NVRAM on both controllers.
- 06 Initialize internals.
- 08 Initialize array scheduler.
- 0A Search for backend devices.
- 0C Backend device discovery complete.
- 0E Enable power supply manager to shutdown if needed.
- 12 Initialize maps and cache via upload from image disks. Attach array to volume set.
- 16 Enable hot plugging.
- 18 Enable warning services.
- 1A Reserved.
- 1C Setup internal data structures based on backend discovery.
- 1D Enable front end ports.
- 1E Initialize array clocks.
- 20 Setup internal data structures.
- 22 Synchronize both controller clocks.
- 24 Startup complete. Enable scheduler. Allow writes to disks.
- 26 Initialization complete.

Virtual Front Panel: Commands

- `vfphelp` Displays syntax for all supported commands
- `vfpssetup` Performs Virtual Front Panel configuration
- `vfpdsp` Displays subsystem status
- `vfpcfg` Subsystem configuration
- `vfpmgr` Subsystem management operations
- `vfpfmt` Performs a subsystem format
- `vfprbld` Performs rebuild operations
- `vfprecover` NVRAM recovery operations
- `vfppsecure` Security operations
- `vfpplog` Displays controller event logs



With Array firmware HP14 or higher, shortened commands are available:

Just remove the prefix „vfp“.

Lab 1: demonstrating the Virtual Front Panel

- Use the vfphelp to find the command reference
- Show configuration details using the vfp
- Create a LUN using the vfp

Performance goals:

- Describe the features and limitations of the vfp



Student Notes:

Backup Slides:

Virtual Front Panel Command Reference



vfpsetup

- Performs VFP configuration.
- vfpsetup [-e {on | off}] [-r <value>] [-c <value>] [-b <int>] [-?]
- Options
 - -e <on | off>
 - Enables or disables character echo. Enabling character echo causes VFP firmware to echo every received character back to the terminal device. This is the default setting.
 - -b <int>
 - Configures the baud rate of the RS232 connection. The default rate is 9600.
 - The numbers 0-3 correspond with the valid baud rate settings: 9600, 19200, 38400, and 57600



vfpsetup -L <language> was an option for setting the local language. This option was available previous to HP14 firmware. The interface is localizable, but until the localization is done the options have been removed.

vfp Usage hints

- HP14 introduced shorthand commands.
 - vfp prefix is not required
 - Setup, dsp, cfg, mgr, fmt, rbld, recover, secure and log
 - vfp is required for commands prior to HP14
 - vfpsetup, vfpdsp, vfpcfg, vfpmgr, vfpfmt, vfprbld, vfprecover, vfpsecure, and vfplog
- Command recall can be done with greater than (>) and less than keys (<)



Student Notes:

vfpsetup – Options Continued

- -r <value>
 - Set the number of rows on the terminal device to the value given. VFP firmware uses this value in formatting output. The default setting is 25. The valid range is from 5 to 1000.
- -c <value>
 - Set the number of columns on the terminal device to the value given. VFP firmware uses this value in formatting output. The default setting is 80. The valid range is from 40 to 8000.
- -?
 - Displays help message. Overrides all other switches.



Student Notes:

vfpsetup Examples

Ready> vfpsetup -e on

Character Echo: Enabled

Ready> vfpsetup -c 50 -r 20

Number of Columns: 50 Number of Rows: 20



Student Notes:

vfpdsp

- Displays Subsystem Status
- vfpdsp [-c <fruLocation>] [-d <fruLocation>]
[e <fruLocation>] [-f] [-L <LUN>] [-s] [-?]
- Options
 - -c [<FruLocation>]
 - Display controller information. If <FruLocation> is not given, display detailed information for each controller installed in the disk array. If <FruLocation> is given, it is in the format of m/c1 or m/c2.
 - -d [<FruLocation>]
 - Display disk information. If <FruLocation> is not given, display detailed information for all disks installed in the array. If <FruLocation> is given, it is in the format of m/d1 – m/d15.



Student Notes:

vfpdsp – Options Continued (1 of 3)

- -e [<FruLocation>]
 - Display enclosure information. If <FruLocation> is not given, display a summary listing of all enclosures. If <FruLocation> is given and specifies an enclosure, display detailed information for the enclosure at <FruLocation> only.
- -f
 - Display a listing of FRUs in the array. Include FRU location, description of hardware, identification, and status.



Student Notes:

vfpdsp – Options Continued (2 of 3)

- -L [<LUN>]
 - Display LUN information. If <LUN> is not specified, display detailed information for all LUNs on the disk array. If <LUN> is specified, display information for only LUN <LUN>.
 - If the LUN is a snapshot, then snapshot information will be displayed.
 - If the LUN has one or more snapshots, list how many and the corresponding LUN ID. Also, its active and its attachment states with its used capacity will be displayed. All listed LUNs will be shown with corresponding world wide names.



Student Notes:

vfpdsp – Options Continued (3 of 3)

- -S
 - Display Array Status, Warnings, and Subsystem Parameter information for the disk array. This includes the configuration settings that control the operation of the entire disk array.
- None
 - Display general information about the disk array. This includes product and vendor information, array state, and capacity usage
- -?
 - Displays help message. Overrides all other switches.



Student Notes:

vfpdsp: Standard Output

Ready > vfpdsp

ARRAY INFORMATION

Array World Wide Name: 0x50060b0000092276
Alias: va7100-4
Array State: Ready

WARNING (if any are given here)

CONTROLLER ENCLOSURE INFORMATION

Serial Number: 00SG10600105
Status Summary: OK

If any warning conditions exist, one or more of the following will be displayed.

WARNINGS

<Warning Message>



Student Notes:

vfpcfg

- Array subsystem configuration
- vfpcfg [-L <LUN> -a <capacity>] [Suffix][-g<RGId>[-?]]
- vfpcfg [-L <LUN> -d] [-?]
- Options
 - -L <LUN>
 - Operation will be performed on LUN <LUN>.
 - -a <capacity> [Suffix]
 - Create the LUN if -L is specified.
 - Default capacity is in Megabytes
 - <capacity> K - is in Kilobytes
 - <capacity> M - is in Megabytes
 - <capacity> G - is in Gigabytes
 - -g <RG Id>
 - Redundancy Group Identifier, 1 or 2.
 - -d
 - Delete LUN if -L is specified.
 - -?
 - Displays help message. Overrides all other switches.



If capacity is followed by 'K' or 'k' the capacity must be a multiple of 512 kilobytes.

The redundancy group identifier will be required in VA7100 with the HPII or higher firmware release.

vfpcfg Examples

Create LUN 0 with 10 MB capacity

Ready > vfpcfg -L 0 -a 10M -g 1

Lun:0

Capacity:10M

Redundancy Group: 1

Generic Example

Ready> vfpcfg -L 4 -d

Delete LUN 4.



Student Notes:

vfpmgr

- Array Subsystem Management Operations
- vfpmgr -s {shut | start}
- vfpmgr -R {full | partial}
- vfpmgr -B {autoraid | raid1}
- vfpmgr -O {on | off}
- vfpmgr -ds {on | off}
- vfpmgr -L <Loop ID> -c <Controller ID>
- vfpmgr -q <queue Full Threshold>
- vfpmgr -os <Host OS> -c <Controller ID>
- vfpmgr -t <Topology> -c <Controller ID>
- vfpmgr -S <speed> -c <Controller ID>
- vfpmgr [-?]



Student Notes:

vfpmgr – Options (1 of 2)

- -S
 - Shutdown or Restart the subsystem
- -R
 - Reset the subsystem
- -L
 - Set the FC Port Loop ID
- -OS
 - Set the Host OS
 - generic -- aix
 - hpux -- netware
 - Nt -- mpe
 - Win2k -- tru64
 - linux -- openVMS
 - solaris -- solarisCluster



Student Notes:

vfpmgr – Options Continued (2 of 2)

- -S { 1 | 2 } -c { 1 | 2 }
 - Set the FC Host Port speed to 1 or 2 GB
- -t { 1 | 2 | 4 }
 - Set the FC Port Topology
 - 1 - Private Loop
 - 2 - Public Loop
 - 4 - Direct Fabric Attach
- -c { 1 | 2 }
 - Controller ID, used to identify controller to be managed
 - 1 – controller in slot 1
 - 2 – controller in slot 2
- -?
 - Displays help message. Overrides all other switches.



Only Private Loop, Public Loop and Direct Fabric Attach are supported.

vfpmgr Examples

Ready> vfpmgr -L 100 -c 1

Set controller M/C1 loop id to 100.

Ready> vfpmgr -os hpux -c 1

Set M/C1 for HP-UX host port behavior.

Ready> vfpmgr -S 1 -c 1

Set M/C1 for 1GB Fibre Channel speed.



Student Notes:

vfpmfmt

- Performs a subsystem format
- vfpmfmt [-?]
- Example

Ready> vfpmfmt

Format Array (Clears ALL data)? {yes/no}

If yes then,

Formatting Array. . .

If any other key,

Array Format Canceled

- Note: Must type "yes" or "no" not "y" or "n"
- This command destroys ALL user data.



-? Displays help message.

vfprbld

- Performs Rebuild Operations
- vfprbld [-p] [-?]
- vfprbld -a {on | off}
- Options
 - -p
 - Display status/progress of an array rebuild
 - -a
 - Auto rebuild on or off
 - -?
 - Displays help message. Overrides all other switches.



Student Notes:

vfprecover

- NVRAM Recovery Operations
- vfprecover [-s [-c]] [-?]
- Options
 - None
 - Begin recovery in interactive mode. The utility will display all recoverable redundancy groups, then prompt for the number of the redundancy group to recover.
 - -s
 - Returns the recoverability status of the disk array, as well as the status of a recovery in progress. The status indicates if recovery is needed, and what percentage of the entire recovery has been completed. When used with the -c option, status will be returned at regular intervals allowing continuous monitoring of the recovery progress.
 - -?
 - Displays help message. Overrides all other switches.



Student Notes:

vfprecover Example 1

Ready> vfprecover

Choice	Volume Set	Serial Number
--------	------------	---------------

0123456789ABCDE1		
------------------	--	--

0123456789ABCDE2		
------------------	--	--

Enter Choice: 1

Attempting to reconstruct NVRAM.

This process can take up to several hours.

Use vfprecover -s to return status and progress of the recovery.



This is a VA 7400 example. The VA 7100 has only one volume set or redundancy group.



vfprecover Example 2



Ready> vfprecover -s

Volume Set Serial Number: 0123456789ABCDE1

Recovery Progress: 50%

(or)

No recovery necessary.



Student Notes:

vfpsecure

- Secure Manager VA operations.
- vfpsecure {[-R] [-e] [-d] [-?]}
- Options
 - -R
 - Resets the LUN security password to AUTORAID.
 - -e
 - Enables security of the device.
 - -d
 - Disables the security functions of the device.
 - -?
 - Displays help message. Overrides all other switches.



Full secure manager va functionality of LUN by WWN is configured with command view sdm, CLUI.

vfpsecure Examples

Ready> vfpsecure -R
Reset the password to AUTORAID

Ready> vfpsecure -e
Enable Lun Security.

Ready> vfpsecure -d
Lun Security Disable.



Student Notes:

vfplog

- Display controller log entries.
- vfplog {[-p <number>] [-?]}
- Options
 - none
 - Displays all events in the buffer.
 - -p <number of last events>
 - Display the most recent number of events.
 - -?
 - Displays help message. Overrides all other switches.



vfplog is available with the HP11 or higher firmware release and the VA7400.

For scheduling reasons, the output will be in hex.

vfpllog Example

```
Ready>vfpllog - p 1
Event Logs Report
Page Code: 1
Revision: 1
Current Controller Time: 0x00241a864da2
Time Basis: 0x000000001805
Page Length: 0x430
```

Number of events: 1

```
Entry 1
Event Type: 0x2
Sequence Number: 0x41911
Time Stamp: 0x0023059c26b1
Module ID: 0x0b6
Controller SN: 00PR00164607
Process ID: 0x1b
Slot: 0x0
Event Length: 0x52
Event Code: 0x0da
Loop Pair: 0x0
Enclosure ID: 0x0
Slot: 0x80
Component: 0xff
Event Unique Information:
```



Student Notes:

Module wrap-up






HP
SurePartner
Training

MNS
Crossupdate

HP Modular Network Storage Solutions

Part 5

Management Software



Management Software Topics

- Command View SDM Overview
- Installation
 - Host
 - Client / Server
 - Web Browser
- Graphical User Interface
- Command Line User Interface



Student Notes:

command view sdm (1 of 4)

- Device management software for hp storage system
- Runs on its own version of JVM (JAVA)
 - Installed with Command View
 - Contained within the directory structure of command view
- Multiple user interfaces (GUI, CLUI, CVUI)
- Multi-device software launcher
- Online help and user documentation
- Licensing support
- Management options
 - Host based
 - Client/server based
 - Web based



GUI – Graphic User Interface

CLUI – Command Log User Interface

CVUI – Command View User Interface

command view sdm (2 of 4)

- Device configuration
 - HW setup and configuration
 - Discovery and mapping
 - Online firmware updates
 - Auto format/ Auto include/ Auto rebuild
 - HP Secure Manager VA
 - HP Business Copy VA



Student Notes:

command view sdm (3 of 4)

- Device diagnostic and fault management
 - Device diagnostics
 - Event monitoring, logging, and notification
 - Log information
 - HP response center link for HP-UX
- Device performance
 - Logging, presentation, and export



Student Notes:

command view sdm (4 of 4)

- Multi-OS support
 - HPUX 11.0 and 11.i
 - Microsoft Windows 2000 Server SP1
 - Microsoft Windows NT 4.0 Advanced Server SP6a
 - Microsoft Windows NT 4.0 Enterprise Edition SP6a
 - Linux Red Hat 6.2 and 7.1
- Solutions
 - SNMP compliant
 - Integration with OpenView, TopTools, and CA Unicenter
 - Qualified with OmniBack and HP OpenView Storage Area Manager



Student Notes:

Command View SDM Minimum System Requirements

- HP-UX 11.0 and 11.i
 - Server class, 256MB memory, 60MB disk (for Command View) and 16MB per 2 months of logs
- Windows 2000 and Windows NT
 - 500 MHz processor speed, 256MB memory, 60MB disk for Command View and 16MB per 2 months of logs
- Linux Redhat 6.2 and 7.1
 - 500 MHz processor speed, 256MB memory, 60MB disk for Command View and 16MB per 2 months of logs



Note that the software is not supported on Windows versions other than the English version; either as a server or client.

Note if you are installing on an alternate drive with Windows (not C:) an additional 45MB is required for the installation process. After installation only 16MB are required.

Memory utilization example (Windows):

Framework: (server)

Real: ~9500 K

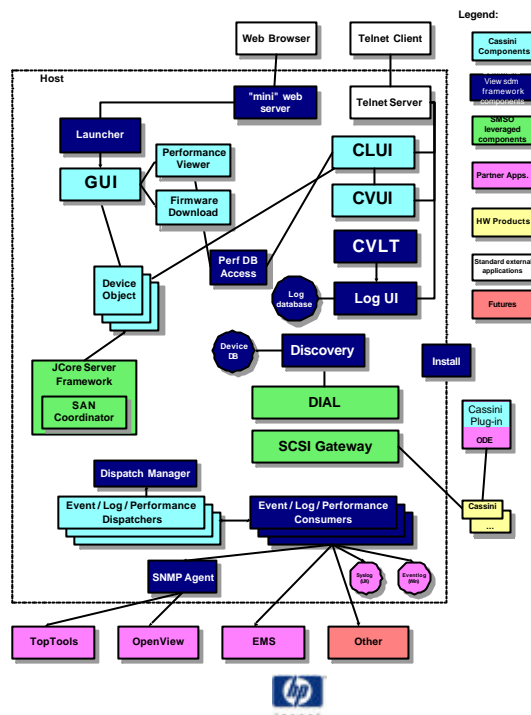
Virtual: ~11000 K

With one device:

Real: ~9700 K

Virtual:~13000 K

Command View sdm: host-based software architecture overview



The Command View sdm software is architected to be extensible: One dimension is to provide a consistent management story for MNS products (in other words, to support multiple types of storage devices with the same infrastructure). Another is to provide consistent management across platforms (easily portable to different operating systems). Integrations with third-party applications are also designed to be easily extensible.

Components

The software is designed to have multiple components: The framework doesn't contain any device-specific or OS-specific functionality. There are OS-dependent components; for example, the HP-UX component contains the HP-UX version of the HostAgent, the HP-UX syslog event consumer, and the EMS monitor).

The VA 7x00 component contains all of the base va 7x00-specific functionality. In future releases, we intend to provide value-add components that may also be device-specific. Other component types include third-party integrations. These integrations, such as the HP OpenView NNM Plug-in for HP-UX, can even be bundled and sold separately

JCore

The Command View sdm software leverages the JCore server framework, including the OpenDIAL and SCSI Gateway components. These are the primary OS dependent pieces of the software (not including the installers). The JCore components are bundled together and installed as the HostAgent. JCore also includes a very lightweight web server, which is enabled by the Command View sdm installer. Other web servers can be easily configured, and the JCore one can be turned off.

Java

- Why Java?
 - Portability, web accessibility, object-oriented, leverage with SAN Manager products
- Bundled JRE 1.3
 - .../sanmgr/jre
 - Will be bundled with HP-UX in the future
- Command View sdm runs in independent JVM instantiation
- Can run other versions of Java on the same system



Why bundle the JRE?

In the Microsoft implementation of their JRE, there are several components that don't exactly follow the Java specifications. In addition, this is the basic test and bundle configuration for the HostAgent.

GUI requirements

The system requires a minimum of Java 1.2.2 support to support the SWING toolkit. We currently use version 1.3.0. The SWING toolkit was used to develop the GUI; in order to support the GUI as a Java application running on the host we need the appropriate version of Java. Just a side note: HP-UX 10.20 doesn't support Java 1.2.2, so the management software can't be ported there (and still include the GUI interfaces).

Remote (web-based) GUI support

Remote GUI support via a web interface requires a minimum of Java 1.2.2 support from the browser used on the remote client system. This is the same basic requirement as for local GUI support; the SWING toolkit dictates the level of Java support required. Again, we are currently using version 1.3.0.

Support of the GUI via a web browser requires Sun's JRE Plug-in. If the plug-in is not available on the system, the browser will be directed to the sun.com website in order to download this plug-in. Note: The software does not work properly with the Microsoft implementation, and is not supported with anything but the Sun plug-in.

Directory Structure

- All files can be found under .../sanmgr/...
- HP-UX and Linux: All aliased to /opt/sanmgr/...
 - Files are split between /opt/sanmgr, /etc/opt/sanmgr, and /var/opt/sanmgr
 - If you are attempting to clean up a suspicious installation, clean up all three directories
 - If you are searching for a file with "find", remember "find" doesn't traverse symbolic links by default
- Windows: Default installation: c:\sanmgr
 - Directory structure is identical to the Unix-based systems
 - Note: This is the "collapsed" directory structure of /opt with symbolic links to /etc/opt and /var/opt



Student Notes:

Host Installation

- Command View SDM must be installed on at least one host system that is connected to the array over fibre channel.
- All installations begin with this step of installing on a host.
- The installation is a straight forward process
 - Setup.exe on the windows platform
 - Software depot on the HP-UX platform



Student Notes:

Client/Server Installation

- This installation begins with the Host installation on the server.
- Modify host's .../sanmgr/hostagent/config/access.dat to add IP information of the client.
- Install Command View SDM on the client.
- Stop the hostagent and dial processes on the client.



Student Notes:

Configuring Web Access

- The Command View SDM software comes with a mini-web server pre-enabled
- Connect from the web browser to either the launcher or a specific array
 - <http://hostname:4096/Launcher.html>
 - <http://hostname:4096/cmdviewVA.html?hostname=arraySSN>
- To set up remote accessibility, remember to configure secure access:
 - Modify .../sanmgr/hostagent/config/access.dat to add IP information
- The mini web server can be disabled if you want to configure and use a different (existing) web server



The first time you connect to the server through the web browser will be led through the JRE installation.

HP
SurePartner
Training

MNS
Crossupdate

Web Access Verification



Student Notes:

Event Notifications

- HP-UX:
 - EMS remote monitor
 - Syslog
- Windows
 - Windows Event Viewer
 - SNMP traps (Windows)
 - Supports integrations into TopTools, OpenView NNM, and CA Unicenter
- Linux:
 - Syslog

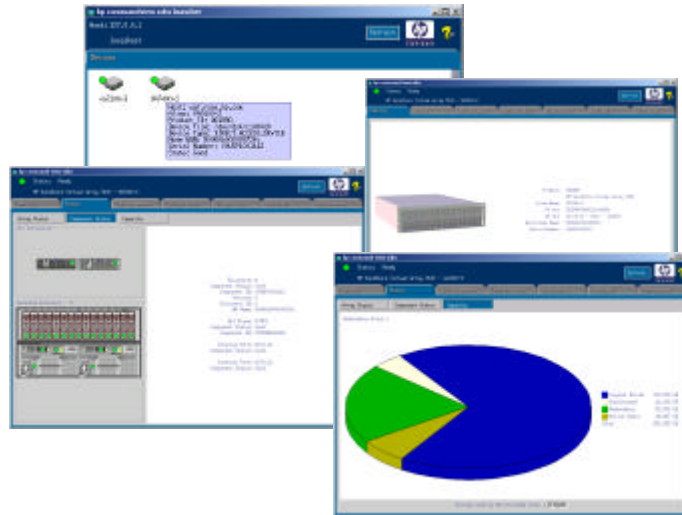


Event notifications are provided through different paths on different operating systems.

HP
SurePartner
Training

MNS
Crossupdate

Command View sdm: Demo



Student Notes:

Lab 2: demonstrating the GUI

- Use the GUI to find out the configuration details of the VA
- Create a LUN
- Delete a LUN
- Show capacity graphs

Performance goals:

- Describe the features of the GUI




Student Notes:

HP
SurePartner
Training

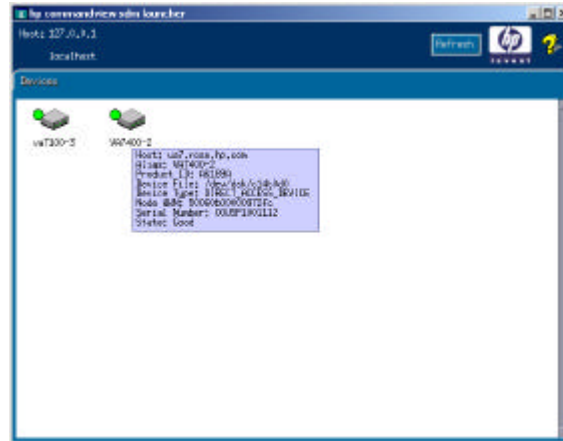
MNS
Crossupdate

Backup Slides – if Demo not Available



Student Notes:

Command View sdm: Launcher



- Give the users a single starting point to select the device to configure and to launch the appropriate management software



The main window area contains:

Serial Number - this area lists the serial number for the selected disk array or device enclosure.

Alias Name - this area lists a user assigned name for the selected disk array or device enclosure. If no alias has been assigned, this line will be blank.

World Wide Name - this area lists the Node World Wide name for the selected disk array or device enclosure.

Status Indicator – The icon to the upper left of the array icons indicates the global status of the array.

Double Click on the referring array icon will launch the device specific management page.

Status Icons



Not Enough Drives



Ready



Shutdown



Shutting Down



Status Up

- Warning



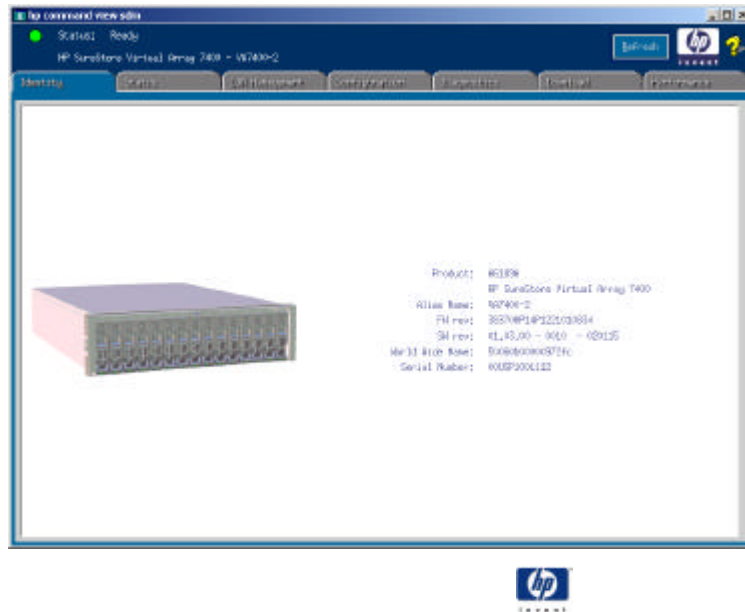
Shutdown Warning

- Controller Microcode Mismatch
- Disk Format Mismatch
- Controller lost microcode
- No Meta Data Available
- Insufficient Volume Set drives
- Controller Revision Mismatch



These icons appear in the banner area of the GUI to indicate status for a component in the array.

Identity Tab



The identity tab contains a single page which lists information about the selected array. There are no buttons to control the array from the Identity tab.

Items listed include:

Product- This area lists the name of the selected array or device enclosure.

Alias Name - This area lists a user-assigned name for the selected array or device enclosure. If no alias has been assigned, this line will be blank.

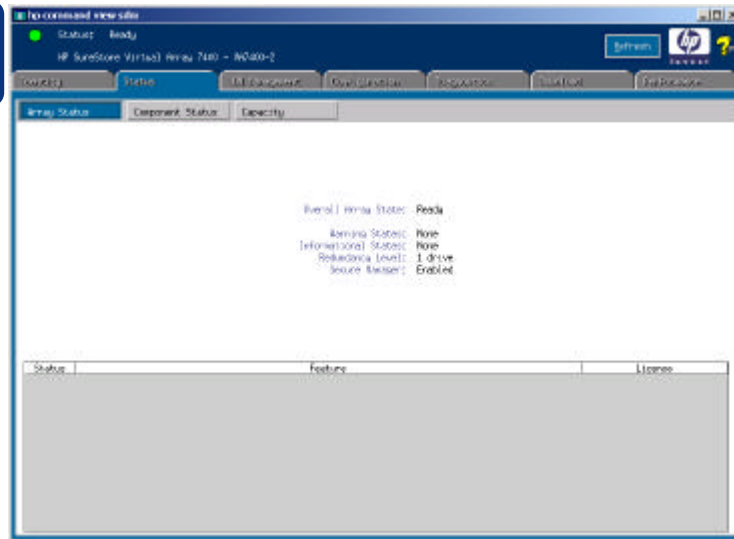
FW Revision - This area lists the revision number of the controller firmware for the selected array or device enclosure. (FW Rev)

SW Revision - This area lists the revision number of the Command View SDM software for the selected array or device enclosure. (SW Rev)

World Wide Name - This area lists the Node World Wide name for the selected array or device enclosure.

Serial Number - This area lists the serial number for the selected array or device enclosure.

Status Tab- Array Page



The Array Status page lists the condition of the selected array.

Status is available for:

Overall Array State - Lists the present state of the array.

Warning States - Lists any warning states the array may be in or entering.

Informational States - Lists any messages to inform the user that the array is performing a task.

Redundancy Level - Indicates the number of disk drives that can fail before a data loss will occur.

Secure Manager - Indicates if LUN security is enabled or disabled.

Status - Lists the status of any options that have been purchased and installed for the array. An icon indicates the present state of the option.

Feature - List any options that have been purchased and installed for the array.

License - Lists the licenses for the options that have been installed on the array. If installed, license serial numbers are shown in the table.

HP
SurePartner
Training

MNS
Crossupdate

Status Tab- Array Page

Static: Warning
HP SureStart Vital Array F100 - N1800v

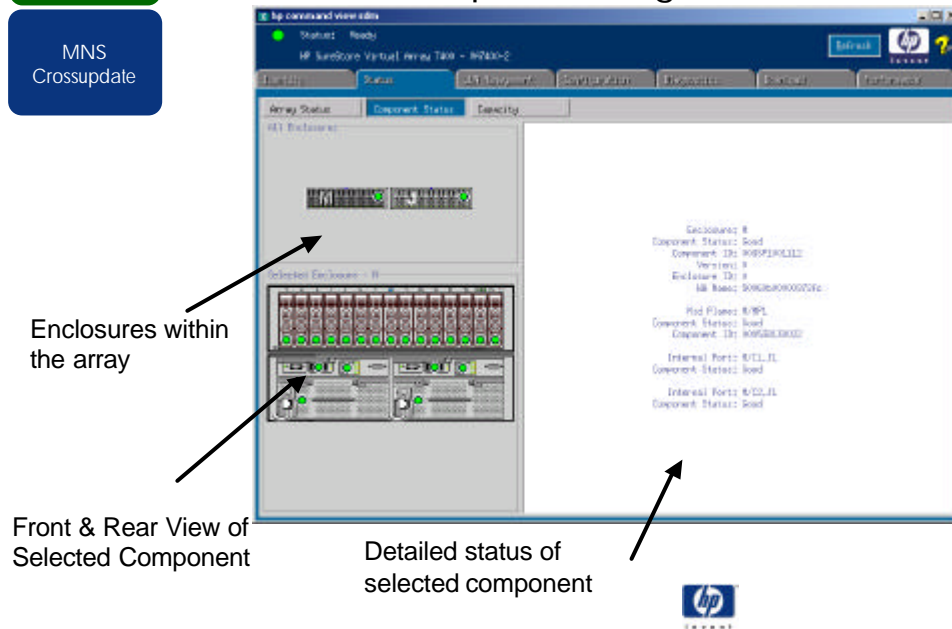
Array Status: Warning
Warning State: Redundancy Loss
Physical Drive Problem
Informational State: None
Redundancy Level: 1 drive
LUN Security: Disabled

Status	Feature	License
✓	LUN-SEC-255	L19A272205D06J8312289
✓	BUS-COPY-3442	L27511eH6008291307



This capture is showing the business copy va and secure manager va licenses installed.

Status tab– Component Page

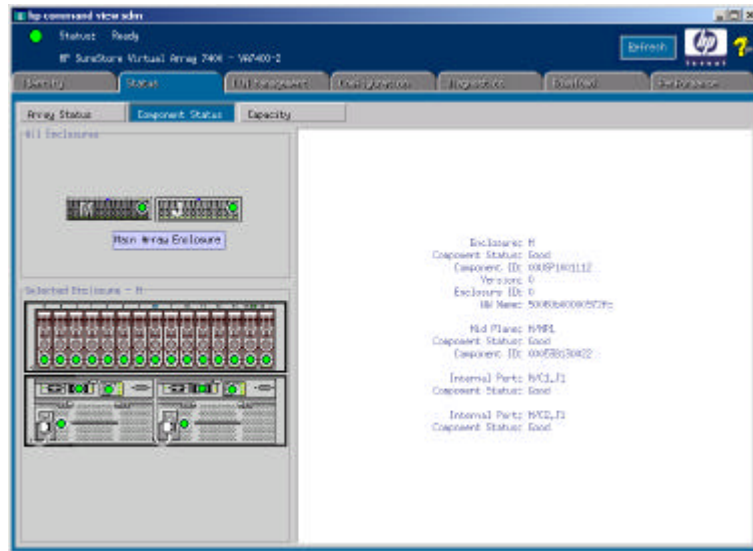


The status of the components are indicated by one of three icons: a green circle indicates good status, a yellow diamond indicates a caution status, and a red triangle indicates an error condition.

Selection of a component in the two left areas is a two-step process. First, select an enclosure by clicking it (in the upper left area). Second, select a component within the enclosure by clicking it (in the lower left area - Front or Rear view). When a component is selected it turns dark gray. Status information is available for the following enclosures and modules:











- Array and JBOD Controllers
- Disk Drives
- Main and JBOD Enclosures
- Fans / Power Supplies
- Front and Back End Ports

Component Page – Main Enclosure



<u>Component</u>	<u>Letter Identifier</u>
Loop Pairs	A or B
Disks	D1 – D15
Power supplies	P1 – P2
Controllers	C1 – C2
Processor	PM1
NVRAM Battery	B1 – B2
DIMM	M1 – M4
Host Port	H1 – H2

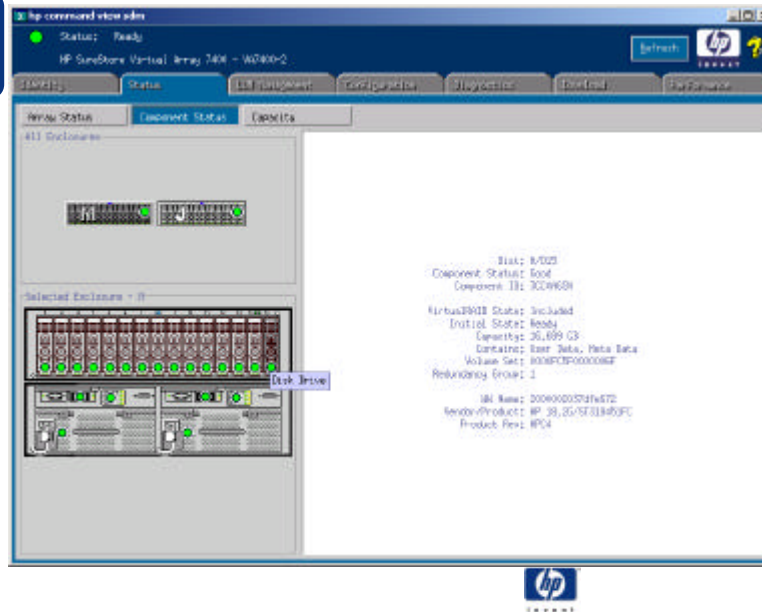
Array Status Icons

-  **Good** - Component is present and no error conditions have been detected.
-  **Warning** - Non-critical error condition detected, component is still operational (blinks).
-  **Critical** - Critical error condition detected, but component is still operational (blinks).
-  **Failed** - Component is present, but has failed (blinks).
-  **Init Failed** - Attempts to initialize the component have failed (blinks).
-  **Downed** - The component has been downed by the management software.
-  **Not Implemented** - Status detection is not implemented for this component.
-  **Unknown State** - The component's presence and/or status cannot be determined.
-  **Unknown State Exists** - Component is present, but status is not available.
-  **Unsupported** - Component is not supported by the firmware.



These icons appear on the Component Status page of the Status tab to indicate status for a component in the array.

Component Page – Disk Status



When a disk module within an enclosure is selected in the upper left side of the Status page, the right side will display the following information. Not all items may appear; the actual status information appearing depends on the state of the module.

Disk - Indicates the location of the component within the array.

Component Status - Indicates the status of the component.

Component ID - Displays the Disk's serial number.

Capacity - Displays the total physical capacity of the disk mechanism in bytes. This value can change due to a hot-plug event.

Contains - Displays the type of data the selected disk contains.

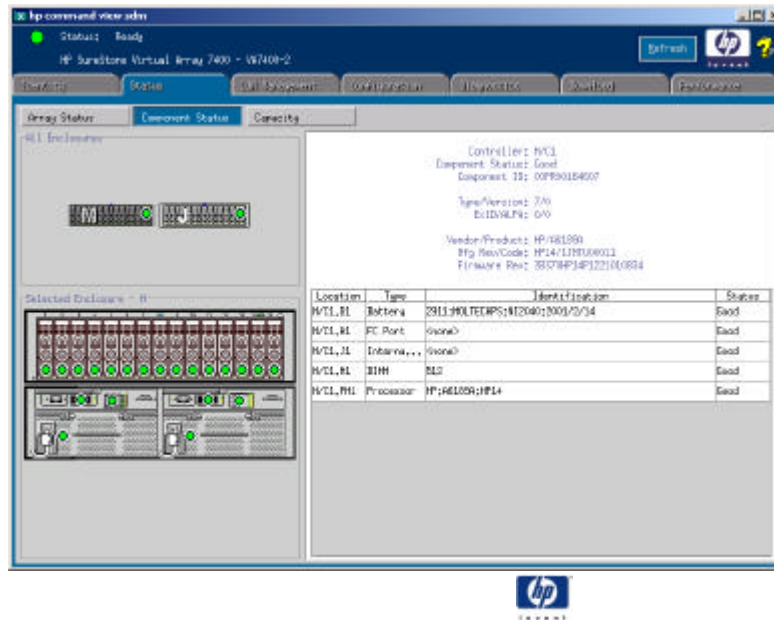
Volume Set - Displays a subsystem assigned unique identifier for the set of disks that make up the current volume set. This identifies the volume set to which the disk last belonged. A volume set is an AutoRAID construct. This value can change due to a hot-plug event.

World Wide Name - Displays the Node World Wide Name for this disk mechanism.

Vendor/Product - Displays a combination of standard SCSI inquiry data for Vendor and standard SCSI inquiry data for Product ID.

Product Revision - Displays standard SCSI inquiry data for Product Revision.

Component Page – Controller Status



When a controller module within an enclosure is selected in the lower left side of the Status page, the right side will display the following information. Not all items may appear; the actual status information appearing depends on the state of the controller module.

Controller - Indicates the selected FRU of the selected enclosure.

Component Status - Indicates the status of the selected component.

Component ID - Displays the Controller serial number of the selected component.

Type/Version - Displays a code for the specific version of the selected enclosure.

ExID/ALPA

ExID- Displays the current setting of the address ID thumb wheel for this controller.

ALPA- Displays the current arbitrated loop physical address being used by this controller.

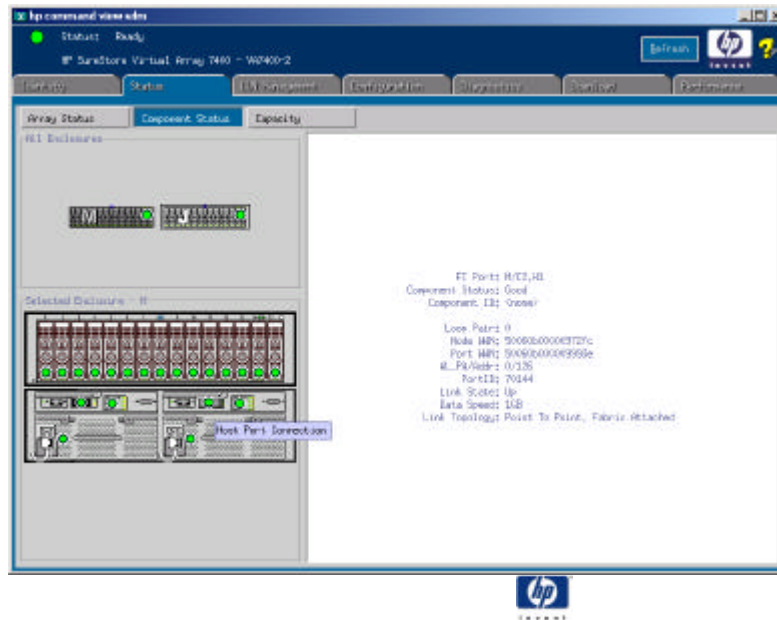
Vendor/Product - Displays standard SCSI inquiry data for Product ID.

Mfg Rev/Code - Displays data for Manufacturing Product Revision and code.

Firmware Revision Number - Displays inquiry data for Firmware Revision

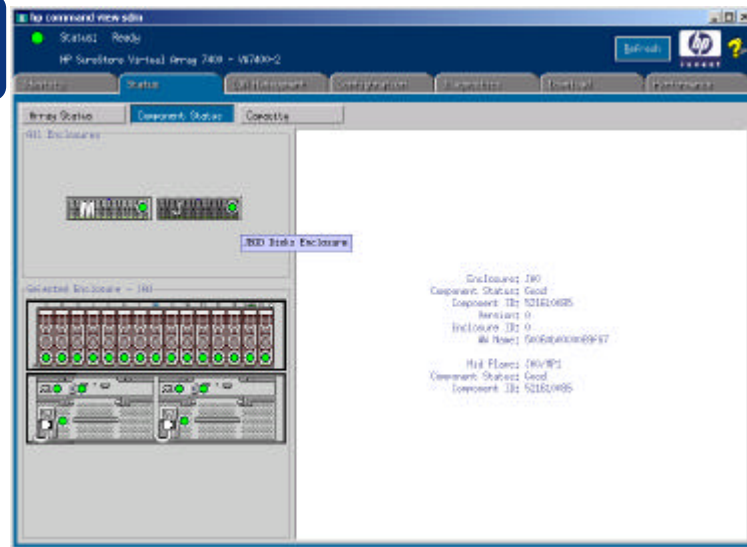
Sub-Component Table - A scrollable table which displays the location, Type, Identification, and status of all sub-components known to this controller. The list will include Controller Batteries, Fibre Channel Ports, NVRAM, Processor, GBICs, and Internal Port.

Component Page – Host Port Status



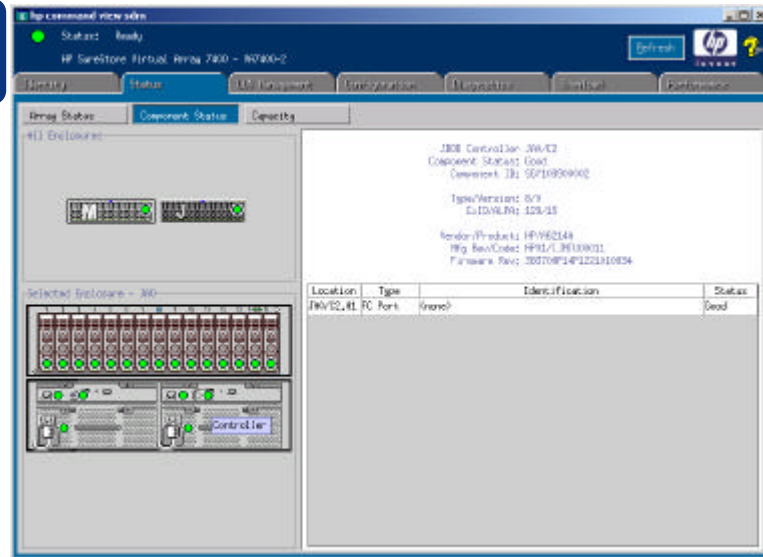
Student Notes:

Component Page – JBOD Status



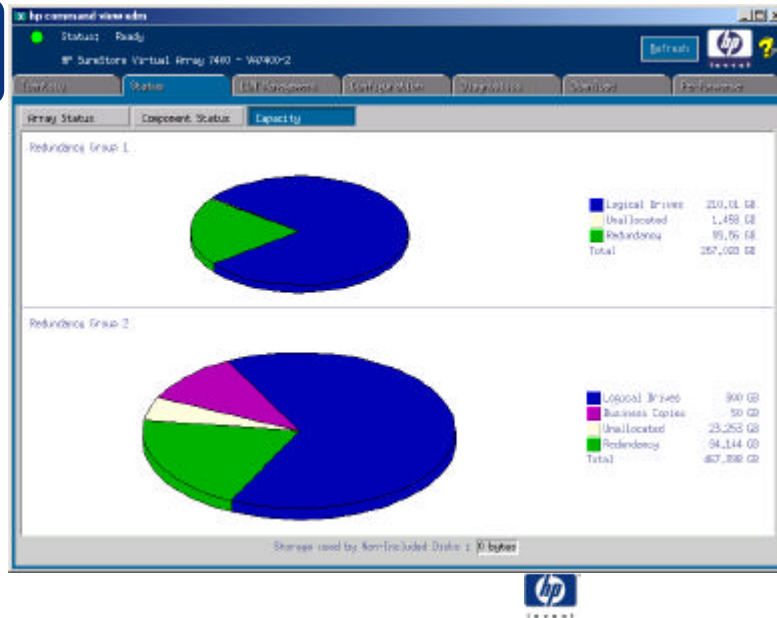
Student Notes:

Component Status – JBOD Controller Status



Student Notes:

Status Tab- Capacity Page



The Capacity page contains a pie chart which lists the size allocations within the array for each redundancy group. Allocation information is shown for:

Logical Drives - The sum of committed space for all LUNs that are not business copies

Business Copies - The sum of committed space for all business copy va LUNs.

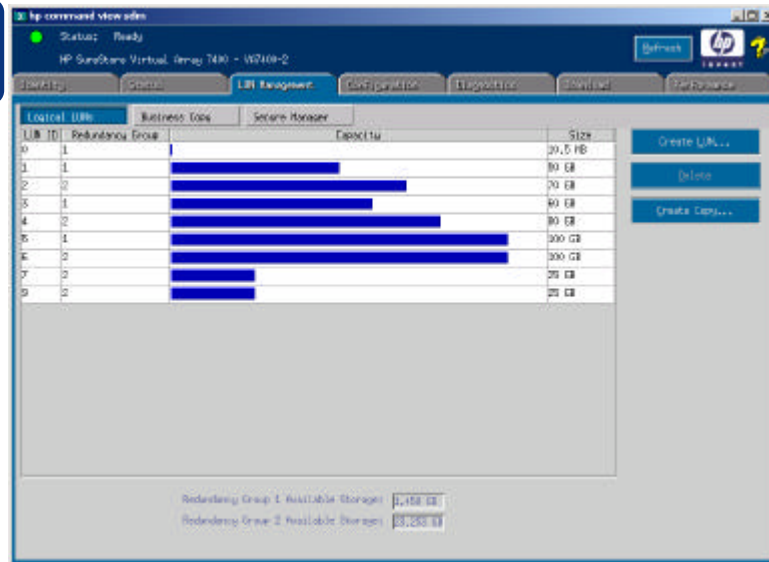
Unallocated - Unallocated is size of the device's uncommitted space. The uncommitted space is the space available for creating new LUNs.

Redundancy - The total space used for redundancy is equal to the sum of all the included drives, minus the unallocated space, minus the sum of all non-business copy va LUNs, minus the sum of the two largest included disks, minus the space allocated for all business copy va LUNs. (Essentially, this wedge accounts for all space not otherwise counted in the other wedges).

Active Spare - The size of the active spare is the sum of the two largest included disks. This wedge is shown only if Active Spare is enabled and available.

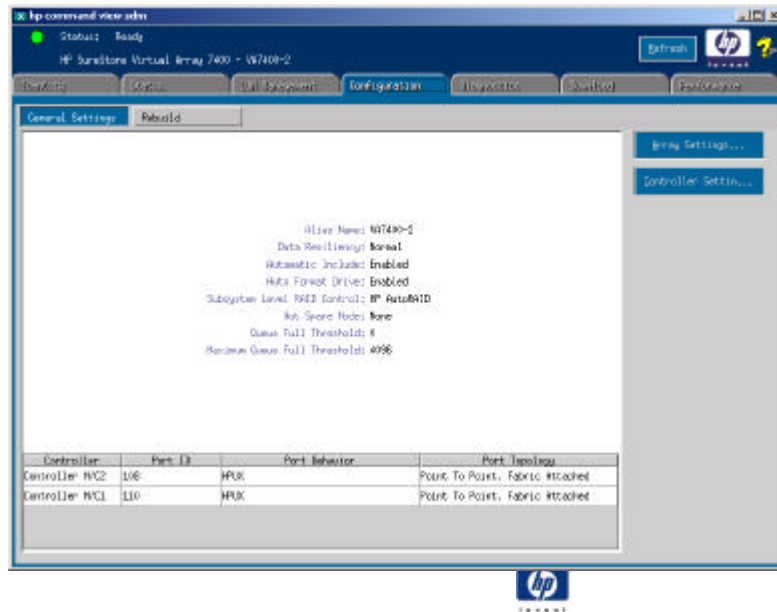
Total Physical Space (Total) - The total capacity of the array is equal to the sum of all the drives physically present: included and un-included.

LUN Management Tab- Logical LUNs Page



Student Notes:

Configuration Tab – General Settings



This page contains settings that apply to the array as a whole. Buttons are provided to make changes for the Array Settings and Controller Settings. The current array settings listed include:

- Alias Name
- Data Resiliency
- Automatic Include
- Auto Format Drive
- Hot Spare Mode
- Subsystem Level RAID Control
- Port IDs
- Port Behavior
- Port Topology

Configuration Tab – Array Settings

High Performance
Restricted Normal
Normal
Secure
Single Controller

HP AutoRAID
RAID 1+0

None
Largest Disk
Largest Two Disks
Automatic



Array Settings Button allows the user to make changes to array settings. Settings do not take effect until the OK button is clicked. Pending changes can be ignored by clicking the Cancel button. Settings that can be changed include:

Alias Name - Enter a user new defined name for the array here.

Data Resiliency

Single Controller, Secure, Normal, Restricted Normal, High Performance

Subsystem Level RAID Control

RAID 0+1 - Provides high performance but less efficiency.

HP AutoRAID - Automatically selects the RAID level.

Hot Spare Mode

None - Do not create a hot spare.

Automatic - Automatically select a hot spare from the available disks.

Largest Disk - Create a hot spare using the largest disk available.

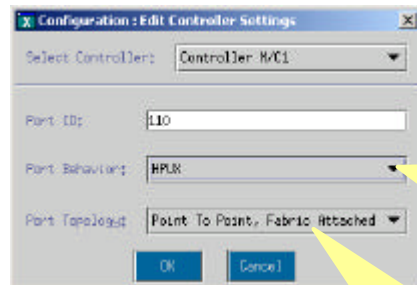
Largest Two Disks - Create a hot spare using the largest two disks available.

Auto Include - Enable or disable auto include.

Auto Format Drive - Enable or disable auto format.

Note! - For some older versions of firmware, Subsystem Level RAID Control will be inactive and Hot Spare Mode will only have ON or OFF selections.

Configuration Tab – Controller Settings



Unsupported OS
HPUX
Windows NT
Windows 2000
Linux
Solaris
AIX
NetWare
MPE
Tru64
OpenVMS
Sun Cluster

Private Loop
Public Loop
Point To Point, Fabric Attached



Controller Settings allows the user to make changes to the array controller(s). A controller is selected using the drop down list box at the top of the dialog box. Settings do not take effect until the OK button is clicked. Pending changes can be ignored by clicking the Cancel button.

Port ID - The Host Port ID is the parameter indicating the fibre channel loop ID that the specified controller and port will attempt to acquire in the hard address phase of the loop initialization.

Port Topology

Public - Select this if the controller is connected to an arbitrated loop connected to fabric.

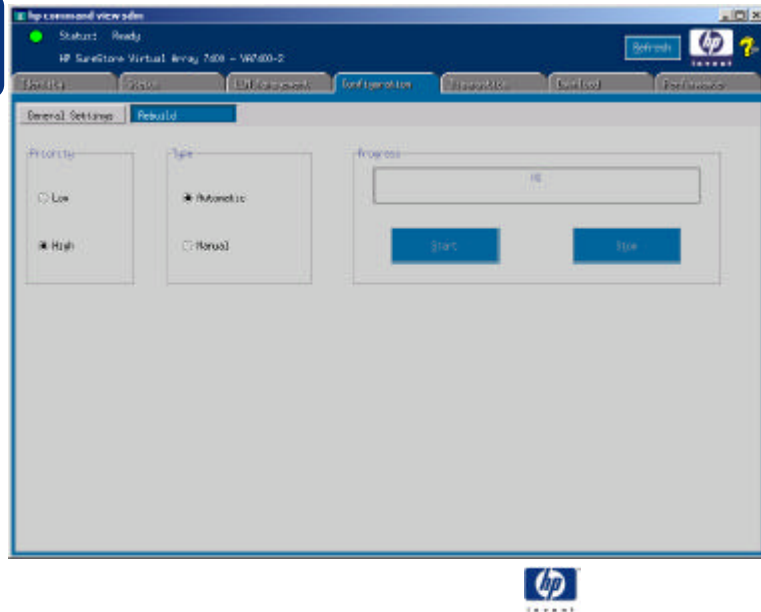
Private - Select this if the controller is connected to an arbitrated loop and not connected to fabric.

Fabric - Select this if the controller connects to many devices in a cross-point switched configuration.

Port Behavior setting implements a behavior that is compatible with the OS's FC driver. Choices include:

HPUX, Windows NT, Windows 2000, Linux, Solaris, AIX, and NetWare.

Configuration Tab – Rebuild



Use this page to view and make settings for the rebuilding the array. Changes take effect immediately. The parameters that can be changed are:

Priority - This feature indicates the priority that rebuild I/O is given with respect to host I/O. High priority indicates that host I/Os should be given a higher priority than rebuild I/Os. Low (the default) priority indicates that rebuild I/Os and host I/Os should be given the same priority.

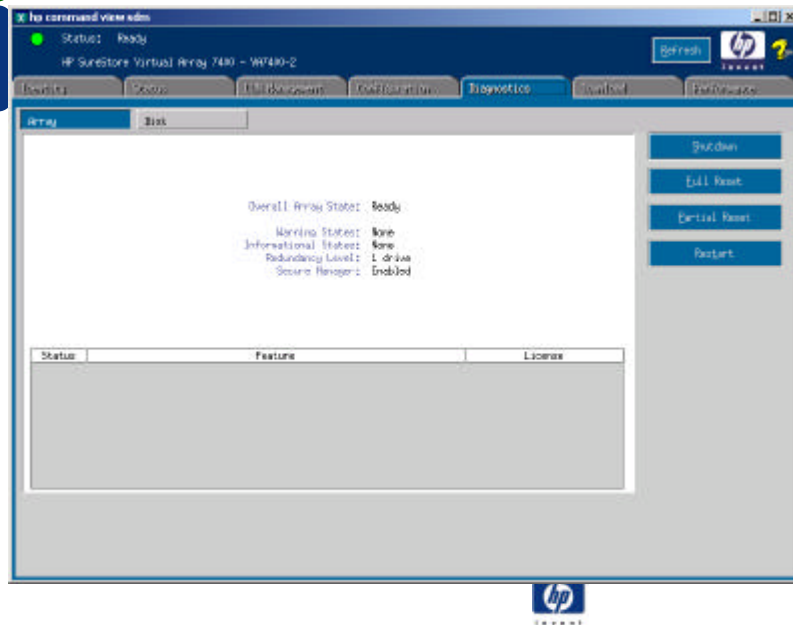
Type - This feature specifies the action to be taken when a drive becomes unavailable. If set to Automatic (the default), redundancy should be rebuilt automatically whenever a drive becomes unavailable. If set to Manual, rebuilds will only occur when explicitly started (the default is Automatic).

Start - Clicking this button will start the rebuilding of missing data redundancy in the array. This button is only active if a rebuild is in progress, and the rebuild type is set to Manual. The Progress bar indicates what percentage of the current rebuild operation has completed (during a rebuild, this progress bar will be updated every 5 minutes).

Stop - Clicking this will stop the rebuild operation. This button is only active if a rebuild is in progress, and the rebuild type is set to Manual. The Progress bar indicates what percentage of the current rebuild operation has completed (during a rebuild, this progress bar will be updated every 5 minutes).

Note! - Clicking Stop during a manual rebuild results in a Rebuild Failed informational state.

Diagnostics Tab – Array Page



This page provides ability to determine where a problem in the array is located. You can also shutdown and restart the array from this page. For convenience, the array diagnostics page duplicates the information Status tab's Array Status page. From the Array page, the array can be:

Shutdown (Shutdown) - Shutdown will stop all background activities in the device. Shutdown should be used in preparation for diagnostic or maintenance functions that may be catastrophic if allowed to occur to an operational device. Shutdown requests the device to stop normal I/O operation and save all map information to physical drive storage.

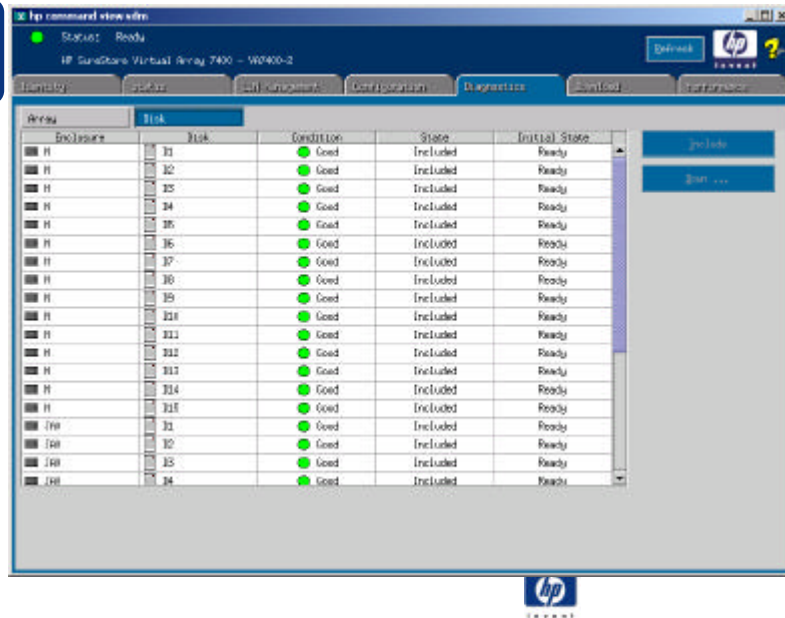
Restarted (Restart) - Clicking this button restarts the normal I/O operation of the array and brings the array back online.

Fully Reset (Full Reset) - Full Reset can be used to cause the device to reset in environments where the Fibre Channel Target Reset task management function is not possible. A Full Reset is used when the device is in an unknown state. A full reset includes memory tests. The button is active when the array is in the: No Code, Shutdown Warning, Shutdown, No Map, No Quorum, or Always Active state.

Partially Reset (Partial Reset) - Partial Reset can be used to cause a device reset in environments where the Fibre Channel Target Reset task management function is not possible. A Partial Reset does not include memory tests. A partial reset is very quick, the entire array is offline for up to 10 sec.

The difference between a partial reset and reset is the memory test. A full reset takes a little less time as power on.

Diagnostics Tab – Disk Page

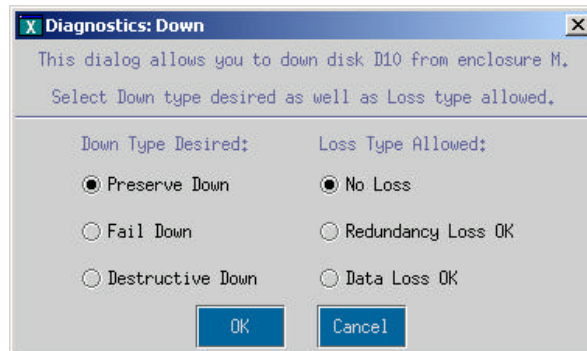


This page provides the ability to determine where a problem disk is located in the array. You can also Include or Down a disk using the buttons and view parameters on all the disks in the array. Parameters include:

- Enclosure
- Disk
- Condition
- State
- Initial State

Note! - To select a disk to modify, highlight (click) the desired LUN.

Diagnostics Tab – Disk Down



Select the down type desired:

Preserve Down - down disk and preserve data. Choose this button if the drive is to be reused and the disk's data will be left intact.

Fail Down - checking this radio button will allow the array to auto-fail the drive. It is possible for either you or the firmware to fail a disk.

Destructive Down -check this radio button if the drive is to be reused and actions that over-write the host data may be taken.

Select the loss type allowed:

No Loss - neither data unavailability nor redundancy loss should be allowed. This would only succeed on disks in the Not Included State.

Redundancy Loss OK - redundancy loss can be tolerated but data unavailability should not be allowed. This would not succeed if a redundancy warning were indicated by the subsystem.

Data Loss OK - both redundancy loss and data unavailability can be tolerated. This will always succeed if the specified disk has been included.

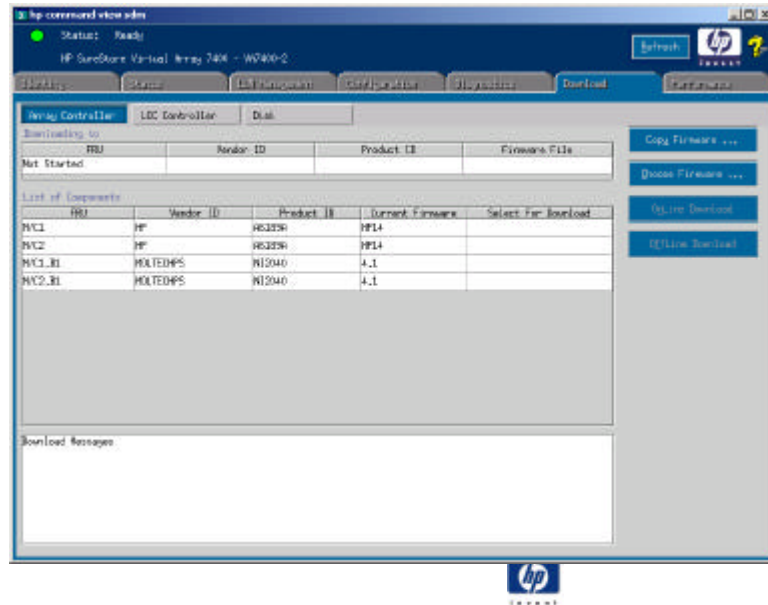
Downloading Firmware

- Backup array data before performing a firmware download.
- Online controller firmware download incurs a 10-15 second pause servicing host I/O (goal is 10 seconds).
- LCC and Disk firmware downloads are done offline.
- If the download is aborted (or fails to complete), the GUI may stop functioning due to an advisory lock being set.
 - break the lock using the `armmgr -b` command.



Student Notes:

Download Tab – Array Controller



The Downloading to... table (upper table) lists information on the selected component.

The List of Components... table (lower table) lists information on the available components in the array controllers

Select For Download - This area will contain a check box and the string "Update Firmware" if the firmware file selected using the Choose Firmware button matches the Product ID of that component. This allows selection of multiple components with matching Product IDs.

Any status messages that are generated during a download will appear in the lower area of the page. Messages will scroll out of the message area with the newest message(s) appearing at the bottom of the screen.

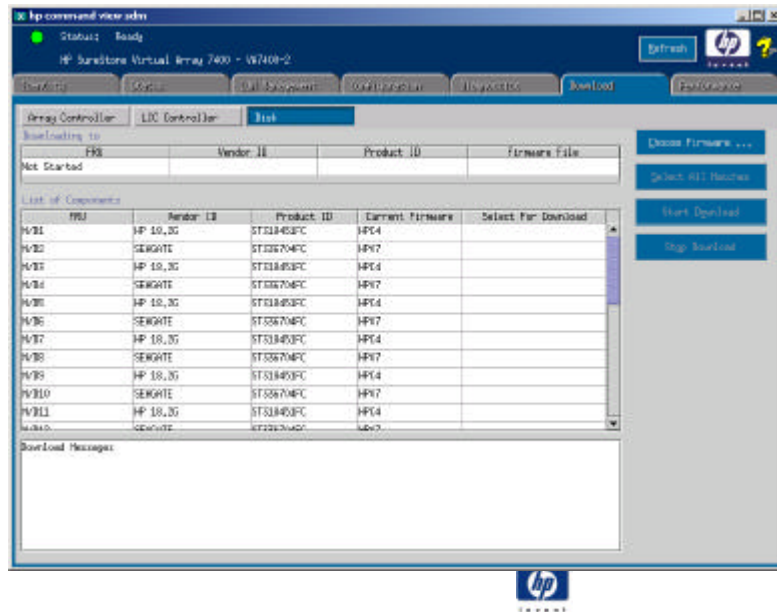
Start OnLine Download - Starts the download of the selected firmware file to the selected component. As the firmware download progresses, progress messages are displayed in the lower area of the page.

Stop OnLine Download - Stops the download of the selected firmware file to the selected component.

Copy Firmware will allow you to select the Source Controller and Target Controller for firmware download.

Clicking the Choose Firmware button brings up the Select Firmware File... dialog box that allows you to browse for a firmware file for downloading. If a firmware file is selected that is "wrapped" (has a .frm extension) then the Update Firmware check box in the Select for Download column will be checked.

Download Tab – Disk



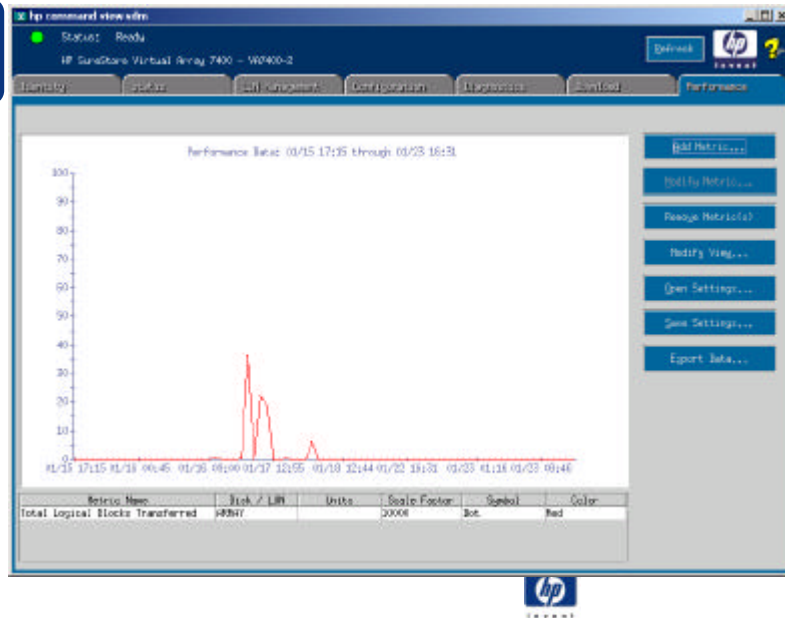
Note! - be sure to backup array data before performing a firmware download. If the download is aborted (or fails to complete), the GUI may stop functioning due to an advisory lock being set. If this happens you will need to break the lock using the `armmgr -b` command.

The Downloading to... table (upper table) lists information on the selected disk.

The List of Disks... table (lower table) lists information on the available disks in the array.

Select For Download - This area will contain a check box and the string "Update Firmware" if the firmware file selected using the Choose Firmware button matches the Product ID of that disk. This allows selection of multiple disks with matching Product IDs.

Performance Tab



The Performance tab contains a single page where the array's performance metrics can be displayed (list of metrics).

Add Metric - Use this button to add metrics to the graph.

Delete Metric - Use this button to select and delete individual metrics from the graph.

New features in this release of command view sdm include opening and saving performance view configurations and exporting data.

Note! - Performance graphs are created from performance statistics (logs) stored on the array. Therefore, the graphs are not "real time." Performance statistics are retrieved from the array approximately every 5 minutes. Graphing a performance parameter consists of retrieving the entries for the desired parameter from the performance log file and displaying the entries on the graph.

Exported Performance Data Example

```
Start Date: Fri May 25 15:05:59 PDT 2001
End Date: Fri Jun 08 15:05:00 PDT 2001
TIMESTAMPS:,ARRAY Host Reads,ARRAY Host Writes
06/05 12:36,2,0,0,0
06/05 12:56,0,0,0,0
.....
06/08 00:42,232891.0,220830.5
06/08 01:43,294276.0,177718.0
06/08 02:44,218536.0,227328.75
06/08 03:45,306243.75,155580.5
06/08 04:46,233980.0,237542.75
06/08 05:48,311168.25,170778.75
06/08 06:48,81890.5,62400.5
06/08 07:48,0,0,0,0
06/08 08:49,0,0,0,0
06/08 09:49,0,0,0,0
06/08 10:49,0,0,0,0
06/08 11:49,0,0,0,0
06/08 12:50,0,0,0,0
06/08 13:50,0,0,0,0
06/08 14:50,0,0,0,0
```



Note there are only 4 columns printed per line. The 5th column of data will be wrapped around and printed on the second line. This results in a mixture of comma and return characters in the data set.

Performance Metrics

- Device-specific performance information
 - Available via command-line and GUI
 - Command-line utility: `armperf`
 - GUI: Performance tab
- Metric calculation uses historical data from `.../sanmgr/commandview/server/logs/performance`
- Purpose:
 - Provide virtual storage architecture performance insight: What is the array doing?



This is not really “new” functionality. Very similar information is provided with the ARM tools for the Model 12 and 12H products. This is the first time a GUI tool available to view trends across the platforms. With the ARM software, the GUI was only available on Windows.

If you are familiar with the ARM tools and are curious about the “recommendations” functionality, we plan to add “smart” performance analysis in a future release of Command View sdm. It isn’t on the list of committed functionality for the VA 7400, but we are actively working on it.

End-user Available Metrics

- Per device:
 - RAID 1 Capacity, RAID 5 DP Capacity, Write Cache Size, Read Cache Size, 1 Hour Write Working Set, 11 Hour Write Working Set, 24 Hour Write Working Set
- Per LUN:
 - RAID 1 Allocation, RAID 5 DP Allocation



The metric list available with Command View sdm is significantly larger than with ARM. The primary reason to provide these metrics is to provide information about the array's activities. As you will notice quickly when looking through the available metrics, many of these will just help characterize the workloads.

To construct a higher-level picture, the most relevant information is provided by throughput and bandwidth metrics. These can be used in conjunction with host and application performance information to develop an accurate picture of the whole-system performance.

Future enhancements include integration of metrics into higher-level applications. Again, the information that is device-independent (i.e., directly applicable to other types of storage devices) will hopefully be used by these applications to develop performance bottleneck indicators.

End-user Available Metrics

- Both device and LUN statistics available:
 - Host Reads, Host Writes, Total I/O's, Host Reads (per second), Host Writes (per second), Total I/O's (per second), Logical Blocks Read, Logical Blocks Written, Total Logical Blocks transferred, Logical Blocks Read (per second), Logical Blocks Written (per second), Total Logical Blocks transferred (per second), Write Command Latency (multiple time intervals), Read Command Latency (multiple time intervals), Cache Pages Read, Cache Pages Written, Total Cache Pages Accessed, Read Cache Hits, Write Cache Hits, Read Cache Hit Percentage, Write Cache Hit Percentage, Transfer Length (multiple sizes), Number of Forced Unit Access Commands



Student Notes:

Command Line User Interface (clui)

- Commands used to
 - display information
 - configure array parameters
 - perform utility functions
- Most feature rich interface
- Scriptable Interface



Student Notes:

Command View User Interface (cvui)

- Text based menu system, used to simplify the command line based management to
 - display information
 - configure array parameters
 - perform utility functions
- Allows access to most of the command line options of the clui
- Not scriptable



Student Notes:

Array ID

- Access to the array is through the device identifiers specified on the command line. The device identifiers can be one of the following:
 - Device File
 - Serial Number
 - World Wide Name
 - Alias Name
- These identifiers will have to be prefixed with the "hostname:" if the array is connected to a remote host
- A list of all the devices known to the host can be obtained using
 - first: the "armdiscover" command, and
 - second: the "armdsp -i" command



Student Notes:

FRU Location

- Each element of the virtual array physical hardware can be identified using the syntax
 - <enclosure>/[<component>].[<subcomponent>]
 - <enclosure> is M for Main or J for JBOD
 - <component> is D1-D15 for disk, C1 for controller/LCC, P1 for power supplies, MP1 for Midplane
 - <subcomponent> is F1-F2 for fans, B1 for NVRAM battery, M1 for DIMM, PM1 for Processor, G1-G2 for GBIC
 - Examples of full string commands:
 - Controller M/C1
 - Disk Slot 1 M/D1
 - Fan M/P1.F1



Student Notes:

Summary of Command Line Utilities (1 of 2)

- **armcfg** Configures major attributes of the array such as LUNs and physical drives.
- **armdsp** Displays the status and operating configuration of the disk array.
- **armfmt** Formats the entire disk array.
- **armlog** Accesses disk array logs.
- **armmgr** Configures minor operating parameters of the array.
- **armrbl** Controls rebuilding of the disk array.
- **armrecover** Reconstructs data mapping and array configuration in the event of the loss of NVRAM contents.



Student Notes:

Summary of Command Line Utilities (2 of 2)

- armsecure Manages all security features of the disk array.
- armcopy Manages snapshot features of the disk array.
- armfeature Queries and modifies the disk array's optional features.
- armhost Queries and modifies the disk array's Host Port Behavior table.
- armdiscover Discovers Virtual Arrays connected to the host.
- armtopology Returns host node & port WWN, product, serial #, RG, LUN, & device path information.
- armdownload Used to download controller or drive firmware.
- logprn Reads the logs stored on the host.
- logdel Deletes the logs stored on the host.



Student Notes:

Lab 3: demonstrating the CLUI/CVUI

- Use the CLUI to find out the configuration details of the VA
- Use the CVUI to access different command line options for viewing and configuration VA properties

Performance goals:

- Describe the features of the CLUI & CVUI
- Be able to describe the available management interfaces



Student Notes:

Backup Slides:

Command View CLUI

Command Reference



armdsp

- Displays array state and status
- armsdp { -a [<r>] | -c [<FruLocation>] | -d [<FruLocation>] |
-e [<FruLocation>] | -f | -L [LUN] [<r>] | -p [<FruLocation>] |
-s | -t } <array-id>
- armsdp { -i [<HostAddr>] }
- armsdp <array-id>
- armsdp -? Display extended help message.



Student Notes:

armdsp Switches

- -a Display the information presented by the -e, -c, -d, -L, and -s options. Specifying -r switch will provide allocated capacity for the LUNs.
- -c Display controller information.
- -d Display disk information.
- -e Display enclosure information.
- -f Display a listing of FRUs in the array.
- -i Display the array-IDs of all disk arrays attached to a given host.
- -L Display LUN information.
- -p Display Port Loop ID of host ports.
- -s Display general configuration information for the disk array.



Student Notes:

armdsp Switches

- -t Display topology of back end loops.
- -? Display this extended usage message.
- With no switch specified, general information about the array is displayed.



Student Notes:

armdsp -i Example

```
#armdsp -i  
Serial Number:00USP1001112  
Alias Name:VA7400-2  
Device Path:/dev/dsk/c14t0d0  
World Wide Name:50060b00000972fc  
Product ID:HP-A6189A  
Unique ID:HPA6189A00USP1001112  
Device Host:ua7.rose.hp.com
```

```
Serial Number:00SG10490112  
Alias Name:va7100-3  
Device Path:/dev/dsk/c19t0d0  
World Wide Name:50060b00000921ef  
Product ID:HP-A6188A  
Unique ID:HPA6188A00SG10490112  
Device Host:ua7.rose.hp.com
```



Student Notes:

armdsp Example

```
#armdsp VA7400-2
Vendor ID:_____HP
Product ID:_____A6189A
Array World Wide Name:_____50060b00000972fc
Array Serial Number:_____00USP1001112
Alias:_____VA7400-2
-----
ARRAY INFORMATION
  Array Status:_____Ready
  Firmware Revision:_____38370HP14P1221010834
  Product Revision:_____HP14
CAPACITY USAGE
  Total Disk Enclosures:_____2
  Redundancy Group:_____1
    Total Disks:_____16
    Total Physical Size:_____267.028 GB
    Allocated to Regular LUNs:_____210.01 GB
    Allocated as Business Copies:_____0 bytes
    Used as Active Hot Spare:_____0 bytes
    Used for Redundancy:_____55.56 GB
    Unallocated (Available for LUNs):_____1.458 GB
  Redundancy Group:_____2
    Total Disks:_____14
    Total Physical Size:_____467.398 GB
    Allocated to Regular LUNs:_____300 GB
    Allocated as Business Copies:_____50 GB
    Used as Active Hot Spare:_____0 bytes
    Used for Redundancy:_____94.144 GB
    Unallocated (Available for LUNs):_____23.253 GB
WARNINGS
  None
```



Student Notes:

armcfg

- LUN Configuration
- Included or Down a disk
- Usage(1) (LUN options):
 - armcfg -L <LUN> { -a <capacity> | -d | -x {true | false}} <array-id>
- Usage(2) (Disk or LCC options):
 - armcfg -D <FruLocation> { -a | -d [-v | -F] [-R | -Z] | -r [true | false]} <array-id>
- Usage(3) (Help option):
 - armcfg -? Display extended help message.



Student Notes:

Usage (1) (armcfg -L <LUN>)

- -L <LUN>
 - The LUN to create or delete.
- -a <capacity>
 - Create LUN of size <capacity>.
 - <capacity> will be rounded up to a multiple of 512K bytes. <capacity> can be a number or a number followed by K for Kilobytes, M for Megabytes, or G for Gigabytes. If there is no suffix, Megabytes is assumed.
- -d
 - Delete LUN.
- -x
 - {true | false} Activates or Deactivates the LUN.



1048576K=1024MB

1024=1024MB

1024M=1024MB

1G=1024MB.

Usage (2) (armcfg -D <FruLocation>)

- -D <FruLocation>
 - The operation will be performed on the Disk or LCC at <FruLocation>.
- The syntax for identifying <FruLocation> is:
 - <enclosure>/<component> where <enclosure> identifies the JBOD or main enclosure, and <component> identifies the disk mechanism.
 - <enclosure> can be JA0-JA9 for the JBOD on loop 0, JB0-JB9 for the JBOD on loop 1, or M for the main enclosure.
 - <component> can be D0-D9 for a JBOD disk mechanism, D0-D1 for a main enclosure disk mechanism, or C0-C1 for a JBOD LCC.



Student Notes:

Usage (2) Continued

- -a
 - Add device.
- -d
 - Down device.
- -F
 - If -d is specified auto-Fail the drive.
- -R
 - If -d is specified down the disk even if it means losing redundancy.
- -Z
 - If -d is specified down the disk even if it means losing data.



Student Notes:

Usage (2) Continued

- -r [true | false]
 - Reset device.
 - If "true", a Loop Initialization reset L_port loop primitive is used to reset the device.
 - If "false", a Target Reset Task Management function is used to reset the device.
- -v
 - Assume that the drive is Valid on reinsertion of the disk.



Student Notes:

armcfg Examples

- `armcfg -L 0 -a 10M ulxr1n:va7100-1`
 - Create a 10 MB LUN 0 on array alias va7100-1 through host ulxr1n
- `armcfg -D JA0/D2 -r false va7400-2`
 - Use a target reset to reset disk 2 in JBOD A0.



Student Notes:

armdiscover

- Scans the host for virtual arrays.
- armdiscover [-silent] [-verbose] [-noscan]
- Switches
 - -silent
 - -verbose
 - -noscan



Student Notes:

armfeature

- Install a software feature license to use.
- armfeature {-r | -a -f <featurestring> -k <key>} <arrayld>
- armfeature -? Display extended help message.
- Switches:
 - -a
 - Add a new enabled feature to the array's feature table
 - -f
 - Feature String of the feature to be enabled
 - -k
 - Key value for the feature to be enabled
 - -r
 - Reads the array's options feature table



Student Notes:

armhost

- Writes and/or Reads the host id table. (host port behavior table)
- Used in heterogeneous configurations
- armhost { -r -f <file> | -w -f <file> | -d <nodewwn> } <arrayld>
- armhost -? Display extended help message.
- Switches:
 - -d - Deletes all the entries for the host Id <nodewwn> from the table



Student Notes:

armhost Switches Continued

- -f
 - Specifies the file which contains (or will contain) the host port behavior table.
 - When used with the -r option, the file will be written from the settings in the array.
 - When used with the -w option, the file will be read and the settings written into the array.
- -r
 - Read host port behavior table from device and write the table from the array into the specified file.
- -w
 - Write host port behavior table to device (read the table from the file specified by the -f option and write it into the array).



Student Notes:

armhost Examples

- Host port behavior table example:
10000060b03e26c8 HPUX
10000060b03e25c3 Win200
- armhost -w -f hosttable.txt va7400-1
 - Writes the information in hosttable.txt to the array alias va7400-1.
- armhost -d 10000060b03e25c3 va7400-1
 - Removes all entries for the named node from the behavior table on array alias va7400-1.



Firmware version HP14, valid host_port_behavior values are HPUX, WinNT, Win2000, Linux, Solaris, AIX, NetWare, MPE, Tru64, OpenVMS, and UnsupportedOS.

Firmware version HP11 and HP13, valid host_port_behavior values are HPUX, WinNT, Win2000, Linux, Solaris, AIX, MPE, or NetWare.

Firmware version HP01 and HP02, valid host_port_behavior values are HpuxFCDriver or Windows/LinuxFCDriver.

Any blank lines or lines beginning with # in the file are ignored. Use with care, the comments are stripped out and not available when read back out.

armmgr

- Basic Management functions for the array.
- `armmgr {-a {on|off}}`
`-b`
`{CreateLun|FwDownload|Security|Select|PassThru|HostPort}`
`-c <FruLocation> | -d | -D <aliasName> | -f {on|off}`
`-h {on|off} | -i {on|off}`
`-J {SingleController | Secure | Normal |`
`RestrictedNormal | HighPerformance}`
`-l <FruLocation> <int> | -n <value> |`
`-p {start|disable} | -r {on|off} | -R [full|partial] |`
`-s {shut|start} | -t <nn%> | -w {on|off} | -x {on|off} |`
`-y {private|public|fabric} <controller> } <array-id>`
- `armmgr -?` Display extended help message.



Switches:

- a Set Auto-Rebuild on or off.
- b Break an advisory lock.
- c Reset one of the internal fibre channel links.
- d Set the fibre channel parameters to default settings.
- D Set the array alias name.
- f Set the Auto Format Drive parameter.
- i Set Auto Include on or off.
- J Set the Resiliency Level.
- l Set the PortLoop ID of a FRU location.
- n Set a limit on the number of LUNs that may be created.
- o Set the Optimize Policy on or off.
- p Set the current state of the scrubbing policy.
- pd Set JBOD port data rate.
- ph Set Host port data rate.
- q Set the Queue Full Threshold.
- r Change the apparent state of read cache disabled.
- R Reset the disk array. A full reset is done if no parameter is supplied.
- s Shutdown (shut) or restart (start) the disk array.
- t Set Capacity Threshold warning to the percentage specified by <percent>.
- w Change the apparent state of write cache.
- x Disable NVRAM on UPS absent.

armmgr

- HP02 firmware or prior
 - B {HpuxFCDriver| Windows/LinuxFCDriver} <controller>
 - h {on|off}
 - y {automatic|private|public|fabric}
- HP11 firmware and greater
 - B {Hpux| WinNT| Win2000| Linux| Solaris| AIX| Netware} <controller>
 - h {None| Automatic| Largest Disk| Largest Two Disks}
 - y {private|public|fabric} <controller>
 - l <FruLocation> {on|off}
 - C {RAID1+0| HPAutoRAID}
- HP 14 firmware and greater
 - o {on | off | true | false}
 - pd {1gb | 2gb}
 - ph {1gb | 2gb} <controller>
 - q <value>



- B Set the Host Port Behavior.
- h Set the Active Hot Spare mode.
- y Set the Host Port Topology.
- l Identify a FRU location.
- C Set the RAID Level

- o Set the Optimize Policy on or off
- pd Set JBOD port data rate.
- ph Set Host port data rate.
- q Set the Queue Full Threshold.

armperf

- Retrieves performance metrics.
- `armperf -c <category> [-u <unit>[,<unit>,...]] [-s <starttime>] [-e <endtime>] [-m "<metric>[,<metric>,...]" [-x COMMA] <array-id>`
- `armperf -i <array-id>`
 - Display all available units and metrics for the array. This switch also displays the start and end times of available data.
- `armperf -?`
 - Display extended usage message.



Student Notes:

armperf Switches (1 of 3)

- -c <category>
 - This switch is used to tell the system which type of metrics are wanted. Required parameter.
- -u <unit>[,<unit>,...]
 - This optional parameter tells the system which devices the user wants data from (which LUNs or controllers). The default is all available units for the chosen category. To get a list of available units see the -i switch.



Student Notes:

armperf Switches (2 of 3)

- -s <starttime>
 - This optional switch is the start time for the retrieval of performance data. The format of the <starttime> is MMddhhmm[yyyy] where :
 - MM = Month
 - dd = Day
 - hh = hour (in 24 hour format)
 - mm = minutes
 - yyyy = year (This is optional. If it is not supplied then the current year will be used)
- -e <endtime>
 - This optional switch is the end time for the retrieval of performance data. The format is the same as <starttime> from above.



Student Notes:

armperf Switches (3 of 3)

- -m "<metric>[,<metric>,...]"
 - This parameter tells the system which metrics the user wants to see data from. The default is all available metrics for the chosen category. To get a list of available metrics see the -i switch.
- -x <exporttype>
 - This switch is used to choose an output type for the data. The default is for console output. The only available export type at this time is comma separated output (COMMA).



Student Notes:

armperf Output

```
C:\>armperf -c ARRAY -m Read_Throughput -s 020211002001 pan-2000:myCassini
Start Fri Feb 02 11:00:00 MST 2001
Finish Fri Feb 02 15:41:27 MST 2001
Array ID: HPA6188A00USF1000521
ARRAY
```

	Read Throughput (per second)
02/02 11:56:24	0.0
02/02 12:11:25	0.0
02/02 12:25:51	3.0
02/02 12:40:54	0.0
02/02 12:55:56	1.0
02/02 13:10:57	0.0
02/02 13:25:59	0.0
02/02 13:41:05	1.0
02/02 13:56:11	25.0
02/02 14:11:13	296.0
02/02 14:26:15	100.0
02/02 14:41:17	239.0
02/02 14:56:19	215.0
02/02 15:11:21	251.0
02/02 15:26:24	155.0
02/02 15:41:27	131.0



Student Notes:

armtopology

- Displays information about the hosts and arrays found in the topology. Displays the hostname, host HBA Node and Port WWNs, and the array product number and array serial number, the Controller/RG/LUN and device path for all the LUNs in the array viewable from the host.
- armtopology [<HostAddr(s)>] [-?]
- None
 - Display the information for all of the discovered hosts in a subnet and the associated arrays.
- <HostAddr(s)>
 - Specify a single or multiple hosts using a space delimiter.
- -? Display extended help message.



Student Notes:

armtopology Example

C:\>armtopology ulxr2n

Host	Host Node WWN	Host Port WWN	Product	Serial Number	C	RG	LUN	Device Path
ulxr2n	50060b00000205f7	50060b00000205f6	A6189A	00USP1001112	0	1	1	\\.\PHYSICALDRIVE3
ulxr2n	50060b00000205f7	50060b00000205f6	A6189A	00USP1001112	0	1	0	\\.\PHYSICALDRIVE2
ulxr2n	50060b000001fbbb	50060b000001fbba	A6189A	00USP1001112	0	0	0	\\.\scsi4.0.0.0
ulxr2n	50060b00000205f7	50060b00000205f6	A6189A	00USP1001112	0	11	11	\\.\PHYSICALDRIVE7
ulxr2n	50060b000001fbbb	50060b000001fbba	A6189A	00USP1001112	0	11	11	\\.\PHYSICALDRIVE9
ulxr2n	50060b000001fbbb	50060b000001fbba	A6188A	00SG10490112	0	0	0	\\.\PHYSICALDRIVE5
ulxr2n	50060b00000205f7	50060b00000205f6	A6188A	00SG10490112	0	0	0	\\.\PHYSICALDRIVE4

- Look for the WWN's being reported by the HBA's
 - Match these with what has been placed in the security and host port behavior tables.
- Look for correct device paths
 - Match these to the expected format for the OS.
 - Incorrect format would indicate an incorrect host port behavior table or setting.



C:\WINNT\SYSTEM32>armtopology ulxr2n

Host	Host Node WWN	Host Port WWN	Product	Serial Number	C	RG	LUN	Device Path
ulxr2n	50060b000001fbbb	50060b000001fbba	A6189A	00USP1001112	1	2	4	\\.\PHYSICALDRIVE19
ulxr2n	50060b000001fbbb	50060b000001fbba	A6189A	00USP1001112	0	1	1	\\.\PHYSICALDRIVE16
ulxr2n	50060b000001fbbb	50060b000001fbba	A6188A	00SG10490112			30	\\.\PHYSICALDRIVE12
ulxr2n	50060b00000205f7	50060b00000205f6	A6188A	00SG10490112			2	\\.\PHYSICALDRIVE8
ulxr2n	50060b00000205f7	50060b00000205f6	A6188A	00SG10490112			1	\\.\PHYSICALDRIVE6
ulxr2n	50060b00000205f7	50060b00000205f6	A6188A	00SG10490112			0	\\.\PHYSICALDRIVE3
ulxr2n	50060b000001fbbb	50060b000001fbba	A6188A	00SG10490112			0	\\.\scsi5.0.0.0
ulxr2n	50060b00000205f7	50060b00000205f6	A6188A	00SG10490112			0	\\.\scsi6.0.0.0

HP
SurePartner
Training

MNS
Crossupdate

Module wrap-up





HP
SurePartner
Training

MNS
Crossupdate

HP Modular Network Storage Solutions

Part 5

Business Copy



Business Copy VA Topics

- Overview
- Positioning
- Cache
- Change LUNs
- Interface
- Backup
- Sample GUI



Student Notes:

surestore Business Copy Virtual Array

- Business Copy VA allows online data replication and backup
- An online copy of LUNs can be dedicated for backup or testing environments



With Business Copy Virtual Array customers can create a nearly instant copy of their data right inside the array. These copies can be used for testing or backup purposes.

Business Copy VA software runs within Command View SDM. Customers can completely configure and manage their data replication environment from one simple point of contact.

And, Business Copy VA works with backup software applications through custom scripting, for all OS environments that are supported on the Virtual Array 7x00.

We have a 50GB demo version currently priced at \$250, and we recommend selling it with every Virtual Array to make sure customers can get there hands on our software and start evaluating it first-hand!

Array Based Data Replication

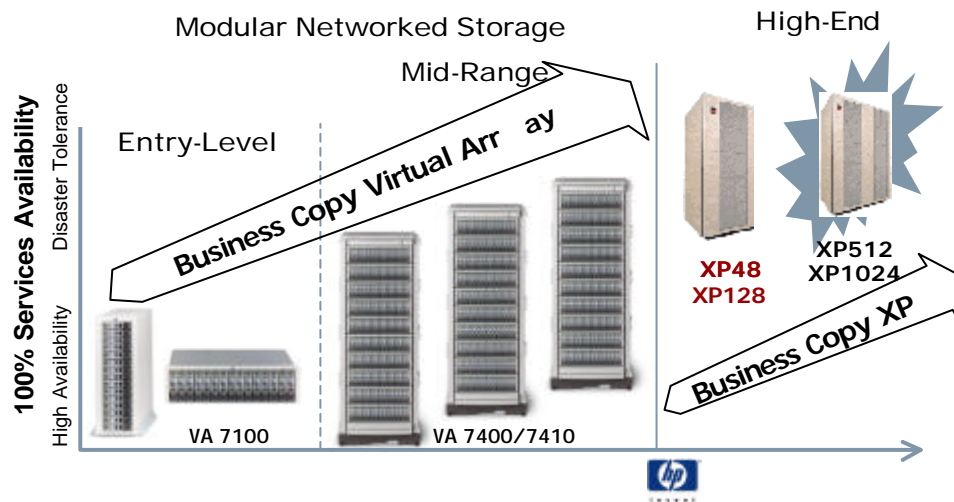
- Makes nearly instantaneous point-in-time online copies of LUNs
 - Uses snapshot technology
- Performs nearly instant online restores of replicated LUNs
- Brings copy and restore functionality to modular storage arrays
- Lets you start with an affordable configuration
 - add capacity licenses as your environment grows
- Command line and graphical user interfaces



Student Notes:

Positioning in hp Product Line

Business Copy Virtual Array allows customers to create online data replication and backup solutions. These are online copies of LUNs that can be dedicated for backup or testing environments.



Student Notes:

Pricing

Business Copy Virtual Array is available as a standalone software product

- License To Use (LTU)
based on the usable capacity of the Disk Array. Customers will buy license capacity to cover only the portion of the array capacity that is being utilized by the software

Business Copy Virtual Array	Upgrade license		
	S/W Kit 50GB DEMO	500GB LTU	1TB LTU
HP Sales Reps (enterprise)	T1007A	T1008A	T1009A
Pricing	\$330	\$4,000	\$7,000
HP Resellers (commercial)	T1026A	T1027A	T1028A
Pricing	\$250	\$3,000	\$5,235



The total size for all Business Copy LUNs cannot exceed the licensed amount for each array. Licensing is based on the total size of all parent LUNs being copied, not the total array capacity.

How Does Business Copy Licensing Work?

- Business Copy licensing is based on the total size of all BC LUNs created. For example, if you're copying 1 terabyte of parent LUN data you must purchase a 1 terabyte Business Copy license.
- Licensing is sold in quantities of 500 Gigabytes and 1 Terabyte. The total size for all Business Copy LUNs cannot exceed the licensed amount for each virtual array.



Student Notes:

Business Copy Upgrade Licenses

- Complete the Software Entitlement Certificate and send to HP by
 - World Wide Web, Phone, Fax, or E-mail
- HP returns a License Key to enable additional capacity based on
 - Certificate registration number
 - Array serial number
- Install the key with armfeature, key, and array SN
armfeature -a -f BUSINESS_COPY_500GB -k
EDC934D901A8 00SG10600107



Student Notes:

Important Features

- Integrates with Command View sdm storage device management software
- Integrates into backup software applications that support a custom scripting implementation
- Has comprehensive user interfaces, Web browser, GUI, CLUI, Command View User Interface (CVUI)
- Supported Operating Systems
 - HP-UX 11.0 and 11.i
 - Red Hat Linux 6.2 and 7.1
 - Windows NT 4.0
 - Windows 2000
- Management Station required for support:
 - Sun Solaris Version 6,7,8
 - AIX 4.3.3
 - NetWare
- Number of LUN Business Copies (with one Parent LUN)
 - Up to 127 BC LUNs with VA7100
 - Up to 1023 BC LUNs with VA7400/7410
- Full RAID protection for the Business Copies



Student Notes:

Business Copy XP is Different

- An XP array has writes for BC going to different drives and can be used for performance
 - Data is mirrored in the XP solution
- Do not use Business Copy VA for performance
 - Data is not actually mirrored
 - Data protection provided by RAID architecture, not Business Copy VA
 - VA solution facilitates backup and testing environments



Student Notes:

Business Copy VA LUNs

- Business Copy VA LUNs are created from Parent LUNs (source of data)
- A Business Copy LUN is created in milliseconds because it copies the map to the data
- Data does not physically move from the Parent LUN to the Business Copy LUN
- A Business Copy LUN is a complete copy of the Parent LUN map at the time it was copied, data is not 'mirrored'
- Each BC reserves the same amount of physical space as its Parent LUN, allowing the BC LUN to change up to 100% of its data



The array controller uses a logical-to-physical data map stored in NVRAM cache.

How to Create a BC LUN

- When you create a BC volume, no data is copied, only pointers to the data (metadata)
- Data will be copied only when the blocks are modified for either the Parent LUN or BC LUN
 - The entire cluster will be copied if a block in the cluster is modified
 - Data copy occurs at the same time the I/O is serviced
- To 'resync' the BC LUN, no data is copied, only the pointers to the data are updated



Student Notes:

Create Business Copy LUNs from Parent LUNs

Virtual Views of Data

Actual Data

Original Copy

Block 1	↔
Block 2	↔
Block 3	↔
Block 4	↔
Block 5	↔
Block 6	↔
Block 7	↔
Block 8	↔
Block 9	↔
Block 10	↔

Business Copy

Block 1	↔
Block 2	↔
Block 3	↔
Block 4	↔
Block 5	↔
Block 6	↔
Block 7	↔
Block 8	↔
Block 9	↔
Block 10 New	↔

Virtual Array
Pointers (Maps)

Block 1
Block 2
Block 3
Block 4
Block 5
Block 6
Block 7
Block 8
Block 9
Block 10
Block 10 New

Reserved Storage
Required For BC LUN





Student Notes:

All of the disks in the virtual
array combine to create a
large physical pool of storage.

Clusters are striped across the
spindles

Physical Pool Of Storage






Student Notes:

[illegible]

Chapter 06 - 16

LUN 0

Virtual Storage Map



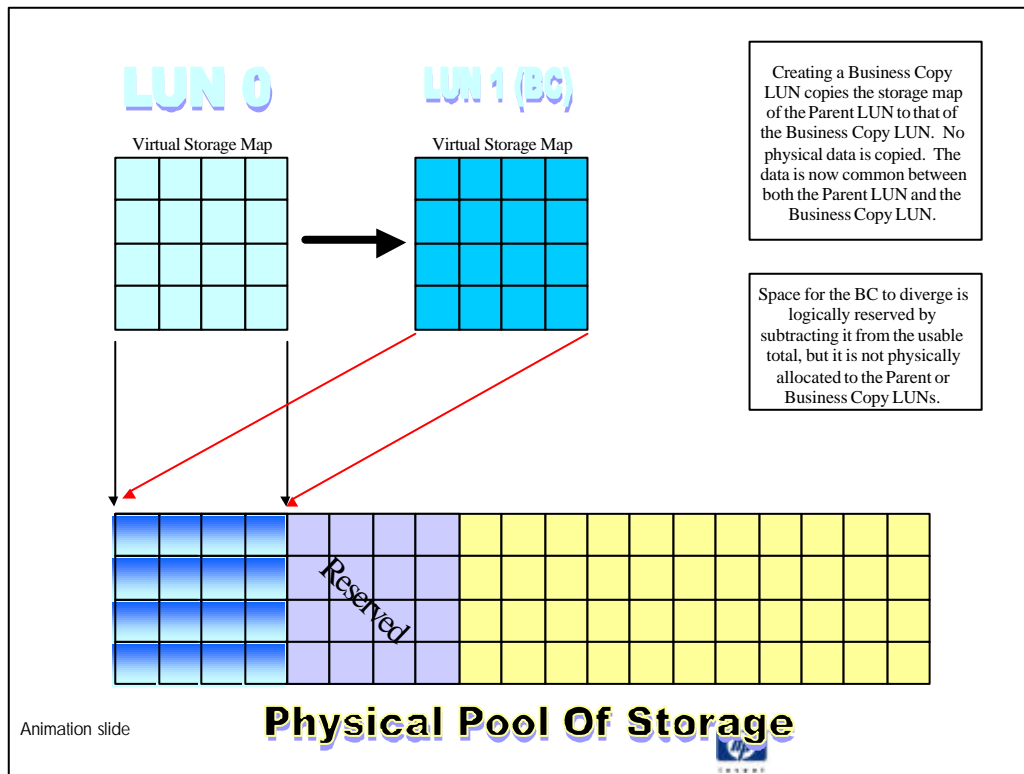
Data that is written to the LUN will result in physical space being allocated to the LUN and pointers to the data being entered into the storage map.

Animation slide

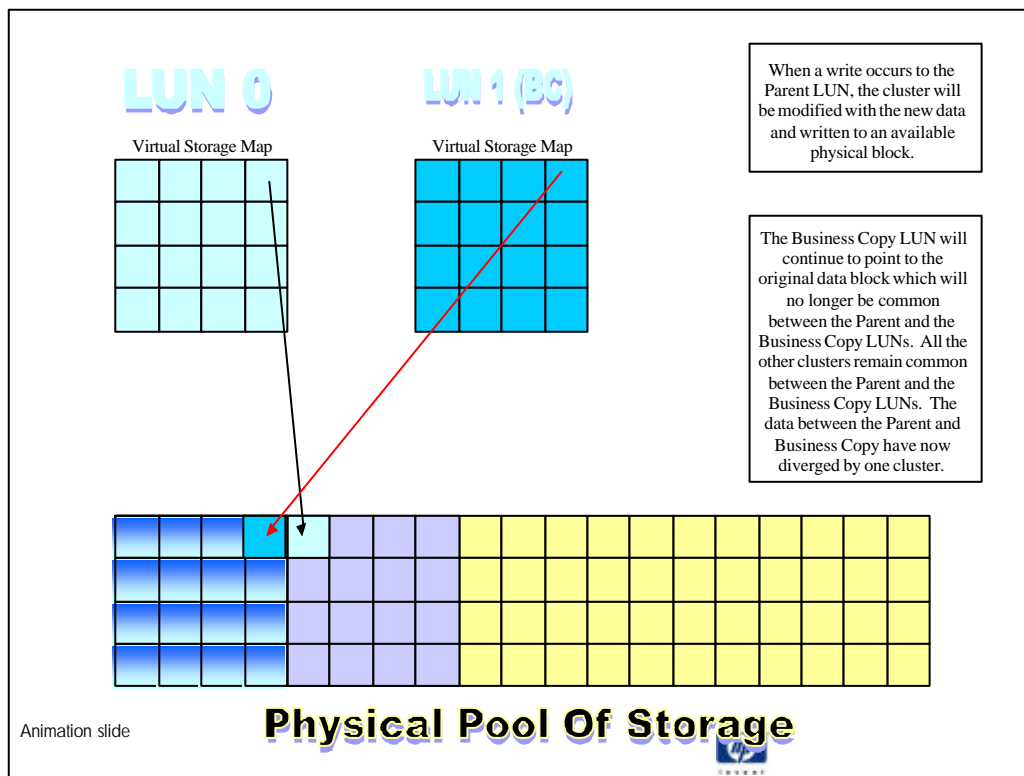
Physical Pool Of Storage



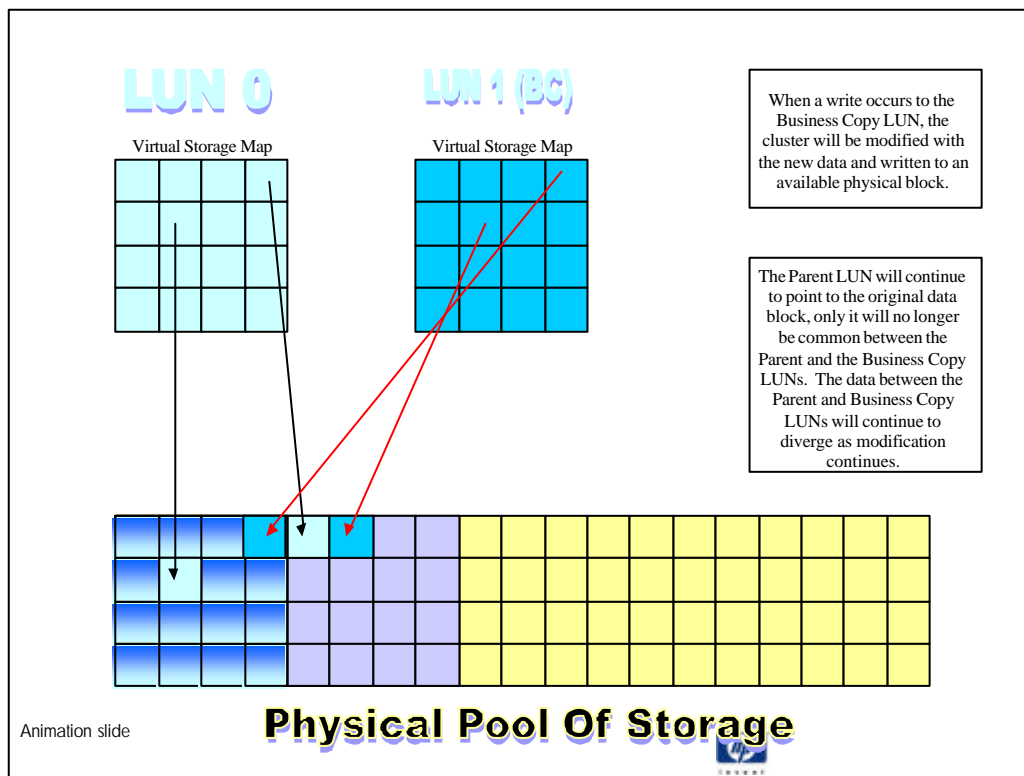
Student Notes:



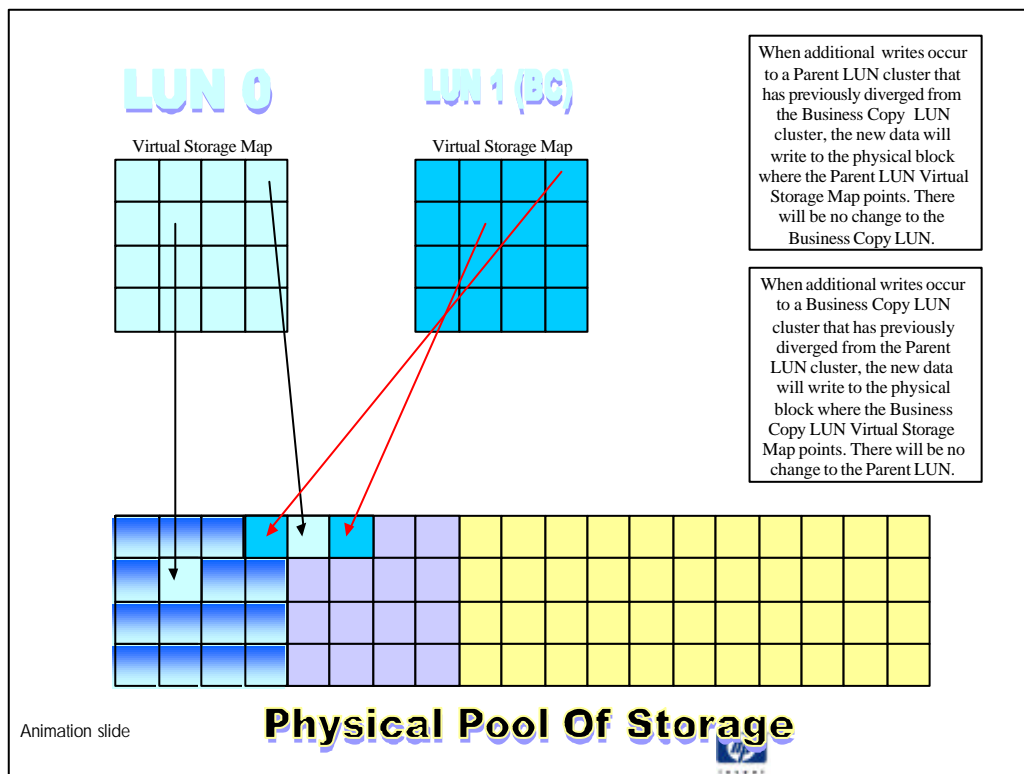
Student Notes:



Student Notes:



Student Notes:

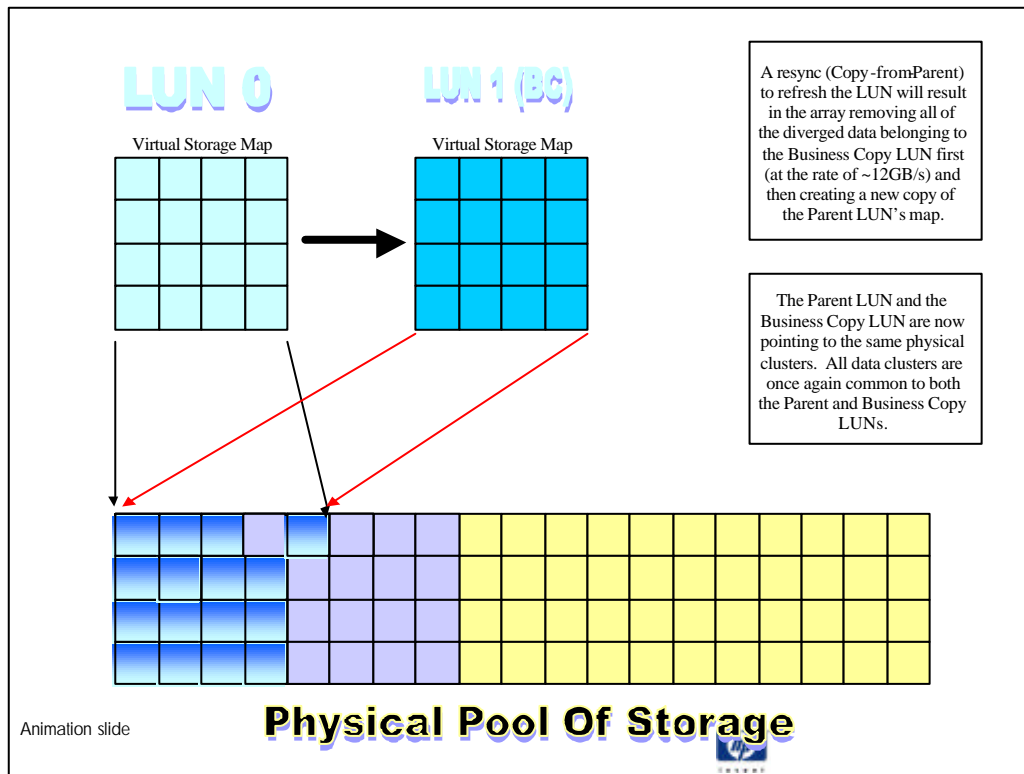


If the data of the Parent LUN blocks is modified for all the clusters in the LUN, then the table that does the logical to physical mapping will point to all new clusters for the Parent LUN. The table for the Business Copy LUN will point to all the original locations. If the data of the Business Copy LUN blocks is then modified in all the clusters of the LUN, then all of the data modifications will take place in the original locations, since the Business Copy map is now uniquely pointing to those clusters.

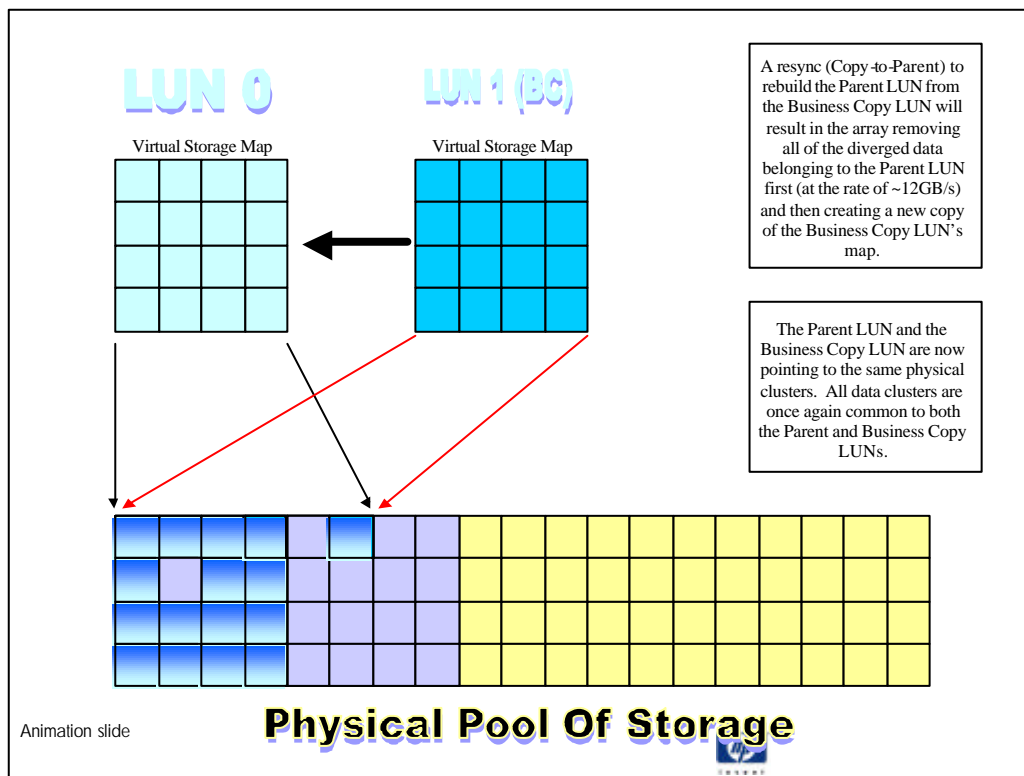
The tables for the pointers are the same size for both the Parent and BC LUNs. At creation of the BC LUN, they both have the exact same number of entries in the tables (which are pointers to where the clusters physically reside on the disks). At the time of creation the tables are both pointing to the exact same locations. After the divergence (blocks of clusters modified in the Parent LUN) they still have the same number of entries. Only now the Parent LUN map is pointing to all new clusters and the Business Copy LUN map is still pointing to the old clusters.

Data is not being stored in the table. The data is stored on clusters and the tables point to those blocks.

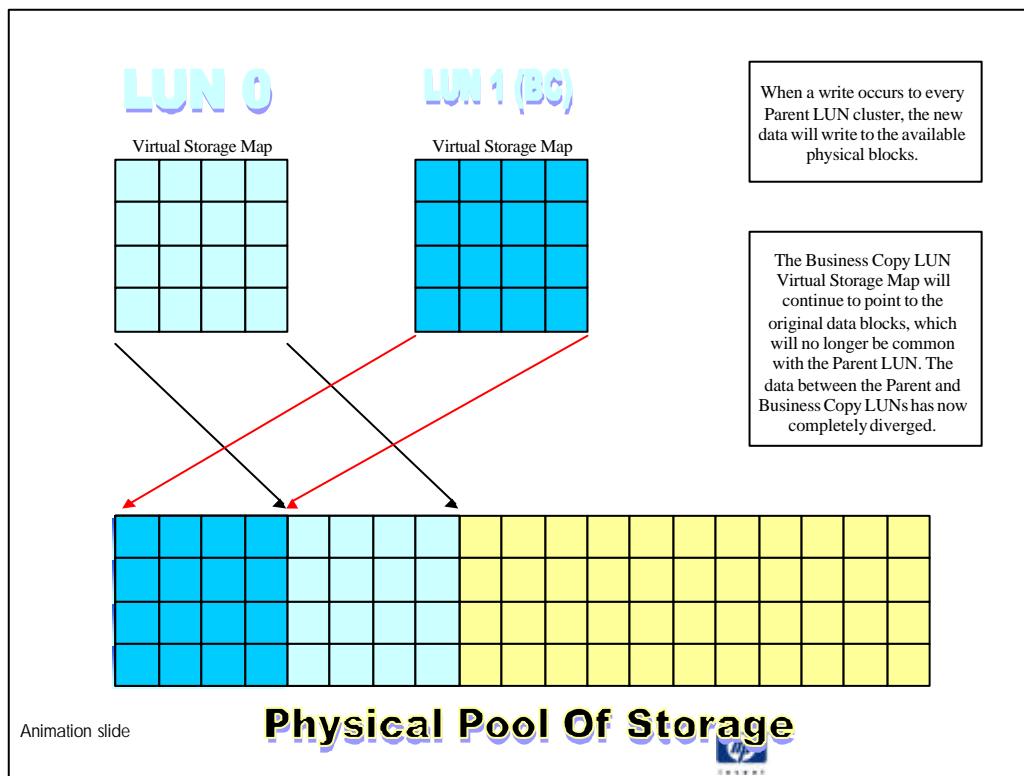
Keep in mind that the LUN is a logical entity on the Virtual Array. The physical mapping will constantly be changing as the array dynamically adjusts the data based upon access (e.g. AutoRAID) and the effects of things like Business Copy and firmware functions such as "promote", "balance", "optimize", error recovery, etc. The backend is continually changing and adapting but all that is masked by the tables that map the logical LUNs to the physical disks.



Student Notes:



Student Notes:



Student Notes:

BC Functionality with Redundancy Groups

- Business Copies are redundancy group dependant
- The 7100 has one RG, the 7400 has 2 RGs
- Business Copy LUN can only be in the same redundancy group as the Parent LUN
 - The copy of a LUN cannot be in a different redundancy group, since it is metadata
 - A customer may want to fill a RG and then create Business Copies. Make certain there is disk capacity in the RG for the BC LUNs



Student Notes:

Backup Applications

- Business Copy VA LUNs are available to the host and can be used like any normal LUN
- Business Copy VA has the potential to work with every backup application that runs on supported operating systems
- Scripted commands to the virtual array using the CLUI can automate the backup application process



Student Notes:

Support for Backup Software

- Goal is to support integration of Business Copy VA with the following backup software applications: Omniback II, Veritas NetBackup, CA-ArcServe, Legato and Tivoli.
- Complete integration with Omniback is released. A zero downtime backup solution is available from HP Consulting.



Student Notes:

No Continuous Access VA

Oracle and Veritas Volume Manager (VxVM) allow remote replication of arrays, performing the Continuous Access function.



Student Notes:

Copy from/to Parent LUNs

- When you copy a LUN in BCVA, the operating system does not see the copied data until you unmount and remount the LUN
 - Use umount and mount for UNIX
 - Use BCopyUtil or BCUtil dismount and mount for Windows
 - The Windows Utility comes with the media kit



Student Notes:

Mounting Business Copy VA LUNs

- BC LUNs copy the Volume Group ID
- Use `vgchgid` to change the VGID for the BC LUNs on the disk array before mounting on the same system
- Once the VGID has been altered with `vgchgid`, the BC LUNs can be imported into a volume group



Student Notes:

Flush Array Cache

- During the snap, cache is flushed to the Parent LUN, and memory to memory copies are done from the Parent map to the Business Copy map.
- If you want to insure data consistency, you can make the Parent LUN inactive (umount or dismount) at the time the Business Copy is created.
- If the Parent LUN remains active during the creation of the Business Copy, open files on the host computer may not transfer (be flushed) to the array
- Unmounting, or making sure that the data bus is inactive, ensures that all data is written to the Parent LUN.
- BCopyUtil, provided with Business Copy VA, can be used to dismount and mount LUNs in Windows.



Student Notes:

Multiple Business Copy VA LUNs

- Any Parent LUN can have one or more Business Copy VA LUNs
- You cannot make a Business Copy of a Business Copy LUN
- Size of a Business Copy VA LUN is defined by the size of the Parent LUN



Student Notes:

Restore

- The Copy to Parent command, used for restoring data, requires that only one Business Copy LUN exists for the corresponding Parent LUN
 - If multiple copies of the Parent LUN exist, all other BC LUNs must be deleted prior to using the Copy to Parent command



Student Notes:

Performance

- Creating a new Business Copy LUN takes less than 100ms, no actual data is copied
- Deleting a BC LUN and recovering the disk space operates at ~12Gbps
 - This time is associated only with the data that has diverged from the parent LUN (i.e., changed BC data saved in BC disk space)
 - The array scans memory and frees any map space used by the deleted LUN



Student Notes:

Is There a Performance Impact?

- Disk reads are the same speed for Business Copy LUNs as they are for regular LUNs
- Disk writes may not be as fast
 - Writes to a new Business Copy LUN that has not yet diverged may be slower for new write allocations
 - Write speed to a copied LUN that has diverged 100% will be very similar to write speed of a parent LUN
- During a backup operation, the backup I/O and the ongoing host I/O use the same set of disks
 - since the copied data is striped across the disk set and the backup process only reads data, the performance impact is minimal



Run backup during off peak hours to avoid a performance impact.

Interface

- User interfaces for hp Business Copy VA

Graphical User Interface (GUI)

Create BC LUNs

Not all BCVA commands are available in the GUI

Command Line User Interface (CLUI)

Scriptable for integration into a backup application

All BCVA commands available in the CLUI

- Create BC LUNs from Parent LUNs
- Copy from BC to Parent LUNs
- Empty the data from BC LUNs (remove access to data but reserve BC LUN disk space)



Student Notes:

Sample GUI screens

The following slides illustrate the GUI
for Business Copy Virtual Array



Student Notes:

Lab: demonstrating business copy VA

- Use the GUI to create a business copy of an existing LUN
- Empty the copy
- Resync with Parent LUN

Performance goals:

- Describe the features of Business Copy VA

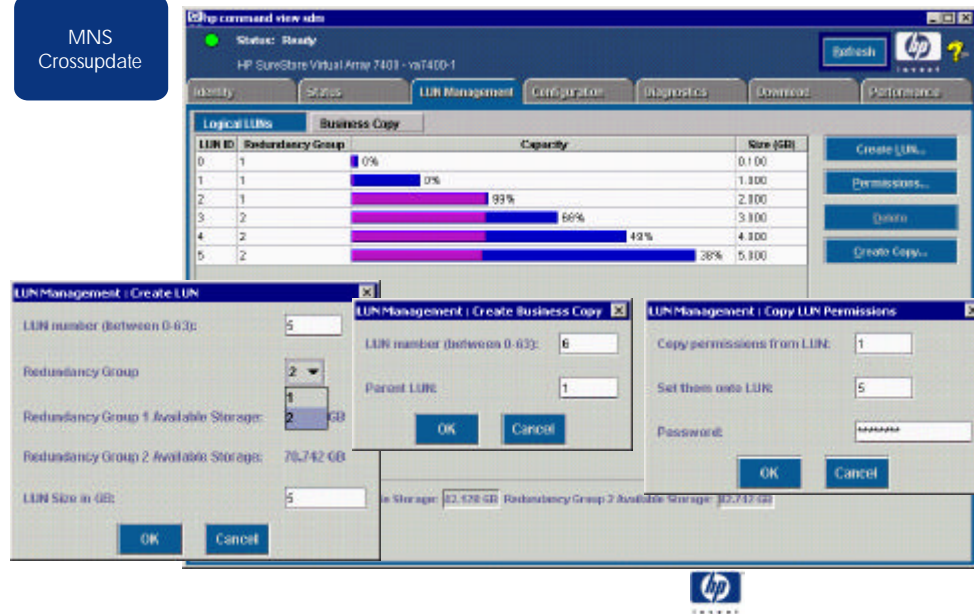


Student Notes:

HP
SurePartner
Training

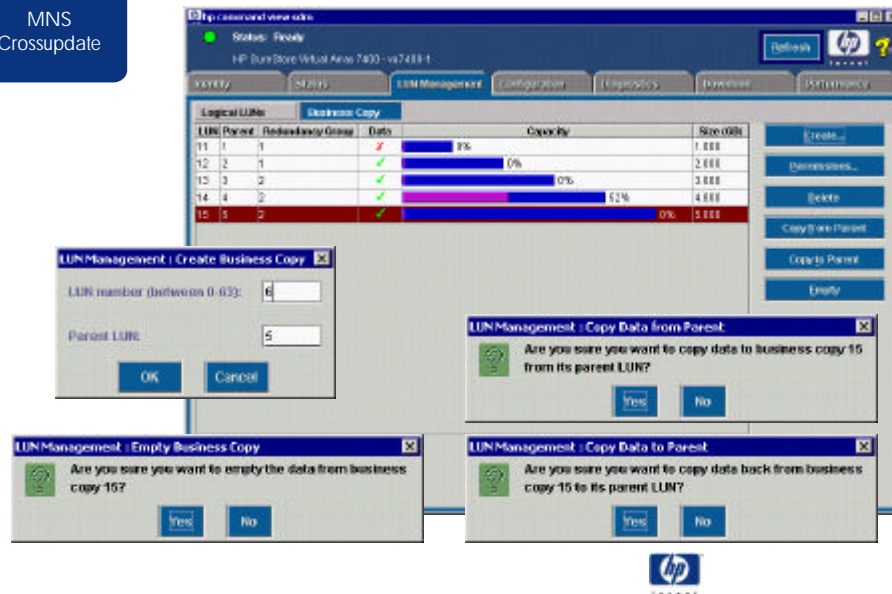
MNS
Crossupdate

Create Parent Logical LUNS



Student Notes:

Create Business Copy VA LUNs

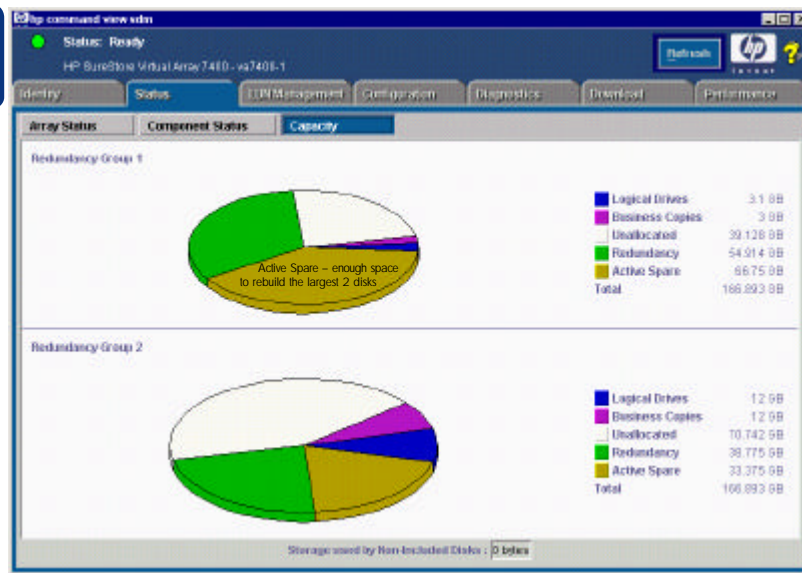


The Windows signature and disk id are copied. The lab will show how vgchgid must be used to mount the Business Copy LUN.

HP
SurePartner
Training

MNS
Crossupdate

Capacity Status



Student Notes:

HP
SurePartner
Training

MNS
Crossupdate

Module wrap-up





HP
SurePartner
Training

MNS
Crossupdate

HP Modular Network Storage Solutions

Part 6

Secure Manager



Student Notes:

Secure Manager VA Topics

- Overview
- Positioning
- Data security
- Security table
- Interface
- Sample GUI



Student Notes:

HP
SurePartner
Training

MNS
Crossupdate

Secure Manager Virtual Array

What is it?

LUN security for shared homogeneous and heterogeneous environments.

What does it do?


Allows LUNs to be locked in a secure shared environment by restricting data availability to specified hosts in a SAN.



Student Notes:

Positioning in hp Product Line

Secure Manager Virtual Array is integrated into hp surestore Command View sdm, a host or web-based device management application that provides a common user interface designed for modular networked storage.



Student Notes:

Pricing

Secure Manager Virtual Array is available as a standalone software product

- One Array License To Use (LTU) is based on usable capacity of the Disk Array

Secure Manager Virtual Array	S/W Kit 50GB Demo	500GB LTU	1TB LTU	5TB LTU
HP Sales Reps (enterprise)	T1003A	T1004A	T1005A	T1006A
Pricing	\$330	\$4,000	\$7,000	
HP Resellers (commercial)	T1022A	T1023A	T1024A	
Pricing	\$250	\$3,000	\$5,250	



A customer must purchase the 50GB LTU to use SMVA. The 50 GB LTU includes the media and user's guide.

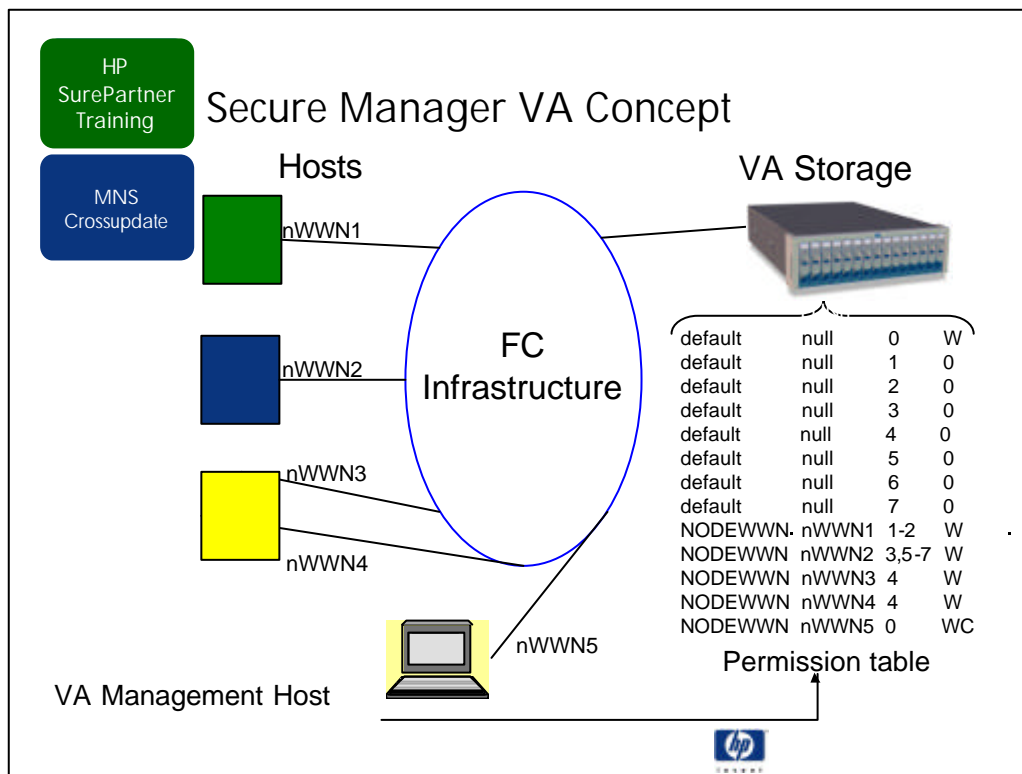
Secure Manger Upgrade Licenses

- LUN security can be enabled/configured for up to 50 GB of usable LUN capacity
 - Licenses are required for additional secured capacity
- Customer completes the Software Entitlement Certificate and sends it to HP by
 - World Wide Web, Phone, Fax, or E-mail
- HP returns a License Key to enable additional capacity based on
 - Certificate registration number
 - Array serial number
- Install the key with armfeature, key, and array SN

```
armfeature -a -f LUN_SECURITY_500GB -k  
EDC934D901A8 00SG10600107
```



Student Notes:



Student Notes:

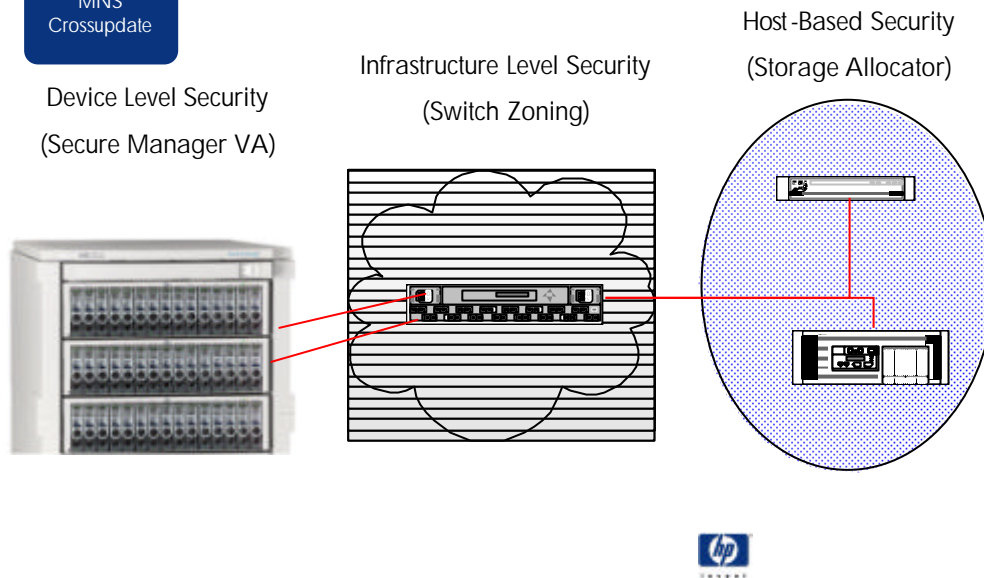
Important Features

- Integrates with Command View sdm
- Comprehensive user interfaces (Web, GUI, CLUI, CVUI)
- User definable security levels by LUN (read/write, write/configure, no access)
- Supported Operating Systems
 - See general OS support
- Secure Participants (WWN principles)
 - 8 secure WWN principles at Initial Release for VA7100
 - 128 secure WWN principles for VA7100 HP11 or greater
 - 128 secure WWN principles for VA7400/VA7410
- LUN Support
 - 64 secure LUNs for each VA7100 at initial release
 - 128 secure LUNs for VA7100 HP11 or greater
 - 1024 secure LUNs for VA7400/VA7410



Secure principle means HBA or basically WWN !

LUN Security Choices – Three Levels



Student Notes:

LUN Security – Storage Allocator


Host Level Security

Advantages:

- Independent of target devices
- Mixed heterogeneous devices
- Ease the management of a storage pool

Disadvantages:

- Management software is host/HBA specific and must be present on all hosts in the SAN to be effective. It must be a co-operated, trusted environment.
- Not a bullet proof security scheme - a host can be plugged into the SAN where it can corrupt data. This scenario is particularly likely in a multiple campus situation.



Student Notes:

LUN Security – Switch Zoning

Infrastructure Level (Switch Zoning)


Devices are segmented based on the ports they connect to a switch

Advantages:

- Independent of Hosts and target devices
- Safeguard SAN against unauthorized hosts

Disadvantages:

- Granularity is at port and node level (Not LUN level)
- Attached nodes on the ports cannot communicate unless they share a common zone
- When multiple, separated SANs are connected to the storage, zoning information may not be consistent
- Most switch vendors use soft zoning when flexibility is required (i.e. moving physical cable connection). This level of security can be breached.



Student Notes:

LUN Security – Secure Manager VA

Device Level (Secure Manager VA)

Advantages:

- Best granularity - LUN level
— Segment nodes by device IDs
- Best safe guarded – from anywhere
- Control LUN access in a SAN



Disadvantages:

- Device dependent - low end array or JBOD may not support this function
- Administration may become cumbersome for a large node count (e.g., 200 NT servers sharing a LUN for a mail database)



Student Notes:

Secure Manager VA is a Data Security Manager

- LUN-level data security
- Set LUN permissions to protect sensitive data
- Guard against LUNs being deleted by unauthorized servers or users
- Control which servers see which LUNs
- Change access permissions online at any time
- Secure Manager VA authenticates access by World Wide Name (WWN), a unique fingerprint
- The Secure Manager VA can use both the nodeWWN and portWWN to enforce LUN security.



Student Notes:

WWN

- NodeWWN is the Fibre Channel Node or Port WWN of the remote port
- WWN is a 64-bit physical address that uses the IEEE 48-bit format with a 12-bit extension and a 4-bit prefix
- WWN is hard-coded to each port
- WWN is a unique name assigned only to that port
- Obtain the WWN from each host port that will be granted access to LUNs



Student Notes:

Overview of Procedures for Managing Security

Manage Security using the CLUI and GUI

- Create and set permissions CLUI, GUI
- View installed permissions CLUI, GUI
- Change permissions CLUI, GUI
- Clear security table CLUI, GUI
- Change password CLUI, VFP, GUI
- Enable and disable array security CLUI, VFP, GUI

The array identifies permissions a host has for a specific LUN using:

Array Security Table – an ASCII text file created by the administrator

Write the LUN security table to array memory with armsecure.



Array security table

When creating the array security table, the following text editors should be used:

- HP-UX use vi
- Windows 2K/NT use WordPad. DO NOT USE NOTEPAD TO CREATE THE ARRAY SECURITY TABLE.

Steps for LUN Security Configuration

- Plan Security
 - Number of hosts
 - Which hosts can access which LUNs
 - Access permissions for LUNs
 - WWNs of HBAs for each host and path
- Generate the LUN Security List
- Add LUN Security List using CLUI command armsecure
 - GUI support to modify access list is a planned future enhancement
- Enable LUN Security feature
- Use the LUNs from the hosts



Student Notes:

LUN Security Table

- Format for a line in the security table file

<participant type> <WWN> <LUN-ID> <permissions>

<participant type> can take one of two values

- | | |
|---------|-------------------------------|
| DEFAULT | – LUN permission for all WWNs |
| NODEWWN | – LUN permission for one WWN |

<WWN> WWN for which permissions are set

<LUN-ID> LUN id for which permissions are set

<permissions>

- W Write permission, allows the specified WWN to perform writes and reads to the specified LUN
- 0 No permission, no information about the specified LUN is available to the specified host (WWN)
- CW Config/Write permission, allows array management commands and write/read to specified LUNs



Array security table

Recommended uses for permissions:

- Do not let the C permission stand alone. Always combine the W (Write) permission with C (Config).
- 0 permission may stand alone.
- W permission may stand alone.
- Give all hosts W access to LUN 0, and only the array management station WC permission to LUN 0.

Example LUN Security File

Entries for a permissions file:

DEFAULT	null	0	CW
DEFAULT	null	1	0
NODEWWN	50060b00092078	1	W
NODEWWN	40134098a44412	1	W
NODEWWN	85c24098344b12	1	W

- LUN 0 has default permission of Configure/Write for all WWNs. LUN 0 can be a small dummy LUN w/no data I/O
- Default is set to no access for LUN 1
- 3 WWNs are given Write permission for LUN 1
- All other HBAs have 0 access (DEFAULT) for LUN 1

This file must be uploaded to the array to set the permissions.

Note: If desired, the permissions file can grant access by default



•HP02 – HP13 defaults:

- The default permission for LUN 0 is CW. The default permission for all other LUN numbers is W.

•HP14 – HP16 / A000 defaults:

- The default permission for LUN 0 is CW. If security is enabled, the default permission for new LUNs is 0. If security is disabled the default permission is W.

•If LUN 0 does not exist (i.e., is deleted or masked from that host) and security is enabled, the software doesn't have any way to actually get the command to the array. The users may see an error message that won't necessarily reflect a LUN security issue.

•When a new LUN is added, it automatically receives DEFAULT participant type with all WWN W access

•Any number of entries may be added to the table as long as

- the number of different WWNs does not exceed 8 for HP02 firmware
- the number of different WWNs does not exceed 128 for the HP11 or greater firmware

•If security is disabled, all hosts have access to all LUNs

•If security is enabled, have at least one LUN with WC permission to allow performance of management operations

•There must be a default entry for every LUN.


CLUI Commands

To view an array's permissions, use the following command format:

```
armsecure -r -f <filename> -p <password> <array-id>
```

Note: If a path is not included as part of the -f <filename> when using the armsecure -r (read) or -w (write) options, the file will be placed in the path specified by the user at the time of Command View SDM software installation. If a path is was not specified during installation or is not included with the -f <filename>, the file will be located in the default directory:

- /opt/sanmgr/client/sbin/< filename> for HP-UX and Linux
- C:\sanmgr\client\sbin\< filename> for Windows

 View the security table using a text editor (Word or Word Pad for Windows, or vi for HP-UX).



Student Notes:

CLUI Commands

To download a newly created security table to an array, use the following format:

```
armsecure -w -c -f <filename> -p <password> <array-id>
```

- The `-w` command writes the security table to the array.
- The `-c` command clears any existing security table on the array.
- When the `armsecure` command is used with the `-c` option, to clear the array security table, the array automatically disables security.

To download a modified security table to an array, adding only the WWN/ LUN# entries that need to be added or modified in the array's security table, use the following format:

```
armsecure -w -f <filename> -p <password> <array-id>
```



Student Notes:

CLUI Commands

To change the array password, use the `armsecure` command with the `-n` option:

```
armsecure -n<newpassword> -p <oldpassword> <array-id>
```

To enable security use the `armsecure` command with the `-e` option:

```
armsecure -e -p <password> <array-id>
```

To disable array security use the following command:

```
armsecure -d -p <password> <array-id>
```



Student Notes:

If LUN Access is Denied

- Be sure to double-check the configuration before you enable LUN security. An error in the security list can prevent access
- If the system becomes inaccessible, use the VFP to disable LUN security

vfpsecure -d



All of the commands that get issued to the array from Command View sdm are passed through the “management” LUN (generally LUN 0). This can be ascertained by looking at the device file or path information reported by “armdsp -i”. In order for the array to work when security is enabled, it has to have LUN 0 and has to be able to send commands to LUN 0 (Write) and LUN 0 has to be able to send them on to the rest of the array (Config). When LUN Security is enabled, the array must have a LUN 0 on it and the permissions on that LUN HAVE to be WC. If the Write permission is there, the user will be able to turn security on and off and see LUNs, but that is it. If LUN 0 does not exist (i.e., is deleted or masked from that host) and security is enabled, the software doesn't have any way to actually get the command to the array. The users may see an error message that won't necessarily reflect a LUN security issue. At this point, neither the host nor the software can see the array. (HP-UX requires a successful inquiry to LUN 0 to determine if the device is a multi-LUN device (i.e. array controller)). Ioscan will report a NO_HW state for the LUN. Although armdiscover (or “armdsp -i”) will appear to see the array, it is actually using the “stale” ioscan information. The only option to the user at this point is to use the virtual front panel and issue “vfpsecure -d” to disable the security on the array. Both the host and the software will immediately be able to see the array again.

LUN 0 in Multi-Host Environments

- LUN security should not be applied to LUN 0
 - For heterogeneous support, recommend a minimum 10 MB size for LUN 0, allowing the Windows driver to see LUNs beyond LUNs 0 through 7
 - HP-UX requires a successful inquiry to LUN 0 to determine if the array is a multi-LUN device
- LUN 0 should always be set with "configuration" (WC) permission for the command view sdm servers to allow the software to manage the array
- If there are no LUNs configured with "C" permissions, you will not be able to create/delete LUNs, or modify other configuration options



Student Notes:

Delete LUN 0

- If LUN 0 was not created or LUN 0 is deleted
 - array automatically creates LUN 0 of zero length
- Command view can get confused if the LUN it uses to communicate with the array is deleted
 - If you have problems talking to the array, run `armdiscover` to recover
- `armdsp -i` will show the user which LUN is being used as the management LUN.



`armdsp -i` displays the serial number, alias, world wide name, device file name and unique name of all arrays connected to the host.
<HostName> denotes the host for which the list is required. If no value is specified, local host will be assumed.

Example `armdsp -i` command:

```
C:\>armdsp -i
Serial Number:00SG10600107
Alias Name:va7100-1
Device Path:\\.\\PHYSICALDRIVE6
World Wide Name:50060b000009227
Product ID:HP-A6188A
Unique ID:HPA6188A00SG10600107
```

HBA Considerations

- Replacing an HBA will require re-configuring LUN security
 - The permissions are based on the HBA's WWN
 - depending on DEFAULT permissions, a different HBA may not be able to see any LUNs
 - change the configuration after you change the HBA, else the LUNs may not be available with the new HBA



Uploading the LUN security table before changing the HBA will not work because the FW checks for valid HBA WWNs.

Password for Security

- A password is required to set security permissions
- Enter the password from the CLUI or GUI
- The default password is AUTORAID (capitals)
 - if the password is unknown or has been lost, it can be reset to the default from the VFP

vfpsecure -R



Student Notes:

HA Usage

- If it is an HA (high-availability) configuration, make sure the permissions are set to allow fail-over



Student Notes:

Demo: Secure Manager GUI

The screenshot displays the NetMiner 4.0.0.0 application. The main window has a top menu bar with 'File', 'Edit', 'View', 'Tools', and 'Help'. Below the menu is a toolbar with icons for 'Discover', 'Refresh', 'Export', and 'Import'. The main area is divided into two panes. The left pane shows a tree view with 'Hosts' selected. The right pane displays a list of hosts with the following columns: Controller, Host ID, Host Name, and Host IP Address. A 'Host Details' window is open, showing details for host 10.10.10.10, including its IP address, MAC address, and various network statistics. The interface is dark-themed with blue and grey accents.

Controller	Host ID	Host Name	Host IP Address
Controller 001	10.10.10.10	Windows 2008 R2 SP7	10.10.10.10
Controller 002	10.10.10.11	Windows 2008 R2 SP7	10.10.10.11

The 'Host Details' window shows the following information for host 10.10.10.10:

- IP Address: 10.10.10.10
- MAC Address: 00-00-00-00-00-00
- OS: Windows 2008 R2 SP7
- Architecture: x64
- Manufacturer: Dell
- Model: PowerEdge R710
- Serial Number: 1234567890
- IP Address: 10.10.10.10
- MAC Address: 00-00-00-00-00-00
- OS: Windows 2008 R2 SP7
- Architecture: x64
- Manufacturer: Dell
- Model: PowerEdge R710
- Serial Number: 1234567890



Student Notes:

Lab 5: demonstrating secure manager VA

- Make sure you have LUN numbers 0, 1, 2 and 3 created
- Use the GUI to restrict access of your management server to only LUN 0 and 3
- Enable secure manager
- Verify LUN isibility on the host

Performance goals:

- Describe the features of Secure Manager VA

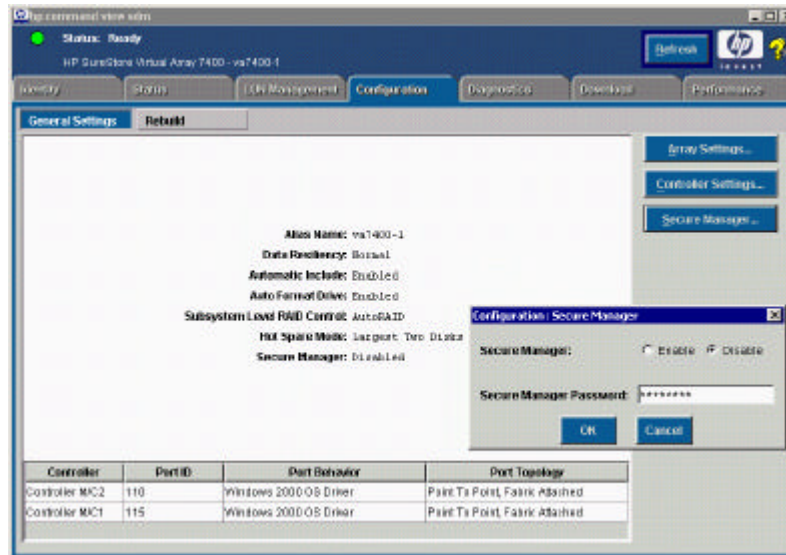


Student Notes:

HP
SurePartner
Training

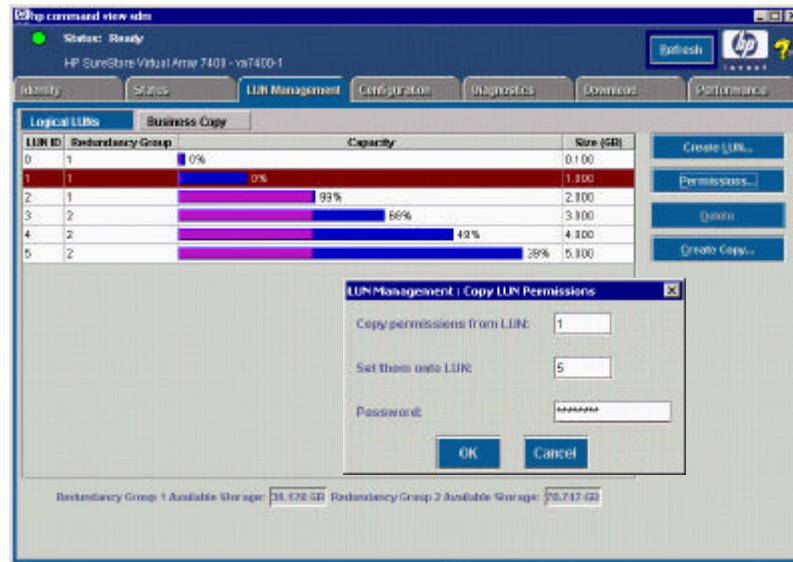
MNS
Crossupdate

Backup Slides (if demo is not available): Secure Manager GUI



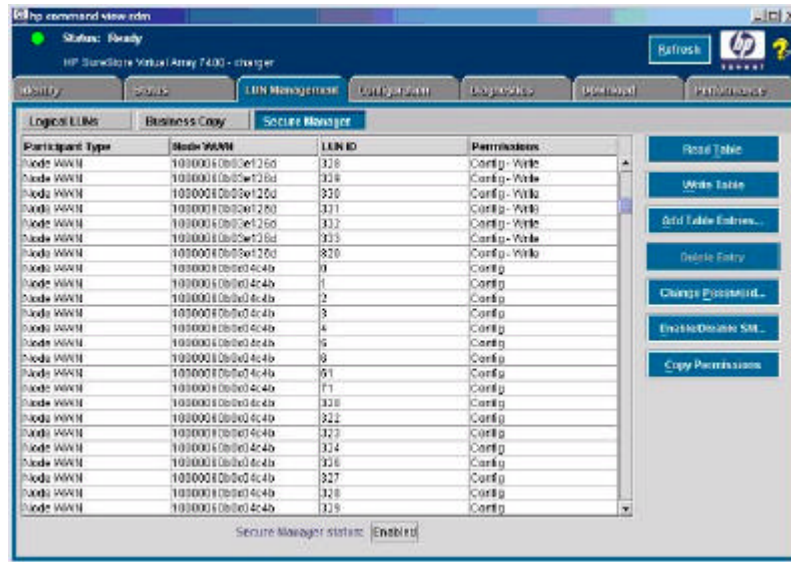
Student Notes:

Copy Permissions



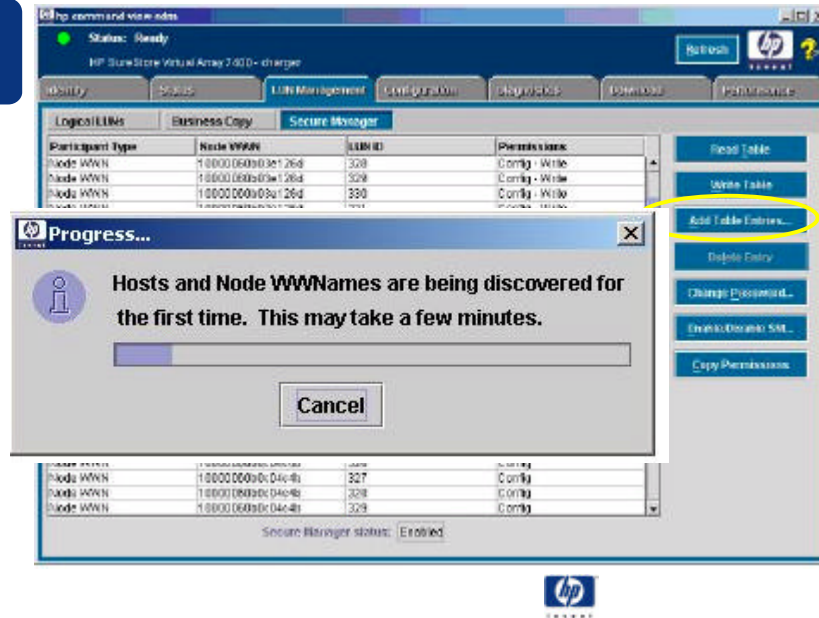
Student Notes:

Secure Manager GUI – v1.03



Student Notes:

Host Discovery



Student Notes:

Security Table Entries

The screenshot shows the 'LUN Management: Add Secure Manager Table Entries' dialog box. It has several sections for configuration:

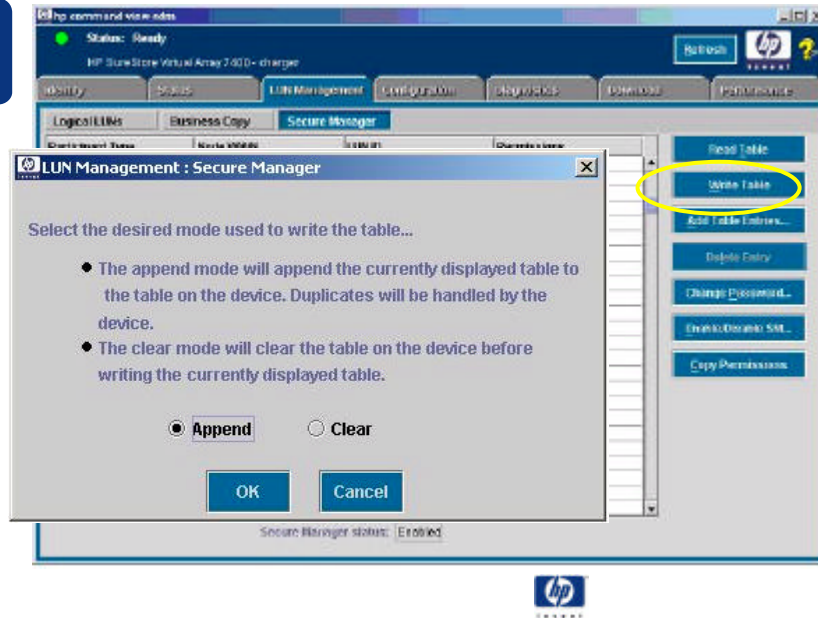
- Participant Type:** A dropdown menu set to 'Node WWN'.
- Host Name:** A list box containing several hostnames, including 'hpbs1712.bol.hp.com' and 'hpbs5577.bol.hp.com'. An 'Add New Host' button is to the right.
- Node World Wide Name:** A list box containing WWN addresses, such as '1000001083fa155:hpbs1712.bol.hp.com'. An 'Add New WWN' button is to the right.
- LINK:** A list box with values 0, 1, 2, 31, and 32.
- Permissions:** A dropdown menu set to 'None'.

At the bottom are 'OK' and 'Cancel' buttons. On the right side of the dialog, there is a vertical toolbar with buttons: 'Refresh', 'Download', 'Permissions', 'Reset Table', 'Write Table', 'Add Table Entries...' (which is circled in yellow), 'Delete Entry', 'Change Password...', 'Enable/Disable SMI...', and 'Copy Permissions'.



Student Notes:

Write Table



Student Notes:

Student Notes:

Student Notes:

HP
SurePartner
Training

MNS
Crossupdate

Module wrap-up





HP
SurePartner
Training

MNS
Crossupdate

HP Modular Network Storage Solutions

Part 8

Auto Path



Student Notes:

Auto Path Objectives

By the end of this module, students will be able to identify the purpose and use of Auto Path. Students will also be able to perform the following tasks:

1. Understand Auto Path's Load Balancing and Path Fail-Over features.
2. Understand Auto Path's supported hardware configurations.



Student Notes:

What is Auto Path?

Auto Path is a device driver that provides automatic input/output (I/O) path fail-over and load balancing for host systems configured with multiple host bus adapters (HBA's) and connections to a disk array. Auto Path is host installed software that utilizes redundant connections between the host server and disk storage in a XP and/or VA disk array subsystem to provide enhanced performance and availability.



Availability: Automatic path recognition and fail-over recovery to an alternate path.

Performance: Definable load balancing policies and device path information.



Auto Path is implemented as a pseudo driver that is installed during boot, layered between the disk manager (Windows) or LVM (HP-UX) and the SCSI device drivers. The layering is achieved by replacing SCSI device driver entry points in the bdevsw and cdevsw device switch tables with the Auto Path driver entry points. The SCSI device driver entry points are invoked by the Auto Path driver. The Auto Path pseudo driver provides the Command Line Interface. Multipath I/O processing

The driver performs the following I/O processing:


- Detects I/O transactions intended for multipath devices
- Passes through transactions intended for single-path devices
- Detects path failures and activates alternate paths
- Detects the “return” of failed paths and activates path recovery
- Implements load balancing policies
- Handles LVM-specific IOCTL's and timeouts
- Records the status of each path

Device Discovery

Auto Path recognizes hardware at system boot and when a failed device is added back on line (recovered). Auto Path detects and records paths to multipath devices. The LibIO Interface is used for detecting the devices and their paths.

Auto Path Path Fail-Over

In the event of a path failure, Auto Path automatically switches to an alternate path, dropping the failed path out of the I/O rotation. This switchover is completely transparent to the host and applications continue without interruption.



Path Fail-over

Auto Path checks the return status of the I/O, which is done in the SCSI I/O callback function. It will be notified at the end of a SCSI read/write, and will take one of several actions based on success or failure. This interrupt handler does error handling by examining the SCSI command pointer's "result" member. The callback function detects I/O failures, and re-routes I/Os to an alternate path. This is done by replacing the `end_scsi_request` function with the Auto Path driver's `my_end_scsi_request` function. It does what the original end request function does, and also updates the Auto Path's "path" structure and overwrites the failed device path with an alternate device path in the "rscsi_disks" structure.


Path Fail-back

Auto Path spawns a kernel thread (which calls the function `scan_for_device_reentry`). The buffer request code activates the path failback detection mechanism if there's a path failure. The kernel thread (`scan_for_device_reentry`) does the following:

- Issues the "Test_unit_ready" SCSI command on the failed device every 5 seconds
- Once the device is recovered, the thread updates the structure "path" and is ready to use by the load balancing policies.

Auto Path Load Balancing

- Auto Path performs dynamic load balancing while monitoring up to 8 HBA's (paths) and 128 LUN's per server to ensure that the I/O is actually completing its transaction. Load balancing is performed using a policy selected by the administrator.



Load balancing is the way the I/O workload can be distributed “equally” across many HBAs in a single system. Load balancing is accomplished by keeping track of the I/O request statistics for each initiator. I/O requests meant for multipath capable devices are trapped in the buffer request code and processed according to the user-established load balancing policies. Load balancing policies are applied here before sending the buffer contents down to the disk driver. The load balancing policies are implemented in the driver code, which is invoked by the function "Autopath_setlb_policy". The path details required by the load balancing algorithms are maintained and updated in the structure “path.” There is a structure for every path to a device. The load balancing algorithm requires the following information:

- The dev_t, for which the I/O is required
- The load balancing policy selected
- All the alternate dev_t physical paths associated with dev_t

This information and other details are maintained in the structure “path”.

Auto Path Features

✍️ Hewlett-Packard designed and developed products for:

Windows NT

Windows 2000

Linux

HPUX 11.00/11i.

The Sun and AIX versions are OEM'ed from Hitachi.

✍️ The load balancing properties are configurable via a GUI and/or Command Line User Interface (CLUI) in Auto Path.

✍️ The Administrator can set a preferred path for data I/O.

✍️ A Server and Client Agent are installed with Auto Path for Windows 2000/NT. A host with only the Client Agent installed can administer configuration settings on a remote host configured with multiple HBA's and an active Server Agent.



Student Notes:

Auto Path VA Product License Pricing

Auto Path Virtual Array	S/W Kit 1 Host LTU	1 Host LTU	5 Host LTU	10 Host LTU	25 Host LTU
HP Sales Reps (enterprise)	T1011A	T1012A	T1013A	T1014A	T1015A
Pricing	\$4,635	\$4,635	\$15,231	\$23,443	\$46,920
HP Resellers (commercial)	T1030A	T1031A	T1032A	T1033A	T1034A
Pricing	\$3,500	\$3,500	\$11,500	\$17,700	\$35,425



Upgrade licenses



Auto Path VA v1.01.00 for NT now supports the VA 7400. However, Auto Path VA v1.01.00 for NT ****WILL NOT**** be supported when used with a VA 7400 in a Microsoft Windows Cluster Environment.

How are XP and VA Autopath different? Fundamentally there is no difference because the bits are the same. It is sold at different price points for XP and VA. The VA version is a little less expensive. For customers who have XP and VA products in their environment, the XP version should be installed to manage both. The XP product will work with the VA, but the VA product **WILL NOT** work with the XP.

Auto Path Virtual Array has a 90 day software media defect warranty.

Auto Path for Windows 2K Known Issues

Known Problems/Limitations

- ✍ Dynamic disks are supported in non-cluster environments. However, there are known limitations of Windows 2000 in supporting dynamic disks, and Microsoft Clustering does not support dynamic disks. Performance impacts may occur when using the dynamic disk striping feature.
- ✍ On any one server, all Host Bus Adapters must be of the same type.
- ✍ In a cluster, all the servers must be of the same type and all Host Bus Adapters on each server must be of the same type.
- ✍ Load balancing is not supported in a cluster environment.
- ✍ It is recommended to install the Auto Path Software before you install the HBA drivers and Microsoft Cluster Service.
- ✍ Parallel SCSI HBA is currently not supported for Microsoft Cluster environment.



Windows 2K Cluster environment load balancing issue defined:

The present SCSI 2 Reserve / Release Implementation only allows exclusive access of a LUN by a single initiator (HBA) on a single path. Host A or Host B can have exclusive access only over Path A or Path B. However, a host with multiple initiators (HBA's) cannot gain exclusive access of a LUN over multiple paths. Host A can gain exclusive access over Path A or Path B, but not simultaneously over both paths. The same scenario applies to Host B over Path C or Path D. Auto Path must use an "Active/Passive" path configuration and cannot support load balancing over multiple paths.

Auto Path for Windows 2K System Requirements

System requirements for installing Auto Path Driver, Remote Access Server, and Command Line User Interface:

- Windows 2000 with Service Pack 1.*
- 40 MB of free disk space.*
- HP XP256/512/48 or VA7100/7400 disk array with formatted logical devices.*
- At least two supported Fibre Channel Host Bus Adapters.*

System requirements for installing Auto Path Remote Access Client:

- Windows 2000/ Windows NT4/ Windows 95/ Windows 98.*
- 30 MB free disk space.*
- A minimum display resolution of 1024x768 with 65536 colors or higher is recommended .*



Student Notes:

Auto Path Supported HBA's

Auto Path VA for Windows 2000 v1.03.00 Supported HBA's:

➤ *Hewlett-Packard D8602A, D8602B with 3.0.4107.0 driver*

Auto Path VA for Windows NT v1.03.01 Supported HBA's:

➤ *Hewlett-Packard D8602A, D8602B with driver versions 3.0.3245.1 and 3.0.4107.0*

➤ *Emulex LP8000 with driver version V4 -4.52a7*

➤ *QLogic QLA2200F with driver version 7.05.05*

Auto Path VA for HP-UX v1.01.00 Supported HBA's:

➤ *HP-UX OS Version 11.00 or 11.11*

➤ *A5158A HBA for V, L& N Class Servers*

➤ *A3404A HBA for K Class Servers*



HP-UX: On systems with PCI A5158A controller cards, the patch PHKL_23939 for HP-UX 11.00 and patch PHKL_23626 for HP-UX 11.11 needs to be applied if not already installed in the system.

Windows NT: Preferred path selection is not supported on Emulex LP8000 HBA's when connected to HP Virtual Array, in a cluster environment.

Lab 6: demonstrating auto path VA

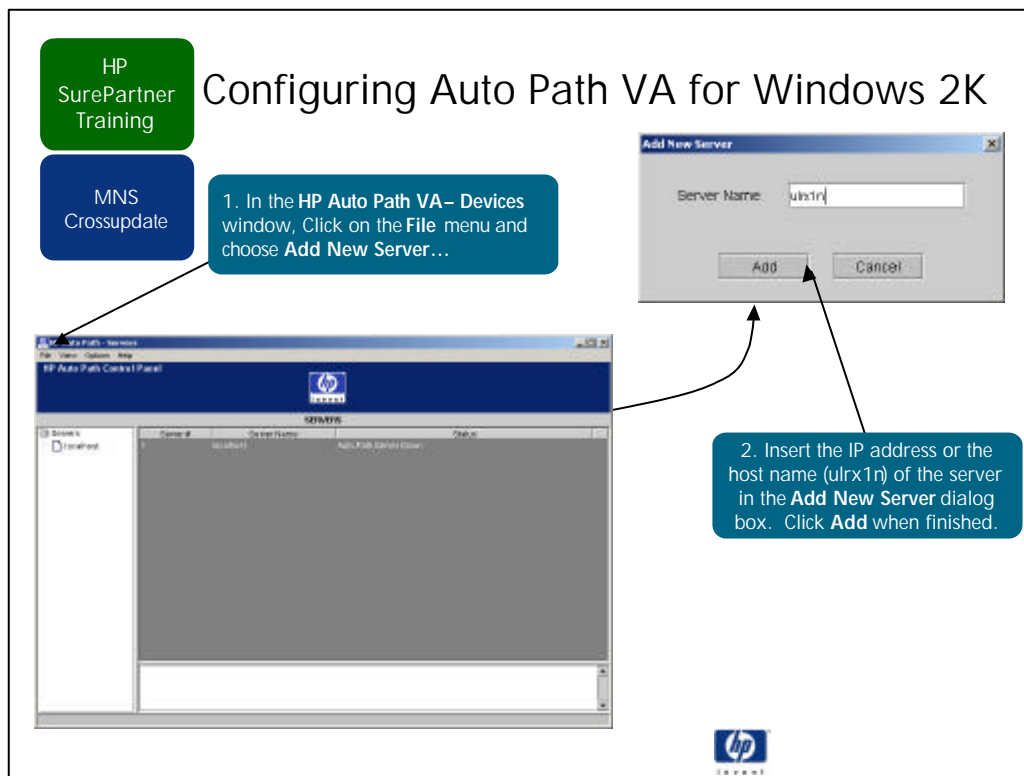
- Use the GUI to display the path connectivity of all visible LUNs
- Change the load balancing policy

Performance goals:

- Describe the features of AutoPath VA



Student Notes:



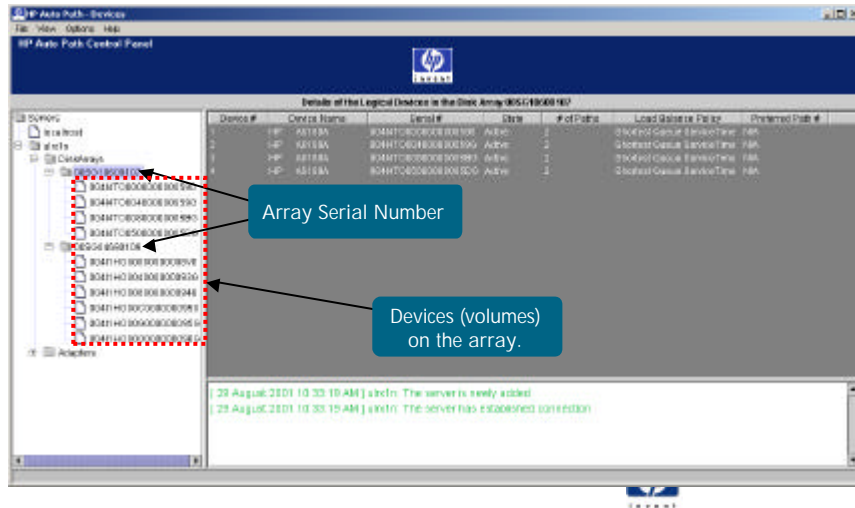
1.To start Auto Path VA, you must first double-click on the **Start HP Auto Path Server** icon on the desktop which causes a DOS window to appear, or simply double-click the Auto Path Client icon on the desktop if you installed Auto Path as a service.

2.Next, double-click on the **HP Auto Path Client** icon located on the desktop. The **HP Auto Path VA – Devices** window opens.

If this is the first time Auto Path has been opened, there will be no Servers or Disk Arrays configured. You must now manually add the Servers to further configure Auto Path VA.

Configuring Auto Path VA for Windows 2K

The Devices window displays the devices for the selected disk array. This is a view- only window; there are no user-configurable functions.



What is now displayed in the **HP Auto Path – Devices** window are all of the devices connected to your host. Below is an explanation of the fields in the Auto Path Devices window:

Device #: An index number used by Auto Path for identifying the device.

Device Name: The device name from firmware.

Serial #: The LUN (Logical Unit Number) serial number of the device.

State: The condition of the device.

Active = device in use.

Failure = device no longer in use or all paths to device have failed.

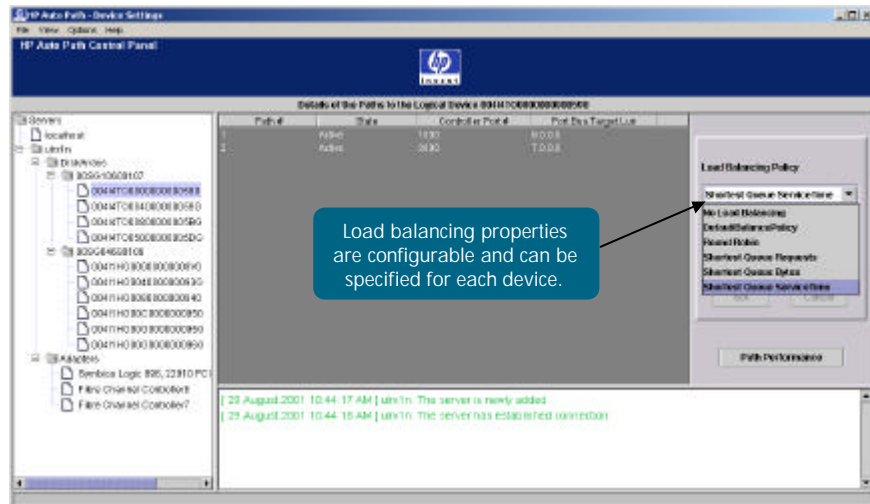
of Paths: The number of paths to the device.

Load Balance Policy: This column displays the Load Balance Policy set for this device.

Preferred Path#: For devices with no load balance policy, this indicates the selected preferred path.

Configuring Auto Path VA for Windows 2K

To change the load balancing properties for a device, you must select a server and a disk array and then click on the device in the **HP Auto Path VA – Devices** window.



Below is an explanation of the fields in the Load Balancing Policy drop-down menu.

No Balance Policy: No load balancing performed. All I/O is routed through the preferred path.

DefaultBalancePolicy: Same as Shortest_Queue_Requests.

Round Robin: The simplest and easiest to understand; No description is necessary other than the above.

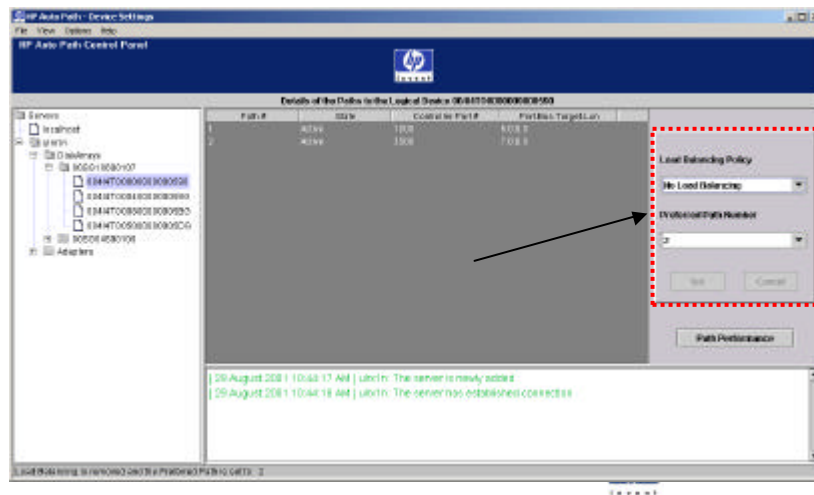
Shortest Queue Requests: Measures the current queue depth by counting the number of requests outstanding for each path.

Shortest Queue Bytes: Measures the current queue depth by the sum of the I/O sizes of the requests outstanding on that path.

Shortest Queue ServiceTime: Measures the current queue depth by the sum of the length of time each I/O request has been outstanding on that path.

Configuring Auto Path VA for Windows 2K

If the **No Load Balancing** policy option is selected, then the **Preferred Path Number** drop-down list becomes available to the user.



Below is an explanation of the fields in the Device Settings window:

Path #: The Auto LUN index number for the path.

State: The condition of the device:

Active = device in use

Failure = device no longer in use or all paths to device have failed.

Port.Bus.Target.LUN: The SCSI address of the path.

Load Balancing Policy: Displays the load balancing policy for the device. The drop-down list allows you to choose from a list of policies.

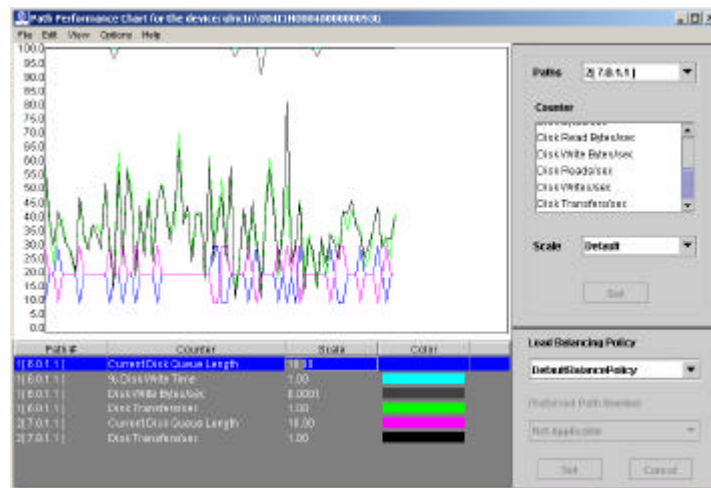
Preferred Path Number: Lets you select the preferred path number to a device.

Note: If the **No Load Balancing Policy** is set for a failed path and the failed path was the designated preferred path, you must reset the recovered path as the preferred path using the **Device Settings** window

Configuring Auto Path VA for Windows 2K

MNS
Crossupdate

The Path Performance window allows the user to create and view performance charts for selected paths of a device. The charts can be used to evaluate the load balancing policy or preferred paths.



Below is an explanation of the fields in the Path Performance Chart window:

Chart parameters

Paths - Drop-down box that is used to select the index number of a path to be charted.

Counter - A list of available counters, measured on the y-axis of each chart.
Choices are:

% Disk Read Time

% Disk Write Time

% Disk Time

% Idle Time

Disk Read/sec

Disk Writes/sec

Disk Transfers/sec

Current Disk Queue Length

Disk Bytes/Sec

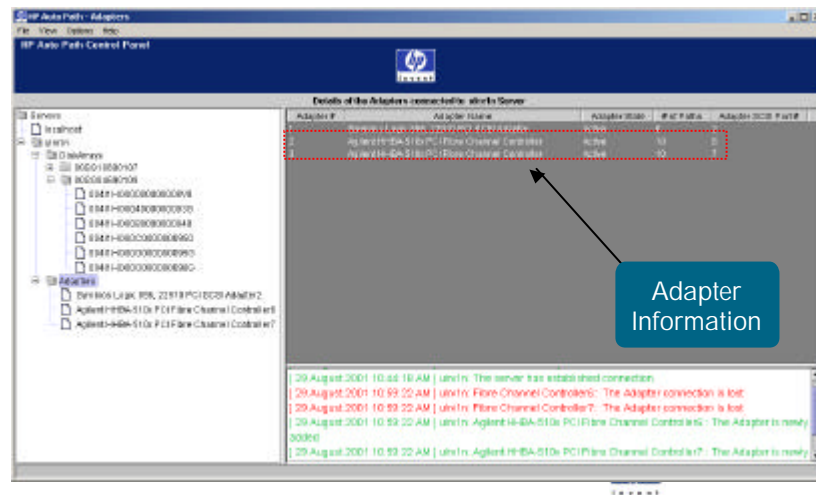
Disk Read Bytes/Sec

Disk Write Bytes/Sec

Configuring Auto Path VA for Windows 2K

HP Auto Path VA – Adapters Window

Displays all of the Auto Path capable adapters (HBA's) on the host.

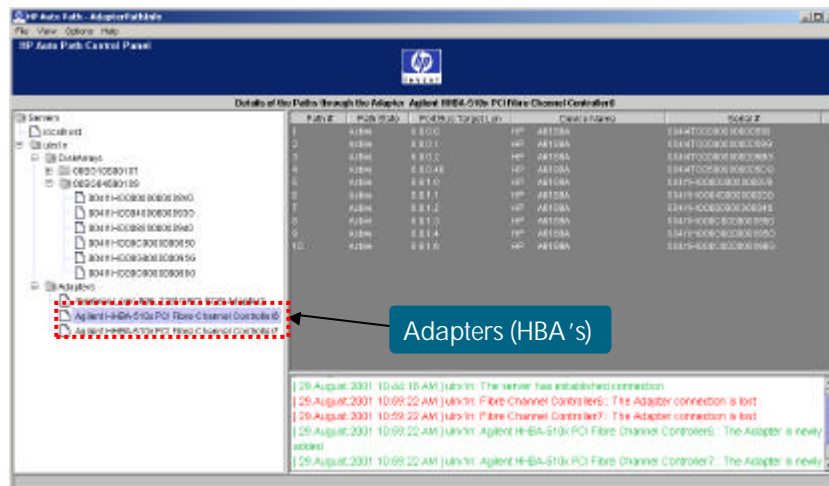


This window displays all of the Auto Path capable adapters present on the host.

1. **Adapter #:** An index used by Auto Path for identifying the adapter.
2. **Adapter Name:** The name of the adapter.
3. **Adapter State:** The condition of the device.
Active = adapter in use.
Failed = adapter not in use.
4. **# of Paths:** The number of paths through the adapter.
5. **Adapter SCSI Port #:** The adapter SCSI port number.

Configuring Auto Path VA for Windows 2K

The AdapterPathInfo window displays the details of the various paths through the adapter. This window is a view-only window; there are no user configurable functions.



Below is an explanation of the fields in the AdapterPathInfo window:

Path #: The Auto Path index number for the path.

Path State: The condition of the device:

Active = device in use

Failure = device no longer in use or all paths to device have failed.

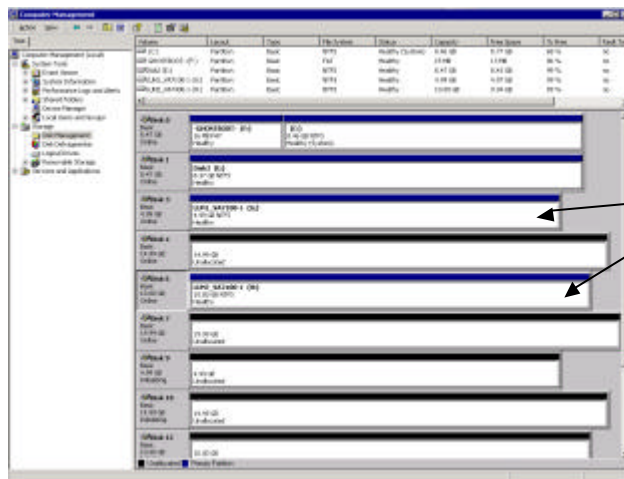
Port.Bus.Target.LUN: The SCSI address of the path.

Device Name: The name of the active disk array.

Serial #: The serial number of the device. Example: **0401EA9D0214**
- **0401EA9D** (Array serial number) **0214** (CU:LDEV).

Device Display from Windows 2K Disk Administrator without Autopath

Without Auto Path, Disk Administrator shows each path to the LUN as a device/disk, even if each path is redundant to the same LUN (i.e.- you could assign > 1 drive letter to the same LUN, format it, and access it).



LUN's as seen in
Windows 2000
Disk Management

Multiple-Paths May Cause Disk Signature to Change

The following information has been taken from a Microsoft Knowledgebase Article:
Problem:

The disk signature and Global Unique Identifier (GUID) for a disk may change unexpectedly on computers that have redundant Host Bus Adapters (HBA) to a common external disk. In this case, programs that depend on these disk signatures in the master boot record (MBR) may fail. Another symptom of changed volume GUID's or the disk signature may be that drive letters may be reassigned (reordered) after a restart. If the same disk is presented to Windows 2000 twice (one representation from each controller), issues can occur. If Ftdisk detects a 0 (zero) or duplicate disk signature, then it writes a new signature without prompting or notifying the user. You must change the signature immediately because Ftdisk cannot report the same unique ID (which is composed of the signature and partition offset) to Mount Manager (Mountmgr). When Mountmgr sees a new unique ID, it creates a new volume GUID and a fresh drive letter assignment for that new volume.

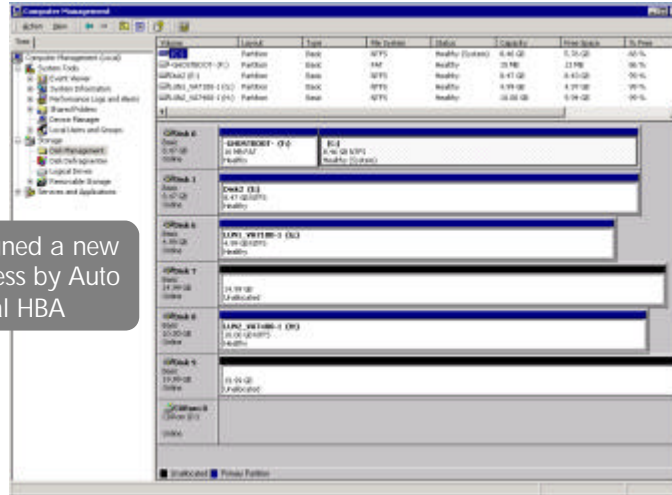
Resolution:

To avoid this issue, use multiple-path software (such as EMC's PowerPath or HP's AutoPath) or restart the computer in Safe mode. When you restart a computer that has more than one controller (HBA) in Safe mode, you may enable multiple-path software. In this case, the controllers may present the disk twice to the operating system, and the disk signature may be overwritten so that these new disks can be addressed.

Device Display from Windows 2K Disk Administrator with Autopath

With Auto Path installed, Auto Path recognizes LUN's that contain multiple paths; in turn, it makes those LUN's unavailable to the operating system.

LUN's are assigned a new
SCSI path address by Auto
Path Virtual HBA

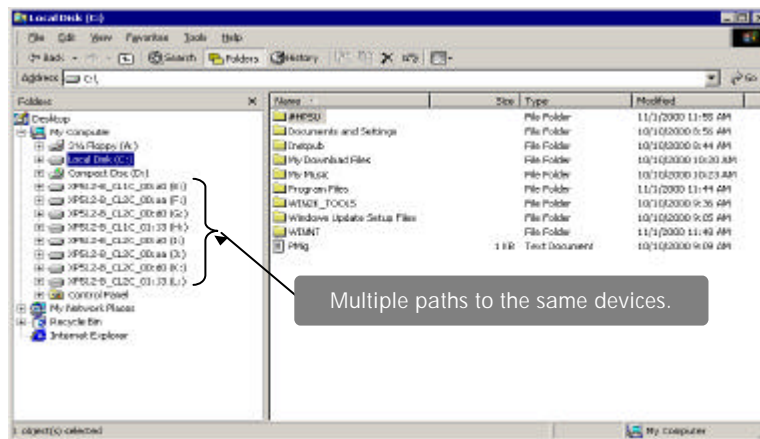


- Q. Why does my system (boot) drive have a new SCSI address (Disk assignment) after Auto Path install?
- A. After Auto Path is installed, the port number field (in the SCSI address) for all non-Auto Path XP-devices present on the system changes. This is because Auto Path creates a virtual HBA on which the Auto Path devices are shown and this installs before the other miniport drivers.
- Q. How does Auto Path decide which path to drop (make unavailable) from the O/S?
- A. Auto Path NT4 & W2K hide all paths to a device. In turn, Auto Path exposes virtual devices to the OS (Auto Path internally maps this device, from the different paths to the end device).

Auto Path: LUN's Seen in Windows 2K Explorer without Autopath

Before Installation of Auto Path:

In Windows Explorer, there are multiple drives mapped to the same LUN on the XP array.

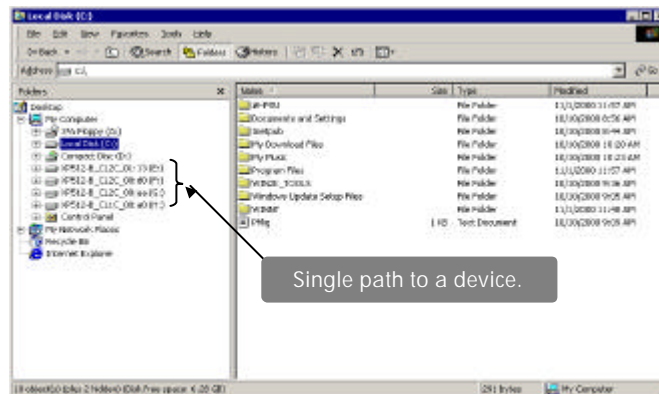


Student Notes:

Auto Path: LUN's Seen in Windows 2K Explorer with Autopath

After Installation of Auto Path:

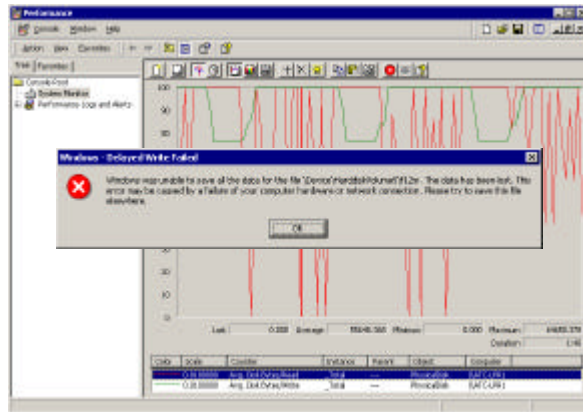
In Windows Explorer, there is now only one path to each LUN on the XP array. Auto Path solves the problem of multiple paths to a device by making the redundant paths unavailable to the operating system.



Student Notes:

Path Failure Seen in Windows 2K Performance Monitor

Without Auto Path installed, a hardware failure will result in a lost data path. When seen through Performance Monitor, a path failure will appear as a failed write operation and will be logged in the Event Viewer.



In this screen, a port was disabled on the switch while I/O was being processed. After 30 seconds, Performance Monitor returned a Delayed Write Failure message.

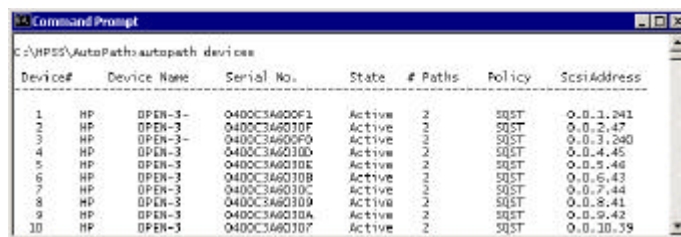


*The 30 second delay is a result of the W2K subsystem and not Auto Path.

Auto Path for Windows 2K CLUI

You can only execute Auto Path command lines from the directory where the installed Auto Path program files reside. If you chose the default directory when you performed the installation, change the directory at the command prompt to the following directory to run your command:

C:\HPSS\AutoPath



Device#	Device Name	Serial No.	State	# Paths	Policy	ScsiAddress
1	HP OPEN-3	0400C3A600F1	Active	2	SQST	0.0.1.241
2	HP OPEN-3	0400C3A6010F	Active	2	SQST	0.0.2.47
3	HP OPEN-3	0400C3A600F0	Active	2	SQST	0.0.3.240
4	HP OPEN-3	0400C3A6030D	Active	2	SQST	0.0.4.45
5	HP OPEN-3	0400C3A6030E	Active	2	SQST	0.0.5.46
6	HP OPEN-3	0400C3A6030B	Active	2	SQST	0.0.6.43
7	HP OPEN-3	0400C3A6030C	Active	2	SQST	0.0.7.44
8	HP OPEN-3	0400C3A60309	Active	2	SQST	0.0.8.41
9	HP OPEN-3	0400C3A6030A	Active	2	SQST	0.0.9.42
10	HP OPEN-3	0400C3A60307	Active	2	SQST	0.0.10.39



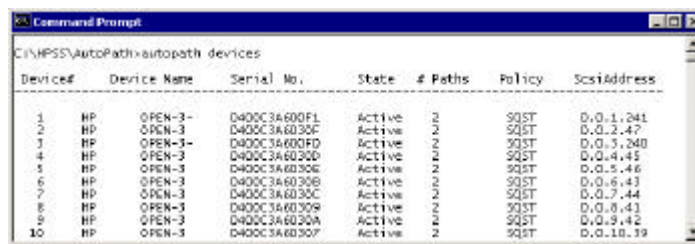
Student Notes:

Auto Path for Windows 2K CLUI

autopath devices

This command displays the details of the multi-path-capable devices in the system.

C:\HPSS\AutoPath>autopath devices



Device#	Device Name	Serial No.	State	# Paths	Policy	ScsiAddress
1	HP OPEN-3	D400C3A600F1	Active	2	SQST	D.O.1.241
2	HP OPEN-3	D400C3A6030C	Active	2	SQST	D.O.2.47
3	HP OPEN-3	D400C3A600FD	Active	2	SQST	D.O.5.240
4	HP OPEN-3	D400C3A6030D	Active	2	SQST	D.O.4.45
5	HP OPEN-3	D400C3A6030E	Active	2	SQST	D.O.5.46
6	HP OPEN-3	D400C3A6030B	Active	2	SQST	D.O.6.43
7	HP OPEN-3	D400C3A6030C	Active	2	SQST	D.O.7.44
8	HP OPEN-3	D400C3A6030A	Active	2	SQST	D.O.8.41
9	HP OPEN-3	D400C3A6030A	Active	2	SQST	D.O.9.42
10	HP OPEN-3	D400C3A6030F	Active	2	SQST	D.O.10.39



Device #: An index number used by Auto Path for identifying the device

Device Name: The name of the active disk array.

Serial Number: The serial number of the disk array.

of Paths: The number of paths to the device.

Policy: This column displays the Load Balance Policy set for this device.

State: The condition of the device:

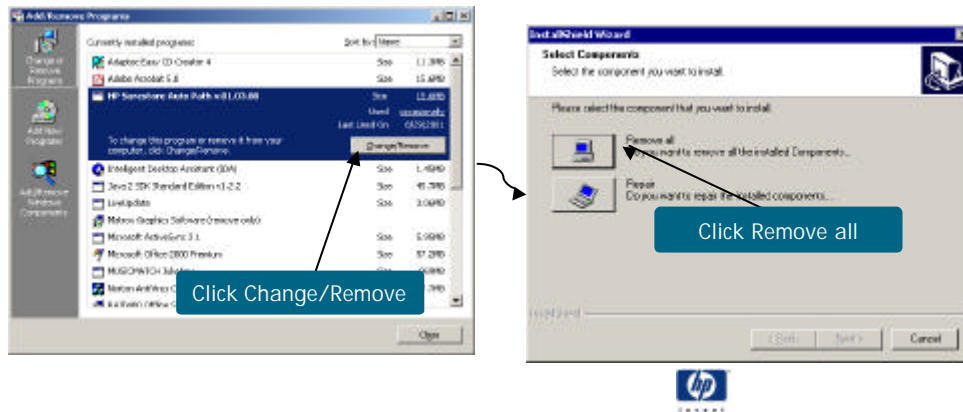
Active = device in use

Failure = device no longer in use or all paths to device have failed.

Uninstalling Auto Path on Windows 2K

To uninstall Auto Path:

1. Go to the Windows 2000 Control Panel and select the Add/Remove Programs icon.
2. Follow the instructions to remove the Auto Path program files.



Student Notes:

Auto Path for HP-UX System Requirements

System requirements for installing HP-UX Auto Path Driver and Command Line User Interface:

To install Auto Path, the system must conform to these requirements:

- ✍ HP-UX 11.00 or HP-UX 11.11
- ✍ Superuser (root) access to the HP-UX system.
- ✍ Devices on the disk array for use by Auto Path, should be part of a volume groups configured using LVM.
- ✍ On systems with PCI A5158A controller cards, the patch PHKL_23939 for HP-UX 11.00 and patch PHKL_23626 for HP-UX 11.11 needs to be applied if not already installed in the system.



Auto Path is implemented as a pseudo driver, layered between the LVM and the SCSI device driver, and supports a clustered environment (MetroCluster/ServiceGuard).

Auto Path for HP-UX Known Issues

Known Problems/Limitations

- ✍ The status of a path will be updated only when the I/O is performed through that path.
- ✍ Only Fibre Channel connectivity is supported.
- ✍ Online addition/removal of devices is not supported. New devices added will be recognized after reboot only.
- ✍ Auto Path supports maximum eight paths to an end device.



Student Notes:

Auto Path HP-UX vs. LVM PVLINKS








Description	AP	PVL
1. <u>Fail Over</u> : automatically fails over to an alternate path when the primary path is no longer available	?	?
2. <u>Fail Back</u> : automatically recognizes the newly available path when a failed path comes back up alive.	?	
3. <u>Load-balancing</u> : automatically and programmatically balances I/O load among all available paths with user selectable load-balancing policies.	?	



- Q. What is the recommendation for Auto Path HP-UX and the coexistence with LVM in regards to setting up primary and alternate links in LVM?
- A. *Auto Path manages alternate links and LVM PVLINKS use is not required.*

Installing Auto Path for HP-UX

To Install AutoPath for HP-UX:

-  Log on to the HP 9000 system as superuser (root).
-  Insert the Auto Path installation CD into the CD-ROM drive.
-  Mount the CD on your file system.
-  From the command line, execute extract.sh script. The script creates an 'Autopath' directory in /opt and copies the files to host.
-  From the command line, execute the install.sh script from the Autopath directory.
-  For systems with PCI A5158A controller cards, the patch PHKL_23939 for HP-UX 11.00 and patch PHKL_23626 for HP-UX 11.11 needs to be applied if not already installed in the system.
-  Reboot the system.



Student Notes:

Configuring Auto Path for HP-UX

No GUI available, completely command line oriented

```
ua2:/root> autopath display /dev/dsk/c12t15d0
```

```
Auto Path Driver Version:    VER 1.01.00
Auto Path CLI Version  :    VER 1.01.00
```

```
Details for path:           /dev/dsk/c12t15d0
```

```
Load Balance Policy      :    No Load Balance
```

```
=====
Device Paths              Device Status
=====
/dev/dsk/c12t15d0         Active
/dev/dsk/c15t15d0         Active
=====
```



Student Notes:

Auto Path - Path Failure in HP-UX

Right after installation, perform tests to confirm that Auto Path is running correctly:

- Run the `autopath display` command
- Set “no load balancing” (“None”) and a preferred path
- Set all paths

⚠ Degraded performance during peak utilization is likely to be the first direct sign to a user that a path has failed

- Run the `autopath display` command to make frequent inspections
- Look for any paths with a “failed” link status or for paths with no load balancing policy in place



Troubleshooting should be devoted primarily to finding and correcting failed path conditions. Some problem resolution can be performed by the customer. For example, customers typically replace failed GBICs on some HBAs. HBA replacement is usually a CE task.

Q. Are path failure events logged anywhere in HP-UX?

A. *No log is maintained for v1.01 of Auto Path HP-UX. Path failure events will be logged in the Auto Path HP-UX v2.0 release.*

Auto Path - Path Failure in HP-UX

Recovering from a path failure

- ✎ Determine that a path has failed
- ✎ Diagnose and correct the failure
- ✎ Run "autopath display" to be sure the path is active
- ✎ If a path was the preferred path, you must reset it "manually"



If there is a policy of "None" (no load balancing) for the failed path and the path was not the preferred path, no user action is necessary. If the failed path was a designated preferred path, the system will fail-over to another path. However, after you recover the failed path, you should then reset it as the preferred path.

Adding New Devices – HP-UX

Adding new devices and paths

- ✍ When you add a new HBA, HP-UX will “autodiscover” it on reboot
- ✍ Use ioscan to find the device file of the new HBA
- ✍ Use Auto Path CLI commands to assign a load balancing policy (or preferred path) to the device.



Student Notes:

Uninstalling Auto Path on HP-UX

To uninstall Auto Path:

1. Login as root.
2. Run the uninstall command:
autopath uninstall
3. Reboot the system.

```
ua2:/root> autopath uninstall
rm: /usr/share/man/cat1.Z/autopath.1m non-existent
Autopath Services stopped Successfully
Compiling conf.c...
Loading the kernel...
Generating kernel symbol table...
PLEASE REBOOT THE SYSTEM
```



There is an immediate work around for an Auto Path HP-UX v1.01 un-installation problem.

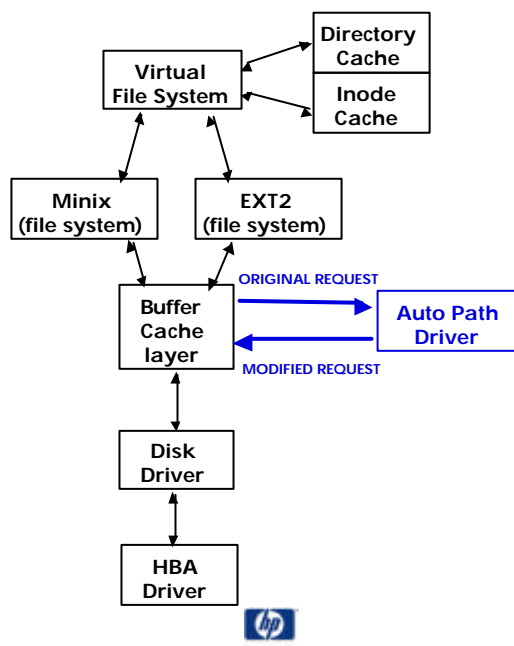
Suggested process for un-installation is:

- 1) Execute the *autopath uninstall* command
- 2) After the kernel is rebuilt, execute the following command:
mk_kernel -o /stand/vmunix
- 3) Reboot the server

After executing the above process a clean kernel without any traces of Auto path. Since Auto Path HP-UX v1.01 is not a DLKM, there is no need to do a kmupdate to rebuild the kernel.

AutoPath for Linux - Auto Path driver

- A block driver (Dynamically Loadable Module)
- The Auto Path driver is layered as shown in the figure



Initialization

The main module of the Auto Path driver registers Auto Path as a block driver with the OS. The Auto Path devices are registered with the OS through the Linux API “register_disk”. The global structure “gendisk” is updated, so Auto Path devices can be partitioned like SCSI disks.

The Autopath driver is a dynamically loadable module, which can be loaded any time after the system bootup and the HBA driver is loaded.

In order to provide load balancing, I/O requests have to be trapped and processed before putting them into a request queue. This is done by Auto Path’s request function “ap_make_request”, which is registered by using the Linux API “blk_queue_make_request(BLK_DEFAULT_QUEUE(MAJOR_NR), ap_make_request)”.

Device Discovery

This module identifies Auto Path capable devices. Discovery is invoked by the CLI command “autopath discover”. The module scans the devices (HBAs) connected to the disk arrays, determines the paths and associated device information (stored in the “device_details” structure) and the “new” Auto Path device is registered with the OS. There’s a structure allocated for each Auto Path device.

Device Discovery Algorithm:

- In the user space, scan the Auto Path capable devices by issuing the SCSI inquiry command to all the device paths.
- Allocate a device_details structure each disk and copy the structure to kernel space. The structure is copied to kernel space through an ioctl “UPDATE_APDEVICES”.
- In the “UPDATE_APDEVICES” ioctl the “device_details” structures is updated and the Auto Path gendisk structure “ap_gendisk” is updated accordingly, so that an alias device points to the original SCSI disks.
- Register the Auto Path new alias disk with the OS by calling “register_disk”.

AutoPath for Linux - Product requirements

- Servers
 - LXR 8500R server
- OS
 - Red Hat Linux 7.1 operating system with 2.4.2-2 kernel
- Clustering
 - HP Service Guard
- HBAs
 - Emulex LP8000 and LP9002 Fibre Channel HBAs
 - HBA driver 4.11e
- Disk Arrays
 - XP256, XP512, XP48, VA7100, VA7400
- Switches
 - Brocade SilkWorm 2800



Student Notes:

Installation components

- Auto Path XP driver
 - Provides failover, failback (recovery), and load balancing
 - Installed on a host machine connected to a supported disk array
- Auto Path XP Command Line Interface



Student Notes:

Installing the software

- You need administrator access to the system
- Log in as superuser (root)
- Follow the installation instructions in the README.TXT file
- See the Student Notes of the presentation for details



The installation procedures in this course are for Auto Path 1.00.00 for Red Hat Linux 7.1. However, be sure to refer to the README.TXT file on the product CD for the latest installation information.

To install Auto Path:

1. Log in as superuser (root).
2. Insert the Auto Path CD into the CD-ROM drive.
3. Mount the CD on your file system.
4. From the command line, execute extract.sh script (on CD).
The script creates the "Autopath"
directory in /home and copies the Auto Path files to the host.
5. Execute the install.sh script from the Autopath directory.

HP
SurePartner
Training

MNS
Crossupdate

Module wrap-up

