

HP SureStore Disk  
Array XP Family  
Technical Pre-Sales  
Training

ESG503SG20303



student  
guide



HP SureStore Disk  
Array XP Family  
Technical Pre-Sales  
Training

ESG503SG20303



training

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

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

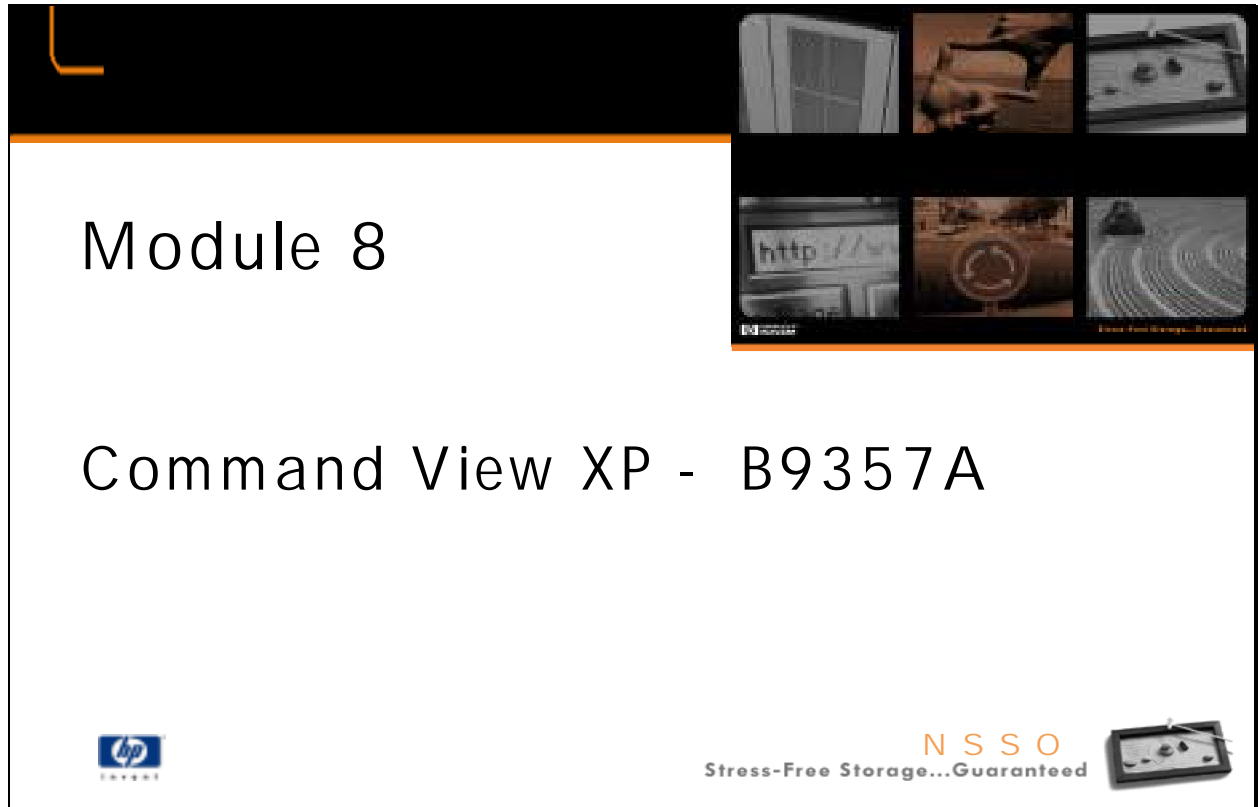
Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for HP products are set forth in the express limited warranty statements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

**HP SureStore Disk Array XP Family Technical Pre-Sales Training**

Student Guide





March 2003





# Module 8

## Command View XP - B9357A


   

- B9357A-001 (XP 256 License To Use)
- B9357A-002 (XP 512 License To Use)
- B9357A-003 (XP 256 License To Use)
- B9357A-004 (XP 512 License To Use)


The HP SureStore E Command View XP is a web-based management tool for the SureStore E Disk Array XP256 and XP512. The HP SureStore E Command View XP will include Remote Control XP software disks or CD. In this training module the name "Command View" refers to the combined software modules which together provides the support for the communication protocol with the client, the presentation layer or user interface (UI), and the low level interaction with the device to perform the management tasks. The Command View can be invoked from a standalone application or from a management platform. The following management platforms are in plan to be supported at release: \* HP SAN Manager DM \* HP OpenView ITO (for Five-Nines program) \* HP TopTools \* HP Web Console \* CA Unicenter \* Tivoli

# Module Agenda

- Command View XP Value Proposition
- XP Array Management Software Defined
  - Command View XP
  - Service Processor (SVP)
- Describe the HP Surestore Command View XP product
- Command View XP Overview
  - Hardware Setup
  - Software Setup
  - SNMP Management
  - HP Storage Map



N S S O  
Stress-Free Storage...Guaranteed



The XP Standalone Management Application provides users with the ability to launch the Command View outside of any management framework. The XP Standalone Management Application will discover all the arrays that are attached to the host (using xpinfo) and register their serial numbers upon startup. However, the user is required to input the IP addresses of the Remote Consoles (RCs) that manage those arrays discovered by the Standalone Management Application. A user interface will be provided to simplify the input of the Remote Console IP setup information. On HP-UX the Command View supports the Netscape 4.04 browser. On Windows NT the Command View supports both the Microsoft Internet Explorer 4.01 and Netscape 4.04 browser. The Command View itself will reside on the Remote Console server and thus only runs on the Windows/NT Operating System. The XP Standalone Management Application will initially be available on HP-UX, Windows/NT, and SUN Solaris. The Command View XP user interface provides the same look and feel as the Top Tools user interface; which include the following operational tabs:  
| Identity, Status, LUN Management, Auto LUN (XP512 only) User Administration, and Support.

## Module Agenda (cont.)



- Install Command View
- Describe the interface and general operations, including Identity, Status, and Administration
- Use LUN Management



- Use Logical Unit Size Expansion
- Use Volume Size Configuration
- Use Cache LUN, Auto LUN & LUN Security



- Use Host Storage Map
- EMC Control Center
- Wrap Up



N S S O  
Stress-Free Storage...Guaranteed







# Storage Management Challenge

Today's Storage Management Challenge

- \* Shortage of qualified IT staff
- \* Service Level Agreements
- \* Managing business continuance
- \* Maintaining flexible storage structure
- \* Capacity Planning
- \* Massive Storage Growth
- \* Storage & Management Consolidation



*XP Management software needed to advance to meet the changing needs of IT administrators*


As the need for information continues to increase, systems professionals are faced with a series of challenges, shortage of qualified IT staff - global businesses, multiple operating systems and extended data centers have created the need for qualified IT staffs to handle these complex environments -- demand has exceeded supply. Service Level Agreements - demands on the limited staffs to maintain availability and performance expectations continue to rise. Business Continuance - Companies simply cannot afford downtime -- resulting in shrinking timeframes for backup and disaster recovery. Flexible storage structure - companies need storage systems that are flexible with the ability to dynamically reallocate resources. Capacity Planning - growth in volumes of required data mean IT professionals must constantly analyze trends to stay ahead. Combined with the 50% growth rate each year for storage and the need for storage consolidation, its clear that a fragmented grouping of tools is not sufficient!

## XP Array Management Choices

- **Remote Control XP (RMC or RC):** XP array management software developed by Hitachi (private LAN). This software was the only end-user XP array management choice prior to Command View XP
- **Command View XP:** XP array management software developed by HP. This software will be the Remote Control XP replacement by 2001. Feature benefit is web based access (public LAN).
- **Service Processor (SVP):** The laptop mounted inside of each DKC is not to be used by customers. The SVP offers the same functionality as RMC and CV plus other initial setup functionality



Stress-Free Storage...Guaranteed



Command View XP started shipping in July 2000. At the product launch Command View XP did not have feature parity with Remote Control XP, therefore each version of Command View sold also shipped with Remote Control. Once Command View XP achieves feature parity with RMC, HP will no longer offer Remote Control. However, each customer that already received Remote Control will continue to receive Remote Control software updates via support contracts.

Is Remote Control XP a requirement for Command View XP to manage an XP array?

-XP256 array management from Command View= YES

-XP512/XP48 array management from Command View= NO

\*please see following slides for more information on this topic

## Hitachi's Remote Control XP

Two versions exists:

- XP256= RMC v14.52.00 or higher requiring microcode v52.47.52/00 or higher
  - RMC is a requirement for Command View to run against the XP256
- XP512= RMC v11.27.00 or higher requiring microcode v1.11.23 or higher
  - RMC is NOT a requirement for Command View to run against the XP512/48

*Can both versions of RMC be loaded onto the same machine?*

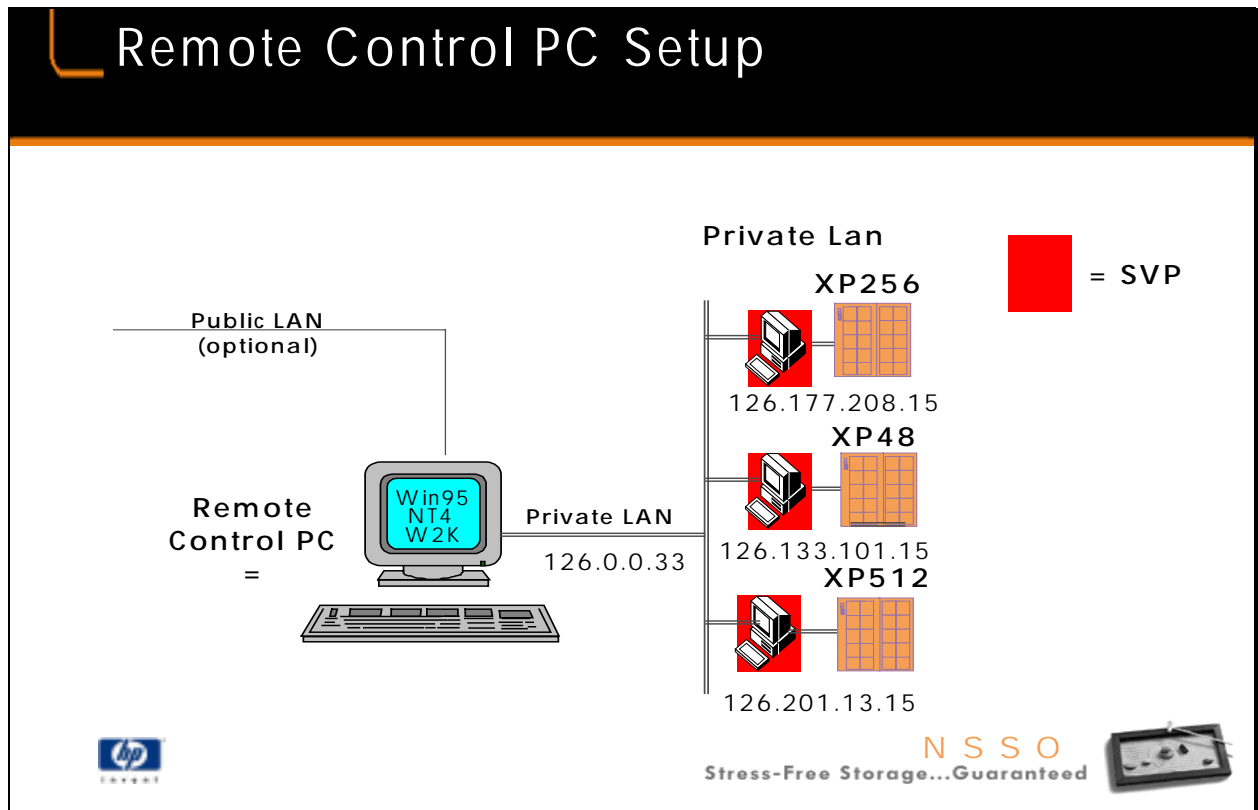
YES. This is fully supported



N S S O  
Stress-Free Storage...Guaranteed



Explained in detail later in this slide set, but RMC is a requirement for Command View to manage an XP256 but Remote Control XP is not a requirement for Command View to manage an XP512. Why? Command View uses SNMP to manage the arrays. For this reason, an SNMP agent must exist in the Command View environment for Cmd View to do SNMP gets, sets, etc. On the XP256 the SNMP agent exist on the Remote Control PC; however on the XP512/48, the SNMP agent exist on the SVP. For this reason, the Remote Control PC (ie- the Remote Control XP software) is not needed in the XP512/48 environment- Command View uses the SNMP agent found on the SVP (served up via the USB connection, which is tied to the external LAN's TCPIP subnet).



The illustration above shows the setup of the Remote Control XP PC. The Remote Control XP software running on this PC is connected must have a private LAN (requirement for management of the XP arrays with RMC) and a public LAN (optional). For the XP256, this private LAN connection into the XP256 is made possible by a 10BASE-T connection to a PCMCIA card inside the SVP laptop. For the XP512/XP48, this private LAN connection is made possible by a 100BASE-T connection to an on-board ethernet adapter with a RJ45 connection.

For more information, please access slides "Remote Control XP PC Setup"

\*Older versions of RMC were supported on Windows 95; however, most current releases only support NT4 or W2K (W2K support began around Q1 of 2001).

## What is HP Command View XP?

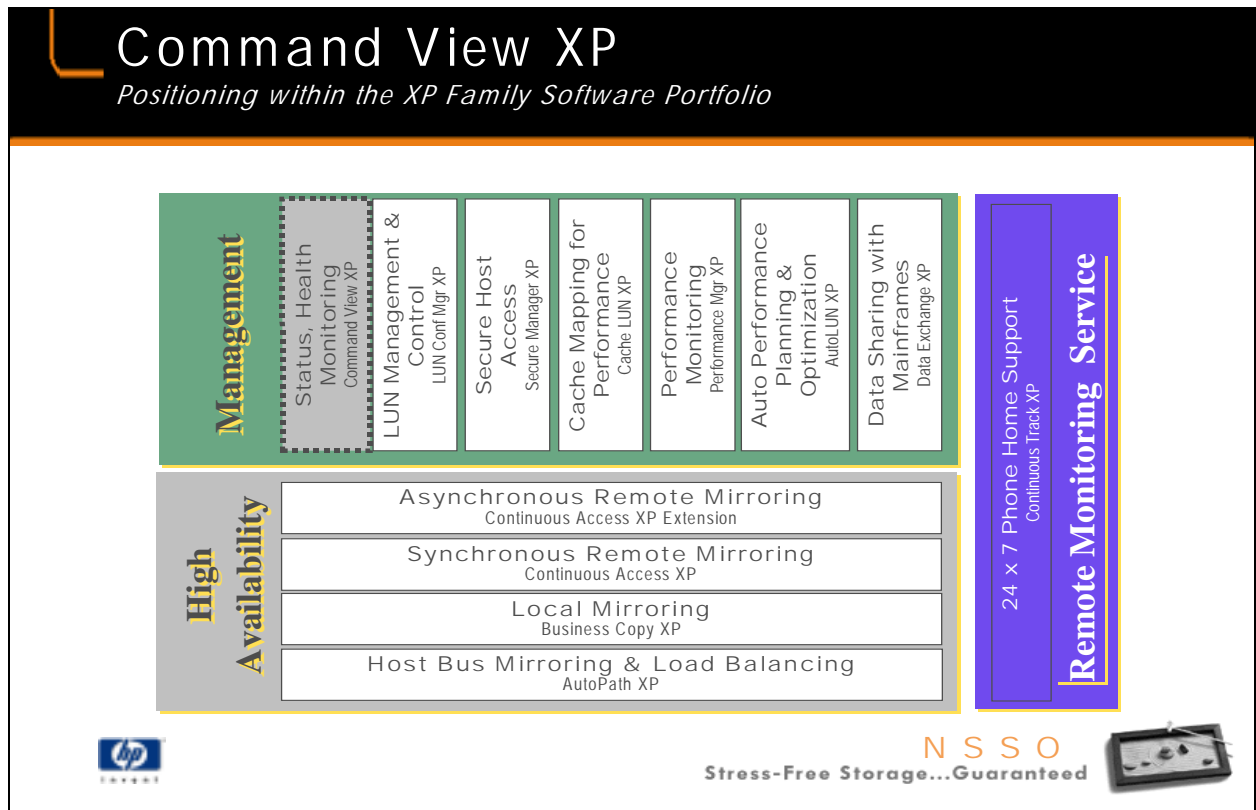
- A web-based platform that manages the XP disk array family while integrating with other applications
- A tool to monitor and configure XP family storage resources from anywhere on the Internet
- A tool with visual representations of host and storage resources
- An array management platform that can be integrated with other network and storage area management platforms



N S S O  
Stress-Free Storage...Guaranteed







Command View Software Package-The Command View product media includes the Remote Control product also. The Remote Control product is installed on the remote console PC that is directly connected to the disk array private LAN. Command View can be installed on any PC that has an Internet connection.

Standard To manage the HP SureStore E Disk Array XP family of products, Command View provides these standard functions:

- Main
- Device Administration
- Session Administration
- User Administration
- Support
- Identity
- Status
- Trap Distribution

Optional You can purchase optional software applications to extend the basic features of Command View. Optional applications are fully integrated with the Command View user interface. Contact your HP account representative for information about these optional products:

- HP SureStore E LUN Configuration Manager XP
- HP SureStore E Secure LUN XP

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Command View XP

- HP SureStore E Cache LUN XP
- HP SureStore E Auto LUN XP

Related information For information about to install and configure the optional products, see the documentation provided with each product.

# Command View XP

## Example Screens

Mapping reports for Ports, LUNs and LDEVs

Overview of Capacity Utilization

Graphical view of XP configuration & health

SSO Guaranteed

Port Displays the port number associated with the LUN. SCSI ID Displays the target number for the LUN.

LUN Displays the LUN ID within the target for the LUN.

CU:Ldev Displays the first logical device associated with the LUN.

ACP Pair Displays the ACP Pair number associated with the LUN.

### Command View Installation Procedures

HP SureStore E Command View XP runs on the remote console PC and communicates with the attached disk arrays through the Internet.

To install HP SureStore E Command View XP, you must install the software components in this order.

1. Install Remote Control on the remote console PC.
2. Install Command View.

To install HP SureStore E Command View XP, follow these procedures:

1. Install Apache Web Server.
2. Install Java Servlet.
3. Install HP SureStore E Command View XP.
4. Run Command View.
5. Start Trap Distributor.

Installation of the Command View software requires approximately 20 minutes.

## SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training Command View XP

You can log on to HP SureStore E Command View XP after software installation is complete. Command View can be started:

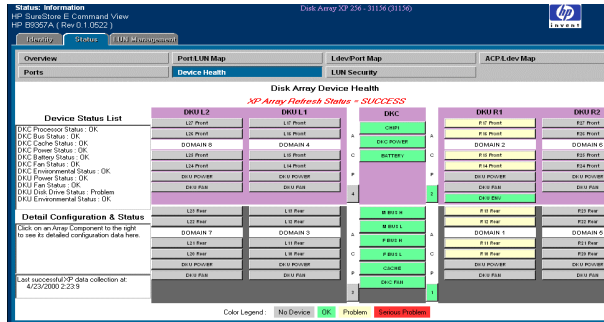
- directly from a browser
- through SAN Manager DM
- through the Host Storage Map application


## HP SureStore E Command View XP:

### A Modular Management Paradigm

*A visual, integrated management platform for the XP family*


- **Ease of Use**
  - Common U/I across XP platforms
  - Quick Installation via InstallShield
  
- **Manage from Anywhere**
  - Remote expert can participate in problem resolution across multiple installations without traveling
  - Excellent training mechanism for junior staff
  
- **Quicker Access to Data**
  - Visual representation of host & storage resources
  - Graphical view of status & health of storage resources enables quicker troubleshooting & problem resolution





NSSO

Stress-Free Storage...Guaranteed



What is Command Center XP? Command Center is a web-browser application that provides a common U/I for the XP storage family reducing IT training, increases IT effectiveness. A common U/I across all XP software solutions means that the solution will be easy to learn -- eliminating the risk of the software becoming shelf-ware. The graphical web-browser interface enables customers to monitor their storage resources from any location and more specifically enables the remote expert to support multiple locations without having to travel into the office. Command Center enables multiple customers to view a single screen making Command Center an excellent training mechanism for junior staff. Our plans are to integrate Command Center into the leading enterprise management products. At first release, Command Center will have event-level integration into OpenView Vantagepoint ITO. We are currently investigating event level integration into Computer Associates Unicenter, Tivoli NetView and BMC Patrol. Device Health Page

The Device Health button under the Status tab displays a graphical representation of the array configuration and the status of critical hardware components.

## When to sell Command View XP...

- When customer purchases **ANY** optional XP256 or XP512 management software
- When customer wants to manage XP arrays themselves
  - Perform LUN configuration & mapping
  - Manage SAN security
  - Do performance monitoring & management
- If customer is interested in SAN Manager DM and XP

**Command View XP should be a part of ANY XP sale!**



N S S O  
Stress-Free Storage...Guaranteed



Target Customers-All XP256 and XP512 customers planning on purchasing optional XP management and disaster recovery products.

## Key features

### ➤ Device Management

- Provides the ability to monitor and configure disk array operation, security and array resources

### ➤ Event Management/Trap Distribution

- Archives and distributes SNMP traps to configured storage area network (SAN) management stations
- SNMP trap messages are converted into readable error messages and archived in log file

### ➤ Host Storage Map

- Displays a graphic representation of disk storage allocation.
- Executed independently on the host system to which the disk arrays are connected
- Link to CommandView facilitates disk array management

### ➤ Remote Control XP

- Provides management functions similar to Command View
- Manages one array at a time



- RC for the XP512/XP48 has a different revision level than RC for the XP256

Stress-Free Storage...Guaranteed



## Key Features (cont.)

- A management platform with a browser-based interface for the XP disk array family. You can:
  - View the status and health of XP resources
  - Configure and manage LUNs
  - Analyze I/O performance
  - Migrate volumes
  - Establish LUN security
- Command View can co-exist on the same server as Performance Advisor
- Command View can share SNMP trap information with any SMNP-based management station, such as HP OpenView Storage Node Manager, CA Unicenter, Tivoli, etc.



N S S O  
Stress-Free Storage...Guaranteed



### Student Notes

#### Device Management

Command View provides the ability to monitor and configure disk array operation, security and array resources.

#### Event Management/Trap Distribution

Command View archives and distributes SNMP traps to configured storage area network (SAN) management stations. Command View can share SNMP trap information with any SMNP-based management station, such as HP OpenView, CA Unicenter, Tivoli, etc. SNMP trap messages are converted into readable error messages and archived in log file.

#### Host Storage Map

Command View displays a graphic representation of disk storage allocation. Host Storage Map is executed independently on the host system to which the disk arrays are connected, and a link to CommandView facilitates disk array management.

#### I/O Performance Monitoring

You can monitor I/O performance by means of Command View's AutoLUN feature and through Performance Advisor integration.

#### Multi-Platform Integration

Command View can be integrated into HP OpenView, CA Unicenter, Tivoli, and HP OpenView Storage Node Manager.

Similar to Remote Control XP



SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Command View XP

Remote Control provides management functions similar to Command View. RC manages one array at a time.

## Features and benefits

### FEATURES

**Web-browser based interface**

**Common U/I for XP management applications**

**Visual and graphical representation of host & storage status and health**

### BENEFITS

Manage your storage resources from any location

Remote expert can participate in problems resolution without traveling

Excellent training tool for junior staff members

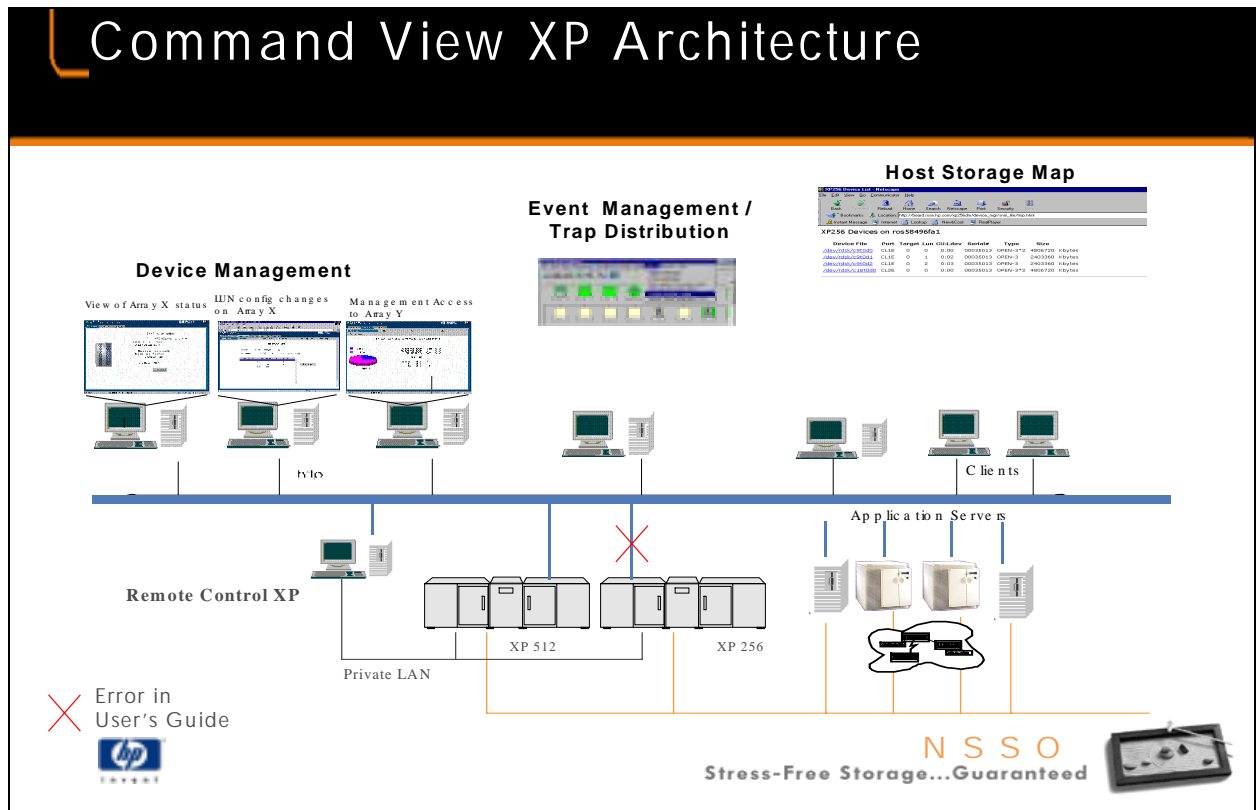
Focus your efforts on managing your environment rather than learning individual tools

Increase the efficiency of your limited IT resources reducing the total cost of storage ownership



**N S S O**  
Stress-Free Storage...Guaranteed





Command View is a web-based management application that uses SNMP standard protocol to manage the XP family of disk arrays. CV can manage any XP disk array installed on the WW. Remote Control is a Hitachi proprietary management application. There are unique versions of Remote Control to manage unique disk arrays. For example, in order to manage an XP256 and XP512 using Remote Control software, you must install the XP256 version of Remote Control AND the XP512 version on the Remote Console. The Remote Console is a PC which MUST be physically connected to a Hitachi-defined private LAN. Residents on the Hitachi Private LAN are usually physically located near each other. The XP512 and XP256 connect to the Hitachi Private LAN in exactly the same manner. The Remote Control software uses a Hitachi proprietary communication protocol to manage the arrays. SNMP is NOT a factor when using Remote Control. SNMP becomes a requirement when you wish to use Command View to manage your XP disk array family from the www. In the case of the XP512, the SNMP agent resides on the XP512 disk array. The XP256 does NOT support an internal SNMP agent. In order for Command View to manage a XP256, it uses an SNMP agent installed on the remote control as a proxy.

#### Command View Concepts and Terminology

The user interface for Command View uses these terms:

**array group**-The number of physical disk drives contained in a RAID group. This number depends on the RAID configuration.

**Controller-**The disk array has one controller (DKC) that controls all data access and storage operations. To organize the storage space attached to the controller, you can group similarly configured logical devices (LDEVs) with unique control unit images (CUs).

**CU (control unit image)-**To organize the storage space attached to the controller, you can group similarly configured logical devices (LDEVs) with unique control unit images (CUs). CUs are numbered sequentially.

**logical device -(LDEV)** An LDEV is created when a RAID group is carved into pieces according to the host emulation mode (OPEN-3/8/9/E). The number of resulting LDEVs depends on the selected emulation mode. The term LDEV is often used synonymously with the term volume.

**LUN -** A LUN results from a mapping of the logical unit number (LU), port ID, and LDEV ID.

**parity group -** A parity group is a mode of disk operation and configuration. A parity group is the same as a RAID group.

**RAID group -** A RAID group is a mode of disk operation and configuration. RAID groups 1 through 5 vary in the number of bits used in data and parity, mirroring, and

striping features. **Command View Description volume** A volume is created when a RAID group is carved into pieces according to the host emulation mode (OPEN-3/8/9/E). The number of resulting volumes depends on the selected emulation mode. The term volume is often used synonymously with the term

**LDEV.** HP SureStore E Command View XP runs on the remote console PC and communicates with the attached disk arrays through the Internet. To install HP SureStore E Command View XP, you must install the software components in this order.

1. Install Remote Control on the remote console PC.
2. Install Command View.

Web based XP256 and/or XP512 management server which contains in-band connected HP Host Storage Map Ships with Remote Control XP. The architecture of the Command View is such that the Web Server, Device Manager, Device Objects and the SNMP agent all reside on the same server; namely the Remote Console server managing the target array. Apache HTTP server must be installed in order to use Command View. Apache is included on Command View installation CD. More information about Apache can be found here: <http://www.apache.org/docs/windows.html>

This slide was taken from the most current Command View XP User's Guide (pg 12, June 2000)

**Why Command View?**

Provides an integrated management platform for the XP family of products

Platform independent

Web-GUI interface common across XP storage products

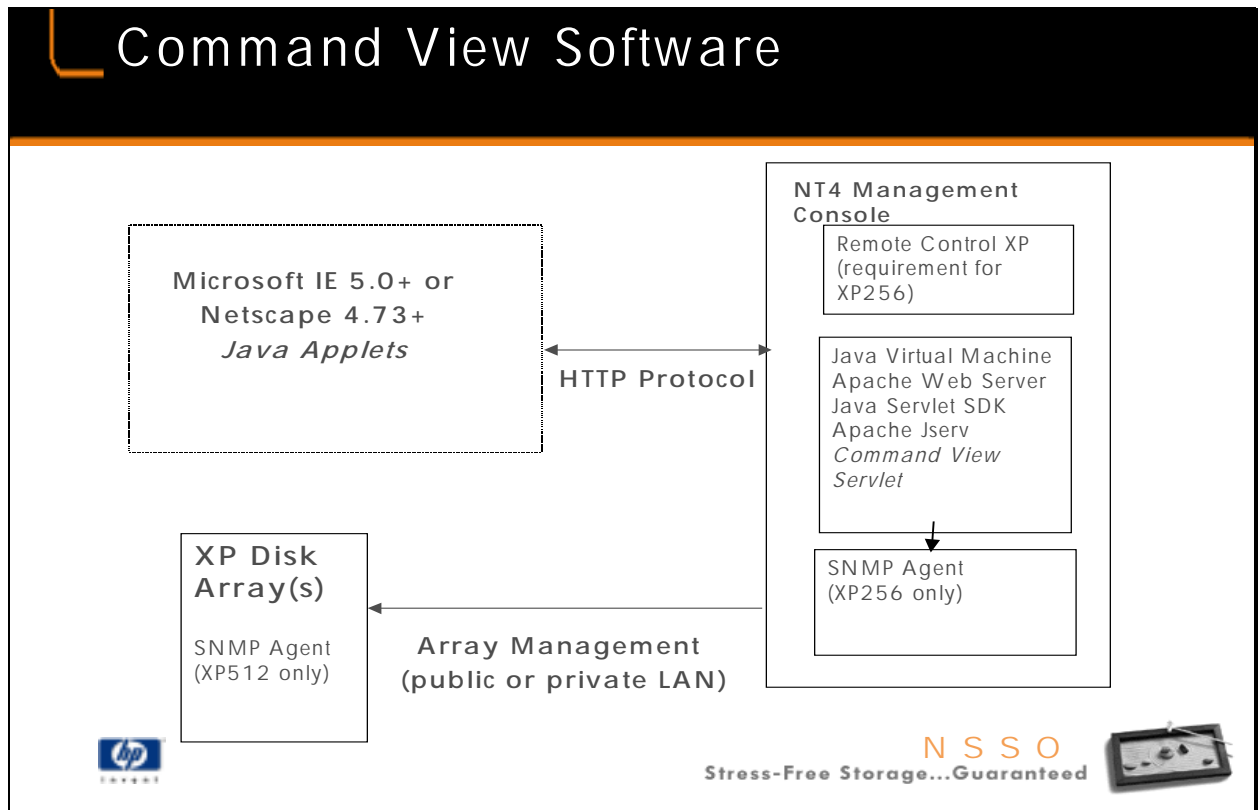
### Ability to manage multiple XP products

#### What is Command View XP?

Command View is a web-browser application that provides a common U/I for the XP storage family reducing IT training, increases IT effectiveness. A common U/I across all XP software solutions means that the solution will be easy to learn -  
- eliminating the risk of the software becoming shelf-ware.

The graphical web-browser interface enables customers to monitor their storage resources from any location and more specifically enables the remote expert to support multiple locations without having to travel into the office.

Command View enables multiple customers to view a single screen making Command View an excellent training mechanism for junior staff.



HP XP Management Console: This marketing concept will support the following software packages all on the same machine

- Command View XP
- Remote Control XP for the XP256
- Remote Control XP for the XP512
- SAN Manager DM
- Performance Advisor
- Performance Manager XP

The machine requirements defined for the "management console" are what defines the machine requirements for the Command View XP software. The NT4 management console can be either W2K or NT4 Workstation or NT4 Server running service pack 4 or higher.

Command View XP requires several individual packages of software. Below is software required:

1. Remote Control XP (RMC): HP now has two separate versions of the Remote Control software. The same RMC software package that manages the XP256 will NOT work on the XP512 (and vice versa). Each version can run on the same PC (illustrated in slides to follow). Technically there is no dependency between Command View and RMC, but Command View in its first release does not contain the full suite of functionality in RMC.

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Command View XP

2. Java Virtual Machine
3. Apache Web Server
4. Java Servlet SDK
5. Apache JServ
6. Command View XP- When you install Command View, items 2 through 6 get installed automatically (ie- you only have to launch one install process).
7. SNMP Agent- only necessary to install if using RMC for the XP256 for SNMP management.

# Command View XP System Requirements




## HP SureStore E Command View XP System Requirements

Pg 20 of the most current Command View User's Guide (June 23, 2000)

See following slide for HP's current "minimum" recommended machine

Before installing HP SureStore E Command View XP, ensure that the system environment conforms to these requirements:

- one or more HP SureStore E Disk Array XP products
- HP SureStore E Command View XP software product
- browsers:  
(HP-UX/Sun Solaris) Netscape Communicator 4.73 or higher  
(Windows NT) Internet Explorer 4.0 or higher
- PC with these features:
  - operating system: Windows NT 4.0 (Service Pack SP5)
  - processor: 400 MHz or better
  - memory (ram): 128 MB or more
  - available hard drive space: 40 MB or more
  - 50 MB of free space is recommended before installation
  - monitor: VGA with 256 colors or better
  - LAN adapter: At least one


Stress-Free Storage...Guaranteed



Command View Hardware  
"Recommended"


A6034A	HP Visualize Personal Workstation P Class
A6024A	800 MHz Pentium III 133 FSB CPU
A6012A#0D1	Factory integrated (\$-200)
A2185A	256 MB DIMM
A5028A#0D1	Factory integrated (\$-100)
A1270A	*13 Gb IDE disk
A5994	9 Gbyte 10K SCSI disk
A1270A#0D1	Factory integrated (\$0)
A5988A	48X CD-ROM drive
A5988A#0D1	Factory integrated (\$0)
A1290A	Elsa Synergy II AGP graphics card
A1290A#0D1	Factory integrated (\$0)
D8910W	19" color monitor
A4199B	Visualize C13 Power Cord Kit
H3462A	3 years next day HW service

Total List Price: \$3,864.00 \$3,946.00



N S S O

Stress-Free Storage...Guaranteed



Listed above is the HP "recommended" hardware for the Command View PC available to SF21 without creating channel conflicts. The updated recommendation has several justifications:

1. This recommendation is centered around making this PC the HP Surestore management console (the vision is to put OVSAM, Command View, Performance Manager, Performance Advisor, OpenView SPI etc on one PC). To do this effectively, the machine will need the horse power to make it usable. With each subsequent release of Command View starting with v1.02 to the most current 1.11 (June 2001) the added functionality of Command View coupled with the tight integration with other XP applications (ie- Performance Advisor) has set a higher workload on the machine. As with most WIntel applications, increasing the processing power of the PC coupled with more RAM usually yields greater performance.
3. Add more horsepower today means you are ready for upgrades tomorrow.

**RECOMMENDATION:** Put as much PC power behind the application as you can at first implementation. This will help improve performance and make for a much simpler upgrade path as new versions are released.

\*If a customer does not choose to use this HP PC hardware solution, that is fine. Just make sure the customers own machine meets the processor, RAM and HDD requirements documented above.





## Command View Hardware "Second NIC Order Options"

Remote Control still uses the private LAN, thus two NICs are necessary. Choose from the following NICs (all supported on the Visual Age workstation PC):

<b>D6937A</b>	HP NIC
<b>J2585B</b>	HP NIC
<b>3C905-DX</b>	3COM NIC

Do I need Remote Control XP if I have Command View XP?

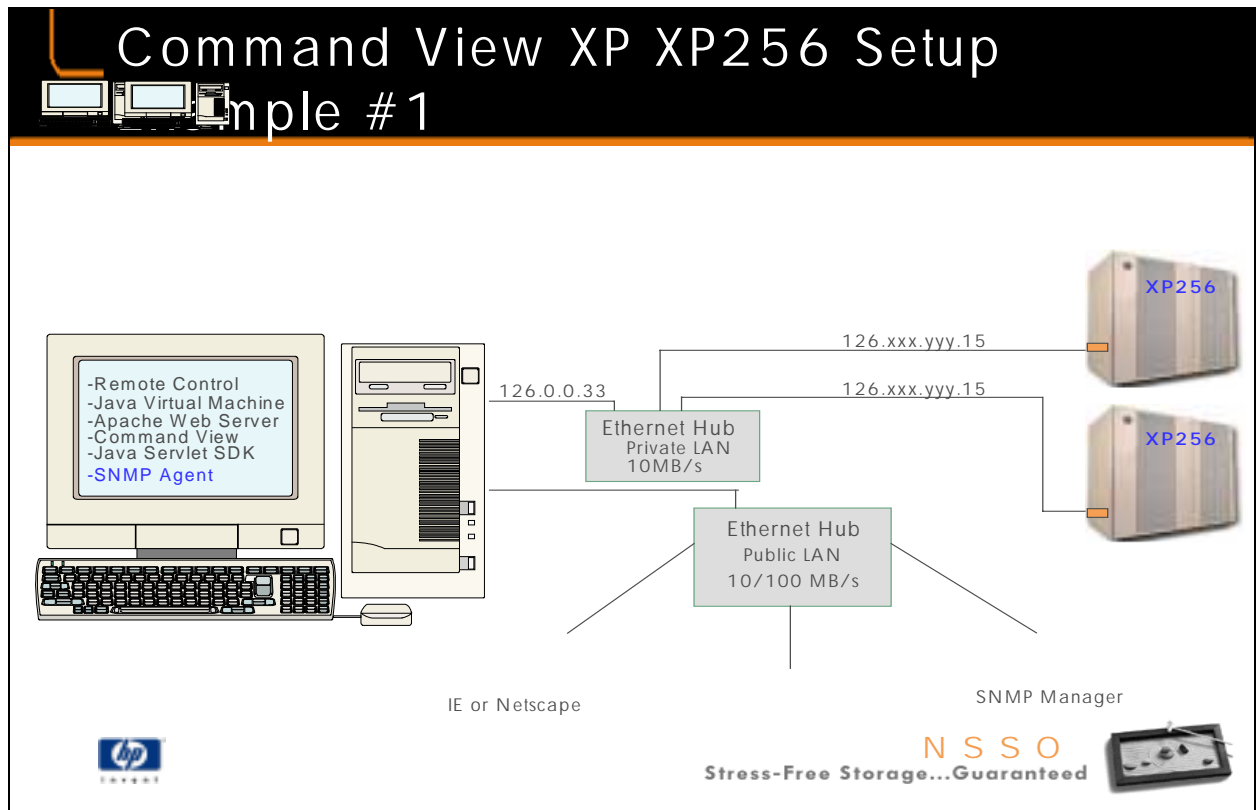
- XP256= Yes. All communication to the SVP occurs via the private LAN. In order for Command View to utilize this private LAN connection, Remote Control XP with the SNMP key enabled is a requirement.
- XP512= No. All communication by Command View is done via the public LAN (the SNMP agent now resides on the SVP and not the RMC PC).

If the customer uses both Remote Control and Command View on the same PC, the Remote Control XP software will still require a private LAN connection into the XP256 or XP512's Service Processor (SVP).

Do I really not need Remote Control XP if I only have an XP512/48 and I currently run Command View?

Technically speaking, as far as setup and supportability you do not need RMC. However, Command View XP currently does not have feature parity with Remote Control XP. For example, as of June 2001 Command View v1.11 is lacking Business Copy, Continuous Access, remote key access, plus a few other features found in RMC. However, Cmd View offers much more than RMC in respects to total array management solutions.



Remote Control XP is a requirement for Command View XP when connecting to an XP256. Command View XP relies on the SNMP module found inside of Hitachi's Remote Control XP software (this means the LUN Management software key disk #4 must be purchased and activated on the Remote Control PC). As this example illustrates, RMC and Command View XP are on the same PC.

### HOW DO I SETUP THIS CONFIGURATION TO WORK WITH CMD VIEW?

Install RMC on the PC that contains both a private LAN connection the XP256 and public LAN connection

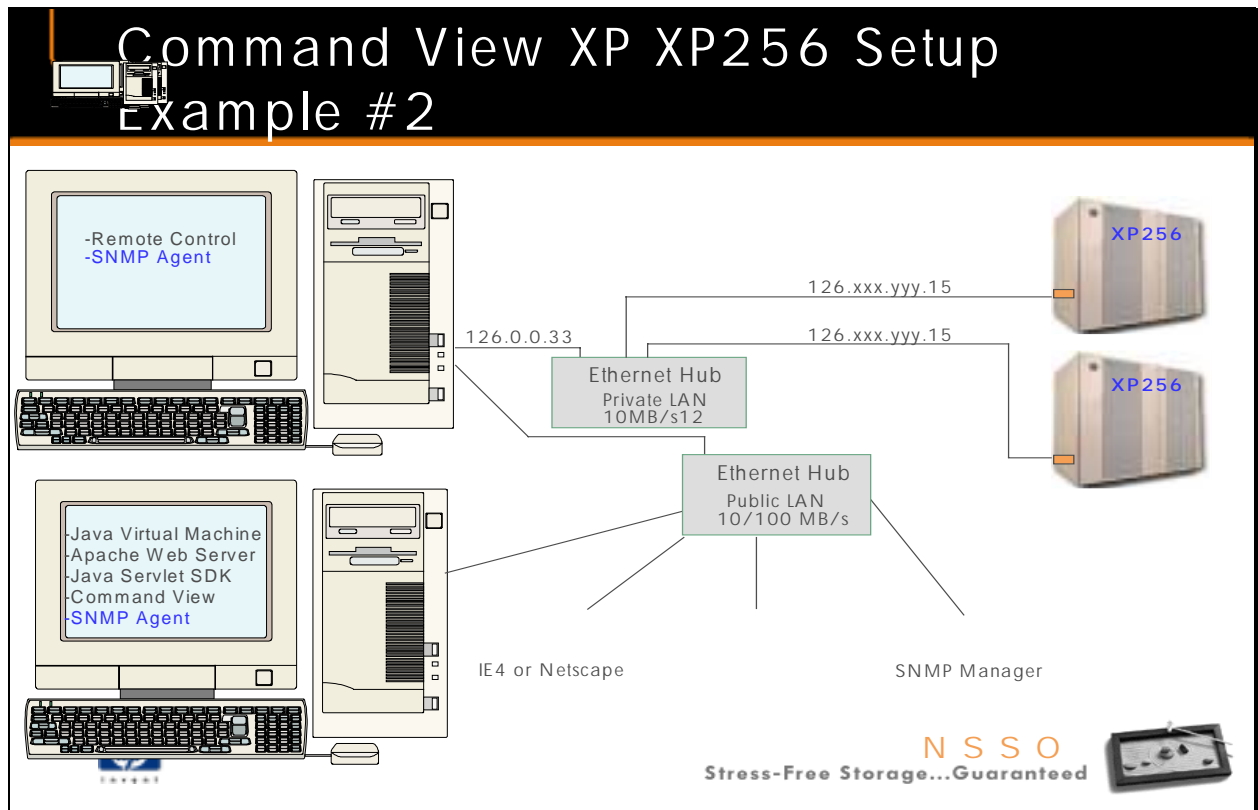
Enable the SNMP option inside of Remote Control XP software. Enable "Extension SNMP for HP RAID Series"

Install Command View XP on the same PC

On the Command View XP PC, install the SNMP Services for NT (click on Start-Control Panel-Network and install the SNMP service from the Microsoft NT CDROM). Following the install of the service, it is not necessary to define any settings for the SNMP service (ie- Trap Tap etc). Remember to reapply any service packs (if you don't you will see an error at NT's startup).

5. On the Command View XP PC, Control Panel-Services-SNMP, set to "interact with desktop"

6. Inside of Command View, click on Device Administration, then define the Remote Control XP's TCPIP address inside box labeled "Agent IP Address" (the box that contains the SNMP agent).



The above example shows Command View XP and Remote Control XP on different PCs. As stated previous in example #1, Remote Control XP is still a requirement for Command View XP when connecting to an XP256, but it can be installed and running on a different machine than the Cmd View PC.

The previous slide showing XP256 Example #1 will be the most common and logical Command View XP installation scenario. Rational is the following: RMC is a requirement for Command View XP when managing the XP256. It makes sense to place both software packages on the same machine. Command View XP against the XP256 will never have Auto LUN, BC or CA support, therefore if the customer wishes to use any of these applications they will always need RMC installed. Example #1 is one less piece of hardware required in the data center to purchase and manage.

#### SETUP NOTES:

Install RMC on the PC that contains both a private LAN connection the XP256 and public LAN connection

Enable the SNMP option inside of Remote Control XP software. Enable "Extension SNMP for HP RAID Series"

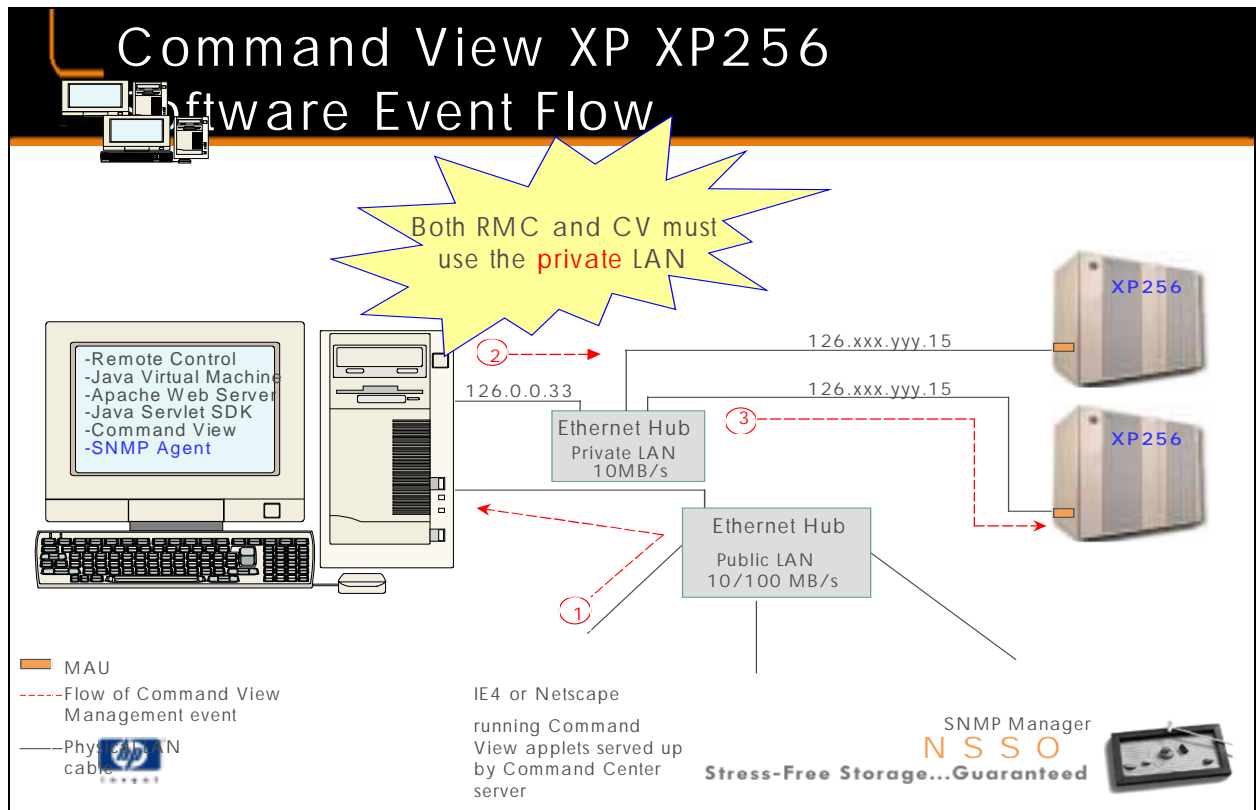
From Control Panel on the Remote Control PC, click on Services-SNMP, and check set to "interact with desktop"

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Command View XP

On a separate PC, install Command View XP.

On the Command View XP PC, install the SNMP Services for NT (click on Start-Control Panel-Network and install the SNMP service from the Microsoft NT CDROM). Following the install of the service, it is not necessary to define any settings for the SNMP service (ie- Trap Tap etc). Remember to reapply any service packs.

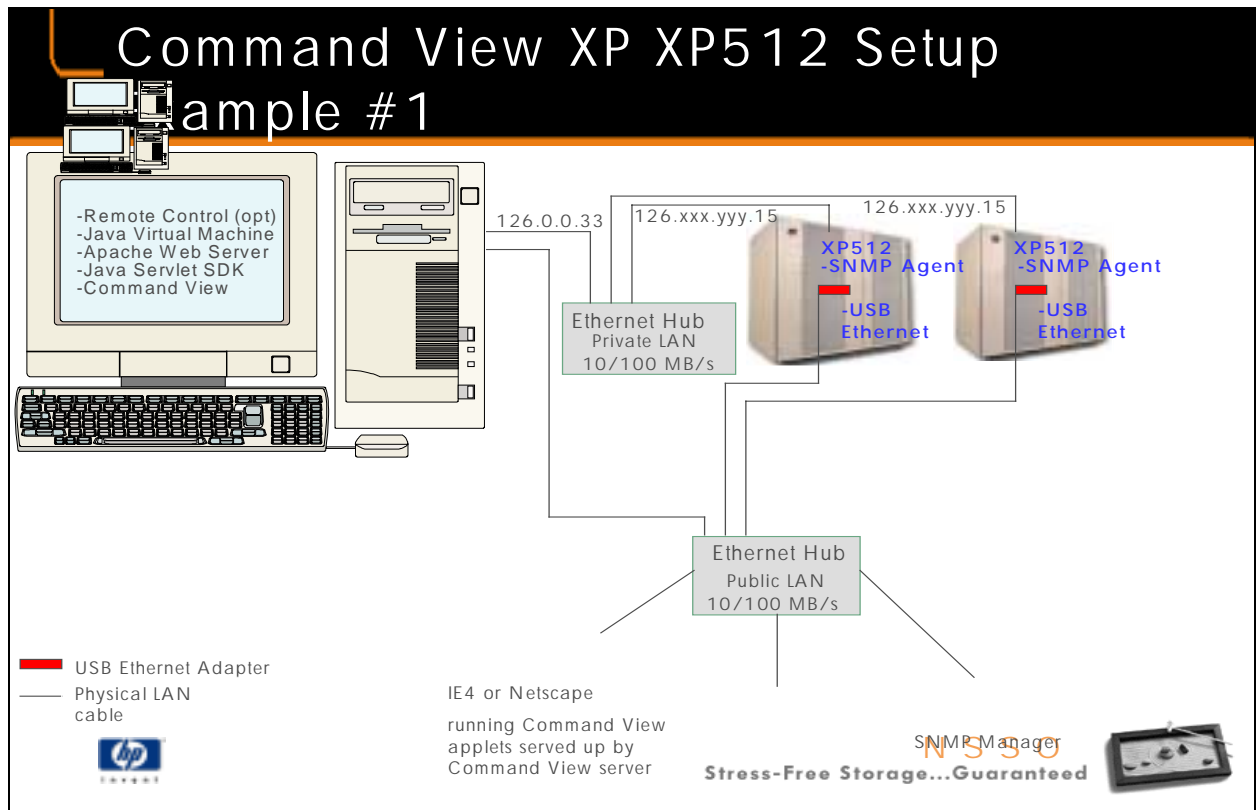
Inside of Command View, click on Device Administration, then define the Remote Control XP's TCPIP address inside box labeled "Agent IP Address"



Command View XP software event flow (from the public LAN browser to the SVP) for XP256 connected arrays:

1. Running either Netscape or Internet Explorer from a public LAN connected host (it is not a requirement that the browser host be on a "public" LAN- it can be run from the private LAN PC) the workstation makes an information read or change request from Command View XP. This request gets sent to the Remote Control PC running the Apache server with Command View.
2. Command View receives and process the request. The request is forwarded off to the SVP via the SNMP agent packaged inside of the Remote Control XP software via the private LAN as an SNMP MIB.
3. The SVP process the MIB and returns result output back via the private LAN to the Apache Web Server running the Command View XP software; from here the result is relayed to the public LAN and displayed back to the Netscape or IE browser connected host.

Just like the Remote Control XP software, all Command View software connectivity to the XP256 is only possible via the private LAN; however with Command View installed a user has public to private LAN XP array management access (by taking the http public LAN connection and ricocheting off the Remote Control PC to gain access to the private XP LAN). As this slide presentation will illustrate, the Command View XP connectivity to the XP512/XP48 is done exclusively via the public LAN.

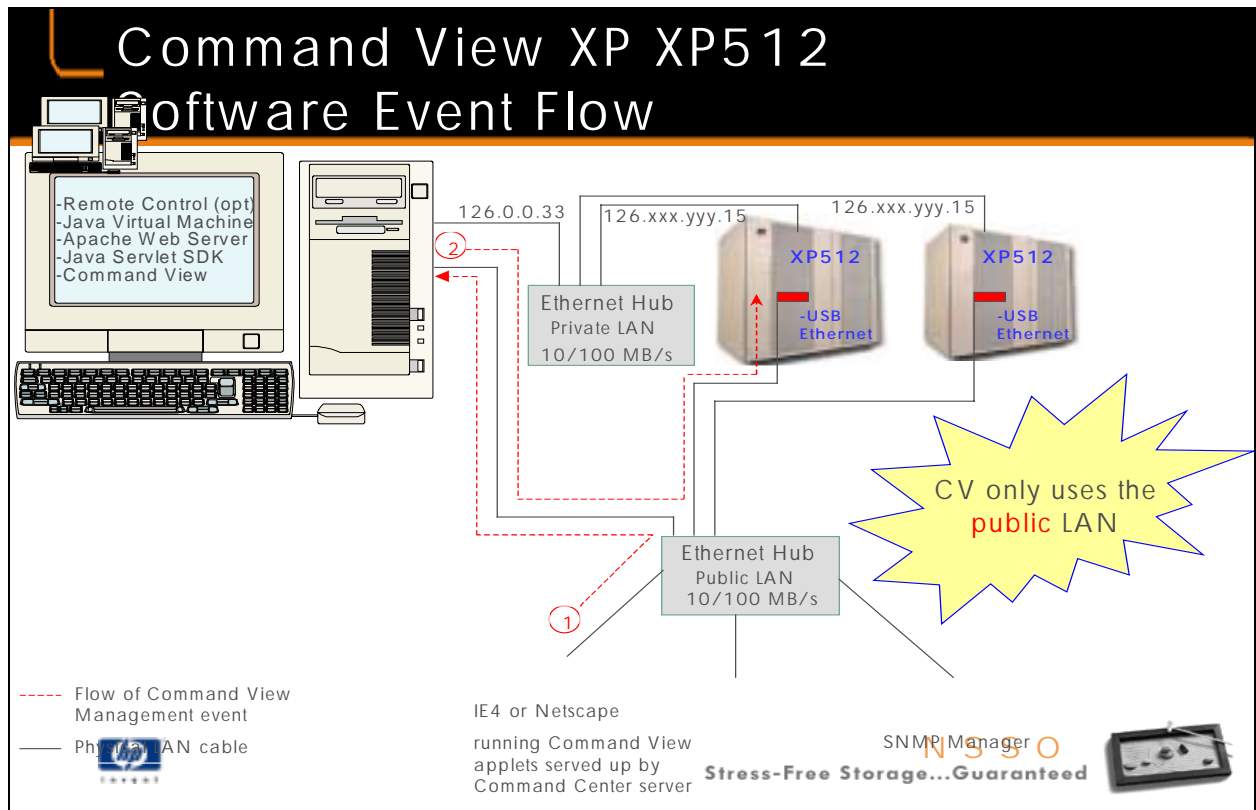


One primary difference to observe is the SNMP agent now exist on the SVP of the 512 and not on the Remote Console PC as it did for the XP256 RMC software. For this setup, we continue to have the SVP on the private LAN and we open up the SVP to the public LAN for SNMP extended SNMP functionality (ie- forwarding TRAPS to a defined SNMP event manager; allowing SNMP management access via Command View).

As the Command View software matures, it will no longer be necessary to have the Remote Control XP software installed on the management console for the XP512; this would also eliminate the need for the private LAN connection between the Remote Control PC and the XP 512 SVP. Why is RMC necessary for the XP512? It isn't, unless you need Business Copy or Continuous Access support via the GUI (you could always use BC/CA via Raid Manager's command line).

The SNMP agent is no longer a requirement on the Remote Control PC- it has been moved to the SVP. For this reason, all SNMP trap information is forwarded directly from the SVP via the public LAN to the SNMP management station (HP OpenView, Tivoli etc) or to Command View.






Command View XP software event flow (from the browser host to the SVP) for XP512 connected arrays:

1. Running either Netscape or Internet Explorer from a public LAN connected host the workstation makes a information read or change request from Command View XP. This request gets sent to the Remote Control PC running the Apache server with Command View. This puts a permanent lock on the SVP (explained in more detail in slides to follow).
2. Command View receives and process the request. The request is forwarded off to the SVP via the public LAN as an SNMP MIB.
3. The SVP process the MIB and returns result output to the Apache Web Server running the Command View XP software.
4. The result of the request is displayed via Netscape or IE on the browser connected host.

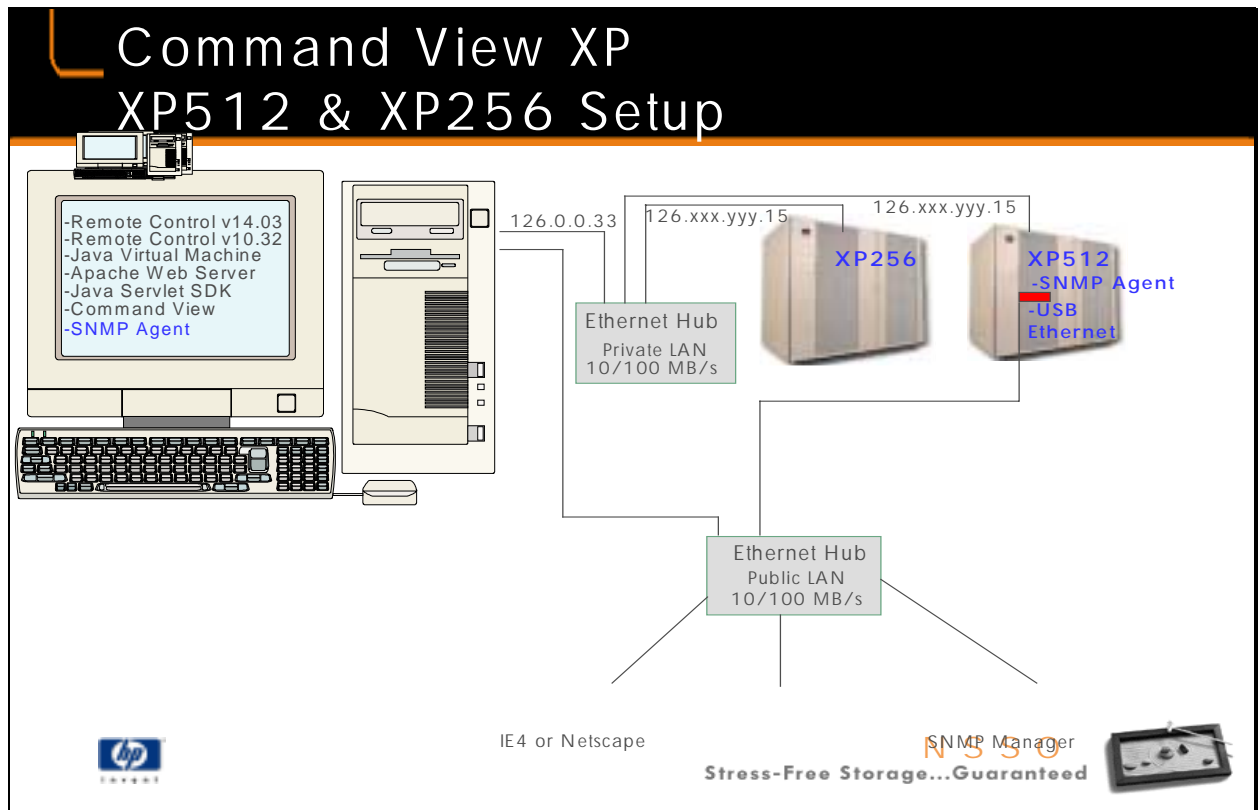
NOTE: Primary difference from previous slide showing software event flow for the XP256 is the software request displayed above is the public LAN. The XP256 being managed by the same Command View software all software event flow is via the private LAN, thus it will always require the Remote Control XP software to be installed.



Command View XP software event flow (from the browser host to the SVP) for XP512 connected arrays:

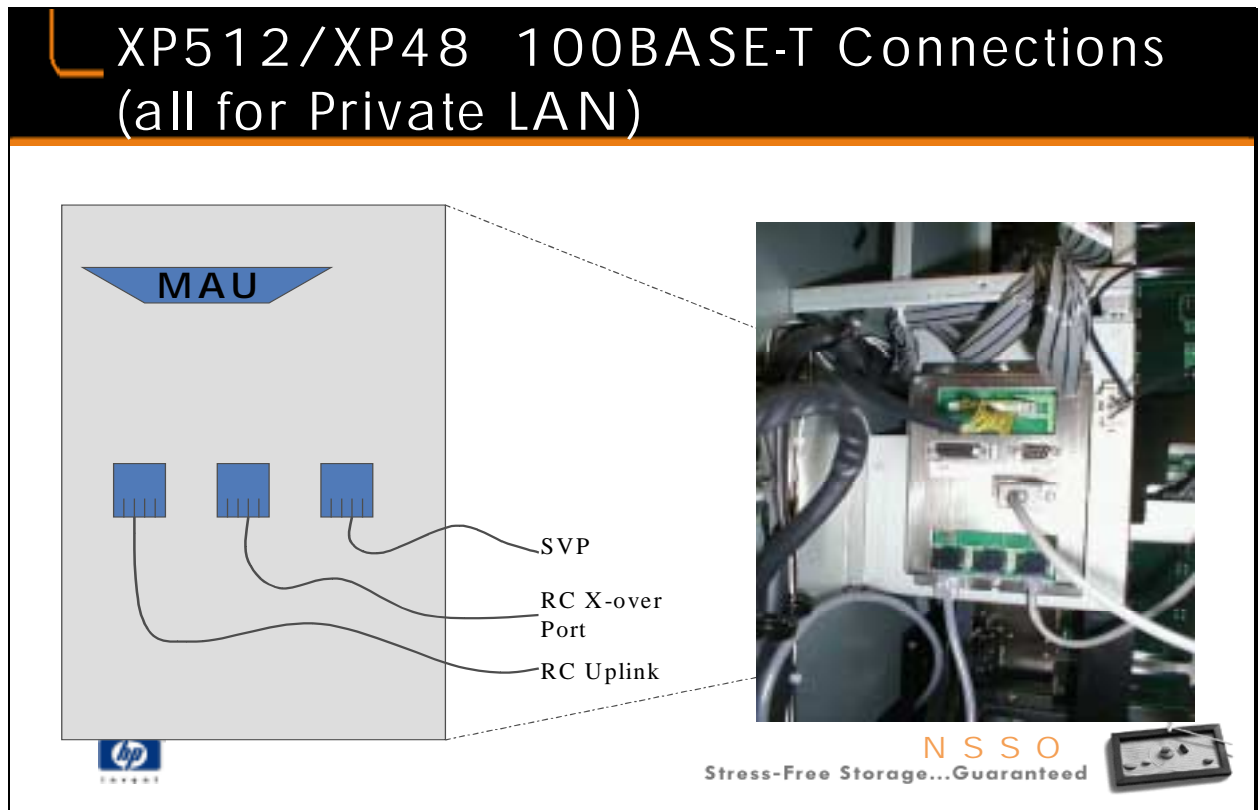
1. Running either Netscape or Internet Explorer from a public LAN connected host the workstation makes a information read or change request from Command View XP. This request gets sent to the Remote Control PC running the Apache server with Command View. This puts a permanent lock on the SVP (explained in more detail in slides to follow).
2. Command View receives and process the request. The request is forwarded off to the SVP via the public LAN as an SNMP MIB.
3. The SVP process the MIB and returns result output to the Apache Web Server running the Command View XP software.
4. The result of the request is displayed via Netscape or IE on the browser connected host.

NOTE: Primary difference from previous slide showing software event flow for the XP256 is the software request displayed above is the public LAN. The XP256 being managed by the same Command View software all software event flow is via the private LAN, thus it will always require the Remote Control XP software to be installed.



The XP256 and XP512's SVP still only allows one modify mode connection at a time. However, Command View running on only one PC will allow us to login via modify mode to two or more arrays (in this example we could have a modify mode connection to both the XP256 and the XP512 at the same time).


Command View will also allow us to login via modify mode to array A and also allow one or more View mode logins to array A (view logins will receive a screen refresh every 15 seconds on certain screens in the event of a change of display information). In the scope of RMC, SVP and Cmd View, modify access is also referred to as a "lock" on the array, or "modify lock."



The XP512 has three direct 100Base-2 ports into an on-board hub for XP512 LAN connections. These three available ports are for connecting the Remote Control XP PC to the SVP via the private LAN. Only two of the three connections will ever be utilized when the machine is in production (dependent upon if a standard cable is plugged into the far left port or the crossed-over center port coming directly from the Remote Control PC). This is to say that the left hand port OR the middle port will be used to connect to the RMC (not both), and the right hand port will ALWAYS connect to the SVP's onboard RJ45 port (not to the SVP's USB RJ45 dongle- illustrated in slides to come).

## XP512/XP48 Physical LAN Connections

- External LAN = USB Adapter with USB Ethernet RJ45 Connector (12Mbps)
- Internal LAN = On-board RJ45 Ethernet Adapter (100Mbps)



To the right of the USB input is the laptops internal or on-board RJ45 ethernet connector. This port is not supported as a replacement for the USB ethernet connection for the public LAN. The SVP has hard coded options for this port; this means if you have altered the settings for the port (added a protocol, modified a protocol, added an adapter etc) the SVP will overwrite those settings on boot or become error prone (ie- the SVP's inability to properly communicate with the DKC's internal components).

USB RJ45: Connects to the external LAN. Cmd View and outgoing SNMP traps will use this ethernet connection.

OnBoard RJ45: Connects to the XP private LAN. Remote Control will use this ethernet connection.

The laptop's PCMCIA ports now house an PCMCIA modem and a 64MB flash memory card (used by the SNMP code), thus one PCMCIA port is for the asynchronous modem, the second is for the flash memory.

## Obtaining "Modify" Rights

- Each SVP allows for only one MODIFY connection to the array's SVP for array management functions (aka "lock")
  - This MODIFY connection can come from...
    - ♦ Remote Control XP software
    - ♦ Command View XP software
    - ♦ Directly on the SVP.
  - If one agent obtains a lock, only that agent can release their lock (ie- From the SVP, a user via the software could not kick a Remote Control XP session off via the user interface).



N S S O  
Stress-Free Storage...Guaranteed



Although the SVP's user interface does not allow an administrator to release the lock, if physical access to the SVP is possible then rebooting the SVP would release a lock obtained from Remote Control or from the SVP itself, but this would not work for a lock acquired by Command View.

## Obtaining "Modify" Rights

- Each SVP allows for only one MODIFY connection to the array's SVP for array management functions (aka "lock")
  - This MODIFY connection can come from...
    - ♦ Remote Control XP software
    - ♦ Command View XP software
    - ♦ Directly on the SVP.
  - If one agent obtains a lock, only that agent can release their lock (ie- From the SVP, a user via the software could not kick a Remote Control XP session off via the user interface).



N S S O  
Stress-Free Storage...Guaranteed



Although the SVP's user interface does not allow an administrator to release the lock, if physical access to the SVP is possible then rebooting the SVP would release a lock obtained from Remote Control or from the SVP itself, but this would not work for a lock acquired by Command View.

## Command View "Locks" on the SVP

- Just like Remote Control XP, Command View must lock the SVP thus preventing change events occurring from more than one location. However, once Command View has obtained a lock, it never releases that lock (even after an SVP reboot).
  - How to release the lock? Logged in to Command View as "administrator" via a browser, from the MAIN tab go to Session Administration and hit "Unlock Array" and wait 3 minutes
    - ♦ NOTE: REFRESH already in progress will be allowed to complete
  - When must I think about doing this? Anytime there is a need to use Remote Control or require direct access to the SVP following a Command View connection



N S S O  
Stress-Free Storage...Guaranteed



Why does Command View maintain this constant lock on the SVP?

The purpose is to prevent Command View from having to go through a new refresh. Once Cmd View has established a lock, this means we have gone through the refresh time (amount of time varies per array) it takes for Command View to grab all array data and download it in a single MIB to the Command View server to display to the connected browser. From this point forward, Command View will only update most recently changed data, and not have to re-grab all data again with each new connection or until a change has been implemented via array management (ie- LUN configuration). There is no way to turn this feature off. Unlike RMC, Cmd View connection state to the array assumes constant, perpetual and never ending (thus it re-locks the SVP even after SVP reboots).

Implications of a static lock:

1. If we are performing a microcode update and we are instructed to reboot the SVP during that update, Command View will attempt to lock the SVP before and after the reboot. One possible workaround for a CE to take during a microcode update is to disconnect the external LAN cable during the entire upgrade, then replace the cable upon completion of the upgrade (invasive, but effective. This is not recommended unless the CE knows the exact outcomes of forcefully unlocking the Cmd View server).



SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Command View XP

If we need to make array management changes not supported by Cmd View (ie- BC or CA), we must first login in with a browser to Command View, unlock the array, then open up Remote Control XP and enter BC or CA.

Why wait 3 minutes? It may release the lock sooner than 3 minutes, but the code must allow for 3 minutes for the operation to either complete or fail the unlock request (this is not documented). Within HPSU's lab testing, waiting the entire three minutes is not generally the norm, but it has happened.

\*FYI: If you disconnect the external LAN cable at the SVP the lock should be free in 7 minutes. However, the proper way to 'unlock' is described above on this slide.

## “Lock” Failures?

- Issue: When a user clicks on an array to manage via Cmd View, the Cmd View server initiates a LOCK on the SVP, then it proceeds to collect data from the SVP (this is known as a REFRESH). Some attempts at this operation may fail during the REFRESH. This results in one of the following (*See speaker notes below for resolution*):
  - A second attempt may succeed
  - A second attempt may yield a result of “SVP already locked” when the RMC, SVP or Cmd View are not connected/locked against the SVP



N S S O  
Stress-Free Storage...Guaranteed



### Steps to resolve a failed LOCK/REFRESH attempt:

If the array failed its attempt on a refresh after displaying the REFRESH screen for 5 minutes or more, it will probably mean you can go straight to step 2 below. However, read step 1 just to make sure all issues are covered.

Read the error screen display by Command View and follow/check its recommendations (ie- is the array locked from another device etc?).

From the MAIN tab, click on Administration-Session Administration and hit UNLOCK ARRAY. Wait 3 minutes and attempt the lock again. If it fails, proceed to step 4.

From the Cmd View server PC....

Open a DOS window and type “net stop apache”

Shut down the Command View server (from the Cmd View server window, hit CTRL-C)

From the DOS window, type “net start apache”

Restart the Cmd View server (icon on the desktop or from the \Start Up folder).

Retry the lock/refresh from Command View. In over 50% of the attempts within the HPSU labs, this works. The other 50% require the next step be executed.

Perform step 4 above, then reboot the SVP. When this step is activated it is what HPSU will term a “Ghosted lock on the SVP.” This means that the SVP is locked, but no locking agent has control of that lock (thus RMC, SVP or Cmd View cannot “unlock” the locked SVP).

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Command View XP

\*these steps are NOT recommended if there is a chance of a SET operation in progress (ie- LUN creation).

## SNMP Management

- **XP256**- SNMP agent resides on the Remote Control PC (RMC)
  - R-SIMS forwarded by the SVP to the RMC; if SNMP agent is installed and configured on the RMC, SNMP traps are generated and sent to the SNMP management station
- **XP512**- SNMP agent resides on the Service Processor (SVP) laptop
  - R-SIMS forwarded by the SVP to the RMC; if SNMP agent is configured on the SVP, SNMP traps are forwarded to the SNMP management station
    - Unlike the RMC for the XP256, RMC for the XP512 does not have SNMP functionality



N S S O  
Stress-Free Storage...Guaranteed



SNMP connectivity into the SVP is made possible and only supported via the Universal Serial Bus (USB) ethernet connector supplied with the XP512 DKC. This kit includes one USB ethernet adapter and a 64MB flash memory placed in the PCMCIA slot. In the event of a R-SIM, the R-SIM will still be forwarded to the Remote Console (if installed) and a SNMP trap will be sent to the defined SNMP management station (ie- Open View or Command View server PC). However, the XP512 versions of the Remote Control XP does not have SNMP functionality like the XP256 Remote Console XP (SNMP has been removed, so the SNMP push button is no longer present on the first Remote Control screen after login).

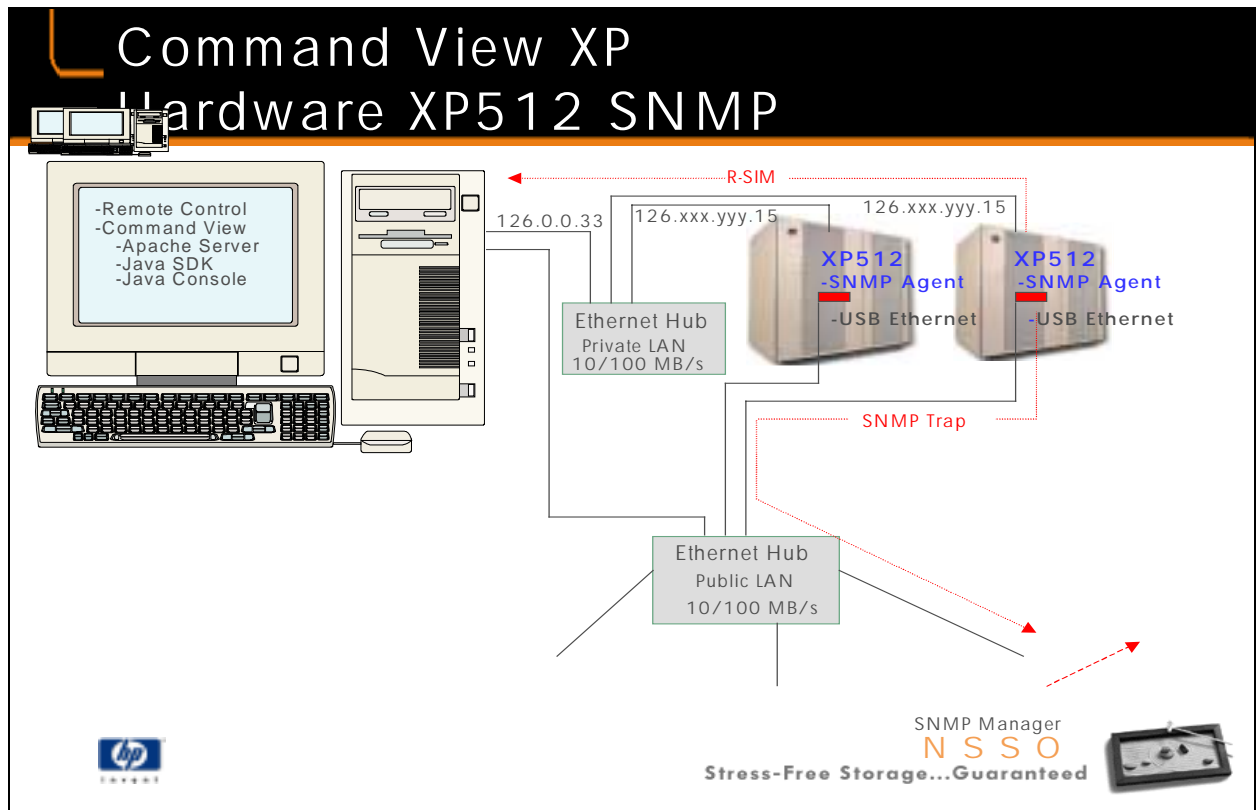
## Event Notification Terminology

- **Simple Network Management Protocol (SNMP)**- network management protocol used in TCP/IP networks (RFC1213)
- **SNMP Trap**- small block of data used to report network management information to one or more host
- **SNMP Message Information Block (MIB)**- the entire set of objects that specifies getting, putting or setting data variables on a host
- **SNMP Transport Protocol**= UDP
- **R-SIM** (Remote Single Information Message)- XP array error & health status messages sent from the SVP to the RMC PC.



N S S O  
Stress-Free Storage...Guaranteed






RSIMS still get sent from the SVP to the connected Remote Control XP via the private LAN. However, unlike the case with the XP256, the XP512 array will forward SNMP traps directly from the SVP to the SNMP management station. Although fully supported, this option may not be desirable. Please see slides to follow illustrating the use of TRAPDIST.BAT with Command View.


## Command View Version Data

- V1.10
  - Shipped only with XP48 units (FW 01.11.59; offered same functionality as v1.11, its predecessor and most current version for all XP arrays as of February 2001).
- V1.11
  - Performance improvement (Cmd View to XP512/48 only). See <http://esp.cup.hp.com>, search keyword COMMANDVIEW
  - XP256: Requires FW 52.47.06 or higher
  - XP48/512: Requires FW 01.12.18 or higher
  - Online operation supported (except for Secure Manager XP)
  - Added support for W2K
  - Supports Japanese browser and installation to Japanese NT box.
  - Apache starts automatically, but Cmd View and TrapDist.bat start only after logging into system (Startup Folder)



N S S O

Stress-Free Storage...Guaranteed



TrapDist.bat and the Command View server batch file are now placed in the Startup folder (in NT/W2K, the startup folder items only launch after a user has logged into the PC).

## Contents of the installation CD

- Command View 1.11
- Apache Web Server
- HP Surestore Remote Control XP
- CD insert
- Electronic software license certificate
- "Try N Buy" manuals on CD



N S S O  
Stress-Free Storage...Guaranteed





## Command View 1.11 has online configuration capabilities

- LUN configuration
  - Add/Delete path, command device, and LUN
- Port configuration
  - Associate LUNs with multiple ports
- Fibre Channel configuration
  - Fibre Channel addressing details



N S S O  
Stress-Free Storage...Guaranteed



### Student Notes

NOTE: Requires a firmware update to 01.12.xx for the XP48 and XP512.

NOTE: Online configuration is a feature of the XP48 and XP512. There is no plan to introduce this feature to the XP256. The XP256 will use Remote Control to perform configuration changes to the disk array.

## Installation – verify or install Remote Control and SNMP (Windows NT 4.0)

- Remote Control must be installed (see the User's Guide for more information)
- Confirm that SNMP is correctly installed and configured
- Confirm or configure the SNMP service to "Allow Service to Interact with Desktop"

NOTE: This procedure is for XP512, XP48, and XP256.  
RC is needed for all three arrays because of firmware key installation.

Over and above that, RC is needed for the XP256 only for the SNMP agent.  
For the XP512/XP48, the SNMP agent resides in the SVP. For the XP256, the SNMP agent resides in Remote Control.



N S S O  
Stress-Free Storage...Guaranteed



### Student Notes

To confirm that the MS SNMP service is installed:

1. Select Start/Settings/Control Panel/Network/Services
2. Verify that SNMP Services is displayed in the Network Services list.

To install SNMP:

1. Select Start/Settings/Control Panel/Network/Services
2. Select Add.
3. Select SNMP Service.
4. Select OK.
5. If Setup requires the i386 folder, enter the location (usually (C:\i386)).
6. If requested, insert the Windows NT CD-ROM and enter its location <drive>:\i386\.
7. Select Continue.
8. Reboot the PC.

To configure the SNMP Service to "Allow Service to Interact with Desktop":

1. Select the service window via Start/Settings/Control Panel.
2. Open the Services icon. Open the SNMP dialog box.
3. Select "Allow Service to Interact with Desktop".
4. Stop the SNMP service. Start the SNMP service.
5. Reinstall the Windows NT Service Pack 5.

## Installation – Command View XP

- Install the Java Runtime Environment 1.2.2.006 (jre-1\_2\_2\_006.exe)
- Install HP Surestore Command View XP (setup.exe)
- Enable the Command View TrapDistributor function
  - Click on the desktop icons, OR
  - Use Start/Run, OR
  - Login to Windows

**NOTE: Microsoft SNMP network services must be installed and configured. If not, TrapDistributor will terminate**



N S S O  
Stress-Free Storage...Guaranteed



### Student Notes

To install the Java Runtime Environment 1.2.2.006:

1. Execute jre-1\_2\_2\_006-win.exe.
2. Click Next at the Welcome screen.
3. Click Yes at the Software License Agreement.
4. Click Next at the Warning screen.
5. Choose the Destination Location for the JRE.

The Default location is recommended, but you can choose a different location for your JRE.

6. Choose Typical Setup and Select Next.
7. Click Next at Select Program Folder window.
8. Click Finish when installation complete.

To install Command View:

1. Insert the Command View XP Installation CD-ROM into the disk drive.
2. Select Start/Run/Browse and navigate to the Command View XP CD.
3. Select setup.exe and click to start the installation program.
4. Select Next at the InstallShield wizard dialog. Select Yes to accept the license agreement.
5. Enter a User Name, Company Name and Serial Number. Select Next.
5. At the Choose Destination Location dialog, either click the Next button to accept the default

Destination Folder (C:\HPSS) or click the Browse button to choose a different folder.

6. At the Setup Type dialog, accept the default Typical selection and click Next.

7. At Select Program Folder, accept the default program folders and click Next.

The setup program transfers application files to destination folder and configures Command View XP.

The setup program will then start the Command View XP server processes, which include

- Apache web server
- The servlet engine
- HP SureStore Command View XP TrapDistributor

8. InstallShield will display "Installation Complete". The installation program will determine if

Windows needs to be rebooted and prompt the user to choose the reboot radio button. Reboot

is usually not required. Click Finish to complete the installation.

To run Command View and TrapDistributor:

1. Click on the desktop shortcuts (HP Command View XP Server and TrapDistributor)

OR select:

Start/Programs/HP Surestore Command View XP/HP Command View XP Server

Start/Programs/HP Surestore Command View XP/HP Command View XP TrapDistributor

OR

Login to Windows. The Windows Common Startup group contains shortcuts to the HP


Command View Server and to the TrapDistributor.

NOTE: The TrapDistributor program requires that the Microsoft SNMP network services are installed and configured.

If the Microsoft SNMP network service is not installed, then the TrapDistributor will terminate and NOTEPAD.exe will open and display the TrapDistributor FAQs of possible reasons and solutions.

## TrapDist.Bat

- Definition: Batch file that is installed by the Command View software that runs in memory on the Command View server. This batch file forwards SNMP traps that Command View has converted to the defined SNMP management station.
  - Benefit: traps forwarded by the Command View server to the SNMP management station are decoded



Just like SERV.BAT, TrapDist.Bat must manually be opened into memory after machine login. This means that if the machine reboots, it is not running by default.

## Starting Command View

- Start the web browser
- Enter the IP address of the Command View XP Management Station in the Address field
- Log in:
  - Enter "administrator" in the User Name field
  - Enter "administrator" in the Password input field
- The main screen appears



N S S O  
Stress-Free Storage...Guaranteed



### Student Notes

To start Command View:

1. Start the web browser.
2. Enter the IP address of the Command View XP server in the Address field.  
The Command View XP Login prompt should appear.
3. Enter "administrator" in the User Name field (Do not include the quotes).
4. Enter "administrator" in the Password field (Do not include the quotes).  
The Main menu appears and displays a prompt to "... add new Array".  
The application is functioning.

Note: The default user ID is "user". The default user password is "user".

Note: After 15 minutes of inactivity, the administrator login times out. This is a security feature. The user login never times out.

## Installing Command View with Performance Advisor

- You may install either product in any order
- The second product should be installed into the same root directory as the first product
- Example:
  - If PA is already installed into D:\HPSS, then you should install CV into D:\HPSS



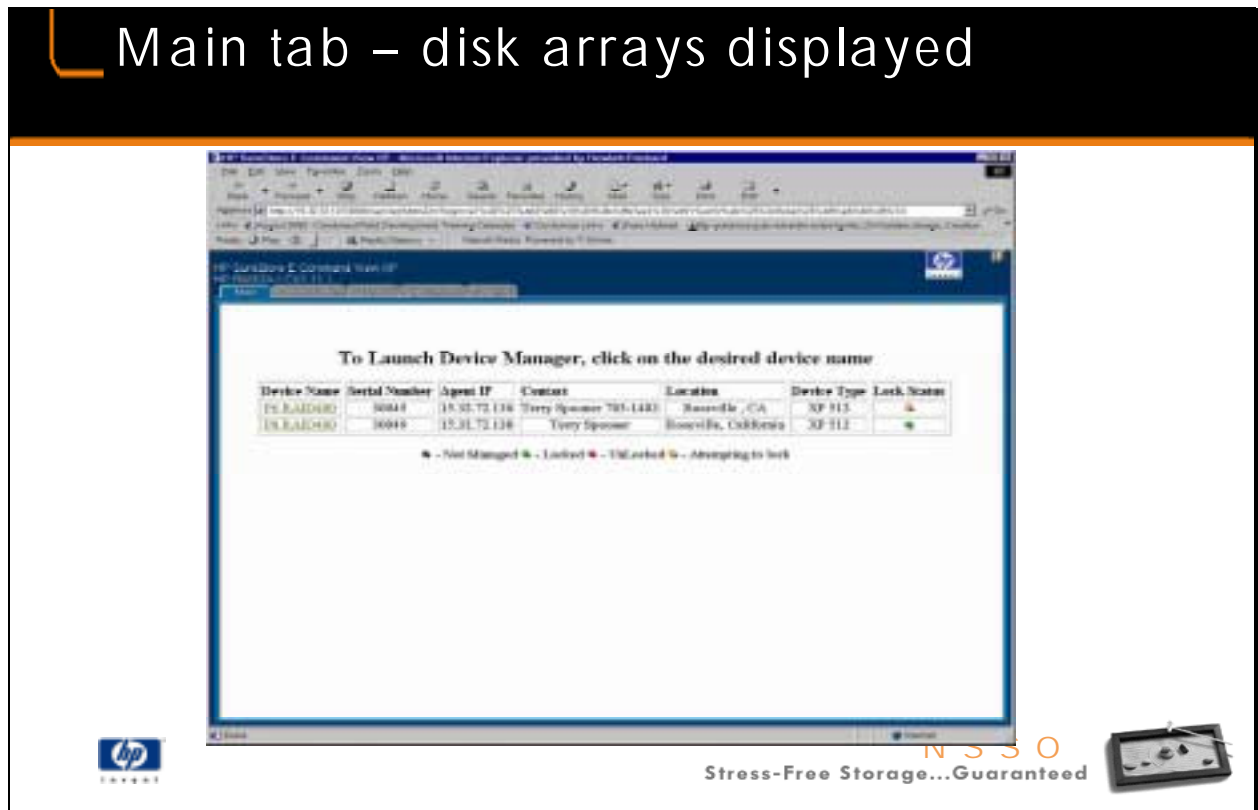
N S S O  
Stress-Free Storage...Guaranteed



### Student Notes:

A note about the HP installed Apache web server service names:

1. If CV is installed first, then the Apache web server as is listed as "Apache" in the Services applet.
2. If PA is installed first, then the Apache web server as is listed as "HPSS Apache" in the Services applet.
3. Both CV and PA work correctly with either service name.



### Student Notes

To start Command View:

1. Start the web browser.
  2. Enter the IP address of the Command View XP Management Station in the Address field. The Command View XP Login prompt should appear.
  3. Enter "administrator" in the User Name field (Do not include the quotes).
  4. Enter "administrator" in the Password field (Do not include the quotes).
- The Main screen appears.

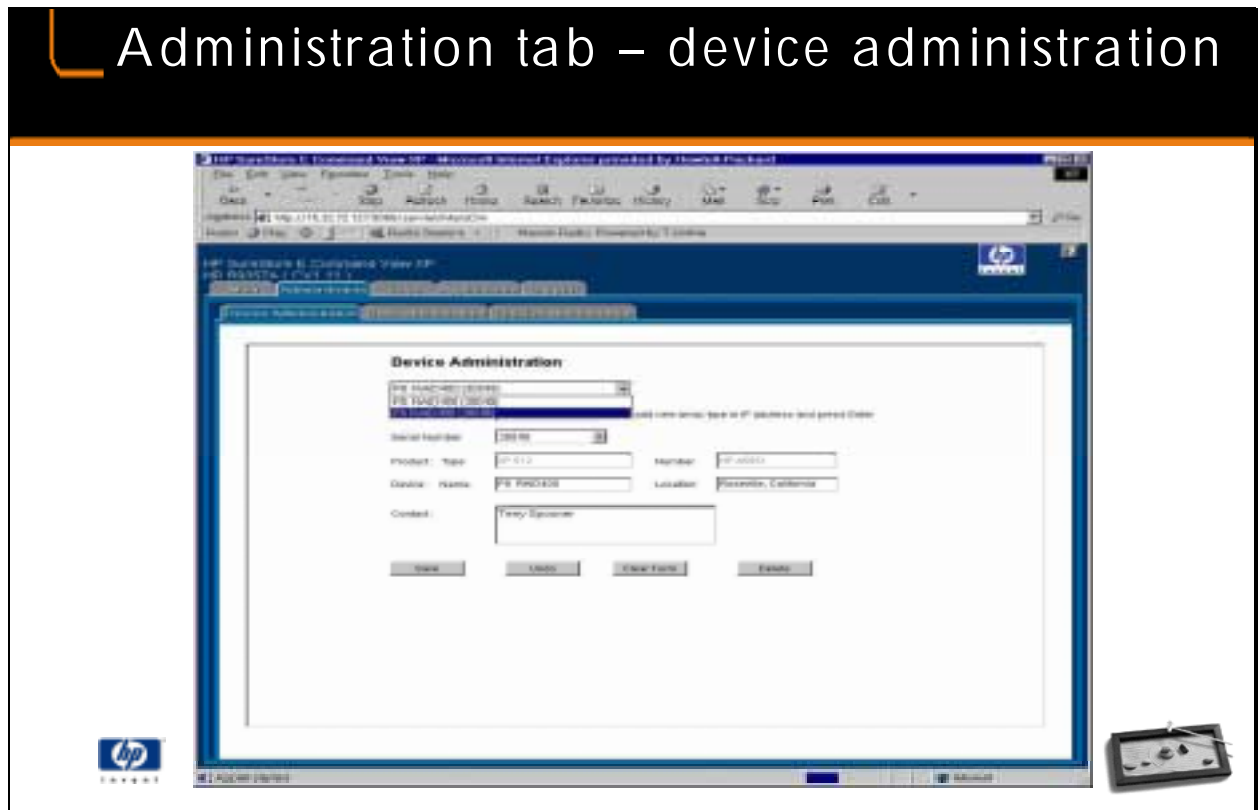
The main screen has tabs for Main, Administration, Traps, Applications, and Support. We'll look at these tabs a little later.

However, before you can enter array management functions, you have to configure devices using the Administration tab.

There are no disk arrays configured, so we'll break the sequence and choose the Administration tab.

Note: The above are the default administrator ID and password. The default user ID is "user". The default user password is "user".





## Student Notes

When you select the Administration tab, the Administration page appears. You can select a tab for Device Administration, User Administration, or Session Administration.

To register a new disk array:

1. Enter the IP address.
2. Enter the a name for the device, a location, and a contact (name).
3. Select Save to register the device.

To unregister a disk array:

1. You cannot unregister a device that is in use. From the Session Administration page, select  
Unlock Array, and the select Stop Managing.
2. From this Device Administration page, select the device.
3. Select Delete.


To change device parameters:

1. Select the device.
2. Modify the parameters you want to change (name, location, or contact name).

3.     **Select Save.**  
      To revert to the previous values, select Undo.


## HP SureStore E Command View XP *User*

### Administration Screens




**Add New Users**

- Set & confirm passwords
- Obtain reports of users logged in & time logged in
- Provides password encryption






**Delete User**



Provides Four Levels of Access:

- Level 1 - read-only function \*
- Level 2 - access to read-write functions
- Level 3 - access to all functions EXCEPT LUN Security
- Level 4 - administrator level, ability to modify passwords, user profiles, view/delete users\*

\* At introduction

Free Storage...Guaranteed

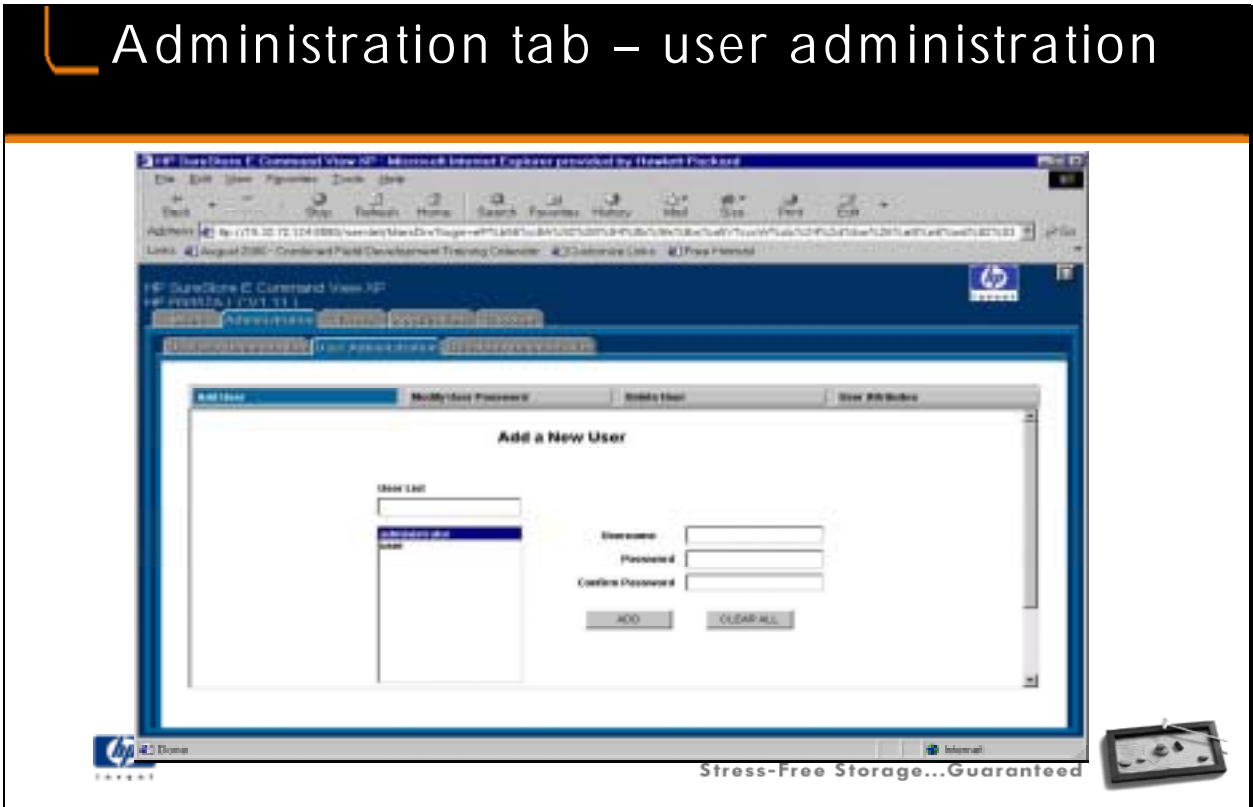
The User Administration tab enables the administrator to add, modify or delete user profiles via the GUI interface. From this interface the administrator can also set and confirm passwords as well as obtain reports of users currently logged in with their time logged in and encryption of password provided by the user. While Command Center will ultimately have four predefined levels of functionality: Level 1, 2, 3, 4. Level 4 will be the highest and Level 1 the lowest. For the first release, we will provide Level 1 and Level 4 functionality. These are defined as:

Level 1 - read-only functions in LUN management screens. Level 1 users can modify passwords and view other users available.

Level 2 - access to read-write functions in LUN management screens. Level 2 users can modify passwords and view other users available.

Level 3 - access to all functions EXCEPT LUN security. The Level 3 user can modify their password and view other users available.

Level 4 - administrator level access to all functions and operate in all modes (online or offline). The Level 4 user can add, modify passwords and users profiles, view and delete other users.



### Student Notes

User Administration lets you add a new user, modify a user password, delete a user, and view or modify user attributes. To perform administrative functions, you must have Modify permission.

To add a new user:

1. Select Add User.
2. Type a name in the Username field.
3. Type a password in the Password field.
4. Type the password again in the Confirm Password field.
5. Select ADD.

To change a user's password:

1. Select Modify User Password.
2. Select a user's name from the list.
3. Type the old password in the Old Password field.
4. Type a new password in the New Password field.
5. Select the MODIFY button.
6. Select OK.

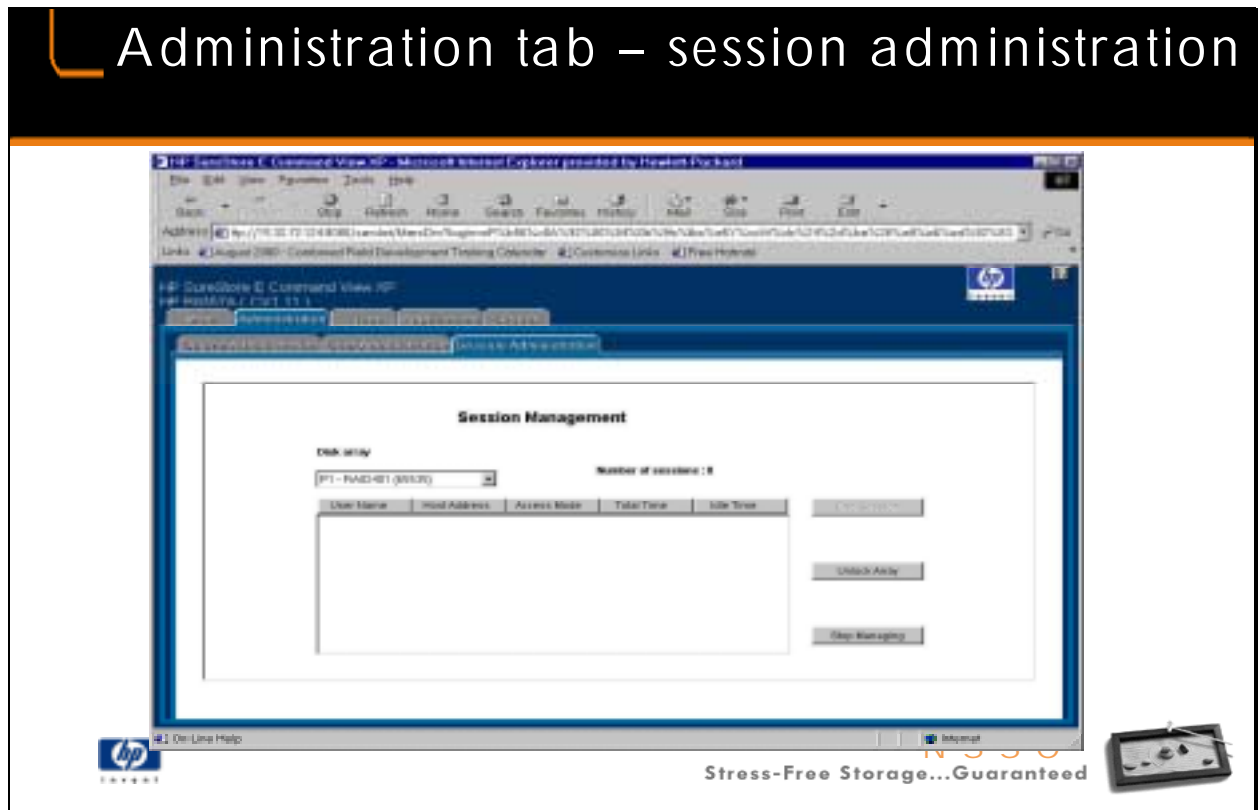
To view or modify user attributes:

1. Select User Attributes.
2. Select a name from the list. You can now view the permissions for that user.
3. To change permissions, highlight each function you want to change, such as "modify" or "view" access.

4. Select Change.
5. When you have finished assigning permissions, choose the SET button.
6. Select OK.

To delete a user:

1. From the User Administration page, select the name to be deleted.
2. Select the Delete button.
3. Select OK.



### Student Notes

Session Administration lets you delete a session, unlock an array, or stop managing an array.

To delete a session:

1. Select a disk array. All sessions for that disk array are displayed.
2. Select the session user name to be deleted.
3. Select End Session.

To unlock a disk array:

To manage a disk array through Remote Control, you must unlock the array. You cannot unlock an array if there is an active modify session.

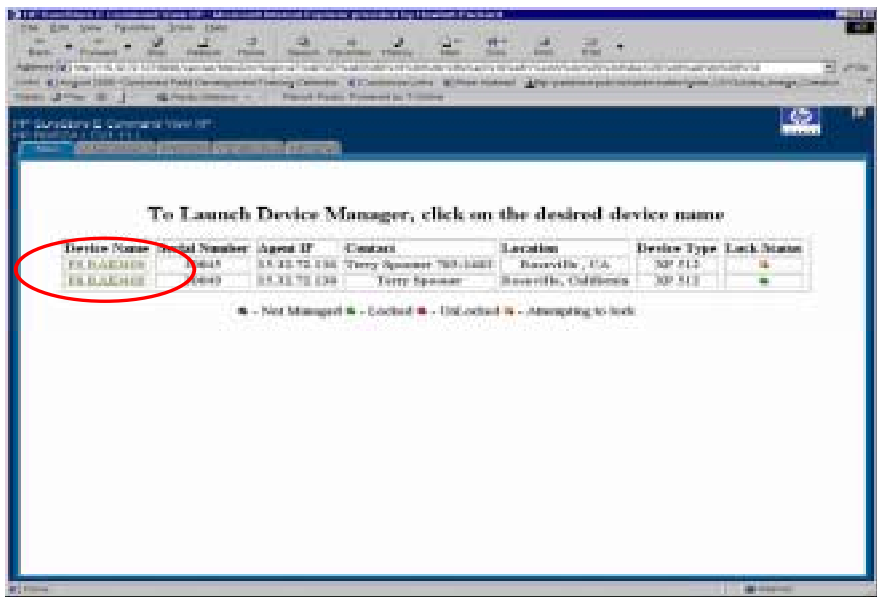
1. Delete the active modify session, if any.
2. Select the disk array. All sessions for that disk array are displayed.
3. Select Unlock Array.

To stop managing an array:

To terminate array management, you must delete all sessions for the array.

1. Delete all active sessions for the array.
2. Select the array you want to stop managing.
3. Select Stop Managing.

## Main tab



To Launch Device Manager, click on the desired device name

Device Name	Serial Number	Agent IP	Contact	Location	Device Type	Lock Status
ES-BLADDER	00000	10.10.10.10	Terry Spencer	Beverly Hills, CA	300 010	Not Managed
ES-BLADDER	00000	10.10.10.10	Terry Spencer	Beverly Hills, California	300 010	Not Managed

Not Managed Locked Unlocked Attempting to lock

N S S O  
Stress-Free Storage...Guaranteed

- We have returned to the Main tab, and now disk arrays are displayed
- Click on an array to see identity and perform management tasks

### Student Notes

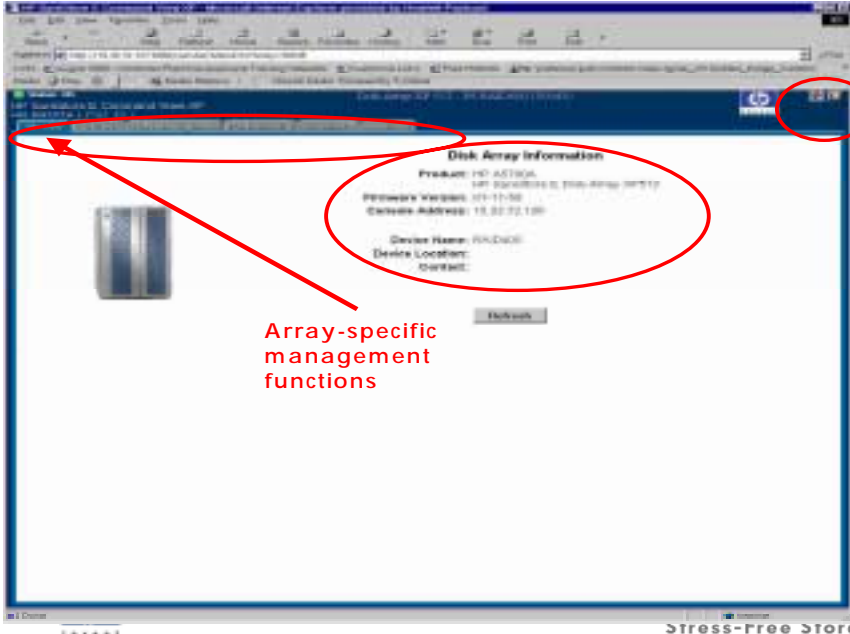
The main screen displays the following array information:

- Device name
- Serial Number
- Agent IP address
- Contact
- Location
- Device type
- Lock status

Lock status may be "Not Managed," "Locked," "UnLocked," or "Attempting to lock"

Click on an array to see more information about it and to go to the management screens for the array.

## Identity tab



Let's look at the Status tab next. We'll look at the management tabs in detail after the tour of the interface

Array-specific management functions

NSSO  
Stress-Free Storage...Guaranteed

### Student Notes

The Identity screen displays the following array general information about a disk array:

- Product number and type of disk array
- Firmware version
- Console address
- Device name
- Device location
- Contact name

### Tabs

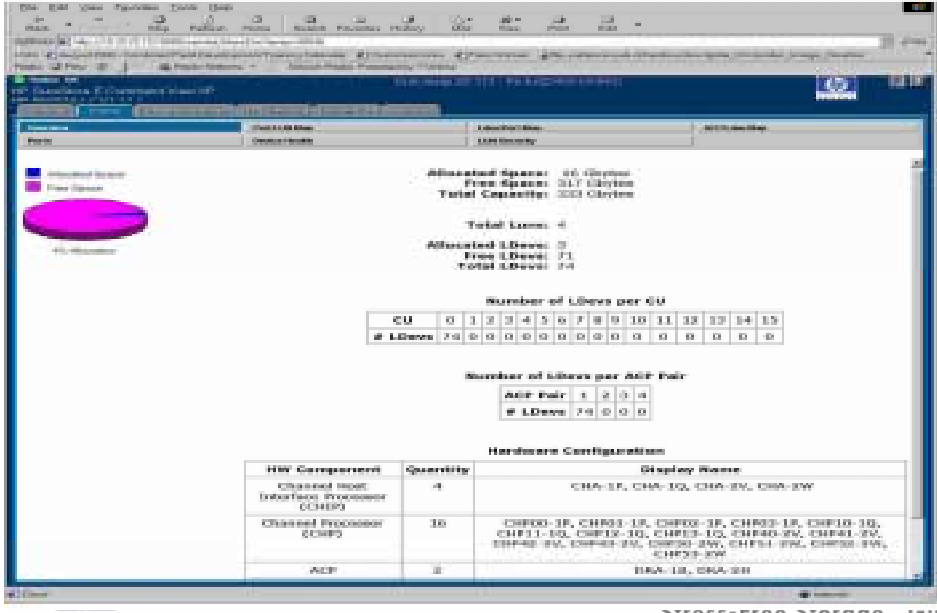
From this screen, you can select tabs for all array-specific management functions:

- Status
- LUN Management
- LUN Security
- Cache LUN
- AutoLun


To return to the main screen, click on the Home icon in the upper right of the screen.



## Status tab – Overview (button 1)



➤ We will now look at all seven status displays



### Student Notes

#### Buttons

The Status screen has buttons for seven information displays:

Overview - Displays an overview of the array and cache configuration information.

Port/LUN Map - Displays array configuration prioritized by port and LUN configuration.

Ldev/Port Map - Displays array configuration prioritized by port and LDEV configuration.

ACP/Ldev Map - Displays array configuration prioritized by ACP and LDEV configuration.

Ports - Displays a summary of port configuration.

Device Health - Displays the overall status of all the components in the disk array.

LUN Security - Displays a list of worldwide names (WWN) for each port in the array and the LUNs that are allowed to communicate with each WWN.

The Overview screen is displayed first. Click on the secondary buttons to see the other displays.

#### Overview screen

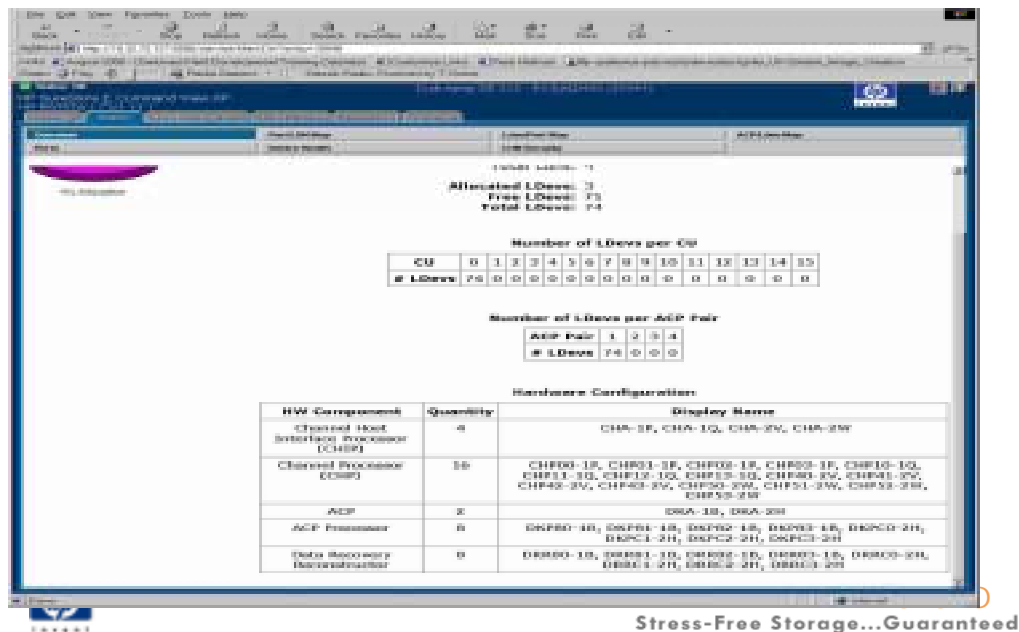
Pie chart % allocation, a graphical representation of the percentage of disk space allocated for use in the disk array.

Allocated Space -Amount of storage allocated in the current LDEV configuration.

Free Space - Amount of storage that is currently unallocated for use.

Total Capacity - Sum of allocated and unallocated disk storage. It represents the total of disk capacity for the selected array.

## Status tab – Overview (button 1) (continued)



### Student Notes

#### Overview screen (continued)

**Allocated Cache** - The amount of cache allocated to cache LUN extents (if Cache LUN is enabled).

**Free Cache** - The amount of cache available for normal I/O use (if Cache LUN is enabled).

**Total Cache Size** - Total cache in the selected array (if Cache LUN is enabled).

**Total LUNs** - The number of LUN paths configured on the disk array.

**Allocated Ldevs** - The number of LDEVs that are associated with LUNs.

**Free Ldevs** - The number of LDEVs not associated with a LUN (a path is not yet assigned to the LDEV).

**Total LDEVs** - The total number of LDEVs configured on the disk array. The number of LDEVs on a system is a function of the selected emulation type.

#### Tables

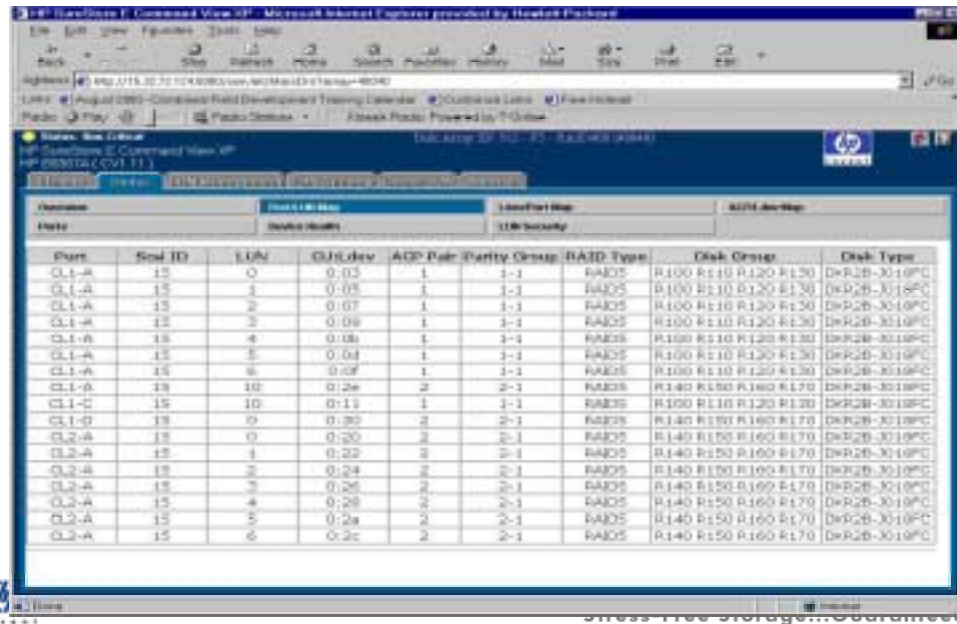
**Number of LDEVs per CU** - The number of LDEVs configured per CU.

**Number of LDEVs per ACP Pair** - The number of LDEVs configured per ACP pair.

**Hardware configuration** - The quantity and names of hardware elements:

Channel Host Interface Processor (CHIP), Channel Processor (CHP), ACP, ACP Processor, Data Recovery Reconstructor

## Status – Port/LUN Map (button 2)



Port	Scsi ID	LUN	CU:LDEV	ACP Pair	Parity Group	RAID Type	Disk Group	Disk Type
CL1-A	15	0	0:03	1	1-1	RAID5	R100 R110 R120 R130	DwR2B-3010FC
CL1-A	15	1	0:05	1	1-1	RAID5	R100 R110 R120 R130	DwR2B-3010FC
CL1-A	15	2	0:07	1	1-1	RAID5	R100 R110 R120 R130	DwR2B-3010FC
CL1-A	15	3	0:09	1	1-1	RAID5	R100 R110 R120 R130	DwR2B-3010FC
CL1-A	15	4	0:0B	1	1-1	RAID5	R100 R110 R120 R130	DwR2B-3010FC
CL1-A	15	5	0:0D	1	1-1	RAID5	R100 R110 R120 R130	DwR2B-3010FC
CL1-A	15	6	0:0F	1	1-1	RAID5	R100 R110 R120 R130	DwR2B-3010FC
CL1-A	15	10	0:1e	2	2-1	RAID5	R140 R150 R160 R170	DwR2B-3010FC
CL1-C	15	10	0:11	1	1-1	RAID5	R100 R110 R120 R130	DwR2B-3010FC
CL1-D	15	0	0:30	2	2-1	RAID5	R140 R150 R160 R170	DwR2B-3010FC
CL2-A	15	0	0:20	2	2-1	RAID5	R140 R150 R160 R170	DwR2B-3010FC
CL2-A	15	1	0:22	2	2-1	RAID5	R140 R150 R160 R170	DwR2B-3010FC
CL2-A	15	2	0:24	2	2-1	RAID5	R140 R150 R160 R170	DwR2B-3010FC
CL2-A	15	3	0:26	2	2-1	RAID5	R140 R150 R160 R170	DwR2B-3010FC
CL2-A	15	4	0:28	2	2-1	RAID5	R140 R150 R160 R170	DwR2B-3010FC
CL2-A	15	5	0:2A	2	2-1	RAID5	R140 R150 R160 R170	DwR2B-3010FC
CL2-A	15	6	0:2C	2	2-1	RAID5	R140 R150 R160 R170	DwR2B-3010FC

### Student Notes

The Port/LUN Map displays a map relating to port and LUN configuration.  
Fields

Port - The port number associated with the LUN.

SCSI ID - The target number for the LUN.

LUN - The LUN ID within the target for the LUN.

CU:LDEV - The first logical device associated with the LUN.

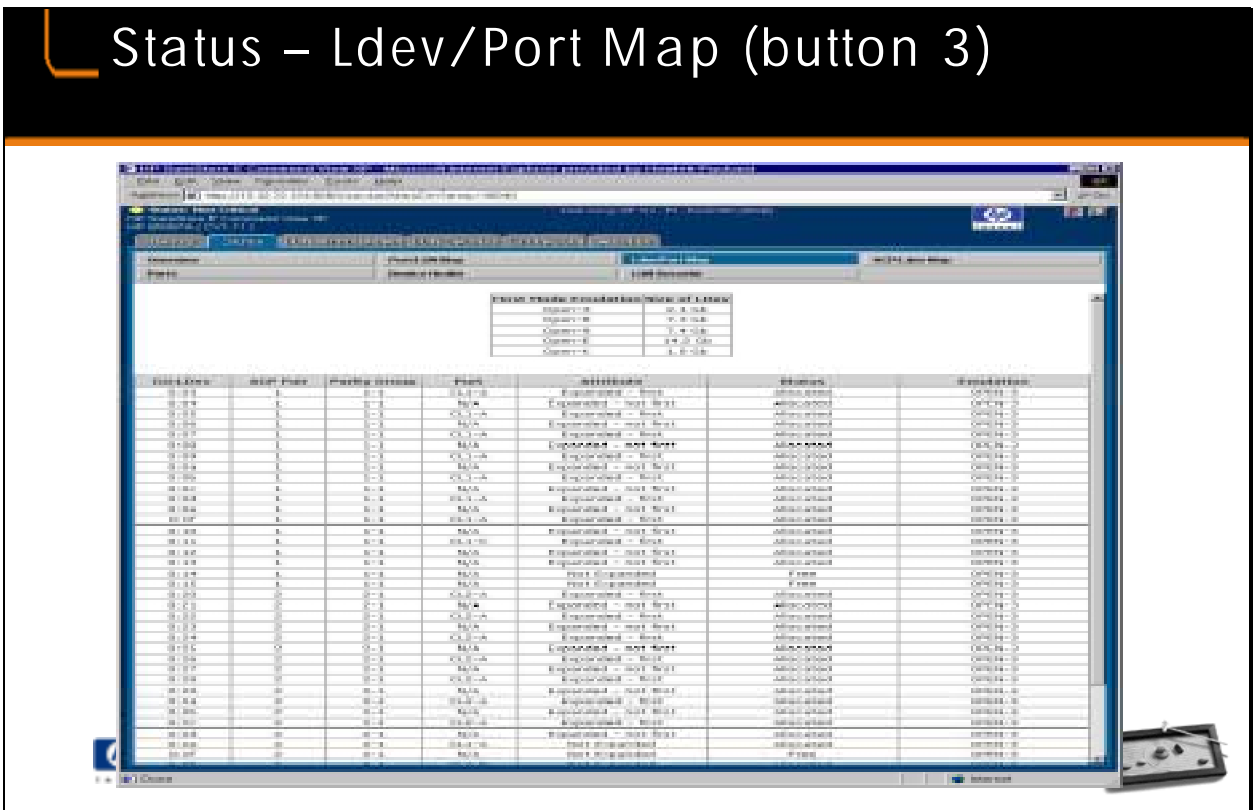
ACP - The ACP Pair number associated with the LUN.

Parity Group - The parity group associated with the LUN.

RAID Type - The RAID type implemented within the parity group.

Disk Group - A list of physical disks contained in the parity group.

Disk Type - The physical characteristics of the disk associated with the LDEV.



### Student Notes

The LDEV/Port Map displays a map relating to port and LDEV configuration.

#### Fields

CU:LDEV - The control unit (0 to F) and LDEV number (hexadecimal) with the control unit. The combination of these numbers identifies the LDEV within the disk array.

ACP Port - The configured ACP for the LDEV.

Parity Group - The parity group associated with the LDEV.

Port - The port number associated with the LDEV.

Attribute - Indicates whether the LDEV is part of an expanded LUN set:

Normal indicates that the LDEV is not part of an expanded LUN set.

Expanded First indicates that the LDEV is the top element in an expanded LUN set.

Expanded Element indicates that the LDEV is participating in an expanded LUN set.

Status - The amount of free space on the LDEV:

Free - indicates that the LDEV is not allocated.

Allocated - indicates that the LDEV is assigned to a port.

Reserved - indicates that the LDEV is not allocated, but reserved for AutoLUN, Continuous Access, and Business Copy applications to use in data migration.

Emulation - The emulation mode of the LDEV, which determines the size and operational characteristics of each logical device.

## Status – ACP/Ldev Map (button 4)

[illegible]

## Stress-Free Storage...Guaranteed

## Student Notes

The ACP/LDEV Map displays a map relating to ACP and LDEV configuration.

## Fields

ACP Pair - The ACP number associated with the LDEV.

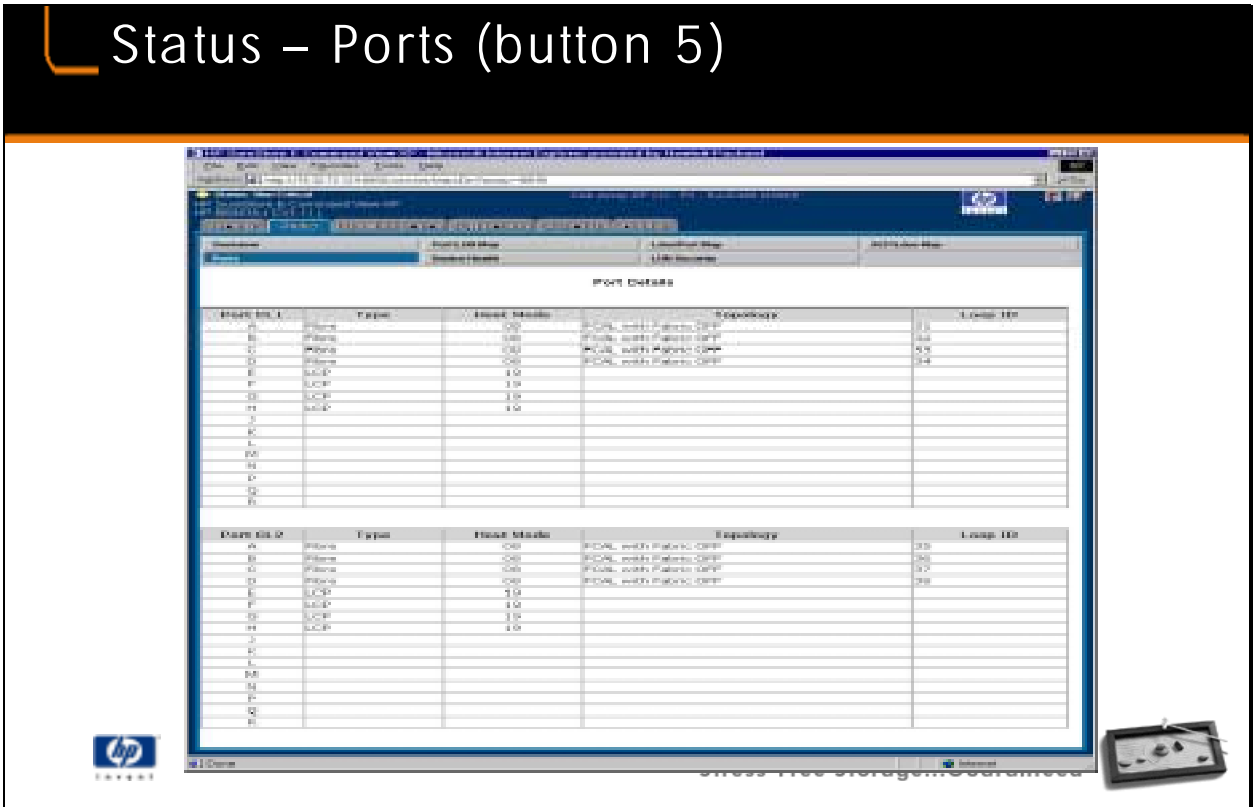
CU:LDEV - The logical device associated with this LUN. The identifier consists of the control unit (CU) and LDEV numbers.

Parity Group - The parity group associated with the ACP.

Disk Group - A list of physical disks contained in the parity group.

**Disk Type** - The physical characteristics of the disk associated with the LDEV.

Port - The port number associated with the ACP.



### Student Notes

The Ports page displays two tables. The first table lists all of the CL1 ports; the second table list all of the CL2 ports.

#### Fields

Port - The port name.

Type - The type of port, which is determined by the I/O card associated with the port:

SCSI Small Computer System Interface

Fibre Fibre Channel

MCP Mainframe port

LCP Mainframe port

Host Mode - The option that can customize each port. A numeric value and a brief description of the mode are listed.

Topology - The Fibre Channel option for the port. The field is blank for other port types.

FCAL with Fabric ON - Fibre Channel Arbitrated Loop with Fabric on

FCAL with Fabric OFF - Fibre Channel Arbitrated Loop with Fabric off

Point-to-Point with Fibre Channel Point-to-Point with Fabric on

Fabric ON

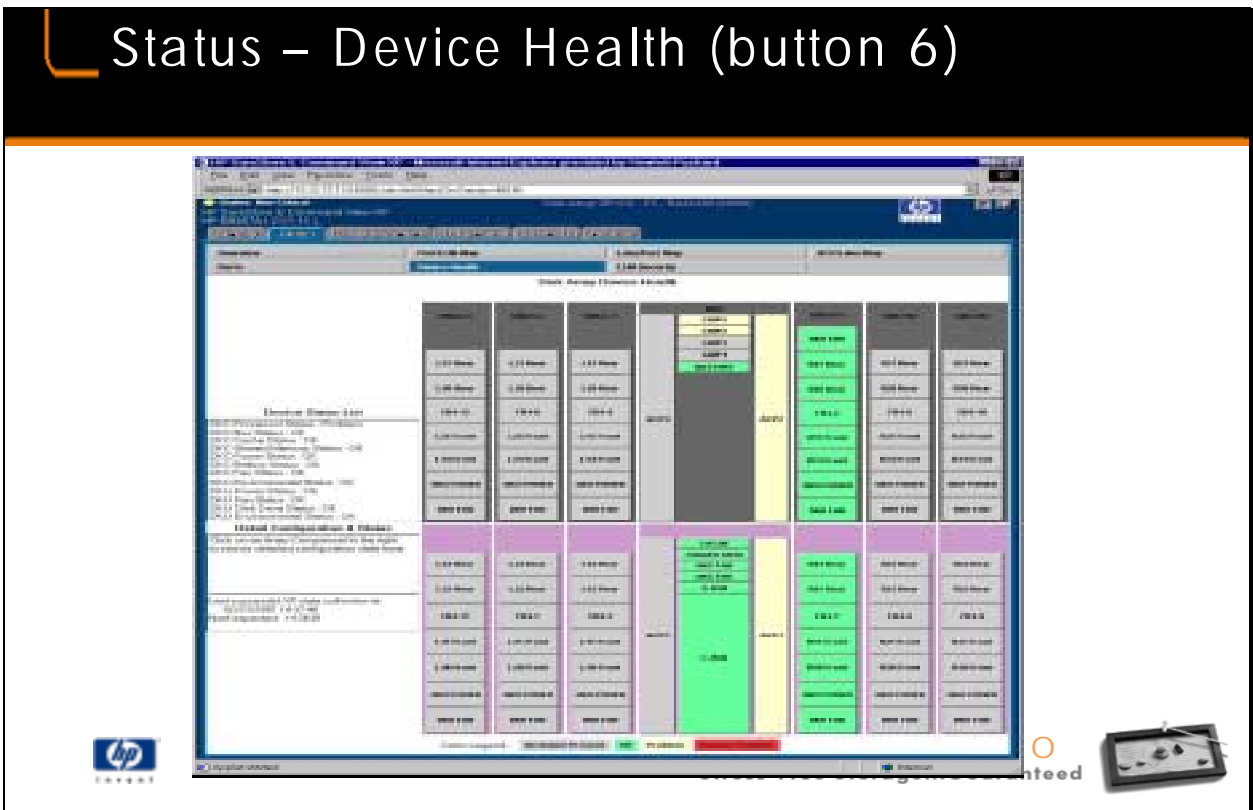
Point-to-Point with Fibre Channel Point-to-Point with Fabric off

Fabric OFF

Not Fibre - Not a Fibre Channel connection

Loop ID - The loop number.





### Student Notes

The Device Health page displays the overall status of all the components in the disk array.

#### Fields

For the DKC, the Disk Array Device Health page displays the status of the processor, bus, cache, shared memory, power, battery, fan, and environment (temperature).

For the DKU, the Disk Array Device Health page displays the status of the power, fan, disk drives, and environment (temperature).

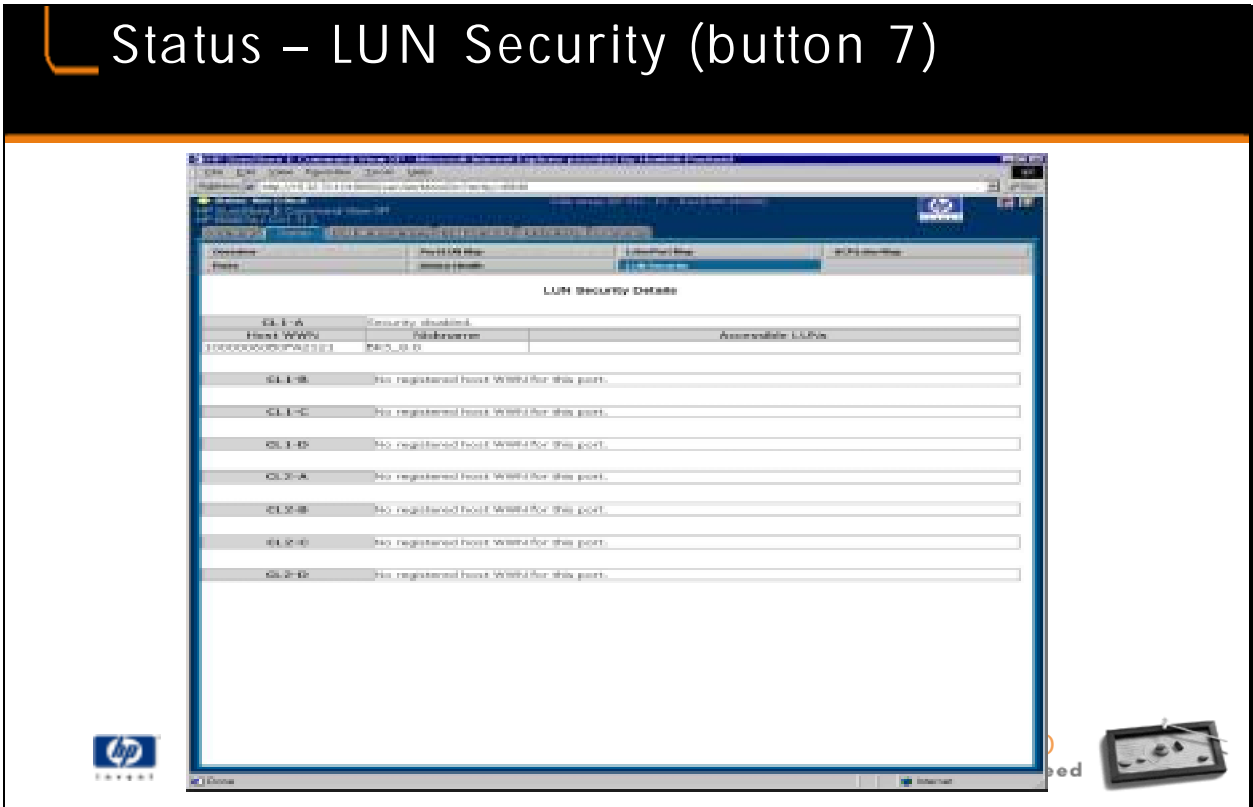
#### Legend

Gray No device present.

Green OK

Yellow Problem.

Red Serious problem.



### Student Notes

The LUN Security page displays a list of worldwide names (WWN) for each port in the array and the LUNs that are allowed to communicate with each WWN.

#### Fields

Port - Provides the port name.

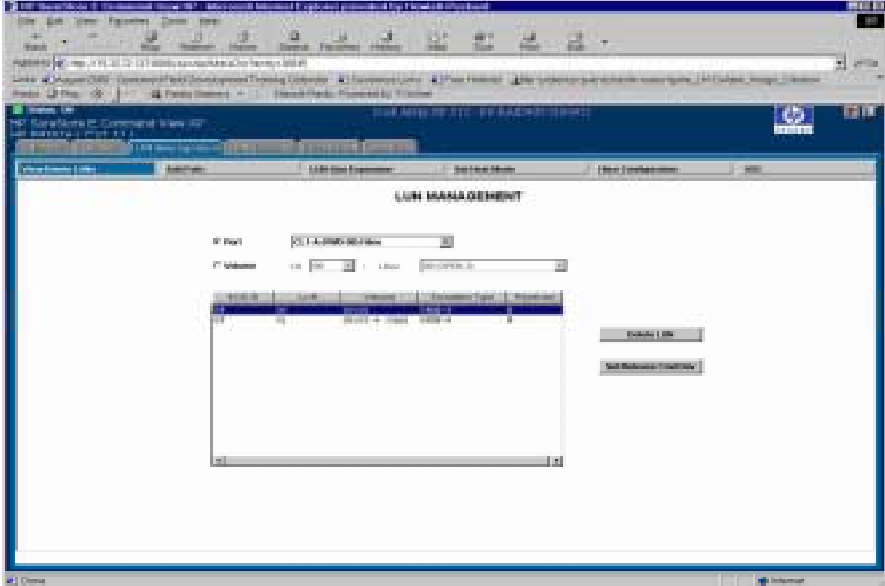
Security Status - Indicates whether security is enabled or disabled for the port.

Host WWN - The host WWN number.

Nickname - The nickname assigned to the host WWN. This field is blank if a nickname has not been defined for the WWN.

Accessible LUNs - The LUN numbers to which the WWN has access.

# LUN Management tab



➤ We'll look at the six functions in detail a little later

SSO  
Stress-Free Storage...Guaranteed

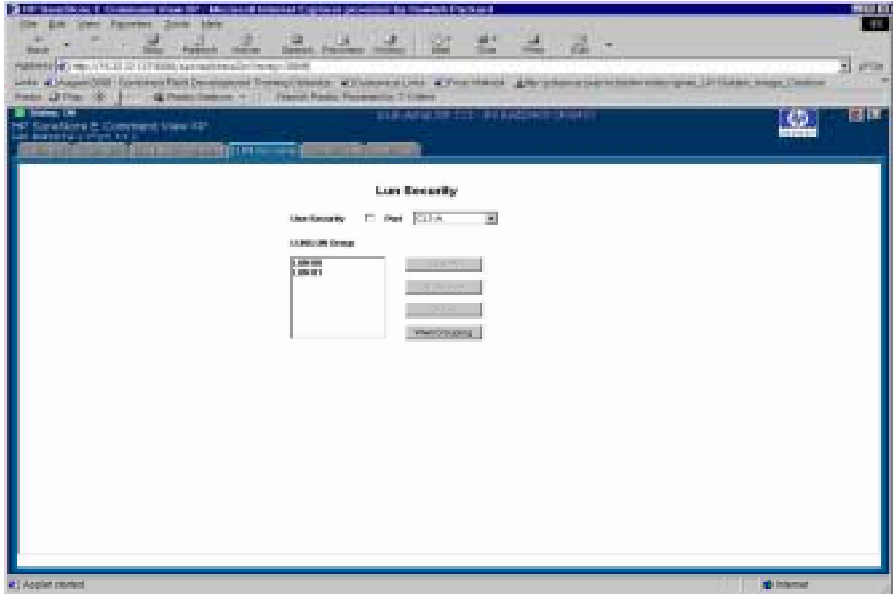
## Student Notes

The LUN Management screen has buttons for six operations:

- View/Delete LUNs
- Add Path
- LUN Size Expansion
- Set Host Mode
- Fibre Configuration
- VSC

The View/Delete LUNs screen is displayed first. Click on the secondary buttons to perform the other functions.

# LUN Security tab



➤ We'll look at LUN Security in detail a little later

SSO  
Stress-Free Storage...Guaranteed

## Student Notes

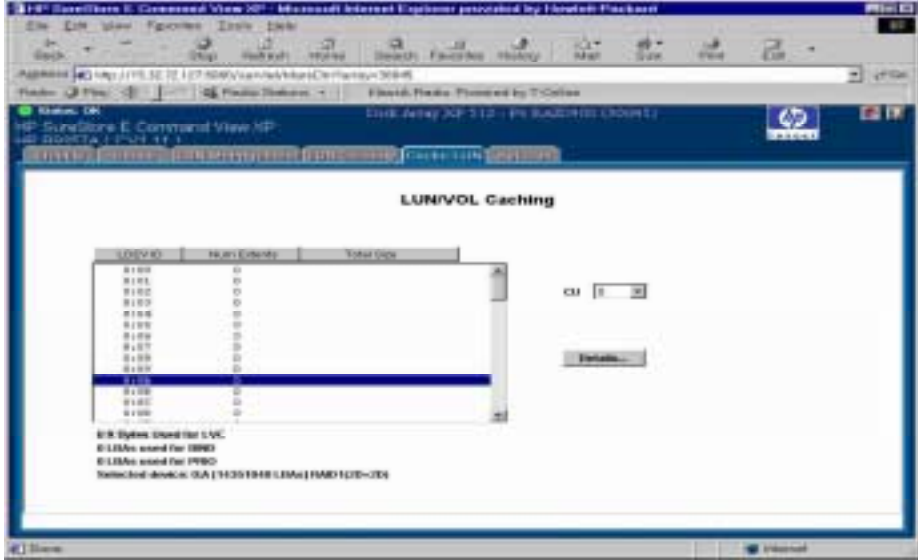
The LUN Security screen is used to manage:

- Security

- LUN groups and nicknames

- Worldwide Name (WWN) names and group names

## Cache LUN tab

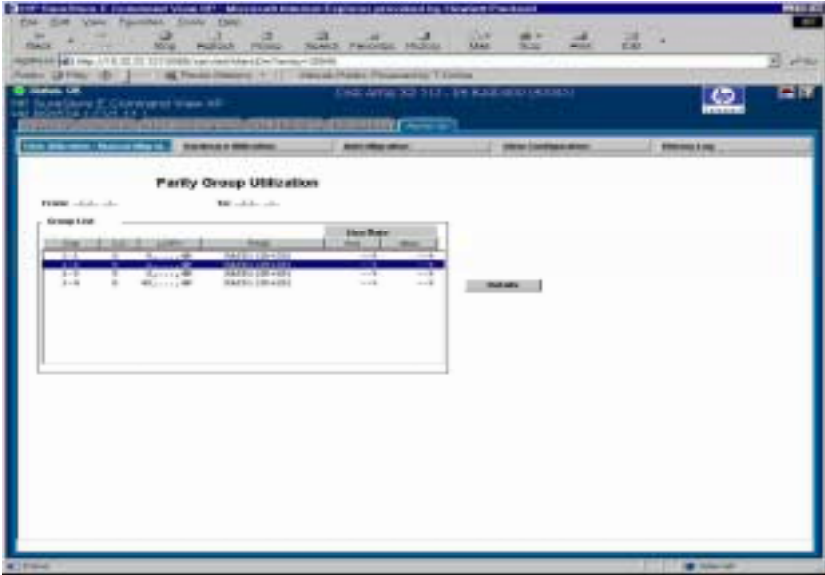


➤ We'll look at Cache LUN in detail a little later

### Student Notes

The Cache LUN function allows you to convert a portion of cache memory for use as very fast “disk space.” As we will see later, you do this by placing extents in cache memory. You can also remove extents from cache memory. The Cache LUN function is enabled when the customer purchases the optional array firmware Cache LUN function. Without Cache LUN enabled in the array firmware, this tab will not be available.

# AutoLun tab



➤ We'll look at the five functions in detail a little later

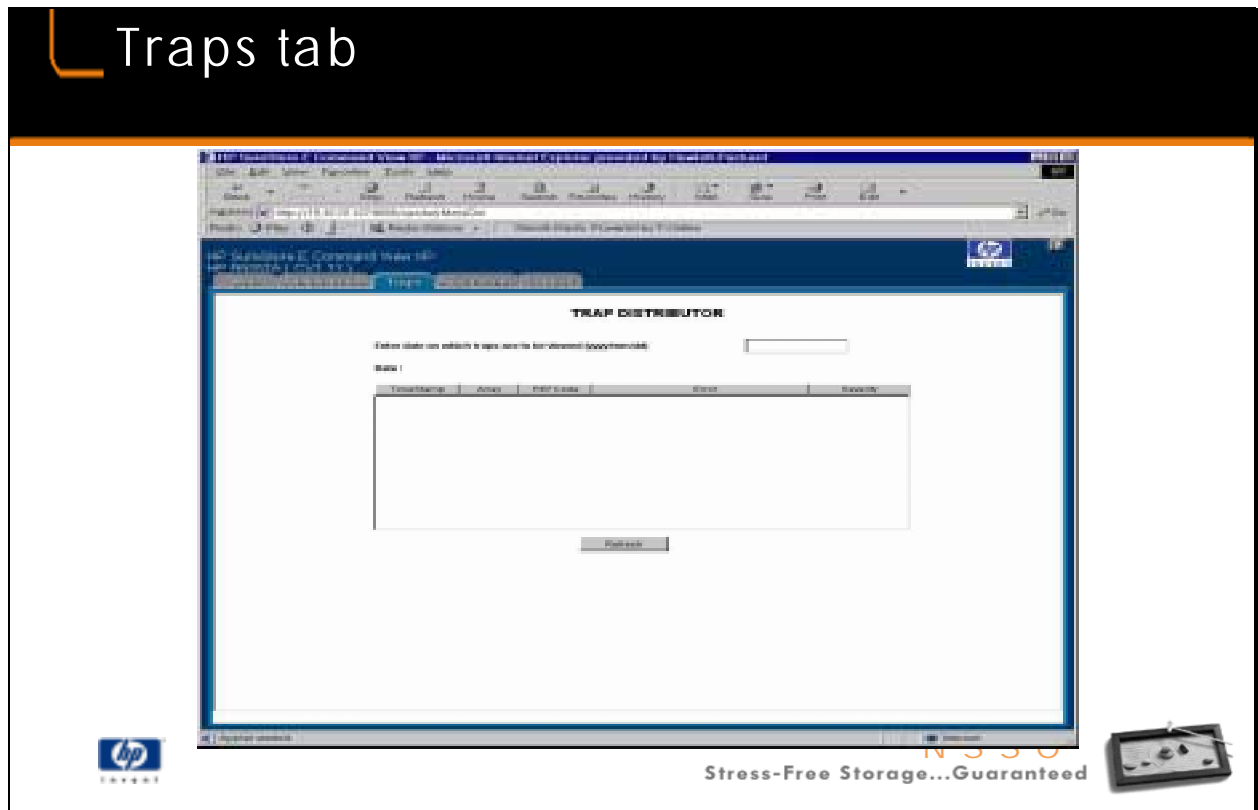
HP  
Stress-Free Storage...Guaranteed

## Student Notes

The AutoLun screen has buttons for five operations:

- Disk Utilization / Manual Migration
- Hardware Utilization
- Auto Migration
- View Configuration
- History Log

The Disk Utilization / Manual Migration screen (Parity Group Utilization screen) is displayed first. Click on the other buttons to perform the other functions. To return to the main screen, click on the Home icon in the upper right of the screen.



### Student Notes

The Traps tab lets you view errors that occurred on a disk array. Command View's Trap Distributor function forwards traps to multi-vendor management stations. The Trap Distributor function also filters out unwanted traps that are generated by the SNMP agent as a result of routine operations. Proprietary R-SIM messages from the array are translated into SNMP traps and forwarded to the Command View management station.

The Trap Distributor uses two configuration files:

xpTrapConfig.conf contains information to forward the traps to the appropriate managers. The file contains entries for the trap object identifier. Each entry specifies whether the trap is sent to the manager, the IP address, and the port number.

RefCodeToErrorMsg.conf contains the mapping information between the R-SIM reference codes and the corresponding error messages. The error messages are appended in the trap PDU along with the severity level of the traps.

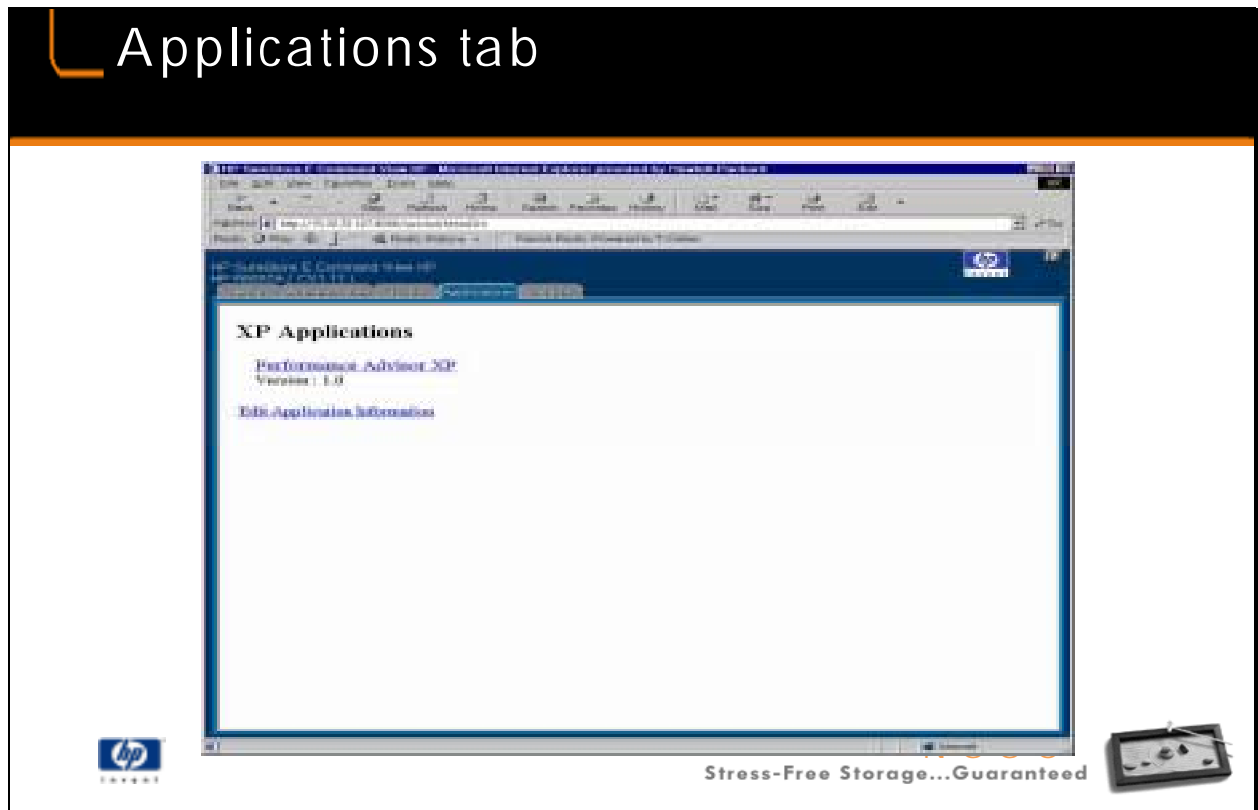
Options:

1. To view the traps for a specific date, enter the date in the format yyyy/mm/dd.
2. When traps are displayed, select Refresh to display the most current traps since the last refresh was done.

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Command View XP

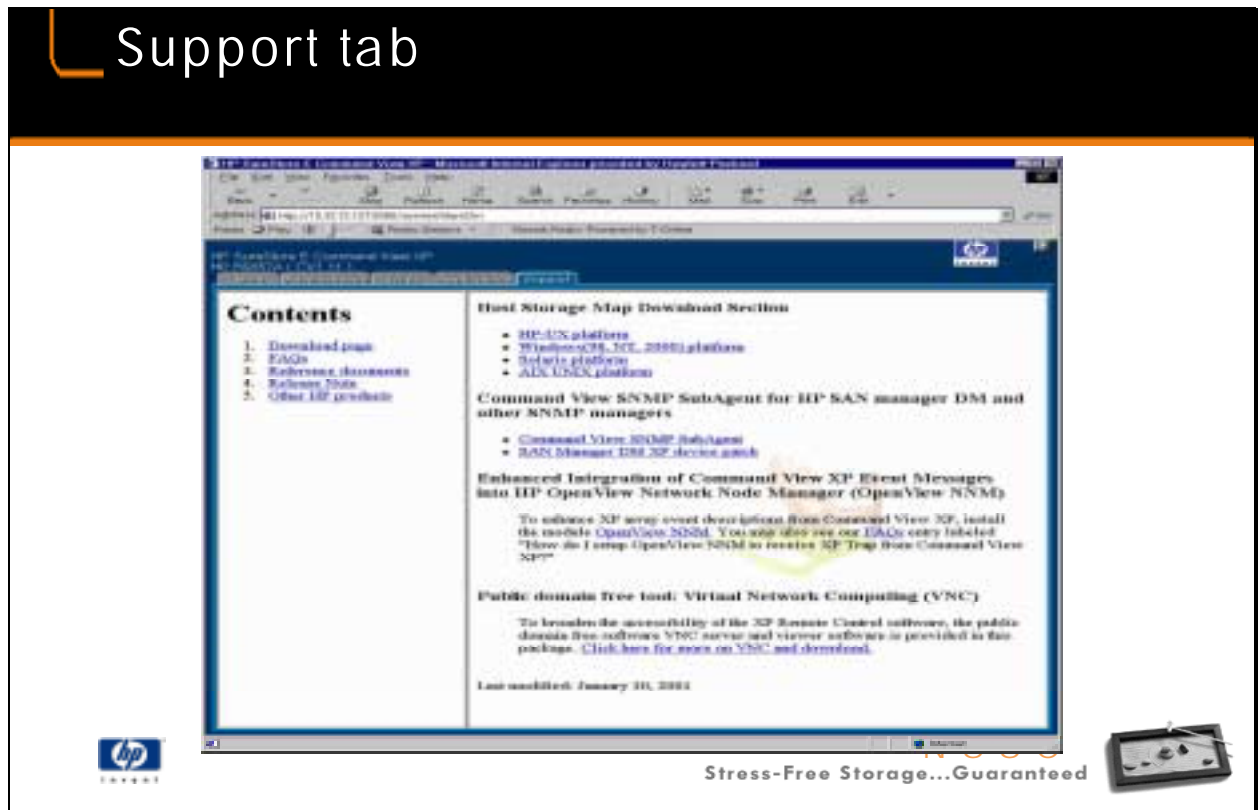
Note that we have returned to the functions on the main page. These are no longer array-specific management functions.





#### Student Notes

Use this screen to incorporate other applications into Command View (for example, integrating Performance Advisor XP into the Command View platform). If you select Edit Application Information, you can add or change application information, including application name, version, and URL.

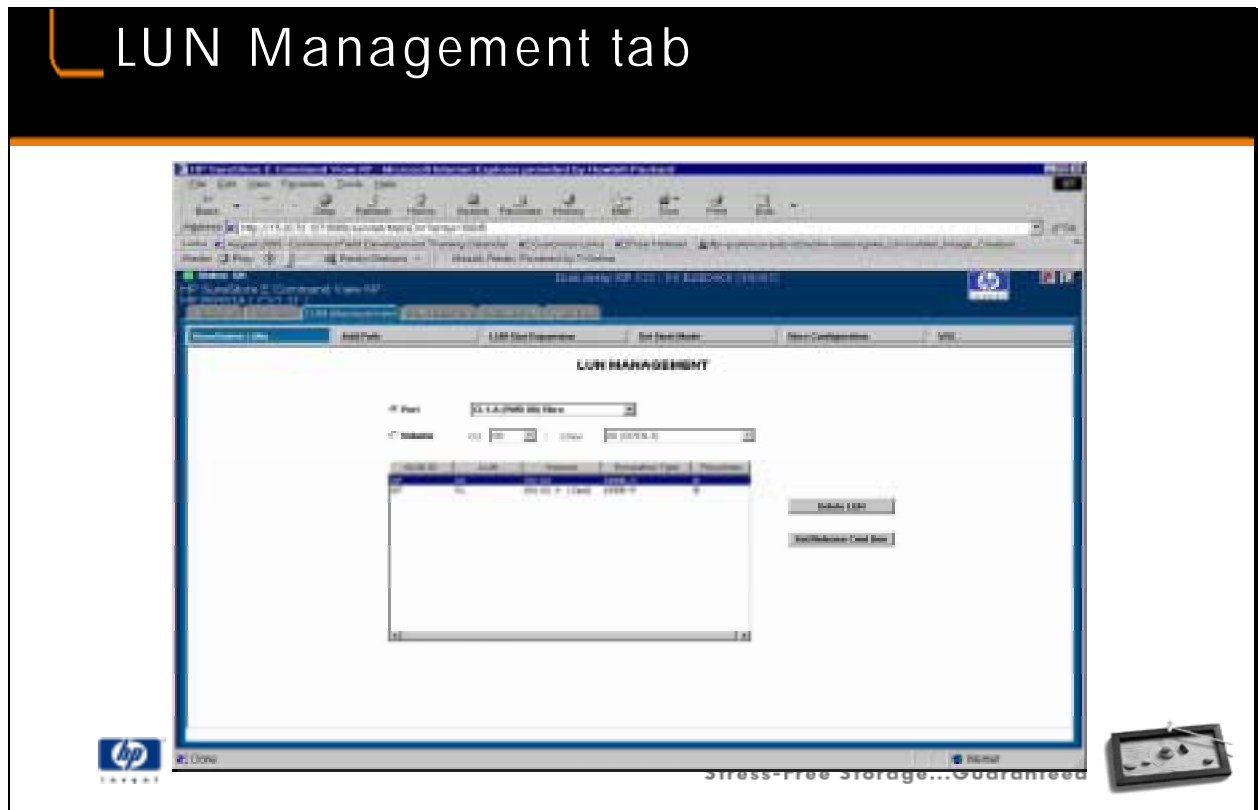


### Student Notes

The Support tab provides you with access to additional information.

You can choose links from the Contents area to access:

1. Downloads. The downloads include:
  - Host Storage Map software
  - A script to auto-integrate Traps into HP OpenView Network Node Manager
  - HP OpenView Storage Node Manager proxy agent (enables Storage Node Manager to autodiscover the Command View management station)
  - miscellaneous manual updates and white papers
  - VNC software (a free version of PC Anywhere)
2. Frequently asked questions
3. Reference documents
4. Release notes
5. Information about other HP products



### Student Notes

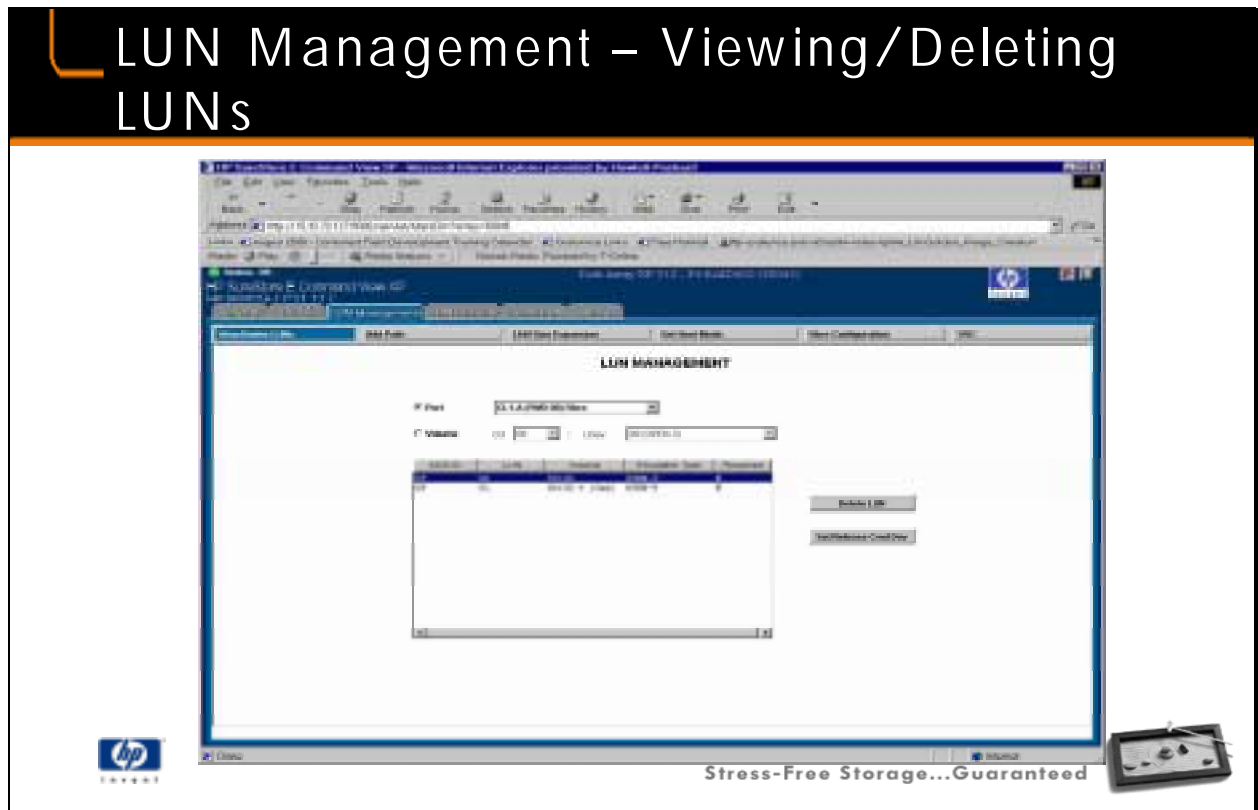
LUN Management enables you to add and delete SCSI paths (XP256 only) or Fibre Channel-to-LUN paths (XP256, XP512, and XP48) for the disk array. Each LUN can be configured for access via multiple ports and multiple target IDs to provide alternate paths for continuous data availability. You can reconfigure the SCSI or Fibre Channel paths at any time to accommodate system configuration changes (for example, adding a host) or to optimize disk array performance. Note: Disks must be formatted into LDEVs before Command View can access data.

The LUN Management screen has buttons for six operations:

- View/Delete LUNs
- Add Path
- LUN Size Expansion (LUSE)
- Set Host Mode
- Fibre Configuration
- Volume Size Configuration (VSC)

Starting LUN Management functions:

1. Start Command View and select a disk array.
2. Select the LUN Management tab. The LUN Management page appears.
3. The View/Delete LUNs screen is displayed first. Click on the secondary buttons to perform the other functions.



### Student Notes

SCSI/FC path configuration for a selected port or volume is displayed on the View/Delete LUNs page. The display shows the SCSI/FC target ID, LUN number, volume (LDEV ID), emulation type, and an indication of whether the path is reserved for AutoLUN, Business Copy or Continuous Access usage.

To view path information for a selected port:

1. Start Command View and select a disk array.
2. Select the LUN Management tab. The View/Delete LUNs screen is displayed.
3. Choose the Port radio button. The Port pulldown selection box is activated.
4. Select the desired port from the Port pull-down selection box. Path information for the selected port is displayed in the list box.

To view path information for a selected volume:

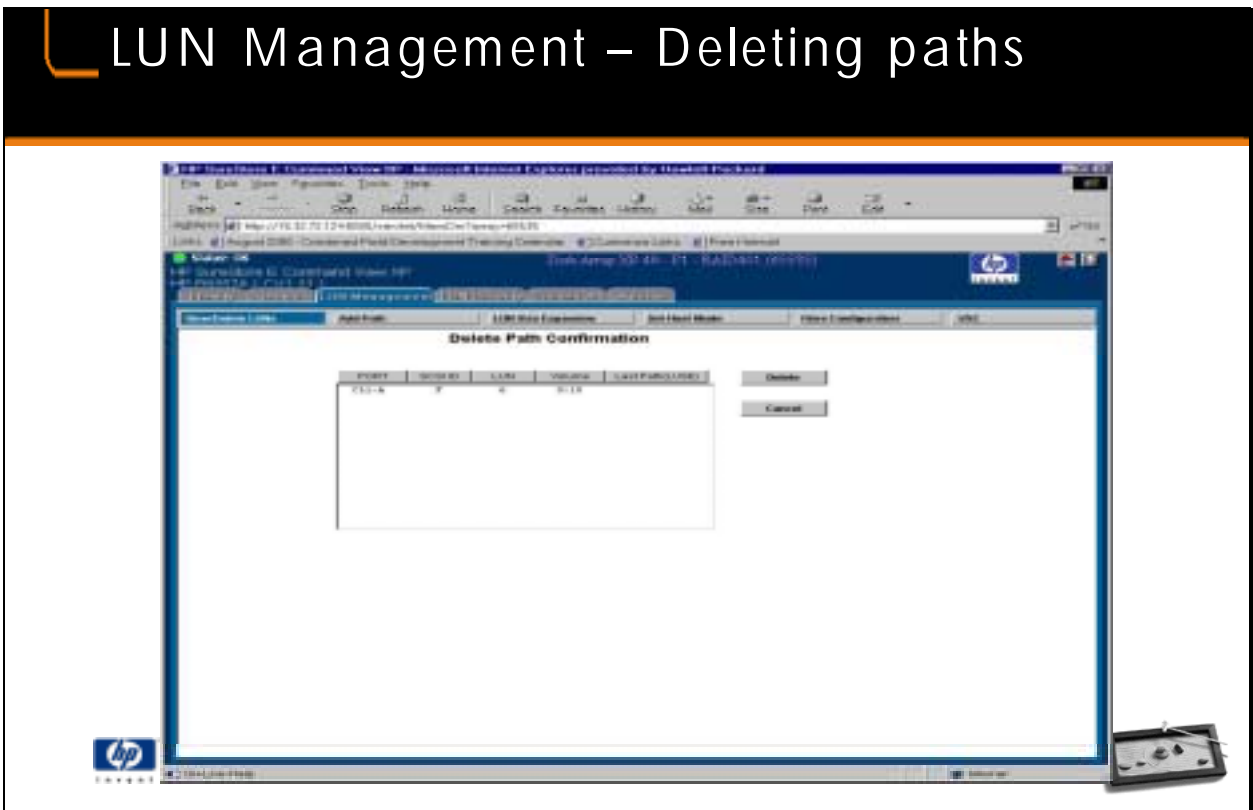
1. Start Command View and select a disk array.
2. Select the LUN Management tab. The View/Delete LUNs screen is displayed.
3. Choose the Volume radio button. The CU and LDEV pulldown selection boxes are activated.
4. Select the desired CU from the CU pulldown selection box.
5. Select the LDEV (volume) from the LDEV pulldown selection box. Path information for the selected volume is displayed in the list box

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Command View XP

**Buttons**

**Delete LUN** - If you select a LUN and click on Delete, you will go to the Delete Path Confirmation screen.

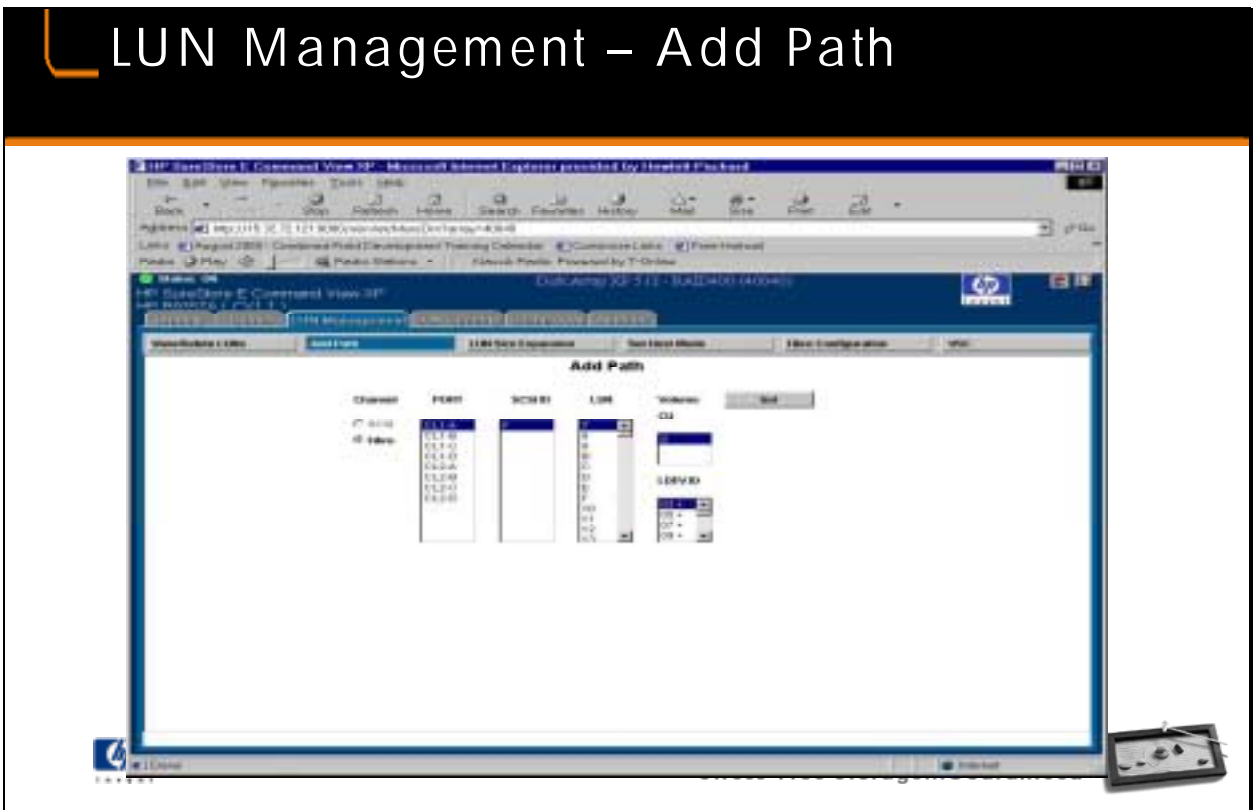
**Set/Release Cmd Dev** - Allows you to specify or release a command device on the array.



### Student Notes

To delete a path:

1. Select the paths you want to delete. You can make multiple selections clicking on the paths. Clicking a path selects it, and clicking it again deselects it.
2. Select the Delete button. The selected paths are displayed in the Delete Path Confirmation page.
3. Verify that the information presented on the Delete Path Confirmation page is correct. If the information is not correct, choose Cancel to return to the LUN Management page.
4. Select the Delete button. The selected paths are deleted.



### Student Notes

Each LDEV can be mapped from multiple ports to provide alternate path high availability. Before assigning a path to an LDEV, make sure that the LDEV is not an Auto LUN reserved device. The Add Path page does not display reserved volumes. If the volume has been reserved, it is write-protected and any writes to that volume will fail. Use Auto LUN to verify the reserve status of the LDEVs.

To add paths:

1. Start Command View and select a disk array.
2. Select the LUN Management tab. The View/Delete LUNs screen is displayed.
3. Select the Add Path button. The Add Path screen is displayed.

Channel Radio buttons for selecting the type of path you are adding, SCSI or Fibre Channel.

PORT - Displays all of the assigned SCSI or Fibre Channel ports.

SCSI ID - If the SCSI radio button is selected, displays all of the allowed values for SCSI IDs. If Fibre is selected, the SCSI ID box displays only one (F) SCSI ID.

LUN - If the SCSI radio button is selected, displays all of the allowed values for LUN number. If the Fibre radio button is selected, displays LUN values from 0 through 119 (0 to 77 in hexadecimal).

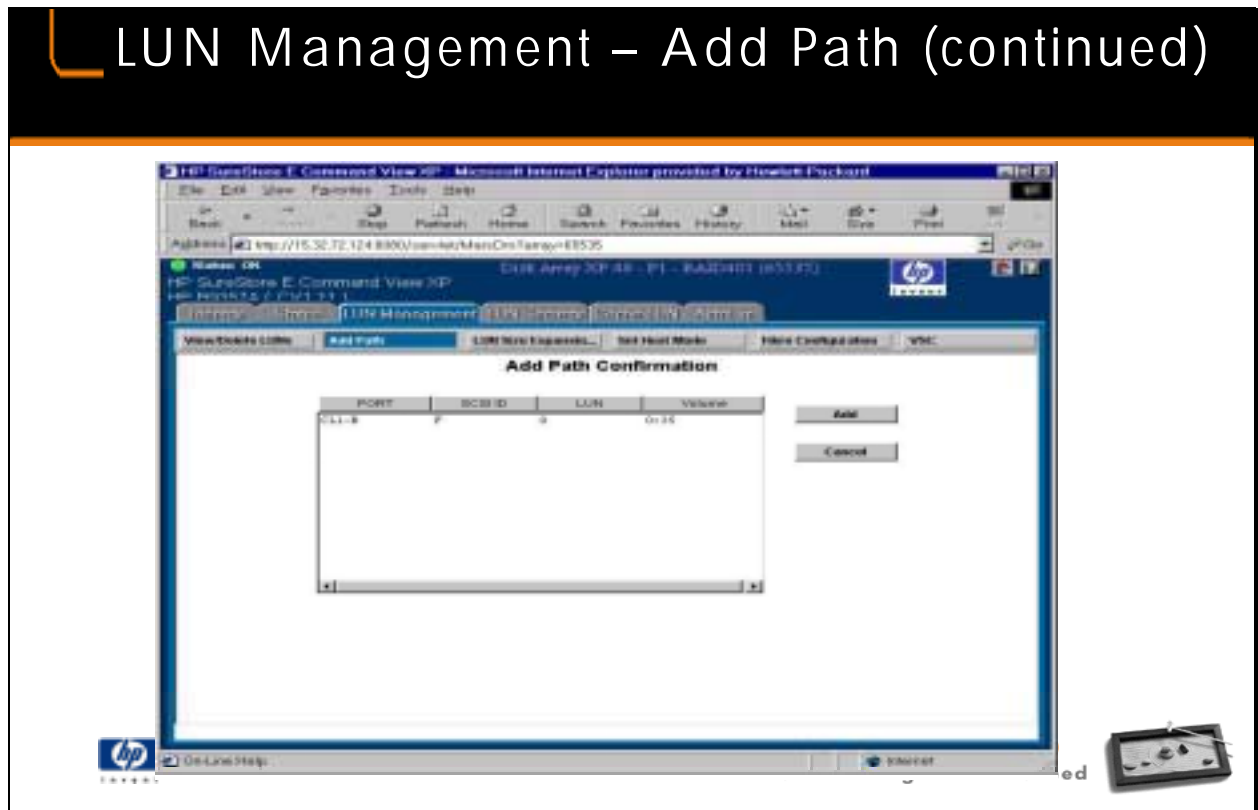
CU - Displays all of the CU numbers currently available in the disk array. The first CU number is selected by default.

LDEV ID - Displays all of the LDEVs for the selected CU.

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Command View XP

A single plus sign (+) indicates that one path is defined for that LDEV. Two plus signs (++) indicate that multiple paths are defined for that LDEV.



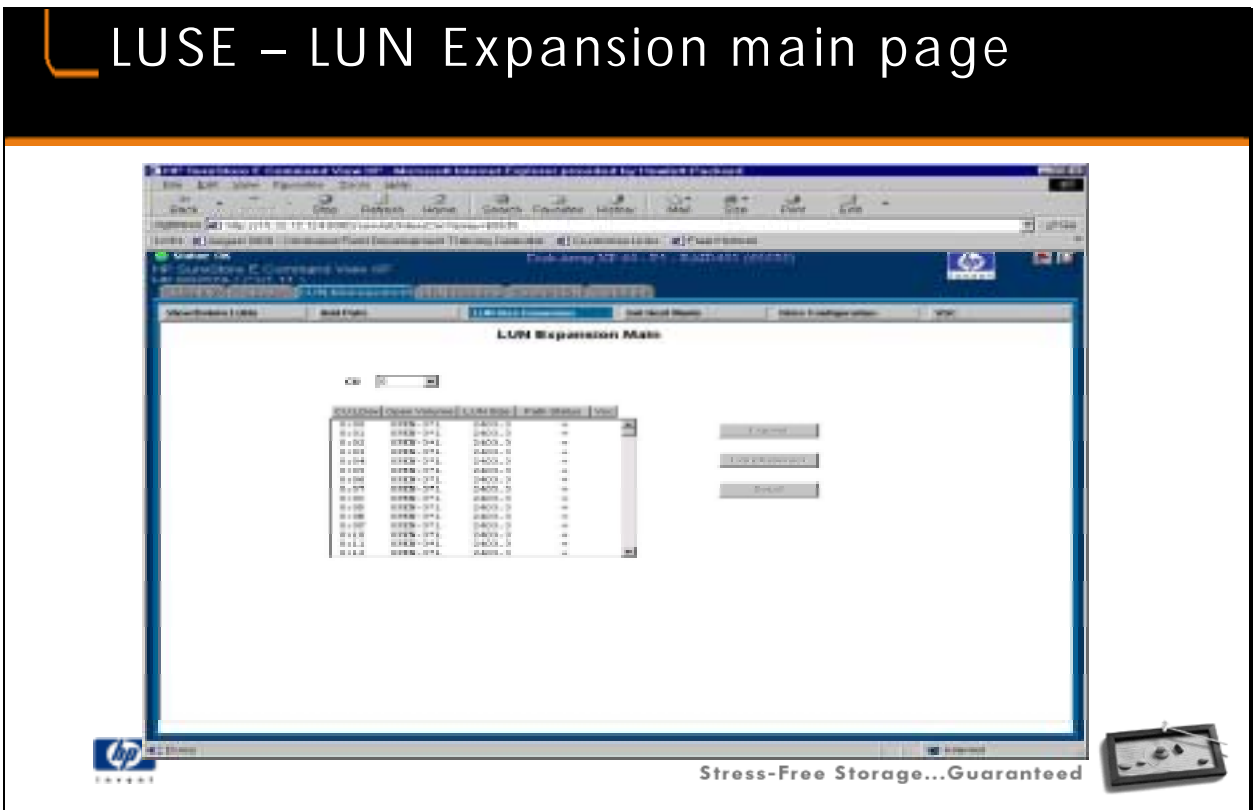


### Student Notes

To add SCSI/FC paths (continued):

4. Choose the SCSI or Fibre radio button under Channel in the Add Path page.
5. Select the SCSI ports, target SCSI IDs, LUN IDs you want to assign, and then select the corresponding CU and LDEV IDs. To calculate the number of LUNs to select, multiply the number of selected SCSI TIDs and the number of selected LUN IDs. For example, if you select four SCSI IDs and two LUN IDs, you must select eight LUNs ( $4 \times 2 = 8$ ). Paths can be defined individually (individual mode), in a group (bulk mode), or randomly (random mode).  
(Individual Mode) - Select PORT, SCSI ID, and LUN one by one for each LDEV ID.  
(Bulk Mode) - Select the first LDEV ID and LUN Manager assigns the remaining paths sequentially, or select the number of LDEV IDs you want to assign. A bulk mode operation fails if the number of LUNs defined exceeds the number of LDEVs available.  
(Random Mode) - Select the LDEV IDs, and the random selection will be assigned numerically.  
For example, if you select LDEV IDs 07, 00, 0F, 01, and 0A, the LDEVs are assigned in the following order: 00, 01, 07, 0A, and 0F, assigned sequentially.

6. Choose the Set button to confirm and save the selected paths. The Add Path Confirmation page appears displaying all of the paths to be added.
7. Verify that the information presented on the Add Path Confirmation page is correct. If the information is not correct, choose Cancel to return to the Add Path page.
8. Select the paths you want to add and choose the Add button to create the selected paths in the disk array.



## Student Notes

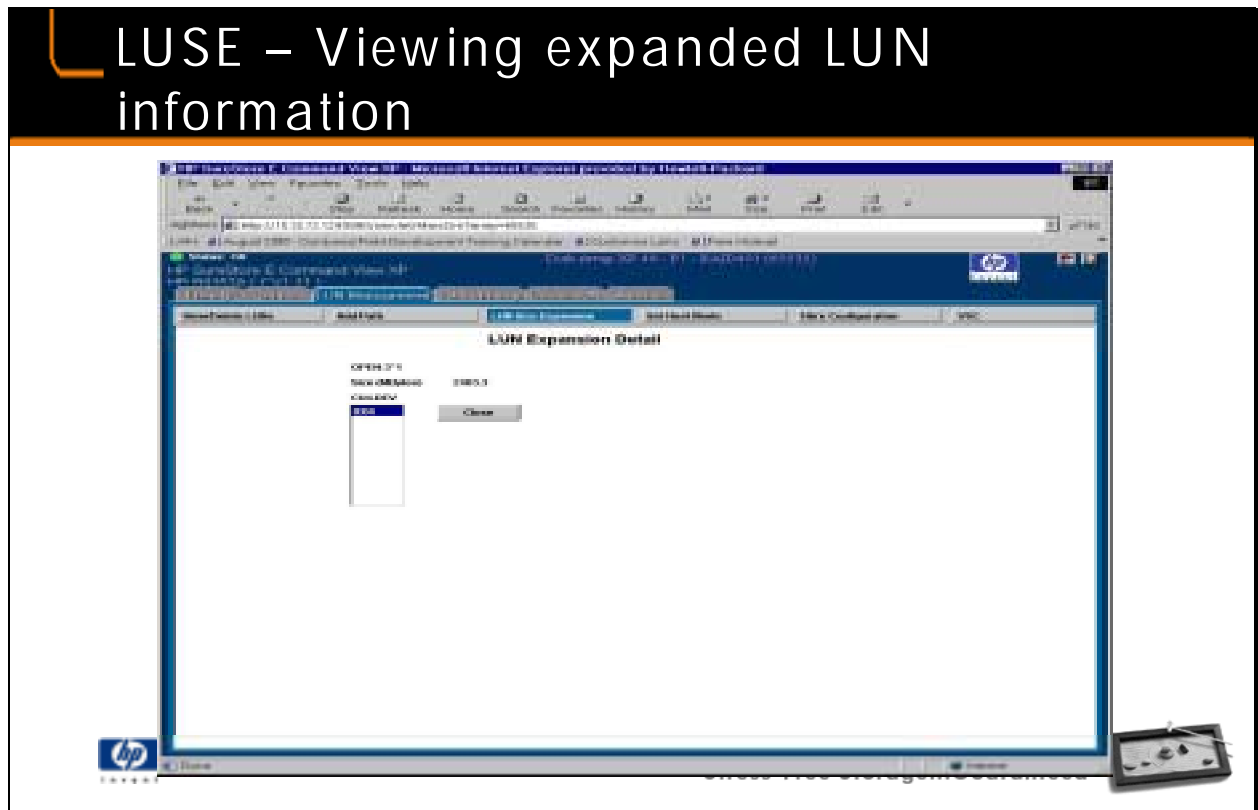
Starting LUN Size Expansion functions:

1. Start Command View and select a disk array.
2. Select the LUN Management tab. The LUN Management page appears.
3. Select the LUN Size Expansion button on the LUN Management page.

The LUN Expansion Main page appears.

The LUN Expansion Main page is used to access the LUN expansion functions.

By default, the Control Unit (CU) box displays 0 and the scrollable list box displays all of the volumes for CU 0. The list box displays all of the volumes corresponding to the selected CU. For each volume, the control unit number, LDEV, open volume type, size, path status, and VSC (CVS) are displayed. If a plus sign (+) appears under the path listing, this indicates that the path has already been set. If there is no +, the path has not been set. Initially, the buttons on the page are disabled. The Expand button is enabled if the selected volume is not already expanded, has no paths allocated to it and has at least one unassigned, numerically sequential LDEV following it. The Exp (Disperse) button is enabled if the selected volume is not already expanded, has no paths allocated to it and there is at least one other non-sequential LDEV which also does not have a path defined. The Detail button is enabled to display information about expanded LUNs.



### Student Notes

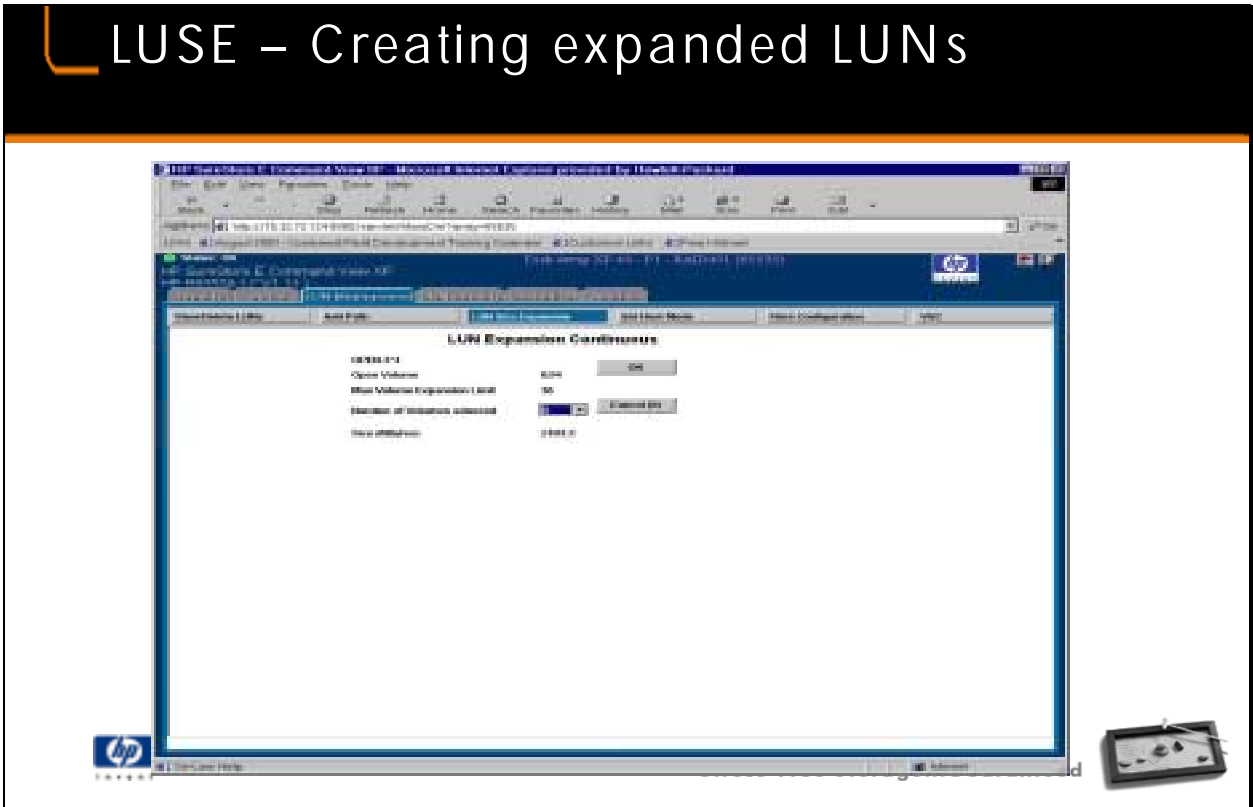
The Detail button on the LU Expansion Main page displays specific information regarding the selected LUN. The list box under CU:LDEV displays the constituent volumes of an expanded LUN. If the LUN has not yet been expanded, there will be only one volume listed.

To view Expanded LUN information:

1. Start Command View and select a disk array.
2. Select the LUN Management tab. The LUN Management page appears.
3. Select the LUN Size Expansion button on the LUN Management page.

The LUN Expansion Main page appears.

4. Select the control unit you want to view from the CU dropdown selection box. The LUNs for the selected control unit are displayed in the list box.
5. Select the LUN you want to view from the list box.
6. Select the Detail button. The LU Expansion Detail page appears.
7. Select the Close button to return to the LUN Expansion Main page.



### Student Notes

LUNs with paths already assigned cannot be expanded. Assigned paths must be deleted before expanding. Available LUNs are displayed as OPEN-X\*1, indicating that only one LUN is present. Volumes that have already been expanded are displayed as OPEN-X\*n, where n is the number of LUNs present. Open Volume displays the CU/LDEV that will be the first LUN after combining. Max Volume Expansion Limit shows number of available LUNs to combine. Number of Volumes selected selects the number of volumes. Size (MBytes) Displays the capacity of the Open volume.

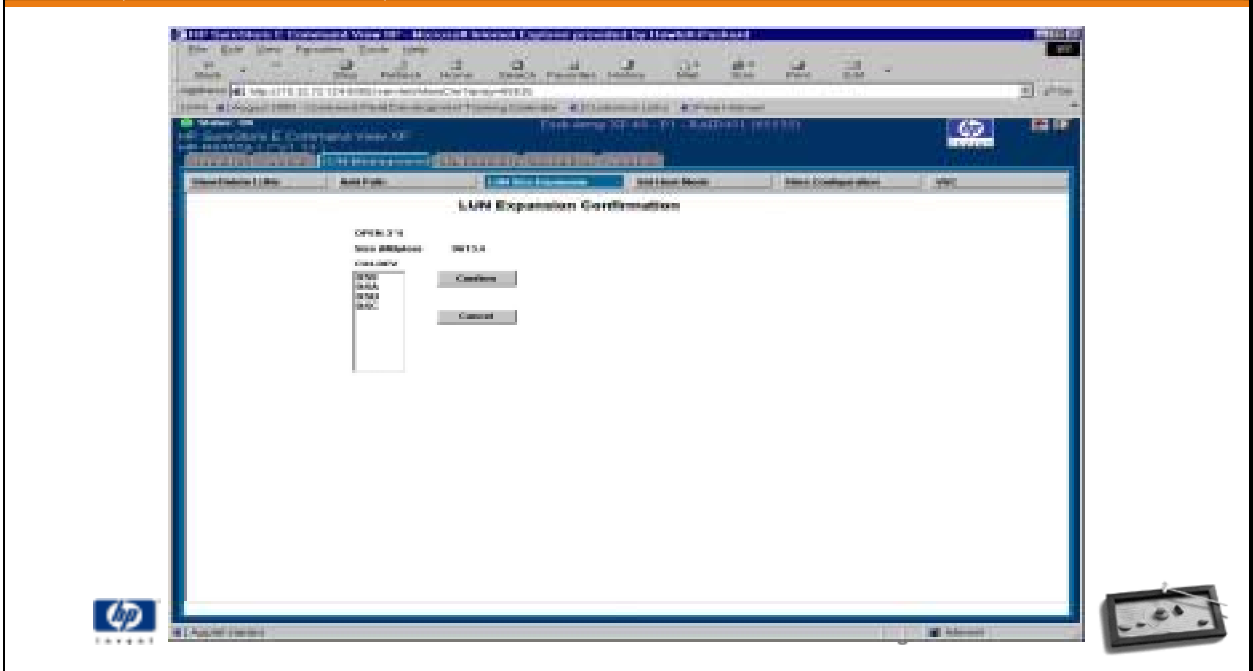
To create expanded LUNs using contiguous LDEVs:

1. Start Command View and select a disk array.
2. Select the LUN Management tab. The LUN Management page appears.
3. Select the LUN Size Expansion button on the LUN Management page. The LUN Expansion Main page appears.
4. Select the control unit you want to view from the CU dropdown selection box. The LUNs for the selected control unit are displayed. If an asterisk (\*) appears under the path listing, this indicates that the path has already been set. If there is no asterisk (\*), the path has not been set.
5. Select the LDEVs you want to expand. The selected volume must not have a path defined.
6. Choose the Expand button. The LUN Expansion Continuous page appears.

7. Select the desired size of the expanded LUN.
8. Choose OK to expand the LUN. The LUN Expansion Confirmation page appears.

Note: You may want to check the Glossary in the Command View online help system for precise definitions of terms such as LUN, LDEV, CU, Parity Group, etc.

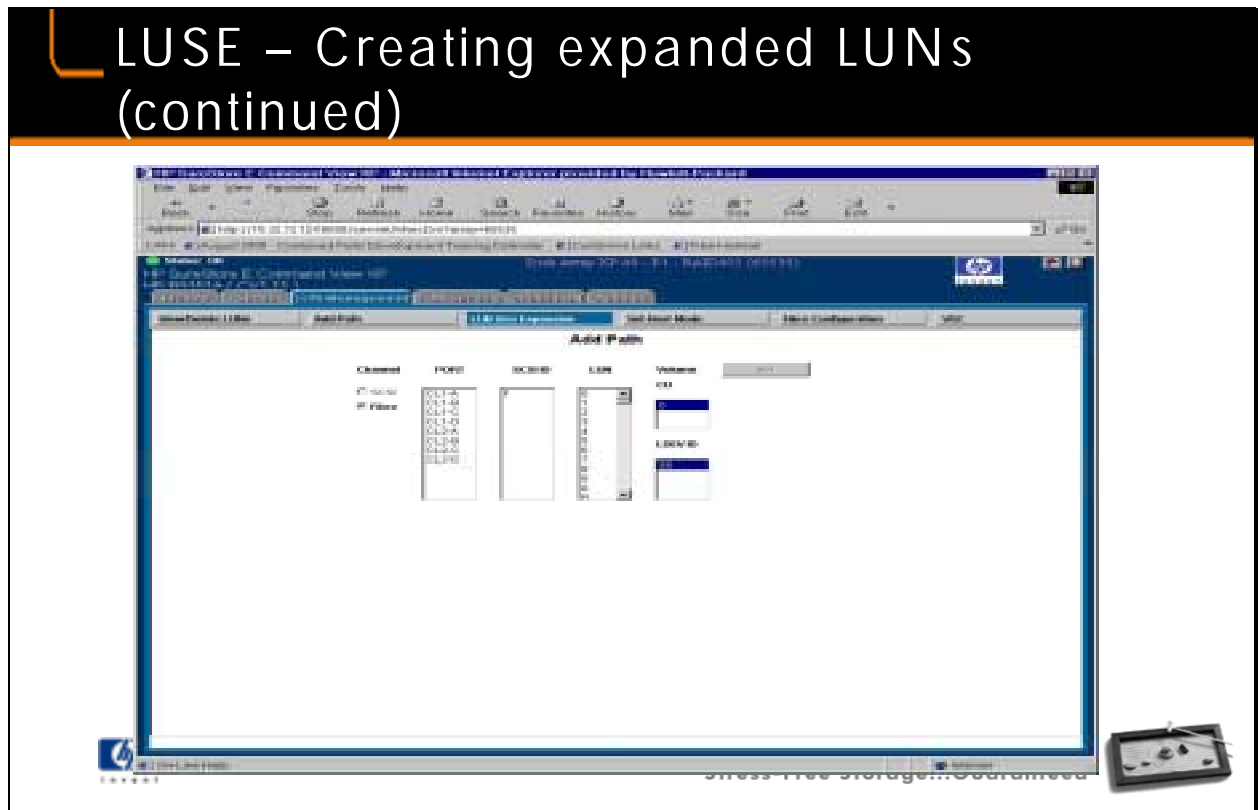
## LUSE – Creating expanded LUNs (continued)



## Student Notes

To create expanded LUNs using contiguous LDEVs (continued):

1. Choose the Confirm button. There will be a warning message.
2. Select OK to confirm. The Add Path screen will then appear. Selecting the Cancel button returns you to the LUN Expansion Continuous page.



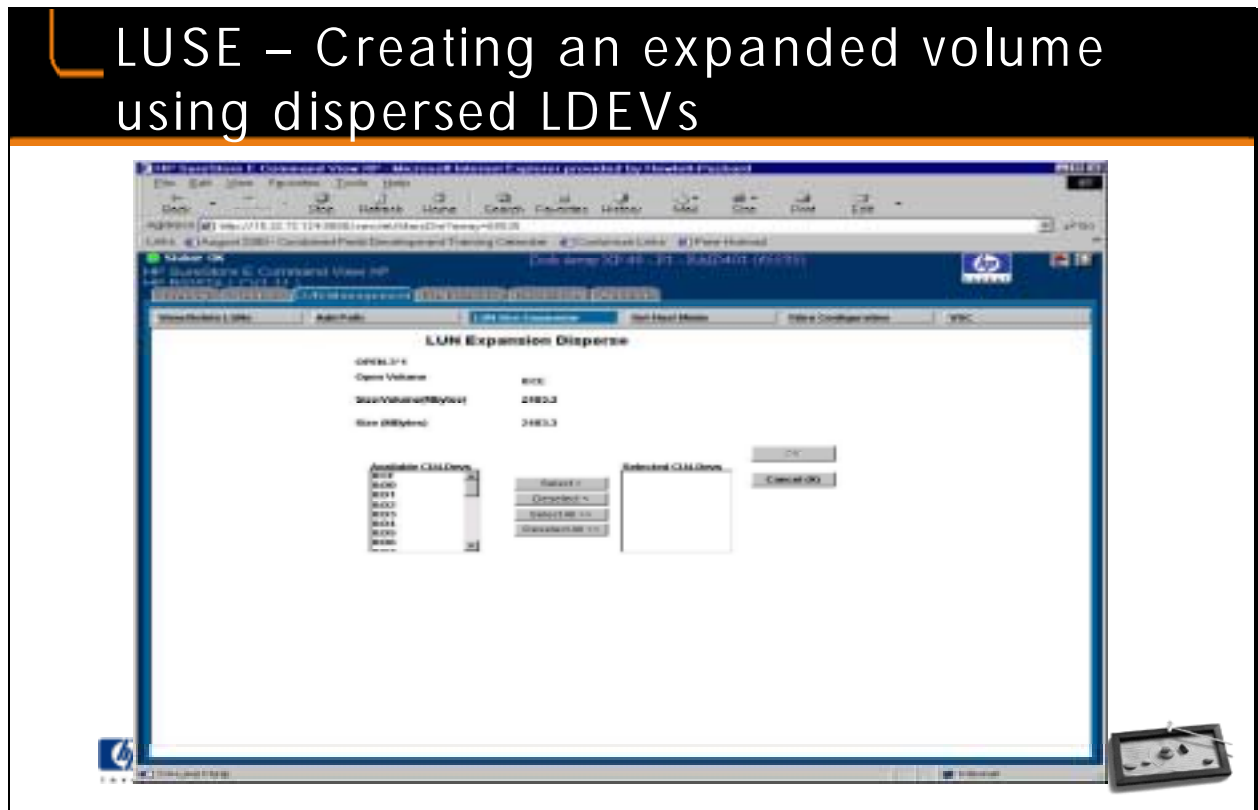
### Student Notes

Use the Add Path page to associate a path with the expanded LUN. If you do not add a path for the expanded LUN, the expansion does not occur. Adding paths was covered earlier in the presentation.

To create expanded LUNs using contiguous LDEVs (continued):

1. Select the SCSI or Fibre radio button.
2. Select a port, SCSI ID, LUN, Volume, and LDEV ID.
3. Select the Set button.



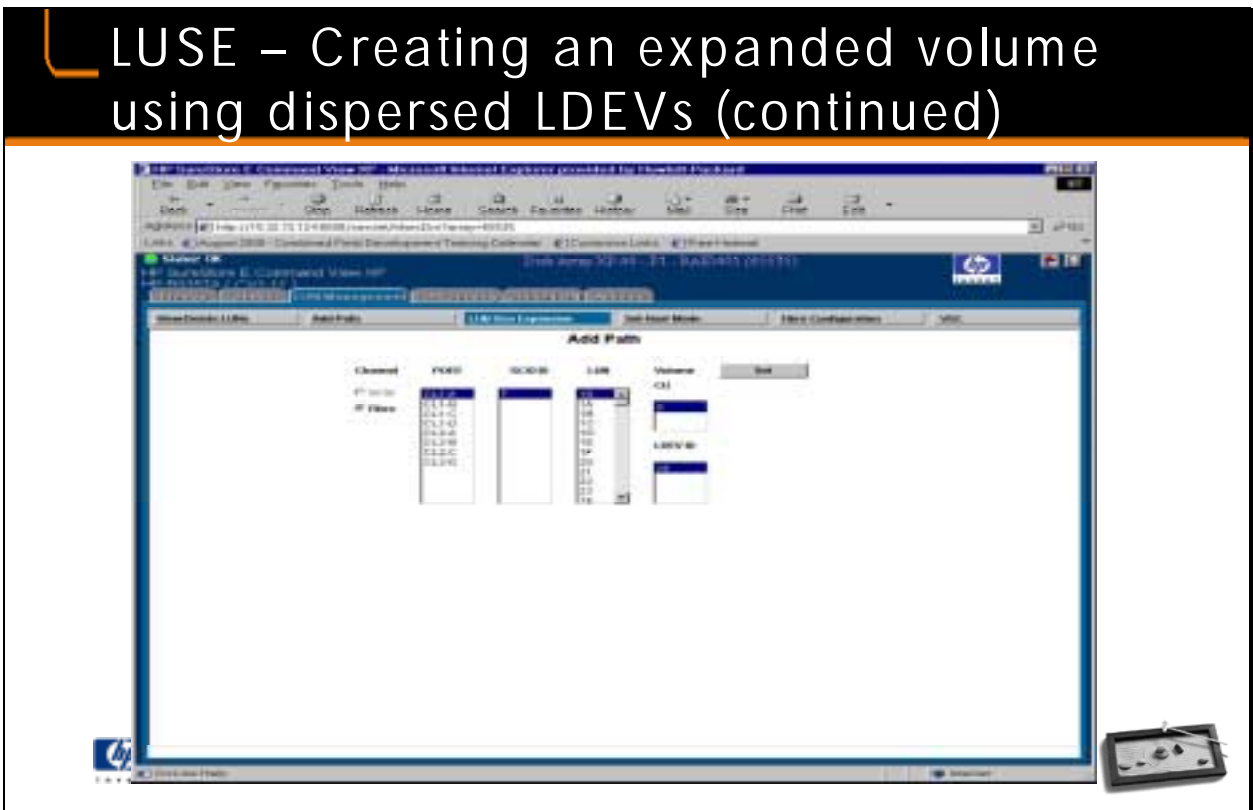


### Student Notes

Use the LUN Expansion Disperse page to create an expanded volume using dispersed LDEVs.

To create an expanded volume using dispersed LDEVs:

1. Start Command View and select a disk array.
2. Select the LUN Management tab. The LUN Management page appears.
3. Select the LUN Size Expansion button on the LUN Management page. The LUN Expansion Main page appears.
4. Select the control unit you want to view from the CU dropdown selection box. The LUNs for the selected control unit are displayed.
5. In the list box, select the first LUN in the series of LUNs you want to expand.
6. Select the Exp(Disperse) button. The selected volume must not have a defined path. The LUN Expansion Disperse page appears.
7. In the Available CU:LDEV scrollable list box, select the volume that is eligible for LUN expansion, and choose the Select button. To select all of the volumes listed, choose the Select All button. The selected volumes are displayed in the Selected CU:LDEV scrollable list box.
8. Choose OK to expand the LUN. The LUN Expansion Confirmation page opens. Respond to the confirmation prompt.
9. Choose the Confirm button. The Add Path page appears.



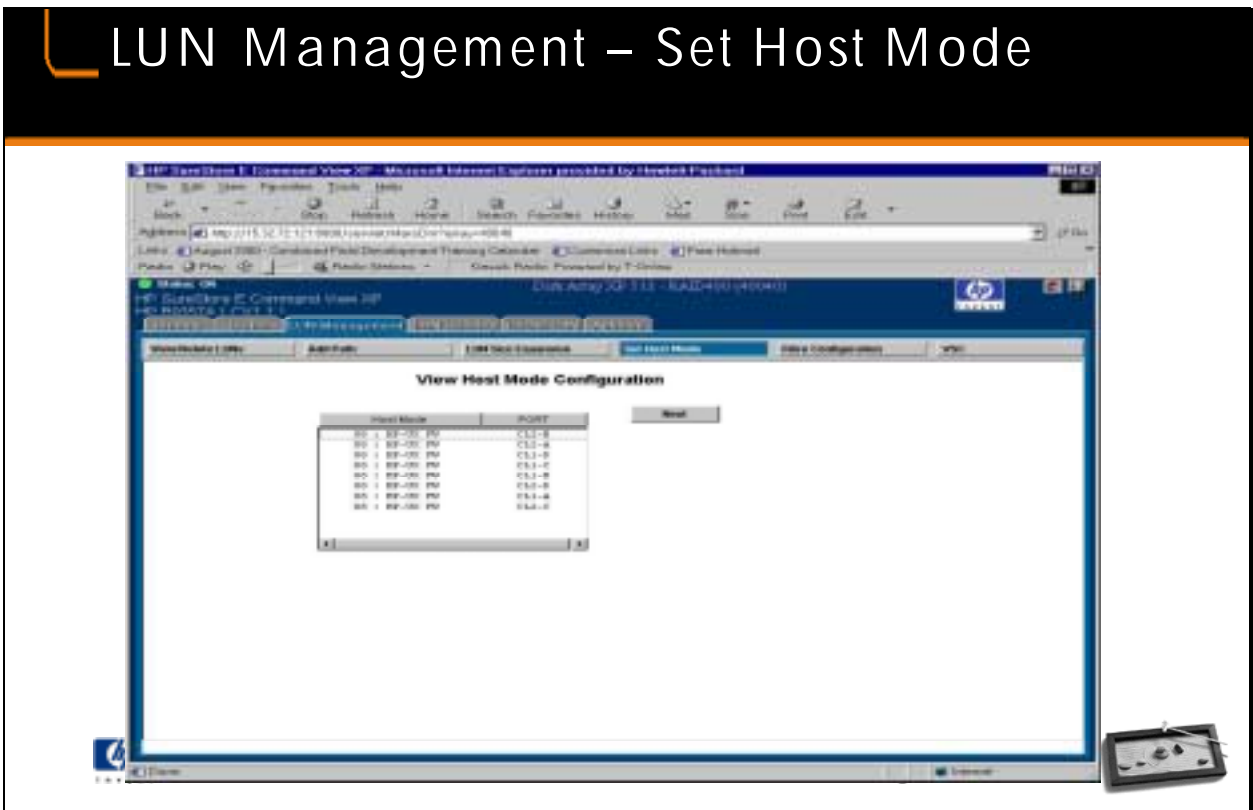
### Student Notes

Use the Add Path page to associate a path with the expanded LUN. If you do not add a path for the expanded LUN, the expansion does not occur.

Adding paths was covered earlier in the presentation.

To create an expanded volume using dispersed LDEVs (continued):

1. Select the SCSI or Fibre radio button.
2. Select a port, SCSI ID, LUN, Volume, and LDEV ID.
3. Select the Set button.



### Student Notes

Set Host Mode allows you to view and set host modes for SCSI and Fibre Channel ports on the disk array.

Refer to the user's guide for a complete list of host modes.

To view host mode configurations:

1. Start Command View and select a disk array.
2. Select the LUN Management tab. The LUN Management page appears.
3. Select the Set Host Mode button on the LUN Management page. The Set Host Mode page appears. When the View Host Mode Configuration page appears, it shows host modes and ports.

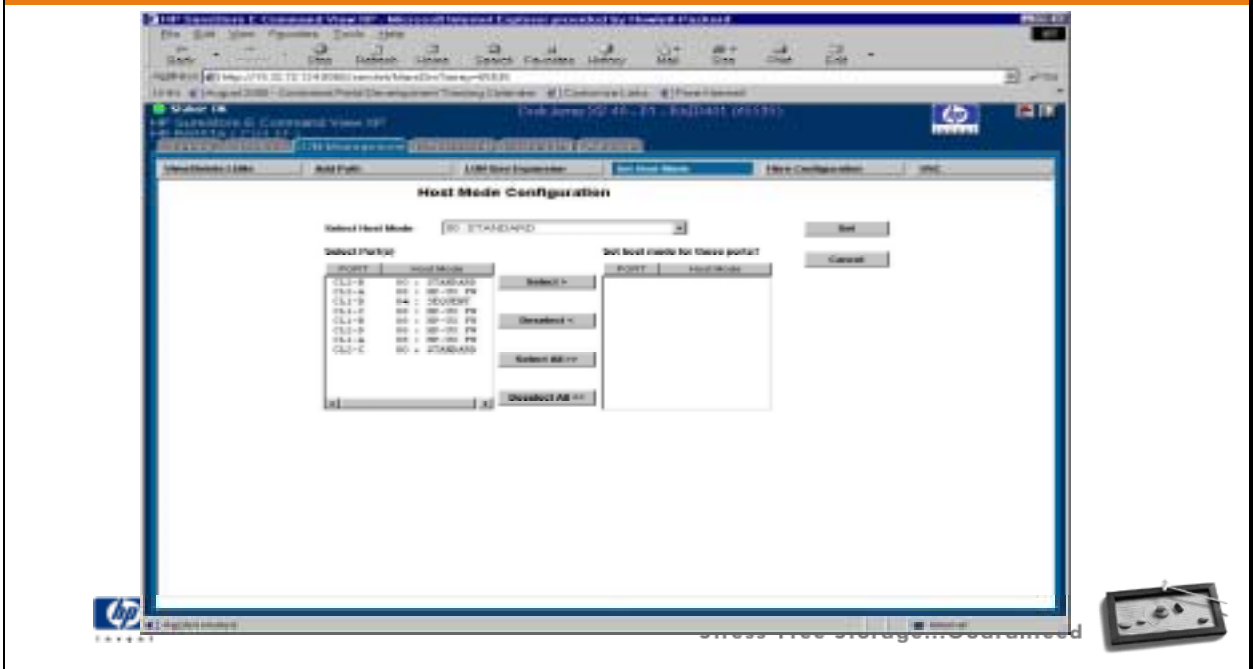
To set or change the host mode configurations:

Caution Back up the data on the port of the host mode you want to change before changing the host mode. The Set Host Mode function is not destructive; however, data integrity cannot be guaranteed without a backup.

1. Choose the Next button on the Host Mode Configuration page appears. The Host Mode Configuration page opens.
2. Select the host mode to be set in the Select Host Mode pulldown selection box. You can select only one host mode at a time. Each host mode is represented by its host mode number and a descriptive name.
3. Select the desired ports in the Select Port(s) box. You can select multiple ports or use the Select All button to select all of the available ports. The

available ports that are listed are all the available ports in the array and the current host mode setting.

## LUN Management – Set Host Mode (continued)

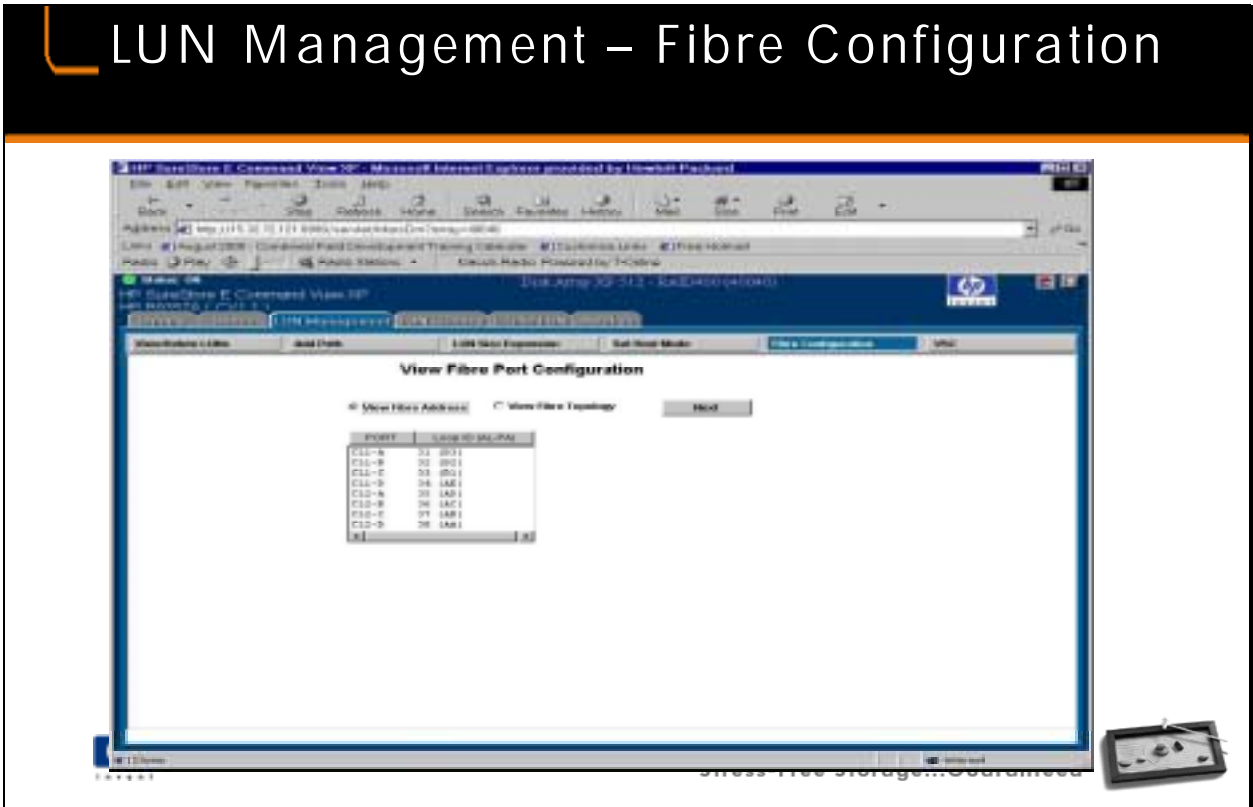


## Student Notes

To set or change the host mode configurations (continued):

4. Choose the Select button to assign the selected host mode to the selected ports. The selected ports for the selected host mode are displayed in the Set host mode for these ports? Box located on the right side of the Host Mode Configuration page.
5. (Optional) If you want to reverse the selected host mode and port combinations, select the combination in the Set host mode for these ports? box and choose the Deselect button. Use the Deselect All button to reverse all of the selected host mode and port combinations. The deselected port names are removed from the Set host mode for these ports? box and reappear in the Select Port(s) box.
6. Choose the Set button if you are satisfied with all of the selected host mode and port combinations. The selected values are set and the data on the page is refreshed.

To cancel any changes and return to the View Host Mode Configuration page, choose the Cancel button.



### Student Notes

Fibre Channel configuration includes viewing and setting the port address and viewing and setting the topology for Fibre Channel ports.

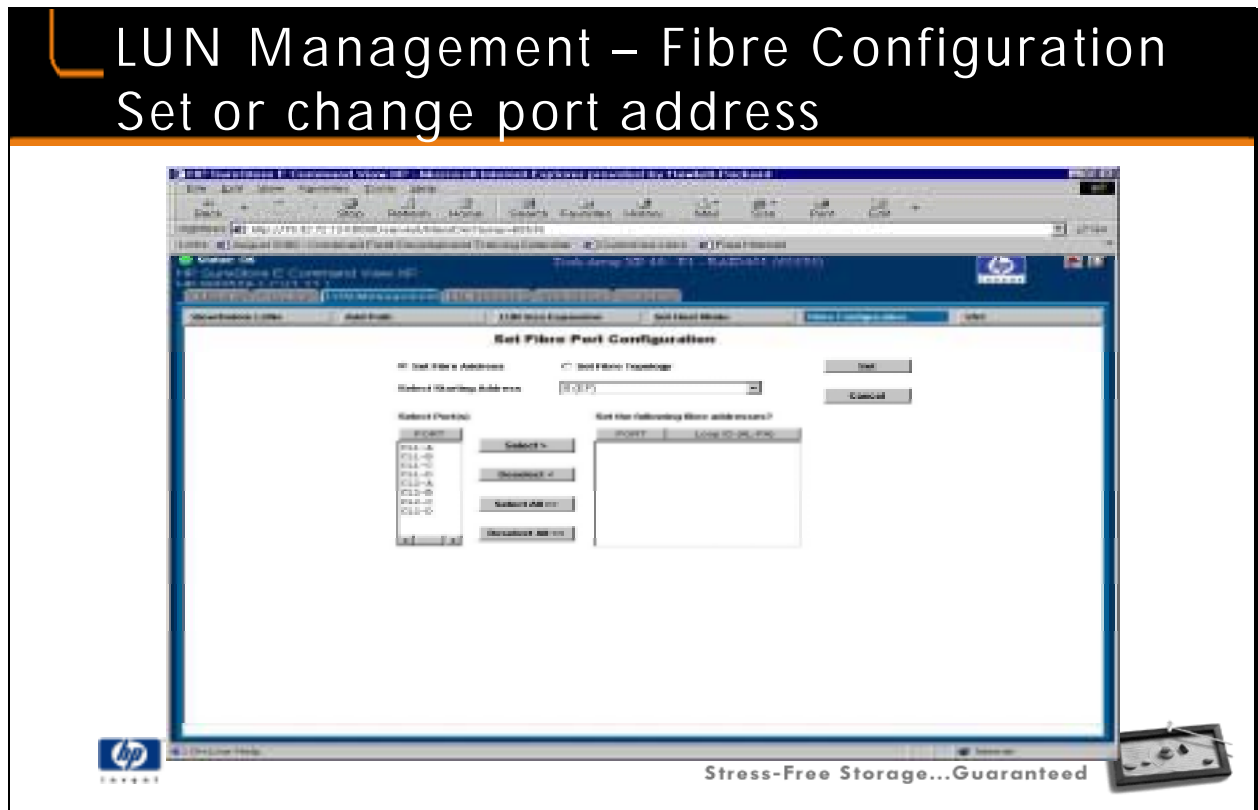
Note: Fibre Channel protocol automatically negotiates loop IDs. If you wish, you can use this function to make permanent loop IDs.

To view Fibre Channel configurations:

1. Start Command View and select a disk array.
2. Select the LUN Management tab. The LUN Management page appears.
3. Select the Fibre Configuration button on the LUN Management page. The Fibre Configuration page appears.
4. Choose the View Fibre Address radio button to view port addresses.
5. Choose the View Fibre Topology radio button to view the ports and their topology.

To set or change a Fibre Channel port address:

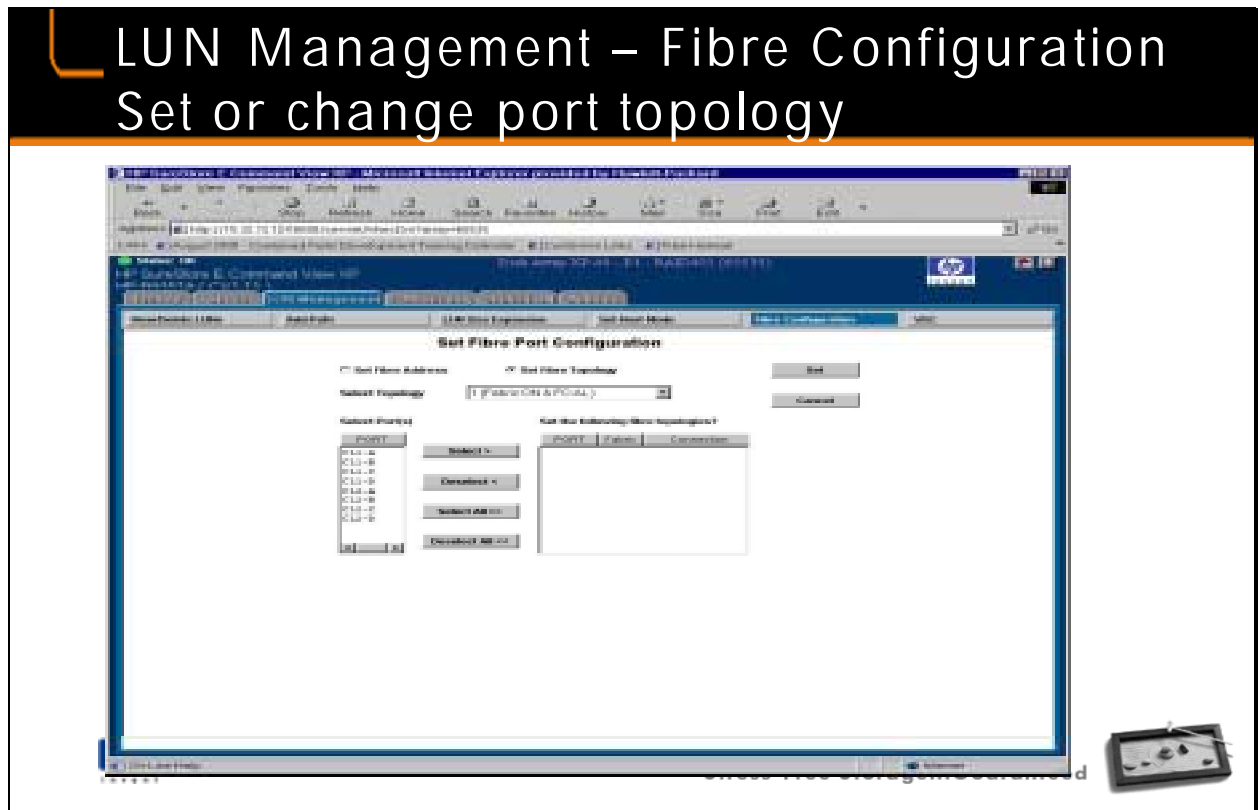
1. Choose the Next button on the View Fibre Port Configuration page. The Set Fibre Port Configuration page opens.



### Student Notes

To set or change a Fibre Channel port address (continued):

3. Choose the Set Fibre Address radio button, if it is not already selected by default.
4. Select the starting address in the "Select Starting Address" dropdown list. You can select only one address at a time.
5. Select the desired port in the "Select Ports(s)" box.
6. Choose the Select button or choose the Select All button to select all of the available ports. The selected ports appear in the "Set the following fibre addresses?" box. The selected ports will have addresses assigned starting with the address selected in the "Select Starting Address" dropdown list.
- 7.(Optional) If you want to reverse your selection of Fibre Channel port addresses, select the ports in the "Set the following fibre addresses?" list box and choose the Deselect button. Use the Deselect All button to reverse all of the selected Fibre Channel port addresses. The deselected port addresses are removed from "the Set the following fibre addresses?" list box and reappear in the "Select Port(s)" box.
8. Choose the Set button if you are satisfied with the port addresses appearing in the "Set the following fibre addresses? List" box. The new address configurations are set in the array, and the data is refreshed on the page. Choose the Cancel button to cancel your selections and return to the View Fibre Port Configuration page.



### Student Notes

To set or change a Fibre Channel port topology (continued):

The procedure starts in the same way as the set port procedure, but here you choose the Set Fibre Topology radio button.

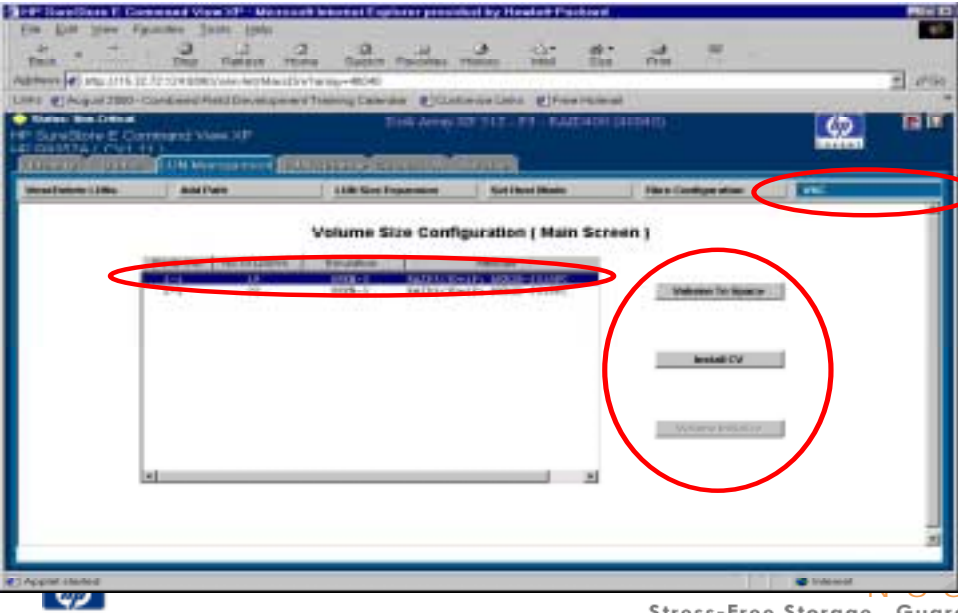
3. Choose the Set Fibre Topology radio button.
4. Select the topology in the Select Topology dropdown list. You can select only one topology at a time.
5. Select the desired port in the Select Ports(s) box.
6. Choose the Select button or choose the Select All button to select all of the available ports. The selected ports appear in the Set the following fibre topologies? List box. The topology selected in the Select Topology dropdown list is assigned to each of the selected ports.
7. (Optional) If you want to reverse the selected Fibre Channel port topologies, select the ports in the Set the following fibre topologies? List box and choose the Deselect button. Use the Deselect All button to reverse all of the selected Fibre Channel port topologies. The deselected port topologies are removed from the Set the following fibre topologies? list box and reappear in the Select Port(s) box.
8. Choose the Set button if you are satisfied with the port topologies appearing in the Set the following fibre topologies? list box. The new topology configurations are set in the array, and the data is refreshed on the page.



SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Command View XP

Choose the Cancel button to cancel your selections and return to the View Fibre Port Configuration page.

## VSC Main screen



- Select the VSC button on the LUN Management screen
- Select a parity group and an operation

Stress-Free Storage...Guaranteed

To perform VSC operations, select the VSC button in the LUN Management screen. The VSC main screen is displayed. All the VSC operations can be done using the following screens:

- VSC Main Screen
- Volume Size Configuration Screen
- Volume to Space Operation Screen
- Volume Initialize Screen
- LDEV ID Configuration Screen

To start VSC operations:

1. Start Command View and select a disk array.
2. Select the LUN Management tab. The LUN Management page appears.
3. Select the VSC button on the LUN Management page. The VSC main screen appears.

Main Screen fields

The screen displays the Parity group, number of LDEVs in the parity group, emulation type, and attribute. RAID information is described in the attribute column as well as information about whether it is a VSC volume or not. The Emulation column will display "Open-\*" if the parity group has volumes of different emulation types.

Main Screen buttons

The three buttons allow you to select the three operations: Volume to Space, Install CV, and Volume Initialize.

The Volume size configuration (VSC) option on the XP48 and XP512 allows you to configure variable-size custom volumes. Creating reduced-size VSC volumes is the opposite of Logical Unit Size Expansion (LUSE). VSC improves data access performance by reducing logical device contention and host I/O queue times, particularly when several frequently accessed files are located on the same volume. VSC enables better utilization of the physical storage capacity of the XP 512 or XP48, and reduces the amount of administrative effort required to balance I/O workloads. Note: VSC (also referred as Custom Volume Size or CVS) used in conjunction with Cache LUN can produce greater performance improvements than when either of these options is used individually. VSC operations can be performed by users with administrator or custom VSC access privileges. Users without administrator or VSC access privileges can only view the VSC configuration information for each attached disk array. Note: VSC is also available for the XP256.

The term "Parity Group" is often used synonymously with "RAID Group." It refers to a number of physical disks associated with the configured RAID level. For example, the RAID Group associated with level RAID1 contains 4 physical disks. The RAID Group associated with level RAID5 contains 4 physical disks. A parity group is called a normal or base volume when all of its member LDEVs are fixed-size. In a normal volume, the size of the member LDEVs is determined by type of device that is being emulated, and the number of LDEVs is determined by the physical size of the disk. On the XP 512 and XP48, VSC formats one or more of the LDEVs into free space. That free space can either be used to install variable-sized custom volumes (CVs), or left as free space for future use.

#### VSC Operations

**Volume to Space** - The first step in performing VSC operations is to convert one or more

normal volumes to free Space by using the Volume to Space operation. When you convert a normal volume to a free space, all of the logical blocks in the volume become available for use in VSC.

**Install CV** - You can add one or more new CVs using the Install CV option. However, you can also use the Volume to Space option to remove a CV and convert it back into free space.

**Volume Initialize** -To deinstall (remove) all the CVs and convert all the free space back into normal volume, use Volume Initialize.

Buttons are enabled or disabled based on the user selection. If the selected parity group has more than one normal volume, then the Volume to Space button is enabled. If the parity group has free space greater than 36 MB, then the Install CV button is enabled. If the selected parity group has any deleted volumes present, the Volume Initialize button is enabled. If there are no custom volumes present in the parity group, the Volume Initialize button is disabled. Multiple selection from the table are not allowed. You can perform VSC operations on the LDEVs of only one parity group at any given time.

Volume to Space operation - Select one of the parity groups and click on the Volume to Space button. This takes you to the screen explained in the Volume To Space operation section.

Install CV operation. Select one of the parity groups and click on the Install CV button. This takes you to the screen explained in Install CV section.

Volume Initialize operation. Select one of the parity groups and click on the Volume Initialize button.

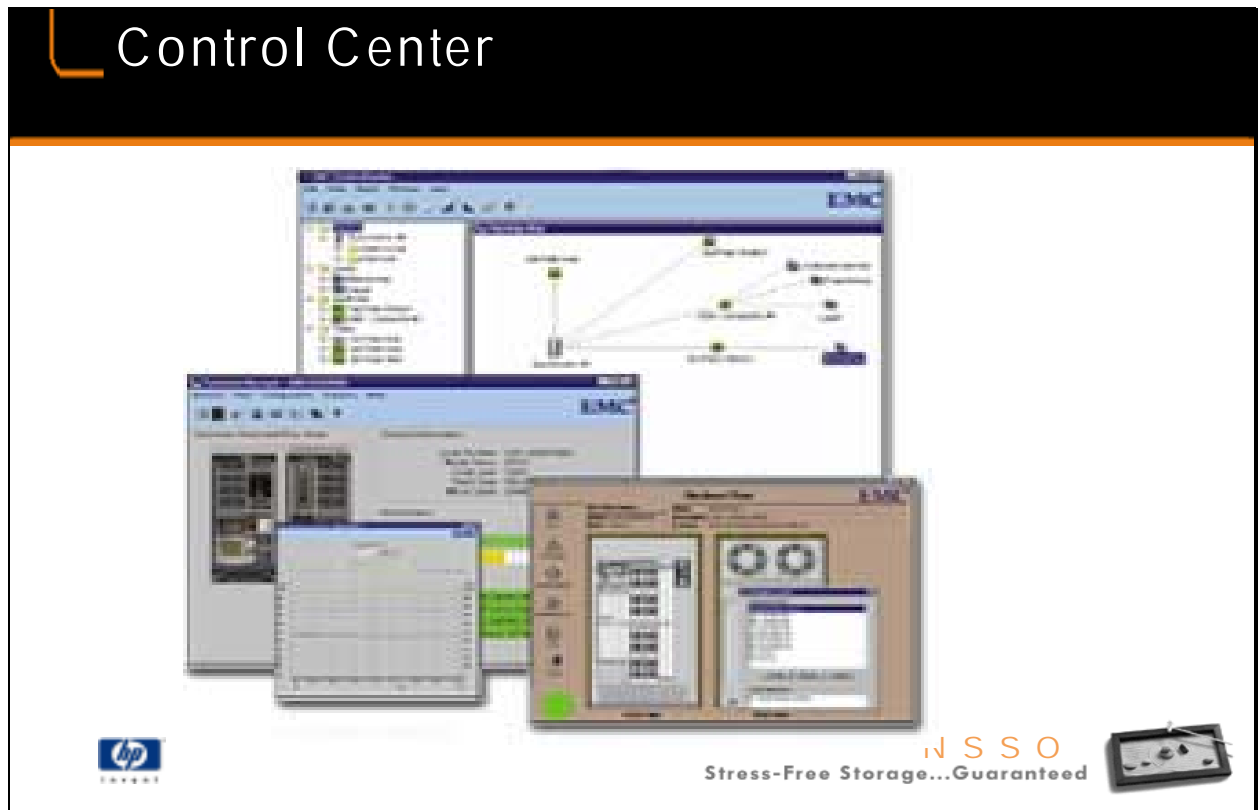
Note: The Volume to Space and Volume Initialize operations are destructive. When you perform these operations, you will get a warning message, which you must acknowledge before proceeding.

EMC Control Center vs. Command View XP	
<u>EMC Control Center</u>	<u>HP Command View XP</u>
<ul style="list-style-type: none"><li>➤ Inconsistent Management Interfaces for tools (some Windows, some X-Windows, some WEB-based)</li><li>➤ Command line interface</li><li>➤ 100% functionality of applications not available<ul style="list-style-type: none"><li>• Example</li><li>• CLI only available for open systems</li><li>• JCL batch utility for OS/390</li><li>• Support only for EMC hardware</li><li>• Optimizer not supported on AS400</li></ul></li><li>➤ EMC directed SAN Management to lock in Symmetrix hardware sales</li><li>➤ Integration into: CA Unicenter, VantagePoint Operations, IBM Tivoli TME</li></ul>	<ul style="list-style-type: none"><li>➤ All Management Applications</li><li>➤ 100% functionality of applications available via GUI</li><li>➤ Integrated Open SAN Management Support with SAN Manager DM</li><li>➤ 60-day trial of Optional XP Management Applications</li><li>➤ Secure login support</li><li>➤ SNMP event integration for leading management platforms (OpenView Vantagepoint IT/O, CA UniCenter TNG, IBM Tivoli TME and others)</li><li>➤ "Deep Integration" with OpenView Vantagepoint ITO and Service Navigator via the XP Smart Plug-In product</li></ul>



Stress-Free Storage...Guaranteed





## EMC ControlCenter Suite

**Centrally manage all Enterprise Storage Network components.**

The EMC ControlCenter family of products enables you to monitor, configure, control, tune, and plan storage across an entire EMC Enterprise Storage Network (ESN) from a single console. EMC ControlCenter greatly simplifies management of your E-Infostructure by providing centralized management for all your:

**Symmetrix Enterprise Storage**

**Connectrix Family of Fibre products**

**Third-party switches and hubs**

**Multi-platform server hosts**

The foundation of EMC ControlCenter is the EMC ControlCenter Console and Symmetrix Manager. Optional software applications leverage that

status, performance, and configuration information to actively manage components in the ESN. These applications include:

**Resource View**— maps host objects to Symmetrix devices with physical, logical, and line-of-business views to help you understand the virtual layers in your IT environment.

More information on Resource View

**Resource Availability**— provides the capabilities to manage logical resources on MVS, UNIX, and Windows servers.

More information on Resource Availability

**SRDF/TimeFinder Manager**— provides local and remote management of business continuity and disaster recovery operations.

**Symmetrix Disk Reallocation**— offers dynamic drag-and drop storage allocation.

**Symmetrix Optimizer**— automatically monitors, analyzes, and moves highly active logical volumes to maintain optimal performance.

More information on Symmetrix Optimizer

**Symmetrix Database Tuner**— monitors and tunes Oracle databases; optimized for Symmetrix.

More information on Symmetrix Database Tuner

**Workload Analyzer**— collects, graphs, and archives performance data for Symmetrix trend analysis, the most critical step of capacity planning.

## EMC Control Center 4.0 Overview


- Centralized graphical management platform for Symmetrix, Connectrix and 3rd party hubs & switches
- Base Product:
  - **ControlCenter Console which includes Symmetrix Manager**
- Optional "Plug-In" Applications for ControlCenter
  - **Resource View** - displays databases, file systems and volume managers and their relationships with Symmetrix
  - **SRDF & TimeFinder** - disaster recovery & business continuance
  - **Disk Reallocation** - ability to reassign volumes, target LUN addresses, adding/deleting host adapter & ports
  - **DataBase Tuner** - enables active monitoring of database enabling storage & database administrators to work together
  - **Workload Analyzer** - graphical representations of performance variables on Symmetrix



N S S O  
Stress-Free Storage...Guaranteed








# EMC ControlCenter 4.0

Software Product Line Comparison

➤ Supports three types of GUIs


- Support is different for each GUI
- Most complete solution is Windows/NT with browser

EMC Software Product	Unix Motif	Windows NT 4.0	Browser	MVS Support
CC Console Symmetrix Manager	SYMM Yes	SYMM Yes	ESN Yes	Yes
Resource View	No	Yes	Future	
SRDF/TimeFinder	Yes – CLI for MVS	Yes – CLI for MVS	Future	No – CLI
Symmetrix Disk Reallocation	Yes	Yes	Future	No
Database Tuner	No	Yes	Future	
Optimizer	No	Yes – CLI for open systems w/CC Base  JCL batch utility for MVS OS/390	Future	
Workload Analyzer	No	Yes	Future	Yes



NSSO

Stress-Free Storage...Guaranteed



## Command View XP

### Resources Available

- Command View XP Sales Guide: **Keyword 'COMMANDVIEW'**
- Command View XP Product Brief: **Keyword 'ESMAYLAUNCH'**
- Software Product Structure Training Slideset: **Keyword 'XPSTOFT'**
- ESP website: **keyword "gohome"**
- Command View XP "demo" slideset:  
**self-paced Command View XP PowerPoint demo**
- XP Configuration/Ordering Guide:
  - XP256: **Keyword 'XP256CONFIG'**
  - XP512: **Keyword 'XP512CONFIG'**
- XP Product Software and Connectivity Summary
  - XP256: **Keyword 'XP256CONNECT'**
  - XP512: **Keyword 'XP512CONNECT'**
- Competitive and Misc. Product Detail: **ESP Search**
- Reference customers: **WINREF Database on ESP**



N S S O  
Stress-Free Storage...Guaranteed



## HP SureStore E Command View

Product Structure

- B9357A - Command View XP
  - Option 01 - For new XP256 installations
    - ♦ List Price: \$24,000
  - Option 02 - For new XP512 installations
    - ♦ List Price: \$24,000
  - Option 03 - For existing X256 installations with Remote Control (Command View bits only)
    - ♦ List Price: \$10,000
  - Option 04 - For existing Command View customers on XP256 that have added XP512
    - ♦ List Price: \$17,500



N S S O  
Stress-Free Storage...Guaranteed



## XP Software - Try 'n Buy

60-day Trial Evaluation of Software

- 60-day Trial Evaluation of four solutions (XP512 only)
  - Lun Configuration Manager
  - Auto LUN XP
  - Cache LUN
  - Secure Manager
- Automatically included with Command View XP
- Software "live" for 60 days from date of installation
  - At end of evaluation period, management of software disabled
  - 10-day emergency key available for CE to reinitiate
- Customer can purchase products individually
- Timeout feature ONLY initialized after installation of Temporary Software Key
- HP provides full support during evaluation timeframe



N S S O  
Stress-Free Storage...Guaranteed




## Module Wrap



N S S O  
Stress-Free Storage...Guaranteed






# Module 9

## Continuous Access XP B9320A B9321A, B9322A, B9323A, and B9324A


&

## Extension – B9325A B9326A, B9327A, B9328A, and B9329A



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



MetroCluster with Continuous Access XP (MetroCluster/CA) is a new addition to Hewlett-Packard's array of high availability products for customers requiring integrated disaster recovery solutions. MetroCluster/CA provides automated reconfiguration of HP SureStore E Family of Disk Arrays using Continuous Access XP in a campus or metropolitan MC/ServiceGuard cluster within the 100 km limit of the FDDI network. This product is being released with the following product number:

- B8109BA license and media

MetroCluster/CA provides automated failover of MC/ServiceGuard packages on local and remote XP Family of Arrays. This provides the means for configuring an MC/ServiceGuard cluster that can be distributed over a range of up to 100 km with automated failover, with a maximum distance of 43 km between the XP256 arrays. The maximum supported range of the FDDI loop is 100 km, which is the total maximum loop length among all three data centers. Each host adapter supports a maximum distance between stations of 2 km. Repeaters must be used to extend beyond 2 km. Integrating MetroCluster/CA with MC/ServiceGuard allows application packages to fail over:

- among local nodes that are attached to the same XP Family array.
- between one system that is attached locally to its XP Family array and another remote node that is attached locally to the other XP Family Array.

## Module Agenda



- Continuous Access XP Overview & Specifications
- Continuous Access XP Performance Considerations
- Continuous Access XP Added Functionality (*Extended*)
- Continuous Access XP Extending the Distance
- Continuous Access XP States of Operation
- Continuous Access XP with Business Copy and OmniBack II
- Continuous Access XP Competitive Landscape
- Cluster Extension XP
- Whats New
- Wrap-Up

HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



MetroCluster/CA Version A.00.04 supports clusters up to 16 nodes on MC/ServiceGuard A.11.xx and clusters up to 8 nodes with MC/ServiceGuard A.10.11 or A.10.12.

### \*\*Summary:

Beginning August 1st, 2001, Continuous Access XP, Continuous Access XP Extension, and Business Copy XP Software for the XP Disk Array will be licensed based upon 'used' capacity.

### \*\*Background:

For 7 XP Software titles (Continuous Access XP, Continuous Access XP Extension, Business Copy XP, LUN Configuration Manager XP, Secure Manager XP, Cache LUN XP, and Auto LUN XP), current license terms dictate that in order for a customer to use these products on an XP Disk Array they must purchase the appropriate number of license-to-use (LTU) products to provide total license capacity to cover the total 'raw' capacity of the disk array in question (Raw capacity is defined as the total nominal capacity of all Disk Array Groups installed on the XP Array). Customers have expressed that licensing based upon the raw capacity of the Array for Business Copy XP, Continuous Access XP, and Continuous Access XP Extension titles is not equitable and does not align with their perceived value of the products owing to the fact that they are most often deployed on a subset of the Disk Array.

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Continuous Access XP

Effective August 1st, 2001, license terms for Continuous Access XP, Continuous Access XP Extension, and Business Copy XP will be changed to be based on the 'used' capacity of the Disk Array. This means that customers will buy license capacity to cover only the portion of the array capacity that is being 'utilized' by the Software.

For these Software titles, used capacity is defined as the total capacity of all volumes that are being managed as a Primary/Source Volume (P-Vol) or a Secondary/Copy Volume (S-Vol) on the Disk Array. This amount represents the total true usable capacity of these volumes corresponding with their configured LDEV sizes (i.e. the amount excludes RAID overhead). For Business Copy XP, this represents all related P-Vol and S-Vol on the Disk Array (every source volume and any associated copy volumes). For Continuous Access XP and/or Continuous Access XP Extension configuration, this represents any P-Vol or S-Vol pairs on a single Disk Array that is being managed in a P-Vol/S-Vol pair with another Disk Array.

Note the following:

- There will be no change to product structure (same products and product numbers as prior to August 1st)
- There will be no change to product pricing for the 1/5/10/25TB LTU products (no price changes in moving to used capacity license terms)
- Licensing for LUN Configuration Manager XP, Secure Manager XP, Auto LUN XP, and Cache LUN XP will continue to be based upon the raw capacity of the Disk Array.

This change results in SW License cost savings for most customers who require these products, especially for those who utilize a relatively small portion of the array for the solution.

\*\*

A CSU Newsflash will soon be posted on ESP which details this change (ESP, keyword CABUSED, )



## Continuous Access XP Description

- Asynchronous Continuous Access XP (B9325A) allows creating and maintaining remote mirror copies of an Open Volume or extended LDEV on a separate XP Family of Arrays.
- Synchronous and asynchronous (*new*) - I/O across ESCON and now 7/01/01, Fibre Channel interfaces
- Performance is optimized to increase throughput
  - Data compression
  - Automatic load balancing with prioritization
- Both the primary and the copy are RAID protected, and both can have Business Copies (BC), typically for disaster recovery and remote backup.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



Volume types supported: Normal, LUSE, VSC and Cache LUN ( DCR )

Continuous Access XP can co-exist with Business Copy XP. Synchronous and asynchronous I/O across Escon are supported.

Continuous Access XP will support all Open systems features of the XP256.

Open systems hosts can be connected to the XP256 via SCSI or Fibre Channel. The HRC link requires ESCON at this time.

Summary of Continuous Access XP, the – these are some of the high points of the product. And for review for some of you people, this is a synchronous, remote copy between two XP256 chassis. This is in contrast to the asynchronous copy that happens on Business Copy. So, this is a synchronous event. People use this for disaster recovery – the example you see here on the slide is actually a MetroCluster, but you're not limited to MetroCluster – and design a larger variety of disaster recovery solutions with Continuous Access. Just to reiterate the – it is a synchronous, right (Iowa), cross (S-CON) links – that is the current link between the chassis. Currently, we have a limit of 43 kilometers using the (S-CON) repeaters or directors and we control the copies – these are the volumes. So, you have a primary volume and a secondary volume. We've been spending some effort to get the language consistent between Business Copy and Continuous Access. In any case, you use RAID Manager or

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Continuous Access XP

Remote Control XP to control, again, the linking and paring and splitting of these volumes between the two arrays. It should be noted at this point that we use in-band communication for this control when we talk about the RAID Manager. This means that we communicate to the host; the host communicates through its channel, whether that be SCSI or Fibre, to give the commands to the array when you're using RAID Manager.

## General CA Concepts/Terminology

- q **Mvol:** the primary volume (which contains [typically] the most current data)
- q **Rvol:** the secondary volume (which is a mirror copy of the mvol)
- q **ESCON:** the uni-directional interconnecting fiber (pair) technology used by CA
- q **LCP:** a Listening ESCON port
- q **RCP:** a talking (initiator) ESCON port
- q **ATM:** a WAN (Wide Area Networking) method subject to outages and non-deterministic routing
- q **T1/T3:** a few WAN transmission method/speeds
- q **CU:** Controller Unit (max 256 ldevs/CU, 4 CUs/array (XP256, 16 CUs per XP512/48))
- q **Sync-CA:** the results of the write are not returned to the mvol host until the outcome of the remote volume write is known (i.e. pass/fail/timeout). Unless noted otherwise, this Module refers to Sync-CA unless noted otherwise. Async-CA: local host write completes immediately (does not wait for data at remote cache). Similar data consistency as fence = NEVER.



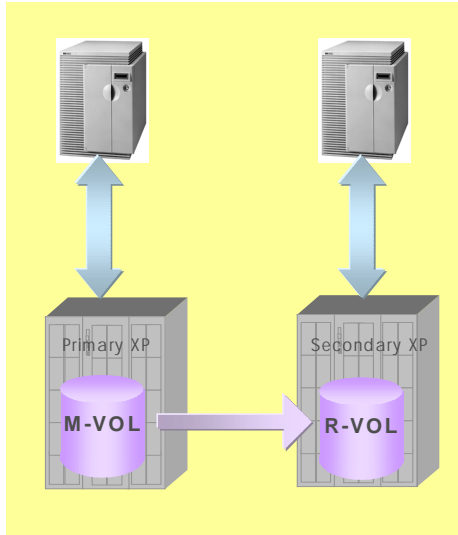
HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



Most of these will be explained in greater detail within this module.

## Continuous Access XP Product Overview



The diagram illustrates the Continuous Access XP architecture. It shows two HP SureStore XP arrays: a Primary XP array on the left and a Secondary XP array on the right. Each array contains a main volume (M-VOL) and a remote volume (R-VOL). A horizontal arrow points from the M-VOL of the Primary XP array to the R-VOL of the Secondary XP array, indicating data replication. Above each array is a server icon connected to the array by a double-headed vertical arrow. The background is yellow.

**HP Restricted**

**What does it do?**

- maintains copies of volumes on a primary XP array on a secondary (remote) XP array
- monitors the status of each primary/secondary volume pair to ensure data consistency and currency
- with RAID Manager, manages failover and failback operations

**How does it work?**

- sends a copy of each write operation executed against a main volume (M-VOL) on a main XP array across a remote link to a remote XP array where it is in turn executed against a paired remote volume/mirror (R-VOL)
- if the remote copy operation is disrupted (e.g. remote link or array failure), system performs failover to R-VOLs while maintaining 'delta' tracking info to allow for quick recovery 'resync' upon failback

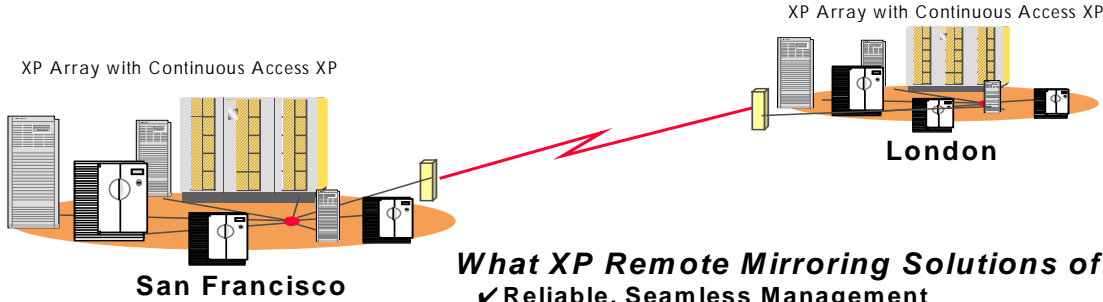
**N S S O**  
Stress-Free Storage...Guaranteed

This is intended to give everyone sort of a basic understanding of what Continuous Access XP is as a product. At a very high level, of course, it's remote mirroring software for the XP family of arrays. Suffice to say, this is the foundation capability that HP is utilizing in turn leveraging to do the IP connectivity solution that we will be talking about later.

## XP Remote Mirroring Solutions

**Continuous Access XP enables a wide range of XP remote mirroring solutions:**

- ✓ **Data Migration:** for migration of data from one array or datacenter to another
- ✓ **Data Replication:** for distributing data between geographically dispersed arrays
- ✓ **Disaster Recovery/Clustering:** for maintaining a 'failover' datacenter



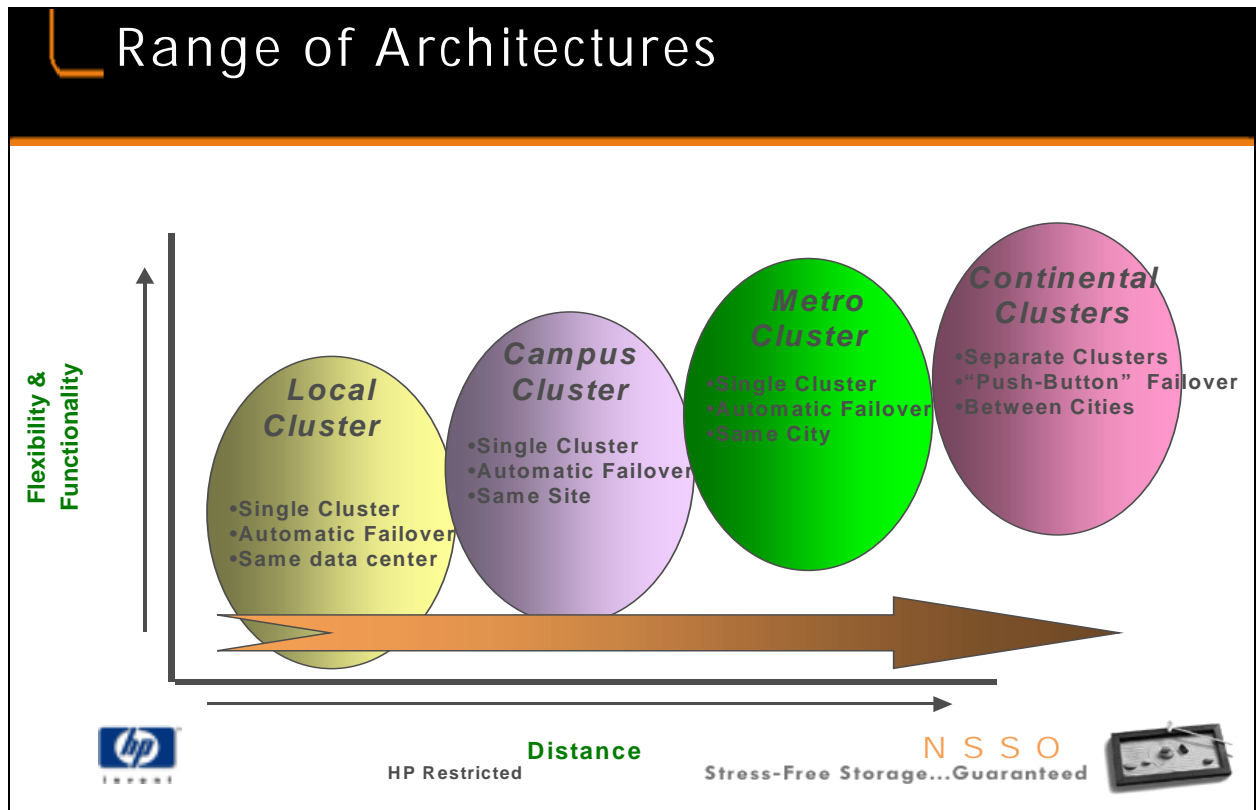
**What XP Remote Mirroring Solutions offer:**

- ✓ **Reliable, Seamless Management**
- ✓ **High Performance, Cost Effective Solutions; cross-town, or across the globe**

HP Restricted

NSSO  
Stress-Free Storage...Guaranteed

The slide title, "XP Remote Mirroring Solutions", you'll see that Continuous Access XP enables a wide range of remote mirroring solutions. HP tends to categorize them as follows: data migration, which is sort of for this one-time or specific event-driven activity of migrating data from one data center to another; data replication, which is sort of more of an ongoing distribution of data between geographically dispersed arrays; and then, last but not least, of course, mirroring for disaster recovery and/or clustering, in which you're maintaining a remote fail-over data center. This type of solution gives you the ability to recover quickly and efficiently in those circumstances. In a fail-over situation, these time-stamped program and data copies are seamlessly restarted with another system. So all in all, what the XP remote mirroring solutions offer are, you know, reliable, seamless management of the disk mirroring operation, and then of course, high-performance cost effectiveness, whether it's for local mirroring or continental distance mirroring.




This slide gives an overview or shows the range of architectures of cluster architectures. HP assumes that most of you are probably familiar with MC/ServiceGuard. And so, as we look at the lower left-hand corner, we're looking at local clusters. And then, as we move to the right, we have Campus Clusters and then Metro Cluster and Continental clusters. And as you can see with the first three, the first two bullets, you'll see single cluster and automatic fail-over, that's all common and the only difference, really, is in the geographic dispersion where with a local cluster, you're looking at the same data center, Campus Clusters; we're looking at the same site, Metro Cluster, the same city. And then as we move to the far right, Continental clusters, where you'll see there is - the first couple of bullets are different. We're talking about separate clusters, not a single cluster. So now we have two clusters involved. They are standard MC/ServiceGuard clusters, but the Continental Clusters product has added additional capabilities on top of those two clusters. What's provided with that product is a push-button fail-over capability. And there really is no distance limit between the clusters. So, and we'll look at some of the requirements for the Continental clusters product.

mc/service guard is the most robust cluster software for open systems, offers reliable and fast failure detection, heartbeat support/cluster reformation over DLPI protocol (single subnet) , supports up to 16 systems per cluster, limited distance up to 10 km

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Continuous Access XP

hp metro cluster needs an arbitrator system to protect against “split-brain” syndromes, supports 7 systems per data center, max. distance up to 50 km




hp continental cluster uses two mc/service guard cluster, no distance limitation, support is limited to up to 31 application packages per continental cluster



# Campus Cluster

Comparative Features

<b>Cluster Topology</b>	<b>Single Cluster up to 4 nodes across 2 data centers or up to 16 nodes across 3 data centers</b>
<b>Geography</b>	<b>Campus, up to 10 km ( Fibre Channel limitations)</b>
<b>Network Subnets</b>	<b>Single IP Subnet</b>
<b>Network Types</b>	<b>Dedicated Ethernet, FDDI or Token Ring</b>
<b>Cluster Lock Disk</b>	<b>Required for 2 nodes, optional for 3-4 nodes, not used with larger clusters</b>
<b>Failover Type</b>	<b>Automatic</b>
<b>Failover Direction</b>	<b>Bi-directional</b>
<b>Data Replication</b>	<b>MirrorDisk/UX</b>




HP Restricted

Stress-Free Storage...Guaranteed

This slides gives you a quick list of some of the features for MC/ServiceGuard for local clusters. And so, you can see by this slide that cluster topology, is basically a single cluster up to 16 nodes, so that's the maximum number of nodes you can have in a MC/ServiceGuard cluster running 11.0 HP-UX. Geography, we're talking data center or campus. We're talking about a single IP subnet, MC/ServiceGuard requires a single IP subnet and here we have shown the network types that are supported, like Ethernet, FDDI and Token Ring. There's also cluster lock discs required for two-node, optional for three and four nodes and not use the larger clusters, fail-over type automatic. Basically, your fail-over direction, you can fail over to any node within the cluster. And as far as data replication goes, typically, we're looking at an environment where you just have shared disks that are accessed by the nodes within the cluster.









# Metro Cluster

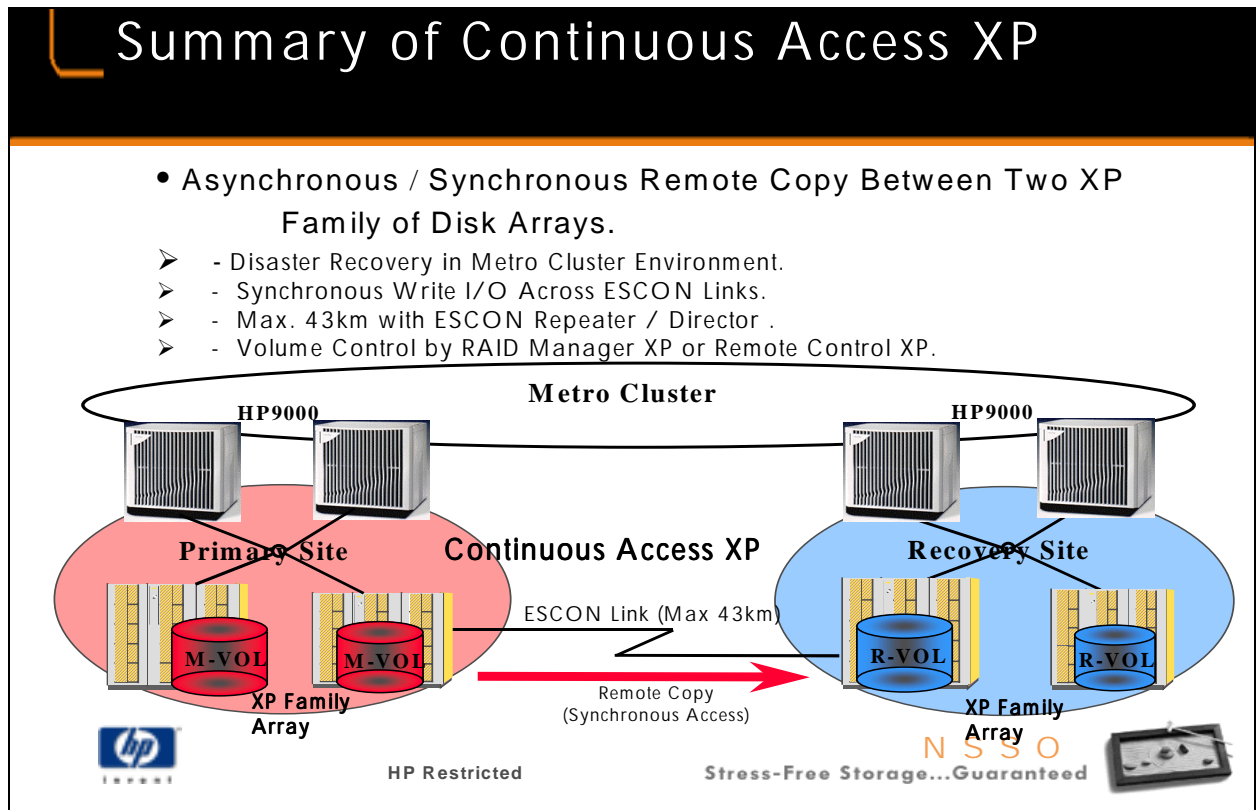
## Comparative Features

<b>Cluster Topology</b>	<b>Single Cluster up to 16 nodes spread across 3 data centers</b>
<b>Geography</b>	<b>Campus or Metropolitan area</b>
<b>Network Subnets</b>	<b>Single IP Subnet</b>
<b>Network Types</b>	<b>Dedicated Ethernet, or FDDI</b>
<b>Cluster Lock Disk</b>	<b>Not Used; 1-2 Arbitrators in third data center act as tie breaker</b>
<b>Failover Type</b>	<b>Automatic</b>
<b>Failover Direction</b>	<b>Bi-directional</b>
<b>Data Replication</b>	<b>Physical, in hardware (XP256 CA or EMC SRDF)</b>



HP Restricted

So in moving from campus to metropolitan or a metro cluster, we basically are still talking about a single cluster, up to 16 nodes and now there's no cluster locks involved. We're basically using three data centers where you have arbitrators at a third data center. It's a campus or metropolitan area. Again, because it's a single cluster, we're looking at a single IP subnet, networking types of Ethernet or FDDI and we have automatic fail-over, bi-directional fail-over. And with Metro Cluster, the data replication is done in hardware and therefore it's a physical replication and we support both the XP256/512/48 Continuous Access, or the EMC SRDF replication facilities.



There are three steps to set up a CA Configuration:

(1) (1) Set up the CA ESCON Links:

Via Remote Control XP, set up the ESCON links:

Change the appropriate ESCON ports from LCP (listener) to RCP (talker) type.

Select the Number of ESCON Links.

Specify the serial number of the remote array

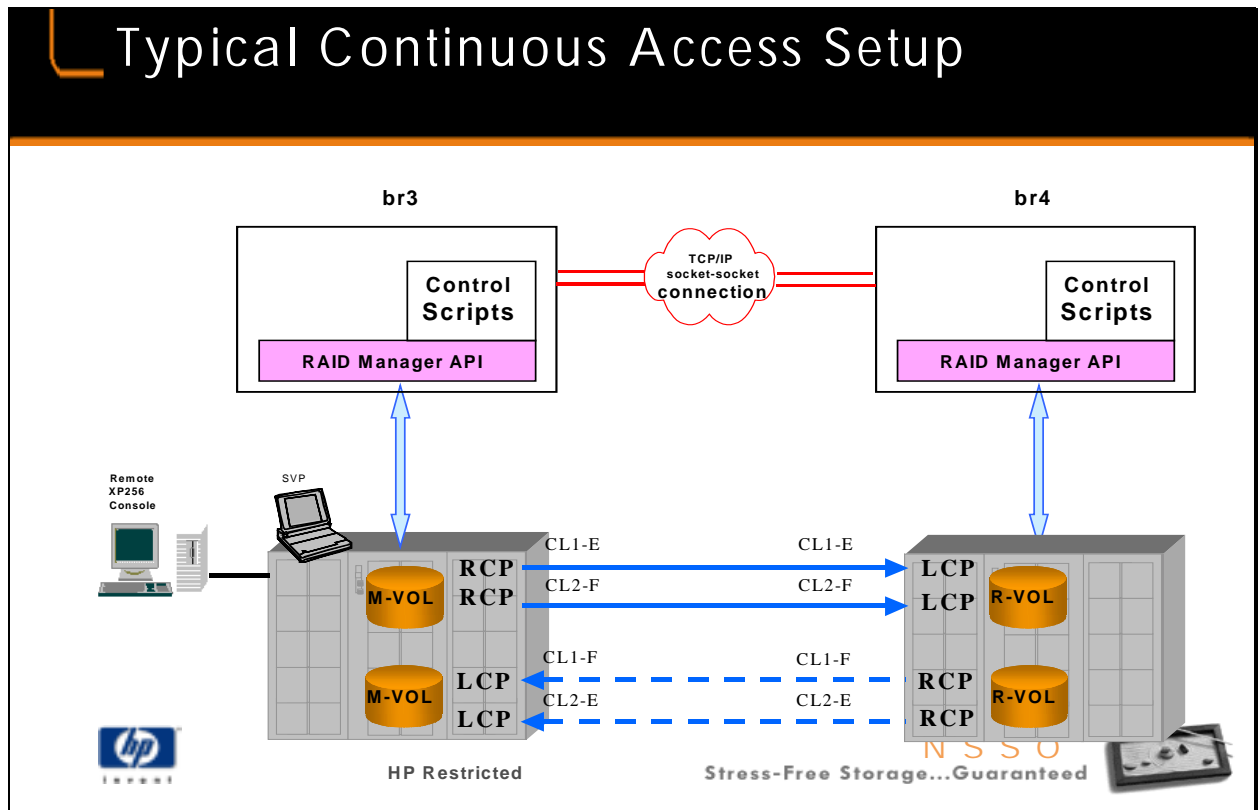
Specify the Port Numbers used by each CU, etc..

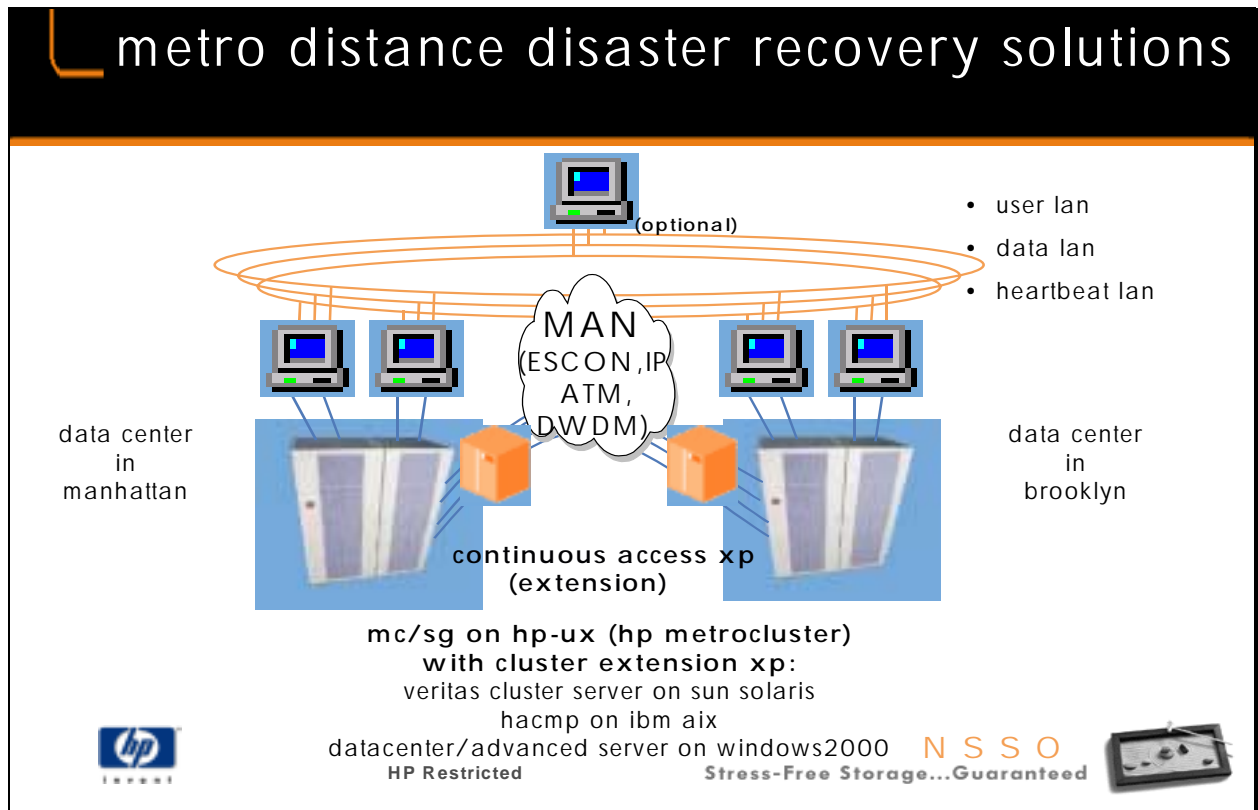
(2) Set up RAID Manager XP:

If you want to use RAID Manager XP, you must create the "Configuration Definition File(s)".

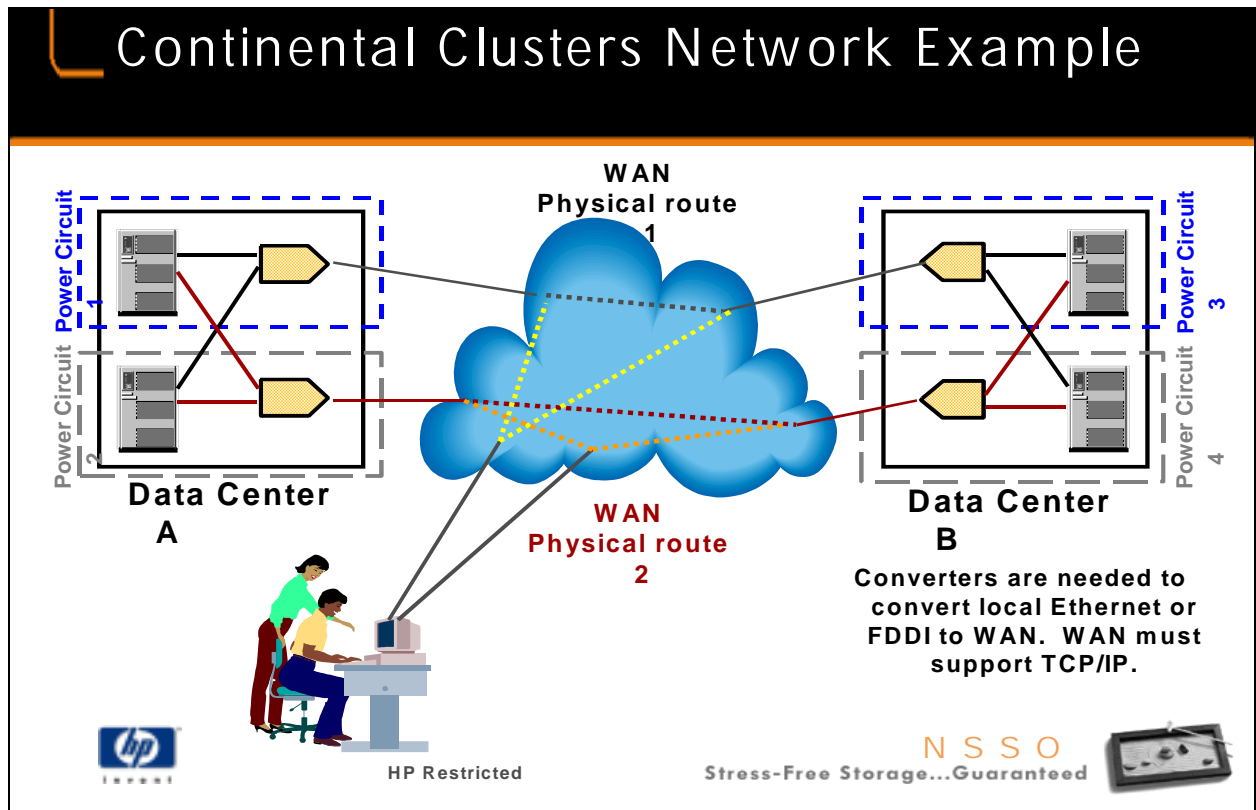
(3) Control of the Volume Pairs;

By operating RAID Manager XP or Remote Control XP; Create Paired Volumes (after defining the LUNs and associating them with ports). In addition, you can control Paired Volumes by using various CA commands (Pair Create, Pair Split, Pair Resynchronization, Pair Display, etc.).





"Metro Distance Disaster Recovery Solutions," this is just a picture of the main set-up of such cluster solutions. A little later we'll explain a little bit more how you disperse data centers over long metropolitan distances. For the HP-UX environment we have offered those customers Metro Cluster and Service Guard for quite some time. And just recently we have extend that offering to our multi-platform OS support model which we will discuss later in this module.



ContinentalClusters is a Hewlett-Packard high availability solution that provides disaster tolerant clustering over long distances. ContinentalClusters is being released with the following product numbers:

- B7659BA license and media
- B7660BA documentation

ContinentalClusters employs semi-automatic failover of MC/ServiceGuard packages from a primary cluster to a recovery cluster following a cluster event that indicates serious disruption of service on the primary cluster. The product consists of a set of configuration tools, a monitor that sends notification of cluster-down events, and a command that moves packages from one cluster to another. In addition, ContinentalClusters includes the following component products:

- B8324BA (Cluster Object Manager)—provides the ability to query ServiceGuard cluster status.
- B6264BA (MetroCluster with EMC SRDF)—provides a special package control script template to implement physical replication between EMC Symmetrix disk arrays.
- B8109BA (MetroCluster with Continuous Access XP)—provides a special package control script template to implement physical replication between HP SureStore XP256 or XP512 disk arrays. Version A.02.00 of ContinentalClusters requires the HP-UX 11.0 operating system and MC/ServiceGuard version A.11.09 or later. The product also requires Event Monitoring Service version

A.03.10, which is included as a component product with MC/ServiceGuard. And as was discussed earlier, we were talking about two separate clusters and we have a primary cluster and a secondary, or recovery cluster and what we can do is configure the monitoring between the recovery and primary cluster and also configure a different event notification that can be generated.

#### HP ContinentalClusters (B7659BA)

A product to automate the failover of MC/ServiceGuard packages among TWO separate clusters

primary and secondary cluster failure notification is configurable

semi-automatic "push button" initiates automated failover

choice of various logical or physical data replication methods

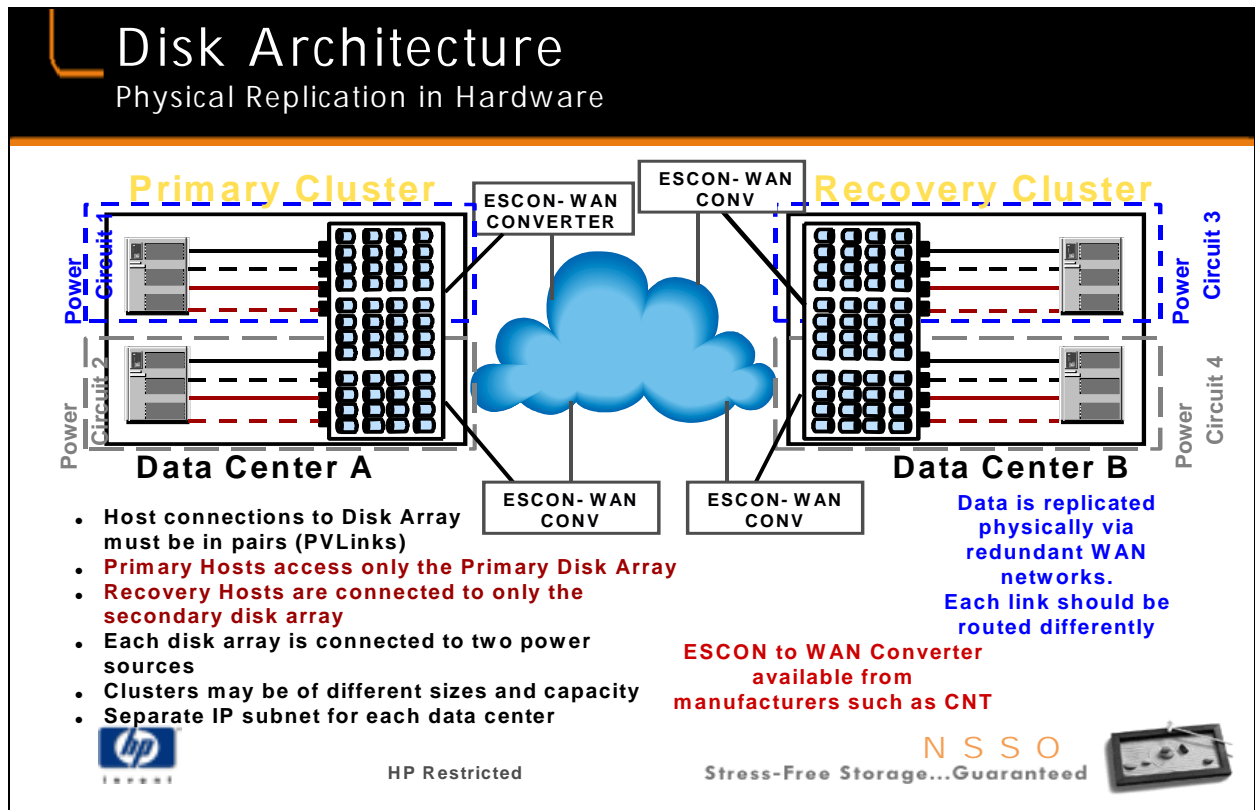
currently, only uni-directional failover is supported

includes scripts for EMC SRDF and HP SureStore E Disk Array XP Family of arrays for physical data replication

local failover still occurs within the primary data center

remote failover is used only when the entire primary cluster fails

network failover and failback defined and implemented by user



Looking at this slide, we see the disk architecture, what we show here is what is occurring at the primary and recovery clusters. The customer has either an EMC Symmetrix array or a XP Family disk array. And basically the data replication is going over a WAN, so this cloud in the middle here is representing the network. ESCON basically is limited, depending on which disk array we're talking about. With the XP Family of arrays, we're limited to 43 kilometers and the EMC Symmetrix is limited to 60 kilometers, and that's based on a number of ESCON repeaters that HP is currently supporting. So to get the additional distance, we basically have to convert from ESCON to a WAN technology and then back. There are some vendors that provide these protocol conversions, one that's in the Bay area is CNT which provides these converters, so that is something that needs to be planned in architecture.

ContinentalClusters supports the choice of physical data replication methods. Integrated data replication solutions that have been tested by HP with ContinentalClusters (HP will SUPPORT the Integration SCRIPTS only) are: HP SureStore E Disk Array XP Family of Arrays with Continuous Access XP or the EMC Symmetrix array with Symmetrix Remote Data Facility (SRDF), no other physical replication solutions are supported for use with ContinentalClusters.

Fence Level Data causes new I/Os to be refused if:

- either the PVOL or SVOL copy of the data becomes unavailable due to hardware failure

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Continuous Access XP

- the CA link hardware fails
- the entire remote XP256 fails

When used in the wide-area environment, a monitor program on local and remote host should run periodically to detect when the CA link is down. When this state is detected on a remote host (SVOL), the BC should be split and fence level data disabled so that the application may continue doing I/Os. When this state is detected on a local host (PVOL), the CA link should be split and the application may then be restarted.



## Port Configuration Detail

- Up to 16 CA ESCON Links (8 RCP & 8 LCP)
  - RCP is Local to Remote Port.
  - LCP is Remote to Local Port.

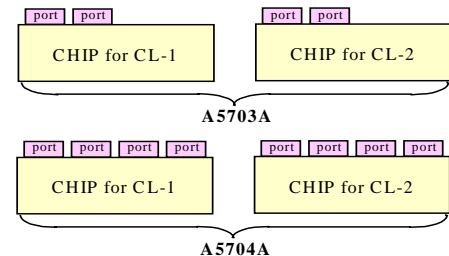
Two Type ESCON Package(CHA).

- A5703A: 2 port CHIP × 2.
- A5704A: 4 port CHIP × 2.

ESCON Link Performance is 17MB/s Maximum.

Up to 1023 CA Paired Volumes supported between two XP256s.

- The LDEVs of a Paired Volume must be the Same Emulation Type (OPEN-3, 8, 9, E).
- In case of LUSE and CVS, the LDEVs must be the Same Capacity and the Same Emulation Type.



HP Restricted

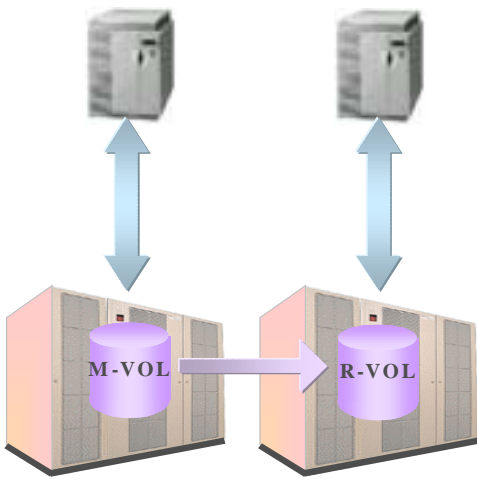
N S S O  
Stress-Free Storage...Guaranteed



This is some port configuration detail about the ESCON ports. These are the chip pairs that you will be using for interconnecting the two arrays with Continuous Access XP. We can have up to 16/64 total ESCON links between the two XP frames. We recommend eight in one direction and eight in the other direction. These ports are unidirectional; so if you define a port with source on the primary volume and destination on the secondary volume, you can only move data in one direction. And then we also recommend having the equal number of lengths which communicate from the remote array to the main array, so that if you're at the fail-back, you have the same kind of performance as you had for fail-over. We have two types of ESCON chip pairs. We have a two-port chip pair, which means that you will get a total of four ESCON connections for the pair. And then we also have a four-port, which will give you eight connections total for the pair. So, you would need two pair, of course, to get up to a number of 16 connections. A person has to be careful with this because these chip pairs compete directly for the same space that your other interface pairs would be using, such as SCSI (i.e. only in the XP256) and Fibre. So, if you need a large number of connections to your host, that'll limit the number of ESCON chip pairs that you can plug in to the XP Family of Disk arrays. ESCON is predominately a mainframe connection. It's performance is similar to Fast and Light SCSI, in that it is 17 megabytes per second, maximum; and that is a maximum number. Typical performance is typically, two-thirds of that, so somewhere in the range of 10 to 15 megabytes per second. All of the (LDEVs) can be paired up between


two XP Family of disk arrays, so there isn't any limitation on the number of paired volumes. In other words, you can go all the way up the number of (LDEVs) supported on the array (i.e. 1024 for the XP256 and 4096 for the XP512). The limitations are pretty simple on the type of pairing you can do. They need to be the same size. So that means open-three to open-three, open-eight to open-eight, and open-nine to open-nine. They can be a mix of RAID 1 or RAID 5. That means the RAID 1 can be continuous access to a RAID 5, and in most cases, I recommend that. You can go the other direction. In other words, you can have a RAID 5 backed up by a RAID 1. The only thing that matters is that they're the same emulation type and same size. In the case of using a CVS or custom volume size, again the rules are the same, the (LDEVs) must be the same capacity and the same emulation type.


## Continuous Access Remote Mirroring



- Provides remote data duplication, remote backup, and disaster recovery
- Allows point-in-time database backup
- Provides restore without latency
- Assures business continuance
- Simplifies workload management
- Remote mirror management via Remote Control XP
- Creates a remote mirror for all logical units in the array up to 43 kilometers in distance

M-VOL = Main Volume  
R-VOL = Remote Volume

 HP Restricted

N S S O  
Stress-Free Storage...Guaranteed 

M-VOL = Main Volume = Primary Volume = P-VOL

R-VOL = Remote Volume = Secondary Volume = S-VOL

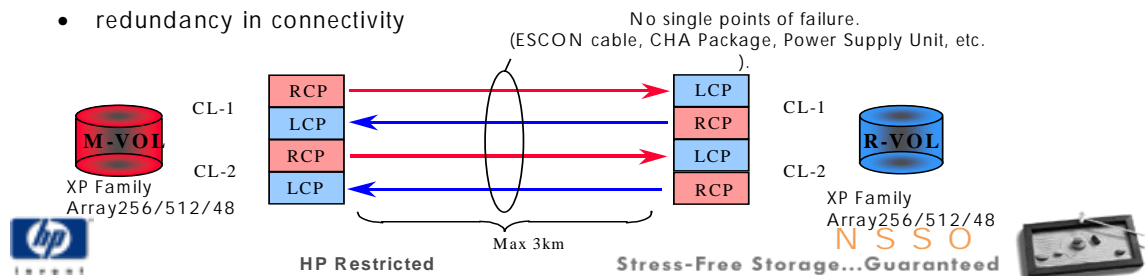
Data protection means not only having fail-safe hardware, but an adequate disaster recovery plan. A disaster recovery plan serves as a critical means of insurance against the intentional or accidental loss or destruction of data. Running HP's SureStore E Continuous Copy XP software on your HP SureStore E Disk Array Family of high-end arrays offer you the best available disaster recovery protection on the market. Continuous Access XP automates the duplicating of current production data from a primary site to a remote site. HP's SureStore E Continuous Access XP allows you to make copies of your data online so that your applications continue uninterrupted. Continuous Access XP operations are nondisruptive and allow the primary Continuous Access XP volumes to remain online to all hosts for both read and write I/O operations. Furthermore, in the event of a system failure or site disaster, the secondary copy of data can be invoked rapidly, allowing applications to be recovered with guaranteed data integrity. That means you can continue running your business without worrying about whether your systems will be unavailable. Continuous Access XP provides continuous access for open systems connected to the arrays. It is available for AIX, HP-UX and Solaris. Continuous Copy XP duplicates data from a primary system to a secondary system, and supports replication of data up to forty-three kilometers over ESCON fiber links. All LUNs can be copied and synchronously maintained. Additionally, copies can be split

for uses by other applications or backup. Split copies can be resynced back to the original primary copies as well. Continuous Access XP supports remote mirroring. For business continuance the mirror acts as business application “hot-site” and to accommodate disaster recovery the solution provides rapid availability with data current to last write. Data center migration capability allows data center move with minimum application outage. Workload management provides service through infrequent scheduled outages. Point-in-time database backup provides required database backup without the downtime associated with off-line backups. This can be accomplished with an application that has an on-line backup utility such as Oracle’s RMAN utility. Source volume for CA is called M-VOL and the copy volume is called R-VOL. Data from a M-VOL is copied to one R-VOL (i.e. 1 to 1 relationship currently). Data from an R-VOL cannot be copied to a M-VOL only re-synced. RCP and LCP, Remote and Local Control Ports. Business Continuance mirror copy acts as a business application “hot backup” Disaster Recovery with remote mirror copy of data rapidly available and current with workload management which provides service through scheduled outages. Point-In-Time Database Backup-backup databases without downtime. Data Migration-migrate data from one XP Family Array to another with no disruption in system operations. CA can be used for device or workload migration from one disk array to another or to load data onto new or upgraded disk array(s).

.

## Continuous Access HW Features

- ESCON links, a mainframe 17 MB/sec optical link
  - 4 or (maximum of) 8 ESCON interfaces per channel pair
  - Additional link extender reference products supported: ESCON repeater, ESCON wide area extenders
  - Automatic load-leveling and hardware failover is supported
  - Each link operates uni-directionally and can be a source or target port
- Installations have a min. of two RCP and two LCP ESCON ports on each XP Family Array
  - data movement in each direction
  - redundancy in connectivity



RAID Manager XP supports 2.4 and 7.3 GB LUN sizes and LUN extension using the LUSE software.

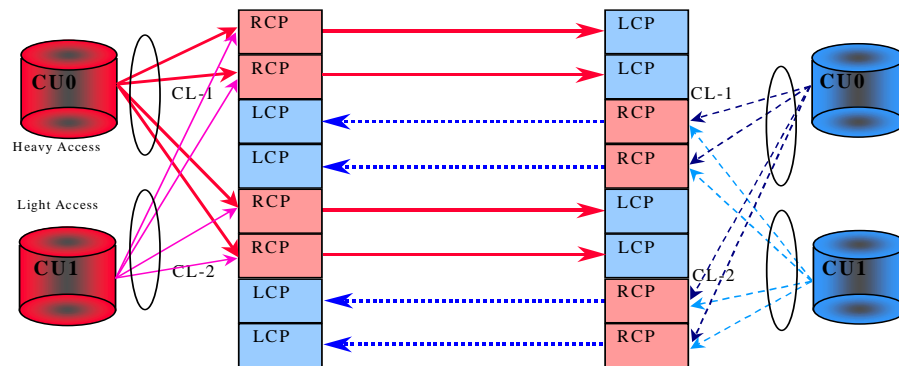
RAID Manager XP Software consists of the RAID Manager XP Manager, RAID Manager XP Commands, configuration definition file defined by users, and logging information. This unit of RAID Manager XP software is called an Instance. It is possible to activate more than one Instance by giving them an instance number. The RAID Manager XP instance runs in cooperation with the software in the UNIX server and the RAID Manager XP control of the HP SureStore E Disk Array. RAID Manager XP for NT and other Oss exist as well. Please check the SW Summary sheet and /or the XP Configuration Guide for more details to supportability. There is no Remote Console support at this time. Most UNIX users will implement this through UNIX commands. The initial setup of the path between the Control Units must be setup via the SVP and bi-directional swap must be enabled between the primary and secondary volumes. However, the RAID Manager XP operation commands can be done through UNIX commands. It is recommended – and the configuration throughout the array all recommend redundant ESCON links. The reason for this is that the array can automatically load balance and can automatically fail-over any link that fails. So, if we have, in this case, four total ESCON links, we have two going toward the secondary and two coming back toward the primary – if we lose one of those links, then it'll automatically fail all of the remaining transactions and links to the remaining link. And in fact, the entire group of 1023 pairs for the

XP256 and 4095 pairs for the XP512 can run on a single ESCON link. Performance, of course, would be fairly poor, due to the performance of the single link. Keep in mind that it is only "write" performance that we are affecting here, reads never have to be posted to the R-VOL. But the important note here is that load balancing is automatic. There is nothing that you need to do to worry about how many volumes are running on each ESCON link. It's completely handled for you and is internal through the array. If you were running less than three kilometers, then you can run regular old ESCON-style cabling (i.e. fiber). No need for repeaters or directors. ESCON will run at fairly high speeds up to three kilometers. This will allow you to approach 17 megabytes a second per connection. What you want to be careful here is, as you notice, because we're running on ESCON pairs, we are fairly resilient, which means if we lose power on the array, one – all of the connections on one half of the chip pair will stay up. Again, load balancing will take over and keep your transactions running back and forth in the case of a power-supply failure. So again, there are no single point of failures in the ESCON connections between the two arrays. If we use ESCON repeaters and/or directors, you can run up to a current supported maximum of 43 kilometers. This is done by simply extending the ESCON port and the ESCON traffic. We have another slide here that shows more detail how this actually wired and connected up. This is an example of the redundant ESCON configuration, so you want to pay special attention to when you install the repeaters and directors, that those are also on separate power supplies, so that a power-supply failure or a power failure on one grid or one area of your computer data center room will not interrupt your traffic between the two arrays. A lot of this is exactly the same as networks, so any of your people who are or have been familiar with connecting up networks will already have the expertise for this type implementation as well.

## Alternate Link and Load Balance

- Automatic Load Balance Across ESCON Links.
- Automatic Alternate Link Control In case of Link Failure.
- Example of a Load Balanced Configuration.

8 ESCON Links should be used for acceptable performance.




HP Restricted

NSSO  
Stress-Free Storage...Guaranteed



In this slide we are showing an example of eight ESCON links connecting two (2) XP256 arrays. This slide is to try and detail the load balancing and the automatic fail-over of a CA implementation. On the primary side, as traffic increases on the red wires, the load will automatically balance across these links, such that we won't be starving out any of the traffic on the lower volume that's in the red group as those I/Os will still continue. But it also means that all of the very heavy traffic off of the top M-VOLs, which is labeled "heavy access" which is distributed across all four of the links so that we make good use of those as well. So, even though we have a lightly-loaded M-VOL and a heavy-loaded M-VOL, we will get maximum utilization out of our ESCON links. The same is true for the response path coming back. We are not forced to have the replies come back on the same wire, so that when you acknowledge an I/O coming back from the secondary array, when it's received it in cache, it simply posts it back onto the ESCON link. The first array gets the acknowledgment that the write on the remote array has happened. Since this is synchronous, it's important for those replies to come back as quickly as possible because that's what's holding up your host-write request – the latency, completely, from front to back and then back to front (i.e. full circle).




This slide is a representation of typical ESCON performance. None of the material in this slide is necessarily proprietary, since it's purely based on ESCON performance. On the left-hand side is the number of minutes. The access across the bottom is the size of the copy and the number of LDEVs that that amount of data is sitting on. The right-hand side is the number of ESCON links. But the top line it appears to go up into the right at a 45-degree angle – is a single ESCON link. Then, the sort of the purplish-pink one is two ESCON links, and yellow is three and so on, for a total number of 16. In each case, they're doubling the number (i.e. 1,2,4,8,& 16). So, what we have at the bottom is a total of 16 links noted by the blue line at the bottom of the slide.

What you can see here is that every time you double the number of ESCON links, you increase the performance, which means that the total time to due the copy operation decreases. So, in the case where you're running the maximum number of links, we're down to about 500 minutes to copy 2.4 terabytes. And these numbers are purely based on the performance of the (S-CON) links. The graph – at the time this test was run, the configuration was 12,000 rpm disks and RAID 5 array groups, with open three and a cache size of 2GB. And in most cases, that configuration is irrelevant to the performance of the ESCON links. In other words, you would've achieved a similar performance with almost any drive configuration.



But the important thing to note is the more ESCON channels, the better performance. And even though they do compete with your chip pairs, you may need to load the machine with ESCON chip pairs if you need a very high aggregate write rate. You need to have at least a total of two pair for the minimum configuration for HA.

One of the things that also applied to Business Copy and EMC applies here, and that is the initial copy is not typically done with the ESCON links. Most of the initial copies are done through tape or through some host connection that is much faster. It is for this reason is why under RAID Manager, we do have the functionality for doing a "pair create, no copy." This means that you can create the images over on the secondary array, make them identical, and then, once that's completed, you can tell the arrays to create the pairs, but don't copy any of the data. And then, from that point on, they will be held exactly in sync. Every I/O will be identical. This would get you around the initial copy time of, say 4.8 terabytes, you would be looking at basically 80,000 minutes to do the copying. And that's purely limited to the performance of the ESCON links.



Continuous Access XP software provides fail over commands and operation commands in order to allow mutual hot standby in cooperation with failover high availability software such as MC/ServiceGuard, First Watch and HACMP. Continuous Access XP supports HA clusters located over ESCON distances for disaster remote clustering. If the primary processor node should fail, the secondary node can do an Continuous Access XP command which transfers the Primary Volume assignment to the copy pair attached to the secondary node. If the P-VOL on the primary node is also lost, the standby node operates with the S-VOL in simplex mode. The takeover commands provided by Continuous Access XP can be activated by shell scripts invoked by the UNIX fail-over middleware software.

Note: The synchronous mode is mandatory to be supported on HA environment. Management:


Host command line interface via RAID Manager XP software GUI on SVP, Remote Console, and SAN Management Console Use Service Processor (SVP) and Remote Console (RMC) Configure mirrored LUN pairs and provide control of remote mirror operation through the out-of-band device LAN link.

Disaster recovery, Continuous Access XP generates sense information when a CA volume pair fails to maintain synchronization. Neither CA nor Remote Control XP provide host fail-over software. This information needs to be transmitted to the remote site by host fail-over software such as MC/ServiceGuard.

## Continuous Access XP

- **Operation (synchronous):**

- each write I/O (1) is confirmed by the primary site array with host (4) after its copy is sent to (2) and confirmed received by (3) the remote site array:



- **What Continuous Access XP (sync copy) offers:**

- the highest level of remote mirror data currency possible

- **Key Solutions:**

- Seamless, reliable Campus/Metropolitan Disaster Recovery/Clustering



HP Restricted

Stress-Free Storage...Guaranteed



This is a simple slide, but a complex set of sequence and events occur here. And this details how the write from the host actually occurs. And this is very important to know when trying to determine the performance and the kind and the number of links you would need between the two arrays because we're running it at ESCON speeds and distances. So, from a rather high-level view, here is how a write process occurs for customer's that are deploying CA. The host will generate a write request – this means that the host has a block of data to write and this could be any size write, let's pick an 8K value for the sake of argument for our example. The host would write that data into the cache on the primary XP256. Once the data is safely in the cache and it's duplexed across both halves, the primary array sends the block of data to its remote XP256 through the ESCON link. The I/O from the host is held up, or literally stopped, until we get this write over to the secondary array and it's safely in the duplexed cache on that array. Once the write is safely in the cache on the remote array, the remote XP256 replies back to the primary that it has received the data. The primary array then sends the write complete back to the host.

This means that for every write, you will be limited by the performance of the interconnect between the primary and secondary XP256. This doesn't sound very complicated, but it can be. For instance, let's take the reads. Reads are unaffected by this operation. So, if you have a database that needs to read, say 100 megabytes a second, there's no impact on reading with Continuous Access.

But let's say that your database would like to write to the database – at 100 megabytes a second. Well, you can see almost right away that that's not possible. If we are running four ESCON links, and let's assume that we're running 10 megabytes a second out of each of those links, the maximum that we could ever imagine our write performance to be would be 40 megabytes a second. So, that would literally throttle or limit the write performance of your database dramatically. And this is critical, because this number slows down over distance. As we approach the full 43 kilometers using ESCON as the interconnect, you can slow down to as much as 3-5 MBs/ ESCON link. So, let's take the five number, just for sake of argument. So, let's take the example of having a distance that limits our ESCON throughput to 5MB/s and we have four (4) links configured, now we are now down to a 20 MB/sec aggregate write rate between the arrays. This would significantly throttle your database even further. So, what you need to do to help improve this performance bottle-neck is either increase the number of ESCON links between the two arrays, or determine what data within the primary array absolutely has to be recoverable (i.e. mirrored on the secondary array). In other words, it has to be mirrored and continuous accessed to the other volume. In the case of running against the database, you probably do not have much choice here in that probably the entire instances file system or raw file system would have to be mirrored.

Because this is synchronous, there are latency concerns here. This means that the database, in most instances, will lock down some region of a table and hold that table, or rows in a table, locked, while it attempts to do this write. This means that the write will bounce from the primary to the secondary, and then back again, which could easily add tens of milliseconds to the overall I/O transaction, which means that you're internal table in the database will remain locked down for a fairly significant period of time. And on a N-Class or V-Class size machine, many processes may block at that point, waiting for the lock to be released. So, when asked to carefully plan the performance between these two frames, the impact on your database may be more significant than you might have implied by, just simply, the throttling of the write.

## Continuous Access XP (Sync-CA)

- The local (M-VOL) and remote (R-VOL) volumes are always synchronized - Secure But Slow
- Write I/O transactions do not complete to the host application until both the local and remote XP Family Disk Arrays subsystems are updated, that is, until the secondary volume acknowledges receipt of data in its cache
- Recovery from site failure is possible because the volumes are always synchronized
- Performance impact: I/O performance is often paced by speed of the inter-subsystem links (ESCON)
- Delay increases as the M-VOL to R-VOL distance increases



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed




Base product provides media, documentation, and license to use on one HP SureStore E Disk Array Family Array with a raw capacity not to exceed 1 TB (terabytes). The base product consists of software keys for the SVP and Remote Console PC, along with a Software Certificate and product manual (order number B7905-90900). The keys will ship on diskette media. HP SureStore E Continuous Access XP is a disaster protection feature of the HP SureStore E Disk Array Family of Arrays. It is an optional software package that provides data replication across two sites in an open systems environment. It allows data on a specific SureStore E Disk Family Array to be copied synchronously to another disk array subsystem up to 43 km (26miles) away via a dedicated ESCON (Enterprise System Connection) physical link. This feature, called remote mirroring, allows for implementation of disaster recovery because the mirror acts as a hot backup and is always synchronized with the primary copy. In a disaster recovery situation, a mirror image allows for rapid availability of data with information current to the last write. In addition, HP SureStore E Continuous Access XP provides fail-over commands and operation commands that allow mutual hot standby in cooperation with fail-over or HA software such as MC/ServiceGuard. The 43 km distance is accomplished using a combination of single mode and multi mode cables, ESCON directors and ESCON repeaters. Continuous Access XP provides for a push copy operation between a primary and a target communication pair. The primary and target pair can exist within one subsystem, however the more common scenario, the copy will take place

between a primary subsystem and remote subsystem target pair. Multiple subsystems copies, up to 16 concurrently for XP256 and 64 for XP512 (16/64 ESCON links), either locally or remote, can be achieved. Both ASCII and EBCDIC system architectures are supported.

## Continuous Access XP Extension

- **Operation (asynchronous):**

- each write I/O (1) is immediately confirmed by the primary site array with the host (2) while its copy is independently sent to (3) and confirmed received by (4) the remote site array:



- **What Continuous Access XP Extension (async copy) offers:**

- Low response time impact on primary-site operations
- Remote sequence-stamping to ensure continuous remote mirror consistency
- Efficient, link-optimized remote copy operations

- **Key Solutions:**

- Cost-effective Continental-distance Data Replication/Migration & Disaster Recovery/Clustering



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



This slide does the description, or goes through the write process for the asynchronous copy mode of Continuous Access XP. And of course, asynchronous differs from synchronous, in that the remote copy operations sort of happen asynchronously with the write and confirm operations happening between a primary host and a primary array. You therefore have this de-coupling effect that's happening with the remote array, that allows for longer-distance solutions, and the opportunity to have sort of a low-impact – performance impact solution, in terms of the primary site operations that you would have with synchronous operations. I/Os going to the remote site, even if out-of-order or lost in flight are corrected for. This is a very BIG point here in the context of HP's CA over IP solution. When we get into why we think HP's CA over IP solution is uniquely positioned in the marketplace to delivery disaster recovery-tolerant solutions – in that, Continuous Access XP Extension, asynchronous mode, does deliver IO consistency at any time. So it's not a question of being exposure in any moment that you would not have recoverable data, let's say, on a remote site. So this is the foundation upon which HP believes a lot of the value inherent in CA over IP is built on.

## Continuous Access XP Extended (Async-CA)

- Async-CA is an optional extension to the XP Family of Array's CA base product
- Async-CA allows longer distance support while guaranteeing the data sequence of writes across ESCON
- Writes on the local side complete immediately, before data reaches the remote array. No write verify is required
  - much less of a performance bottleneck on the local side
  - safeguards are built in to handle data packets that may not arrive at the remote site, or may arrive out of order



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



PRODUCT NAME: HP SureStore E Continuous Access XP-Extension.

KEY PRODUCT NUMBERS: B9325A

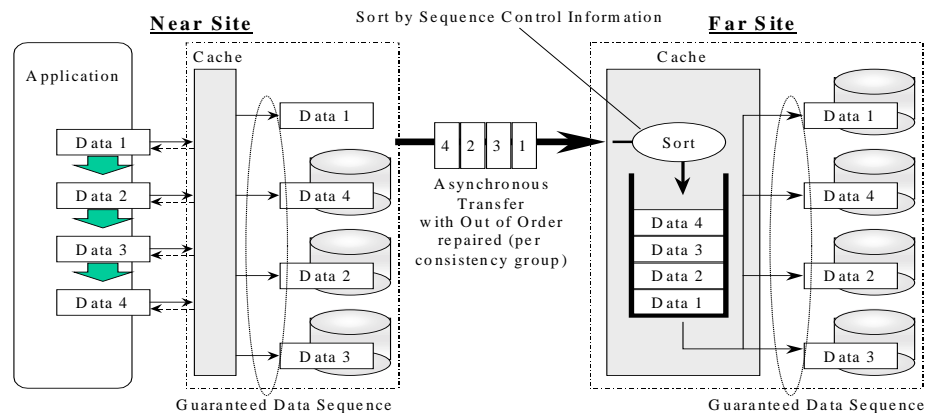
CA asynchronous update copy mode provides high performance data updates to the secondary site. This mode employs sequence numbering information for all "in-flight" data and tracks the order and completeness of the data received on the secondary disk array. It does not depend on the order of receiving write I/O requests from the host. HP SureStore E Continuous Access XP- Extension is an advanced "disaster protection" feature recently made available for the HP SureStore E Disk Array Family of arrays (Feb. '00). It is an optional software package that provides data replication across two sites in an open systems environment. Extending the basic synchronous Copy Mode capabilities of this product, HP SureStore E Continuous Access XP Extension enables Continuous Access XP to operate in asynchronous and/or synchronous Copy Mode to another disk array subsystem up to 43 km (26miles) away via a dedicated ESCON (Enterprise System Connection) physical link. It enables Continuous Access XP/XP256 copy a local volume asynchronously to a paired volume on another, secondary disk array subsystem via a dedicated ESCON (Enterprise System Connection) physical links (using extenders for long distance connections) or via WAN links using ESCON converters. Whereas each write operation must be executed and confirmed at the remote subsystem before final write confirmation is issued by the primary subsystem in a basic (i.e synchronous) Continuous Access XP implementation, Continuous Access XP



Extension enables asynchronous operation. This means that primary subsystem write confirmation occurs independently of secondary subsystem copying functions. Data integrity at the secondary subsystem is insured via Sequence Control Algorithms, which allow data operations sent asynchronously (and in potentially varying sequence) to the secondary subsystem to be 'reassembled' in the correct order. Overall primary subsystem response times are significantly improved due to the asynchronous nature of the Local/Remote data transfers. Both Continuous Access XP (B9320A) and Continuous Access XP Extension (B9325A) are required to achieve full synchronous/asynchronous capability. Continuous Access XP Extension may also be purchased as an upgrade for customers currently using Continuous Access XP.

# Asynchronous-CA Control

## Asynchronous-CA Control



At far site, I/O #3 will be held until I/O #2 shows up



HP Restricted

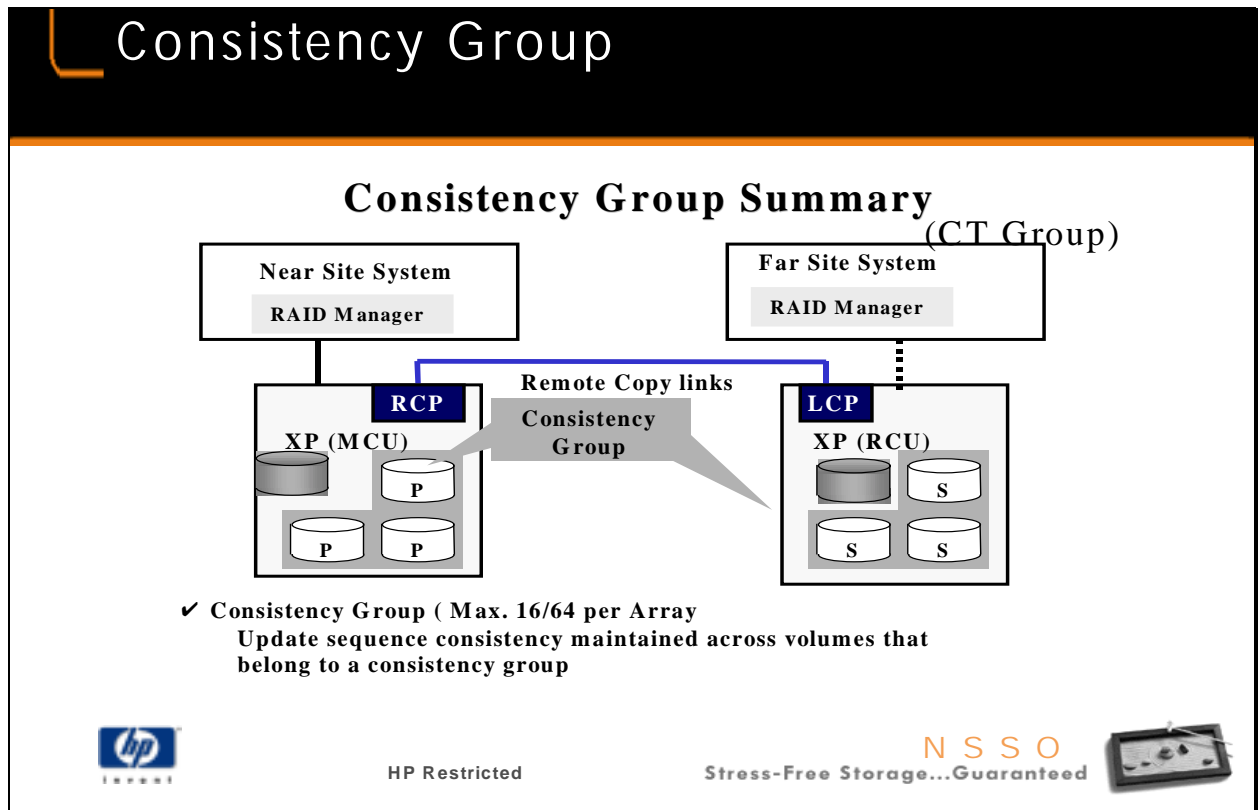
NSSO  
Stress-Free Storage...Guaranteed



This is the basic async-CA architecture showing I/O re-ordering per consistency group (if necessary):

Each update write has a four-byte data sequence number to ensure re-assembly of the data at the remote disk array correctly.

The I/O sequence ordering function of async-CA is very important. Without that, various paths of differing speeds (and routes) could allow I/Os to arrive out of order. The picture illustrates how I/Os are sequenced by number and then re-assembled (if necessary) at the far end. This is a key/strong point. Two concepts which will be discussed are the "sidefile" and "consistency" groups.



The CA asynchronous update copy mode manages application related volume groups in consistency groups. If intermediate data is lost on a transfer, the secondary disk array suspends consistency group pairs after a timeout. Consistency group (CT group) - unit of granularity for I/O ordering assurance. A consistency group is a grouping of LUNs that need to be treated the same from a data consistency (I/O ordering) perspective.

Notice that the unit of granularity for I/O ordering assurance is the Consistency Group (CTG). A Consistency group is a grouping of LUNs that need to be treated the same from a data consistency (I/O ordering) perspective. A consistency group can contain one or more volume groups in the RM config file and can be specified in the pair-create command (RM 01.03.02 and above).

# General Async-CA Concept



## Async-CA Specifications Summary

**Configuration**  
Near Site 1 Unit/Application : Far Site 1 Unit/App (i.e no NxM yet)  
43km via ESCON (Max. 16 Links, 50km via DWDM)  
For Longer Distance use Channel (ESCON) Converters

**Data Sequence Guarantee**  
Asynchronous Data Transfer, and Guaranteed Data Sequence  
Guaranteed Data Sequence per Consistency Group  
Guaranteed Data Sequence at the completion of Resync-Operation

**Asynchronous Volumes**  
Up to 1023/4095 Volumes (1 for "Command Device")  
Open-3, 8, 9, with Extended LUSE /CVS/DCR

**Control**  
RAID Manager XP  
Remote Control XP  
SVP (used by CE only)

 HP Restricted N S S O  
Stress-Free Storage...Guaranteed 

Notice that data consistency is assured only within the granularity of a Consistency group (explained below) which goes from exactly one XP Array to exactly one XP Array. These are some Async-CA specs.

# Async-CA Guarantee

## Async-CA Data Sequence Guarantee Specs Summary

### Policy

- ✓CT Group Data Sequence is managed by both the XP256/512 MCU & RCU

### Data Set Control

- ✓Each ESCON Write has 4 bytes of Data Sequence # information included
- ✓Data Sequence #s are in order (within a CT group)
- ✓Data Sequence #s are added to the Status reply at Write completion
- ✓Transfer of Data to RCU may be out of order
- ✓RCU writes Data in Cache after a Sort via Data Sequence # and CT group

### If intermediate Data is lost in Flight?

- ✓RCU suspends CT group pairs (after a timeout)



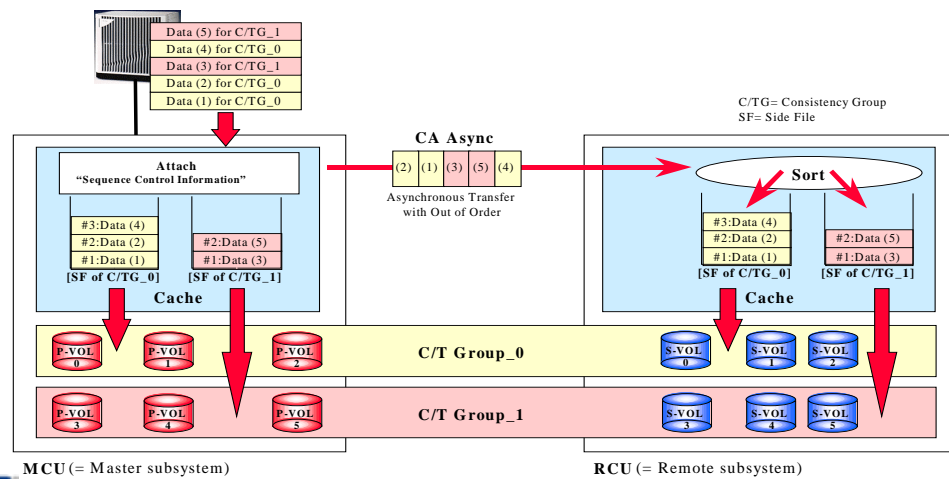
HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Different CTGs Are Handled Differently

### *Guaranteed Updated Data Sequence per CTG*



HP Restricted




NSSO  
Stress-Free Storage...Guaranteed



## CTGs

### *More Specifications for Consistency Groups*

- ✓ C/T Group guarantees the Updated Data Sequence .
- ✓ Each CA Async paired volume must belong to some C/T Group.
- ✓ A CA Async paired volume can belong to only One C/T Group.
- ✓ Up to 16/64 C/T Groups (#0-F) per XP256/512
- ✓ One C/T Group can include up to 1023 CA Async paired volumes.
- ✓ A C/T Group can't belong to multi MCU&RCU pairs. (MCU : RCU = 1 : 1)
- ✓ Every P-VOL in the same C/T Group must belong to the same XP256 (MCU).
- ✓ Every S-VOL in the same C/T Group must belong to the same XP256 (RCU).



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Continuous Access XP Extended

- CA asynchronous mode provides high performance data updates to the remote site at unlimited distances (WAN)
  - ESCON extenders and converters and utilizing common-carrier links allow transfers at distances greater than 43 kilometers
- Consecutive updated records are combined into a single write operation to reduce ESCON interlocks



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



CA asynchronous update copy mode works with ESCON extender and converter technologies to allow transfers at distances greater than 43 kilometers. To reduce the overhead associated with CA remote copy operations, the disk array uses a special write command which is allowed only for CA initial and update copy operations. This CA command transfers the control parameters and the format data, which includes consecutive updated records in each track, using a single write operation. This capability reduces interlocks on the ESCON interface.



## Sidefile Area

(DCR = Cache LUN)

- ✓ Side File % = Side File Area / (Total Cache - DCR area)
- ✓ Side File has control points 30% & Threshold

Diagram labels: Total Cache - DCR, DCR, Side File Area, High water mark, 30%, Side File Threshold, Default 50%, CT=0, CT=1, CT=2.

High water mark (30% , Side File Full)	Start of “Delayed Response”, Host side slow down
Side File Threshold (Timeout)	Slow Down gradually at 30% , 35% , 40% , 45% ... Host Write waits for less than Threshold, “Wait until NOT Full” Monitors time of over Side File Threshold Suspend (Default : 90s)

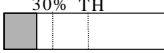
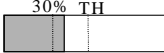

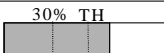
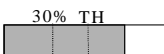
HP Restricted


NSSO

Stress-Free Storage...Guaranteed

CA asynchronous update copy mode employs a side file to maintain the ordering of write I/Os prior to transmission across the ESCON link. The side file is a competitor for cache resources and varies in size automatically, depending on CA asynchronous write activity. What happens when writes to the primary volumes come faster than they can be transmitted? Side file size expands, cutting down on cache available to the primary host I/O. Side file is a competitor for cache resources. Plan for enough cache to handle write I/O peaks and avoid system performance degradation. Initial recommendation for customers using async-CA is to order 1.5-2x the cache for a system without CA. This design allows the side file to be bounded only by the total cache size minus the area dedicated for DCR (cache-based LUN). You must account for this when configuring the array's cache size. Sequence numbering information tracks the order and completeness of data sent to the remote site, each update write has a four-byte data sequence number. Asynchronous writes can be received out of order, with the sequence number tracking the order of the updates. They are assembled in time order at the remote XP Family Array. When a write is done on a primary volume, a time-stamped copy of the record is stored in cache in a side-file. The side-file has control points: High Water Mark (HWM) at 30% of cache, Threshold – Side-file threshold values limit asynchronous remote copy cache from 30-70% (default is 50%). This value can be set from the SVP. Side-file has top priority for taking cache slots.

## Sidefile Behavior and HWM/Threshold


Side File Condition	DKC Control	Raid Manager States
less than 30% 	Normal Async Transfer	PAIR : PAIR
30% - Threshold 	Async Transfer with Host Slow Down, "Delayed Response"	PAIR : PAIR
over Threshold 	Async Transfer with Host waiting until "less than Threshold"	PAIR : PAIR similar to "Sync"
over Threshold after timeout 	Start of Bitmap control after timeout	PSUE : PSUE
over Threshold after timeout 	Suspend with Bitmap	PSUE : PSUE



HP Restricted

NSSO

Stress-Free Storage...Guaranteed




This gives more specifics on how behavior will change inside/outside the 2 control points (HWM/Threshold):

Side file Definition - an area of XP Family Array cache used to store Data Sequence #, Record location, Record length and queued control information before transmit over the ESCON link. If the side file reaches the High Water Mark (30%), write I/Os will have waits added. If the side file continues to fill (e.g., temporary WAN link outage) the I/O that crosses the threshold (50%) will be waited as long as necessary until room for more side file data in cache is avail. If this does not occur within a set timeout (1-255 seconds, default 90) the ESCON link is suspended. After bringing the link up, the volume pair is re-synchronized.

## Example of Link Down or Delayed


<u><b>Event</b></u>	<u><b>Host side Write I/O</b></u>
less than Side File Full 30%	Normal Async Writes
30%	Write I/O with delay 500ms
35%	Write I/O with delay 1 second
40%	Write I/O with delay 1.5 seconds
45%	Write I/O with delay 2 seconds
Side File Threshold (Default 50%) Over Threshold	Write I/O is waited for "less than Threshold" <b><u>"Wait until NOT Full"</u></b>
Monitoring Term	One by One depending on enough space
Timeout (Default 90 sec)	Queued Write I/Os to Bitmap
After Suspending	Suspend status receiving host I/Os

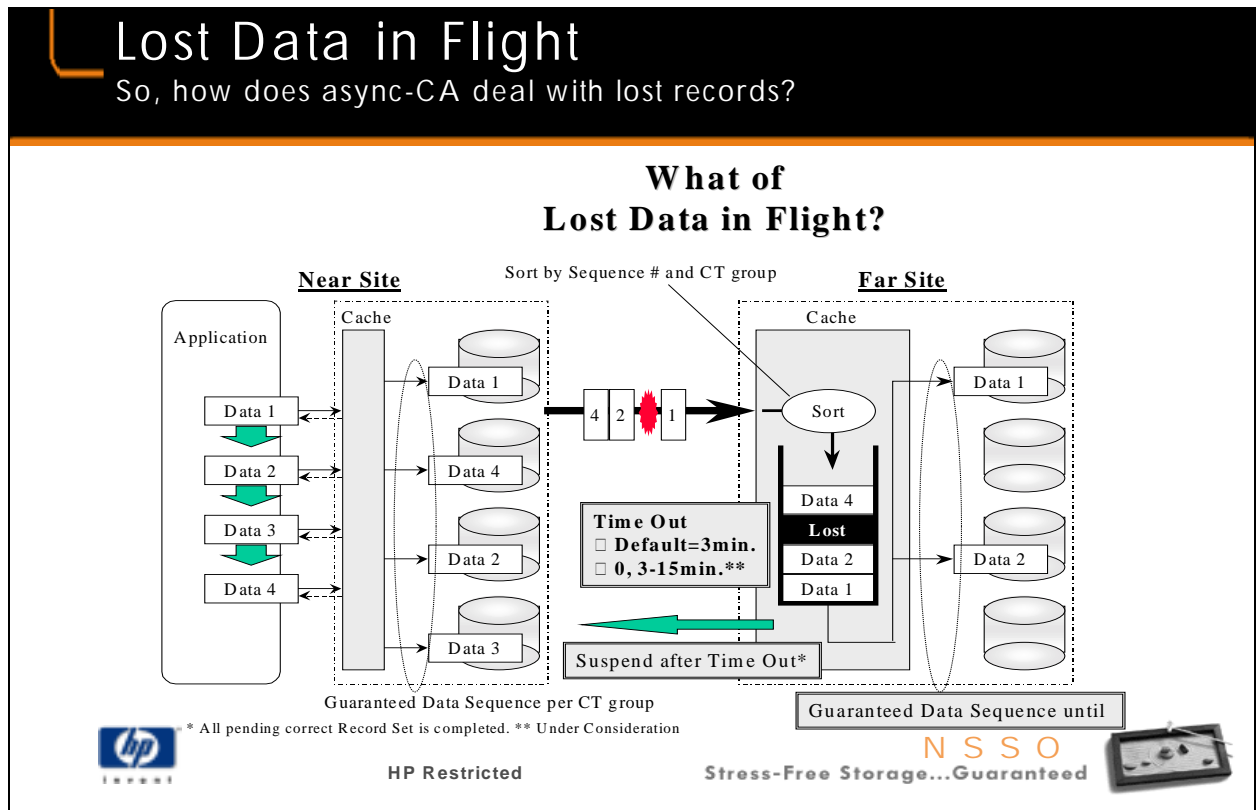
A normal situation should see few if any delays.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed





Writes 1 and 2 make it to the write cache and the remote disk array. Write #3 has a timeout show up. If the timer pops, the link is suspended until operator intervention (pair-resync command) is issued. Converters like CNT do automatic re-transmission which should help out in this type of situation. While the link is suspended, both sides start using a bit map (#4 in remote array goes into a bit map, all new and unsent side file writes go into an primary array bitmap. After the reason for the link problem is rectified (and possibly after a data consistent BC copy is made at the far side) a pair-resync command causes an non-ordered bitmap copy to bring the pairs back into synchronization. The suggested BC mirror is because a site failure during this non-ordered copy time would leave the remote data non-consistent.

## CA Features and Terminology

- A **Main Volume** can have only one **Remote Volume**
- A Main Volume and its Remote Volume must be of the same emulation type, Open-3/9/8/E
- Main Volumes and Remote Volumes can be any combination of RAID-1,0/1 and RAID-5
- Different RAID types must be on different ACP pairs (XP256 only)
- If the Main Volume is a customized (CV) or extended volume (LUSE), then the Remote Volume must be of equal or greater capacity
- Up to 1023 LDEVs in each XP256 and up to 4095 LDEVs in each XP512/48 disk array subsystem can be assigned to continuous access pairs



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



An array subsystem has 1024 LDEVs.  
One LDEV is reserved as the open system command device.  
CA operations involve the main control unit (MCU) and the remote control unit (RCU):

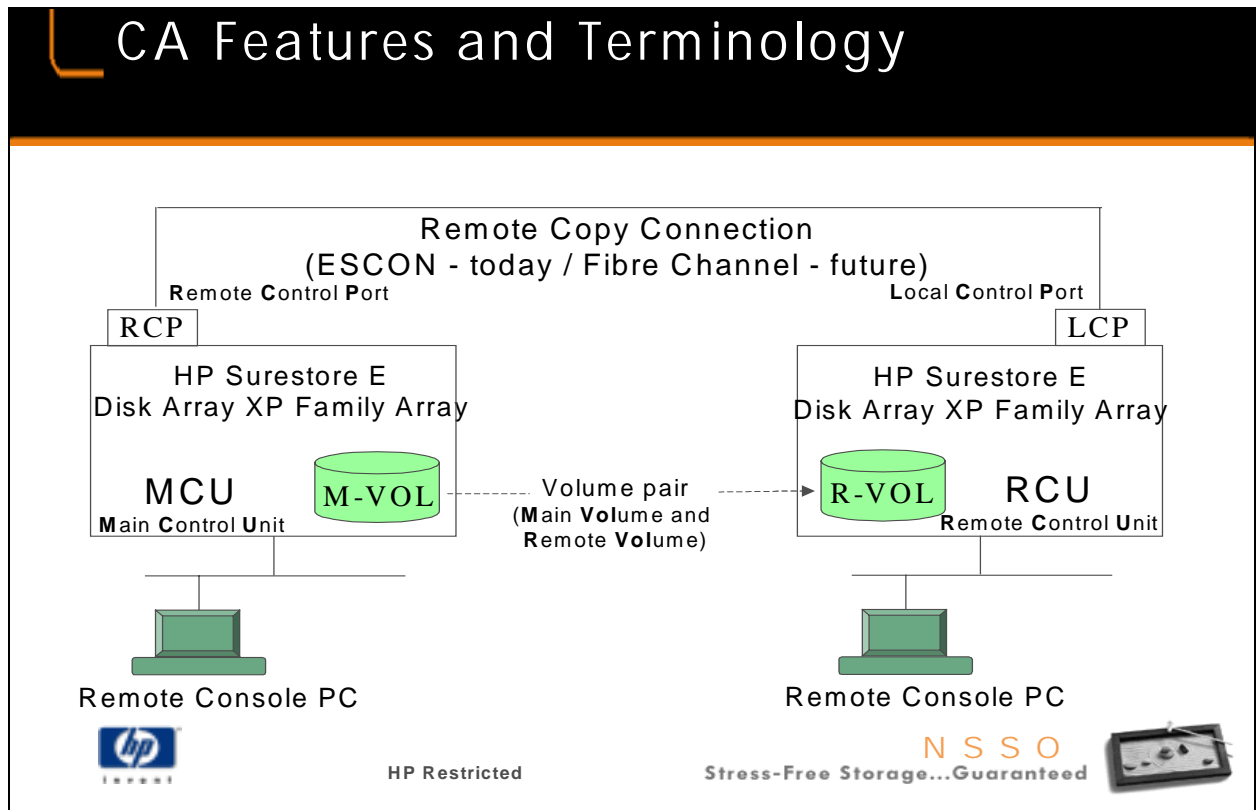
The MCU is the control unit (CU) in the main disk array which controls the P-VOLs of the CA volume pairs. The MCU communicates with the RCU. The MCU controls host I/O operations to the CA P-VOLs, as well as the CA initial copy and remote update copy operations between the P-VOLs and the S\_VOLs. The MCU also manages the CA pair status and configuration.

The RCU is the CU in the remote disk array which controls the S\_VOLs of the CA volume pairs. The RCU assists in managing the CA pair status and configuration (for example, rejects write I/Os to CA S-VOLs). The RCU executes the remote copy operations issued by the MCU. A second remote console PC must be in operation at the remote site. The RCUs must also be attached to an open system host to allow sense information to be reported in case of a problem with a remote volume or disk array.

The MCU and RCU can be defined separately for each CA volume pair. The CU can function simultaneously as a MCU for one or more P-VOLs and as an RCU for one or more S-VOLs, provided the remote copy connections and serial interface connections are properly configured.

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Continuous Access XP

The disk array CU provides four CU images (CU 0,1,2 and 3), with each CU image controlling up to 256 logical devices. CA operations can be performed on 1,023 logical volumes in all four CU images.



#### M-VOL & R-VOL:

M-VOL (the source) and R-VOL (the target) must be of the same emulation type (Open-3, Open-E, Open-8 and Open-9), but need not be of the same RAID type (Note: Different RAID types have to be on different ACP pairs within a single subsystem). If the main volume is customized or extended, then the remote volume must be of equal or greater capacity. An M-VOL may have only one R-VOL.

#### RAID Manager:

Provides a command line interface to control remote mirror operation through the in-band Fibre Channel or SCSI link.

#### SVP and Remote Console:

Configures mirrored LUN pairs and provides control of remote mirror operation through the out-of-band device LAN link.

#### ESCON Links:

4 or 8 ESCON interfaces per 'CHIP' pair (maximum of 16 links per XP256 and 64 links per XP512 array subsystem supported). Additional link extender reference products supported for flexible link configurations: ESCON repeater; ESCON-wide area extenders. Automatic load-leveling and hardware fail-over is supported. Each link operates uni-directionally and can be configured to be either a source or target port.

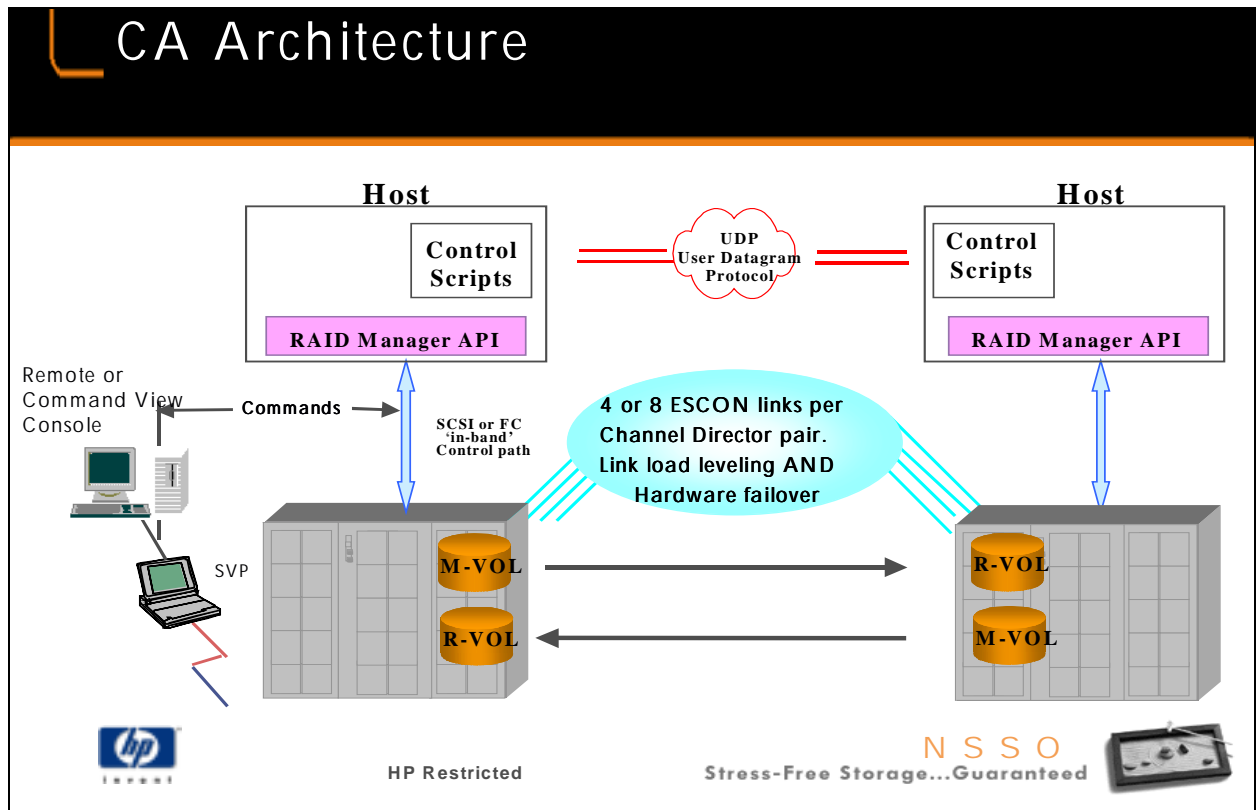
#### RAID Manager to RAID Manager communication:

Any host-host link that supports a TCP/IP socket-socket connection. Used by the


SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Continuous Access XP

RAID manager to communicate with its counterpart on the remote subsystem for control.





Copies a logical volume onto a separate subsystem over a unidirectional ESCON link. Can have R-VOLs and M-VOLs on an XP256 at the same time. ESCON links: 17 MB/s links. Array FW does automatic load-leveling with hardware fail-over (eg. if one ESCON link fails, the FW will migrate all traffic to the remaining links). 4 or 8 links per CHIP pair. Can have up to 1023 mirrored volumes (1023 LDEVs) between the 2 locations with the XP256 and (4095 LDEVs) or mirrored volumes for the XP512. With the XP512 you can now have 65 ESCON paths versus 16 with the XP256. There should be an SVP and Remote Console at each site (picture only shows configuration at one site). Control LAN can be any link capable of carrying TCP/IP traffic. In MC/ServiceGuard/MetroCluster environments, this could be the heartbeat LAN. Used by Continuous Access software for control and synchronization between the 2 sites. S-VOLs & P-VOLs are standard RAID-1 or RAID-5 volumes. Possible to intermix (for example, P-VOL can be RAID-1 and the S-VOL can be RAID-5). Can be combined with Business Copy volumes.



This slide gives detail on the directory and repeater configurations that are currently available and supported. And in this case, these examples are IBM-style directors and repeaters, but any ESCON director or repeater would function similarly. Please see the latest version of the XP256 Ordering & Configuration guide for supported manufacturers of this equipment. The total distance in either case is still going to be 43 kilometers or less. And we have a limit to the distance from the array to the first director, and that same limit has to be maintained from the last director to the array. We can run up to 20 kilometers between a director and a repeater. The other configuration shows two directors with repeaters on the end. Again, the distance is exactly the same. The only reasons you might do this has to do with the configuration of your machine room and what kind of services you have out in the maintenance areas in the infrastructure of the streets, or wherever you're going to be running these fibre optic cables.

It should be noted at this point that these pieces of equipment are typically not cheap. And renting or leasing or installing the dark fiber required to run ESCON over these directors and repeaters is fairly expensive. You can literally spend millions of dollars putting in the long-haul or the dark-fiber connection between your arrays. Because of this, it makes the inter-connection of the arrays the single largest expense your customer will go through in setting up a continuous-access or disaster-recovery site. And in the case of some of these cities that do not have the dark fiber infrastructure, they'll be out there with a backhoe and

digging up the street to lay the dark fiber. If you're in a city, a large metropolitan area that already run the infrastructure, which means that the dark fiber is already in the ground, all you have to is, then, lease it from them, but that can still be fairly expensive.

The maximum distance between primary and target disk arrays is 43 km.

\*\*\* Per Hitachi HORC training class: Only multi mode ESCON cables up to 3 km long can be attached to the SureStore E Disk Array XP256.\*\*\* The information in the rest of this section is from HDS. Please confirm this, as SSP Product Structure shows the following ESCON fiber optic cables: A5752A:

7 m A5752A #001

13 m A5752A #002

22 m A5752A #003

31 m A5752A #004

46 m A5752A #005

61 m A5752A #006

IBM 9032/9033 ESCON directors (ESCDs) or 9036 ESCON repeaters are required with disk array subsystems more than 3 km apart. The IBM 9032/9033 ESCON director supports an extended distance facility (XDF). XDF uses single mode ESCON cables up to 20 km long. The IBM 9036 ESCON repeater supports single mode to single mode connection or single mode to multi mode. In order to locate subsystems further apart than 9 km, the XDF connections provided by the ESCON directors or ESCON repeaters are required. The following diagrams show examples of distance configurations:

Key: Symbol Meaning

-----	-----
P	Primary SureStore E Disk Array XP 256
T	Target SureStore E Disk Array XP 256
=	Multi mode cable up to 3 km long
-	Single mode cable up to 20 km long
D	IBM 9023/9033 ESCON director
R	IBM 9036 ESCON repeater

P===D-----D===T  
P===D-----R---D===T  
P===R-----D-----D===T  
P===R-----D-----R===T  
P===D-----R-----R===T

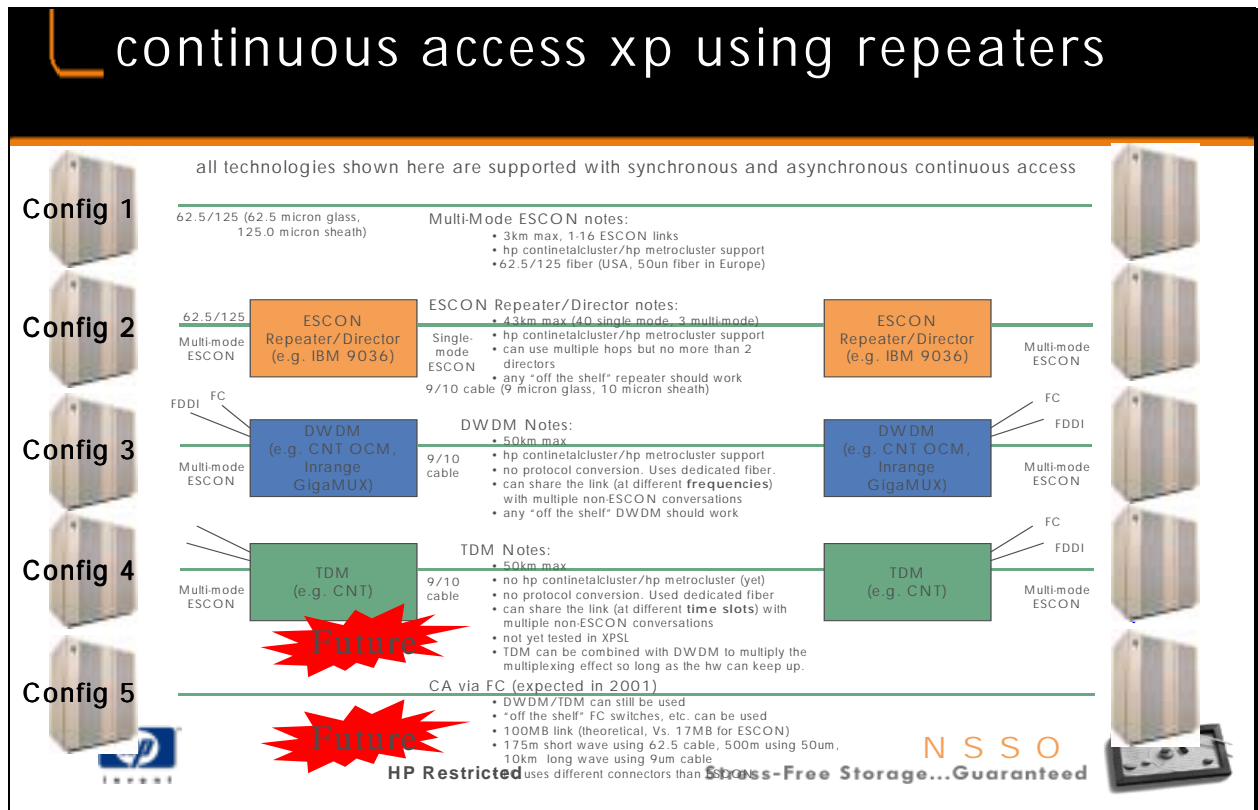
## Existing Remote Interconnect Options

HP Restricted

- For all options: dedicated infrastructure is used exclusively by the remote mirroring activity (for DWDM, shared with other SAN-oriented apps)
  - ESCON: dark fiber
  - WAN/LAN: Switched Circuit Virtual Private Network (deterministic/non-deterministic pathing)
  - DWDM (Dense Wave Division Multiplexing): share infrastructure with other applications (FC, FDDI)
- Pros
  - ESCON & deterministic WAN/LAN: constant reliable bandwidth/latency
  - ESCON: MetroClusters support, works with 'standard' ESCON repeater/director technologies
  - DWDM: works with 'standard' DWDM technologies
- Cons
  - Dedicated infrastructure can be expensive
  - For ESCON & WAN/LAN solutions: interconnect infrastructure cannot be shared with other applications
  - non-deterministic WAN/LAN: potential latency fluctuations
  - Cannot leverage a customer's existing IP networking infrastructure
  - Requires investment in another type of infrastructure

This is a high-level slide, really giving the three broad categories of interconnect options that HP currently offers in our remote mirroring solutions. Option one being an ESCON-only solution, in which you are employing dark fiber or directly connecting to arrays, either directly over a distance of three kilometers, or more generally and probably a more common deployment would be direct connection, but employing standard repeater/director technology to extend the overall distance to 43 kilometers. The takeaway there, of course, is that the overall link is a pure ESCON link. A multi-mode connection from the XP array to the repeater, and then single mode to extend between repeaters and directors onto the fail-over site, or the remote mirroring site. So that is a what you would describe as a "pure" fiber solution. Option two is sort of a cousin of that, employing what we're calling DWDM technology which stands for "dense wave division multiplexing" technology. Differing from a pure ESCON solution in that the ESCON protocol is being multiplexed with several possible other protocols over a common optical networking infrastructure, and then being de-multiplexed on the other side. As far as the XP arrays are concerned, it is a pure ESCON dialog that is happening, so it's really transparent technology in the purest sense of the word. But in fact, the DWDM is a new multiplexing technology employed. So the converter boxes would be different, and the infrastructure connecting the two sites would be different, but effectively, still results in sort of a pure ESCON dialog happening between the two boxes. Option three is an option that differs distinctly from option one and option two, in which the reason is that the ESCON

protocol is being actually converted to some sort of a switch circuit (WAN-LAN) protocol ATM technology – T3, OC3, or something similar. In this particular solution, we are in fact converting the ESCON protocols to some other sort of (WAN-LAN) protocol technology, and because of that are able to extend overall solutions beyond the sort of sub-50-kilometer distance. And with CNT's UltraNet Storage Director platform, in fact, we can theoretically have unlimited distance solutions between sites. HP does support other converter technologies, In-Range primarily. Those solutions – the thing to emphasize there, of course, is that you are utilizing a (switch) circuit network, not an IP or networking-based network. It would be effectively a dedicated switch circuit that would be employed between the two sites.



## continuous access xp using escon converters (non-IP networks)

Config 6

asynchronous continuous access extension strongly recommended




Preferred

Static Switched Circuit (virtual private network [ATM PVC]) WAN converter notes:

- T3 or ATM OC3 (fiber, copper, satellite, microwave link, etc); non-IP (i.e. not switched packet),
- constant path/latency (best if using sync continuous access)
- hp continentalcluster support, but no hp metrocluster support
- converter box provides buffering/compression
- can't use regular customer IP WAN infrastructure
- If CNT Utranet, unlimited distances (sync continuous access latency may become prohibitive)
- If INRRANGE: sync ca up to 31 miles, async ca up to 100 miles.

Config 7



Tolerated

Dynamically Switched Circuit WAN converter notes:

- T3 or ATM OC3 (fiber, copper, satellite, microwave link, etc); non-IP (i.e. not switched packet)
- hp continentalcluster support, but no hp metrocluster support
- path (and latency) may change (converter box provides buffering/compression)
- can't use regular customer IP WAN infrastructure
- async ca may need a higher local timeout and a larger (cache) due to non-deterministic path/latency
- if CNT Utranet: unlimited distances (sync continuous access latency may become prohibitive)
- if INRRANGE: sync ca up to 31 miles, async ca up to 100 miles. may not tolerate varied latencies as well as CNT.



HP Restricted

Stress-Free Storage...Guaranteed



## continuous access xp using escon converters for IP networks

asynchronous data transfer with continuous access extension **only!**

**Config 8**

Shared/switched packet IP-based interconnect

**• A traditional IP network is used for the remote mirroring activity**

- Customers can leverage their existing networking infrastructure
- Requirements: The network must provide a deterministic path and guarantee latency/throughput/ packet-loss

**• Pros**

- **Much more cost effective compared to dedicated/switched circuit options**
- CA/IP combination can be used for a Disaster Recovery solution (Competitive advantage)
- Future Continental Clusters support

**• Cons**

- Requirements for network interconnect (packet-loss, WAN Delay, minimum throughput) limit the type of applications
- CNT is the only Converter with this IP-based capability

**New in November 2000**

HP Restricted

This slide is meant to give you sort of an overall picture of what this solution looks like, and sort of the pros and cons. Below, we mention some of the requirements for this configuration. We've got XP arrays at either end which are employing CNT's UltraNet Storage Director platform. That platform provides the function of converting the ESCON protocol coming from the XP array and converting that to a fast Ethernet networking technology. And then that connection would be routed through a customer's existing switched packets, a (WAN/LAN) infrastructure or a dedicated circuit that they would possibly lease from a telco. The message there, of course, is that it's a switched packet (WAN/LAN) circuit. The bullets state a traditional IP network is used for the remote mirroring activity. The hard requirements that the customer has is that the network that's employed between the two sides must provide a deterministic path and guarantee some level of latency throughput impact if lost to make sure that the ongoing relationship or status of the volume pairs in a CA solution stay viable, and maintain up and running. And the requirements for throughput, of course, are dictated or driven by what sort of write traffic needs to be mirrored across a link. So all of those things work together to driving to specific requirements for the quality of the link, if you will, between the primary and remote sites. So the pros of this sort of solution should be pretty obvious. This could be a much more cost-effective solution compared to dedicated or switch circuit connectivity options. The fact that HP's CA over IP solution combines the asynchronous CA capability that is available for the XP arrays, and the fact that can be used for disaster recovery, is in fact a competitive advantage. We think we have an



SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Continuous Access XP

interesting leading capability in this regard. The customer Network Requirements are as follows:

LAN/WAN operating parameters must stay within a 'supported' envelope (delay, packet loss, minimum throughput) at all times to allow for a reasonable initial copy window and ensure ongoing solution viability

Should be at least 100bT LAN/WAN infrastructure, end-to-end (i.e. not 10bT)

The network must provide a deterministic path

Amount of network traffic must guarantee the minimum throughput

No hubs allowed in IP interconnect between arrays

CNT Ultraset Storage Director (USD) Requirements


Must utilize CNT throttling (to avoid 'overrunning' the network) and compression (to improve performance)


No more than 2 ESCONs per USD, must have 2 100bTs out

## Remote Interconnect Options Comparison

	Pure ESCON	ESCON to Switch Circuit WAN/LAN	ESCON Over TDM/WDM	ESCON To IP Switched Packet WAN/LAN
Supported Distance	(Campus/Metro) 3 Km multi-mode 43Km multi/single-mode	(Campus-Continental) unlimited*	(Campus/Metro) 50 Km	(Campus-Continental) unlimited
Overall Solution Cost	High-Medium	High	Medium-Low	Low
Solution Link Speed	17 MB/sec/link	Lesser of Tx/ATM, or ESCON	17 MB/sec/link	Max 200 Mb/s total (w/ compression)
Vendors supported	'standard' Repeaters/Directors (IBM, etc)	CNT, InRange	'standard' TDM/WDM (Nortel, CNT, InRange)	CNT
Dedicated/Shared Interconnect	Dedicated	Dedicated (VPN)	Shared	Shared
Synchronous CA support?	Yes	Yes	Yes	Yes**
Asynchronous CA support?	Yes	Yes	Yes	Yes
MetroCluster support?	Yes	No	Yes	No
ContinentalClusters support?	Yes	Yes	Yes	Planned
Non-HP-UX Cluster support?	Yes	Yes	Yes	Planned
Typical Solutions	Mirroring & Clustering Metropolitan Distances	Mirroring & Clustering Continental Distances	Mirroring & Clustering Metropolitan Distances	Replication Continental Distances

\* InRange 9801 SNS: 3 miles w/sync CA, 31 miles w/async CA  
\*\* Network latency must be guaranteed


**HP Restricted**

**N S S O**  
**Stress-Free Storage...Guaranteed**


“Remote Interconnect Options Comparison.” This is a little bit more detailed comparison between the three options previously described and this new interconnect option. Initially, we are of course supporting primarily asynchronous CA. With synchronous CA, there are specific issues that need to be addressed, given that synchronous CA has this sort of atomic lockstep remote copy operation. It is especially important that network latency and throughput impact loss, all the things that we’ve specified here really need to be looked at quite carefully, given the fact that this is not the sort of high-bandwidth solution that you could possibly have in some of the other connectivity options. Considering synchronous CA over IP is something that must be looked at quite carefully before considering in any solution, given that you’re much more susceptible, let’s say, or exposed to the inherent network latencies that you would see in these sort of solutions. Of course, we’re supporting asynchronous CA at this time. As far as Metro cluster and Continental Cluster support Metro Cluster, of course that there is a sort of standing requirement that HP’s BCC partners have had, in that they do not support any sort of protocol conversion in a remote link for Continuous Access XP.


The last bullet there, typical solutions, this is really just a high-level view on where HP believes ESCON-to-IP will be very useful for the customer base. In terms of continental distance replication, where, let’s just say, the options for remote interconnect aren’t as plentiful, HP believes that IP provides a viable option, given that many customers, of course, have continental distance

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Continuous Access XP


networking infrastructures already set up, and that the infrastructure can be leveraged for these purposes.

## CNT UltraNet Storage Director (USD)

- USD from CNT performs the transformation from ESCON to IP
- Industry leading product
- Throttling within the USD assures WAN/LAN performance for other users.
- A dedicated 100bT line (100Mb/sec) is about equivalent to the usable bandwidth of an ESCON link.
- USD can be configured with up to two ESCON and two 100bTs.
- Data compression can further reduce the bandwidth used on the 100bT link




*Available Through  
HP Complimentary Products!*



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



As the CNT UltraNet Storage Director relates to CA over IP, much of what's in the slide here we've already talked about, data compression capability and the requirements for the link connectivity for this sort of solution are listed here.

### CNT UltraNet Storage Director (USD)

Data Compression, Load Leveling, Alternate path routing, Pipelining

Compression: allows the USD to efficiently utilize the bandwidth of the remote link

Throttling: allows the USD to be tuned to the bandwidth limitations of the remote link

Buffering: makes the USD tolerant of link latency fluctuations

Emulation: makes the USD more tolerant of latency issues arising from long distance connections

Remote Diagnostics, GUI-based System Mgt

6 & 12 slot chassis

Channel Interfaces: ESCON, SCSI, FC

Network Interfaces: ATM, DS-3 (T3/E3), Ethernet (10/100)

Two USDs required (local and remote) for complete solution

USD Strengths

sophisticated traffic routing/management: more appropriate for extended distance solutions

Ethernet (IP)-capable

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Continuous Access XP

robust System Mgt tools

## HP XP48/256/512 Remote Mirroring with CNT

**HP Continuous Access**

IP  
or E1, E3, ATM, SDH  
 $\lambda$   
or Dark Fibre/ DWDM

HP-UX, NT/W2000,  
Solaris, AIX

HP XP  
Storage

CNT  
UltraNet




Network

CNT  
UltraNet

HP XP  
Storage

### High performance remote mirroring






- HP NSSO qualified and endorsed
- UltraNet Storage Director comes with full HP Support
- Supports ESCON, SCSI and now Fibre Channel
- HP Complementary Products (hpcp) resell CNT UltraNet products

 HP Restricted  Stress-Free Storage...Guaranteed 

2 CNT USD's provide a high-speed switching platform at each site, local and remote, to interconnect HP XP's over unlimited distances to create an effective enterprise-wide SAN. HP has had several successful implementations to date with British Telecom, Yellow Pages, Takas Bank and Dutch Tone. The CNT USD's currently perform Continuous Access (CA) over WAN or over IP. CNT introduce Fibre Channel Interface Cards to allow Continuous Access using FC as an alternative to ESCON. The FC Interface Cards are now available on the HP Price List. For all CNT related information, please contact your local hp complementary products (hpcp) representative or visit the hpcp homepage at: [http://hpcp.grenoble.hp.com/Content/storage/partners\\_and\\_products/](http://hpcp.grenoble.hp.com/Content/storage/partners_and_products/)

## CA Operating States

Simplex . . . . .	Not currently paired
Pending Duplex . . . . .	Initial copy in progress
Duplex . . . . .	Volume pair is in sync, synchronous updates are in progress
Suspended . . . . .	Volume pair is out of sync, synchronous updates have stopped, difference data may be being maintained depending on fence level
Duplex LUSE warning . . . . .	An additional state when an individual LDEV within an expanded volume has failed
Pending LUSE warning . . . . .	An additional state when an individual LDEV within an expanded volume has been detected as failed during initial copy



This slide describes CA volume pair status descriptions. Unlike Business Copy, there is no split state.

In Business Copy, pairs in duplex state are NOT guaranteed synchronized. In Continuous Access, pairs in duplex state are guaranteed synchronized. Please reference page 20 of the CA User's Guide for further definitions.

## CA Fence Levels

- Data - M-VOL fenced if update copy fails
  - data is not allowed to differ between the M-VOL and R-VOL
  - replication link failure causes application to fail when running on the S-VOL. S-VOL inaccessible to applications for updates
  - useful for customers whose business requires no loss of data
- Status - M-VOL fenced if R-VOL cannot be suspended
  - data is allowed to differ between the M-VOL and R-VOL
  - I/Os accepted by M-VOL and MCU tracks updates for re-sync
- Never - non-current data at remote site if ESCON link fails
  - data is allowed to differ between the M-VOL and R-VOL
  - replication link failure does not cause application to fail when running on the S-VOL, but data no longer protected



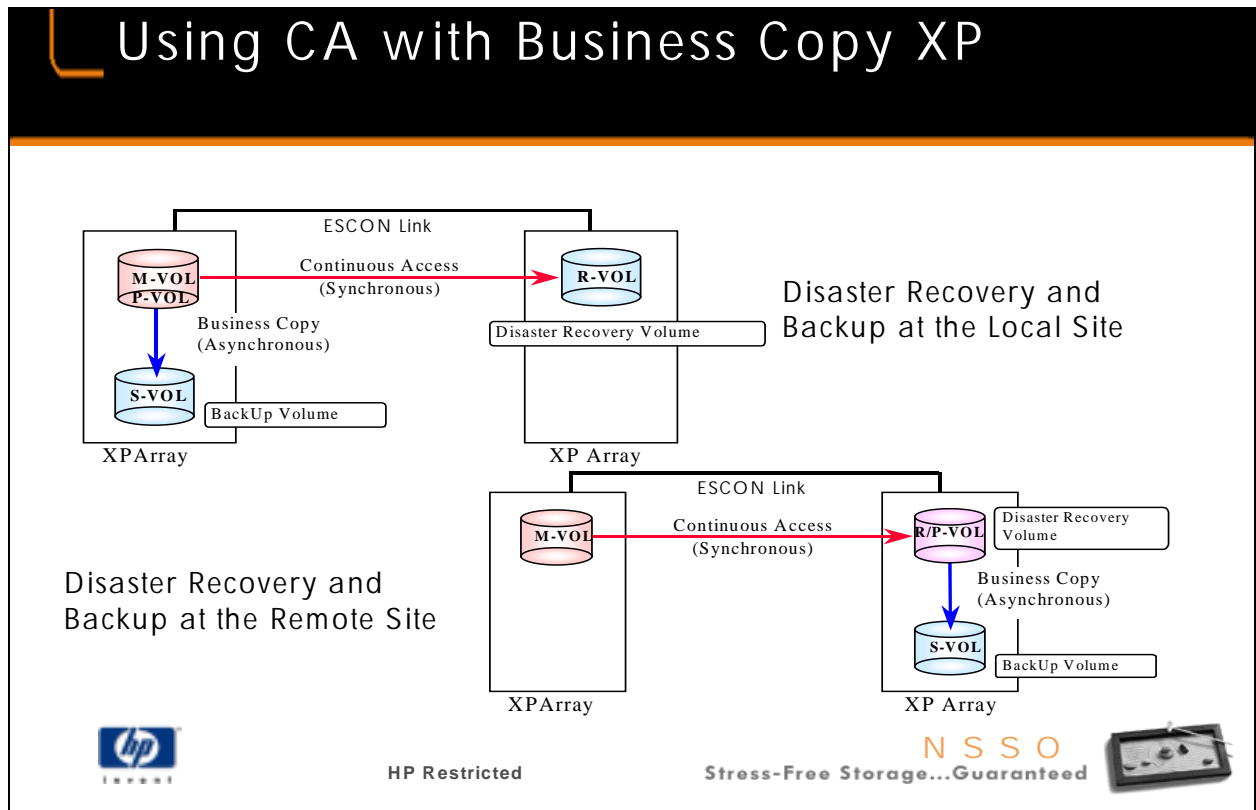
HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



Different fence levels allow different levels of assurance regarding the mirroring consistency of critical paired volumes. It is necessary to determine the data consistency and takeover action according to the pairing status and fence level of the corresponding secondary volume.





The diagram shown here is the solution of choice for HA customers because a synchronous backup at the remote site (R-VOL) can be quickly be put into use for automatic disaster recovery situations since it is in sync with the original data. Backups are streamed from the P-VOL to the S-VOL at the remote location, so the split occurs there. There is never a suspend (or split) between the local and remote sites. Customers will generally not want to stream their backup directly from the R-VOL because this will necessitate a suspend between the M-VOL and the R-VOL. Since the streaming backup operation can take up to several hours, the R-VOL will be several hours out of sync with its M-VOL, disallowing immediate disaster recovery. This is the strategy that will be supported by MC Metrocluster. Streaming backup will occur off of the remote S-VOL.

This solution provides:

- full disaster recovery / zero downtime backup

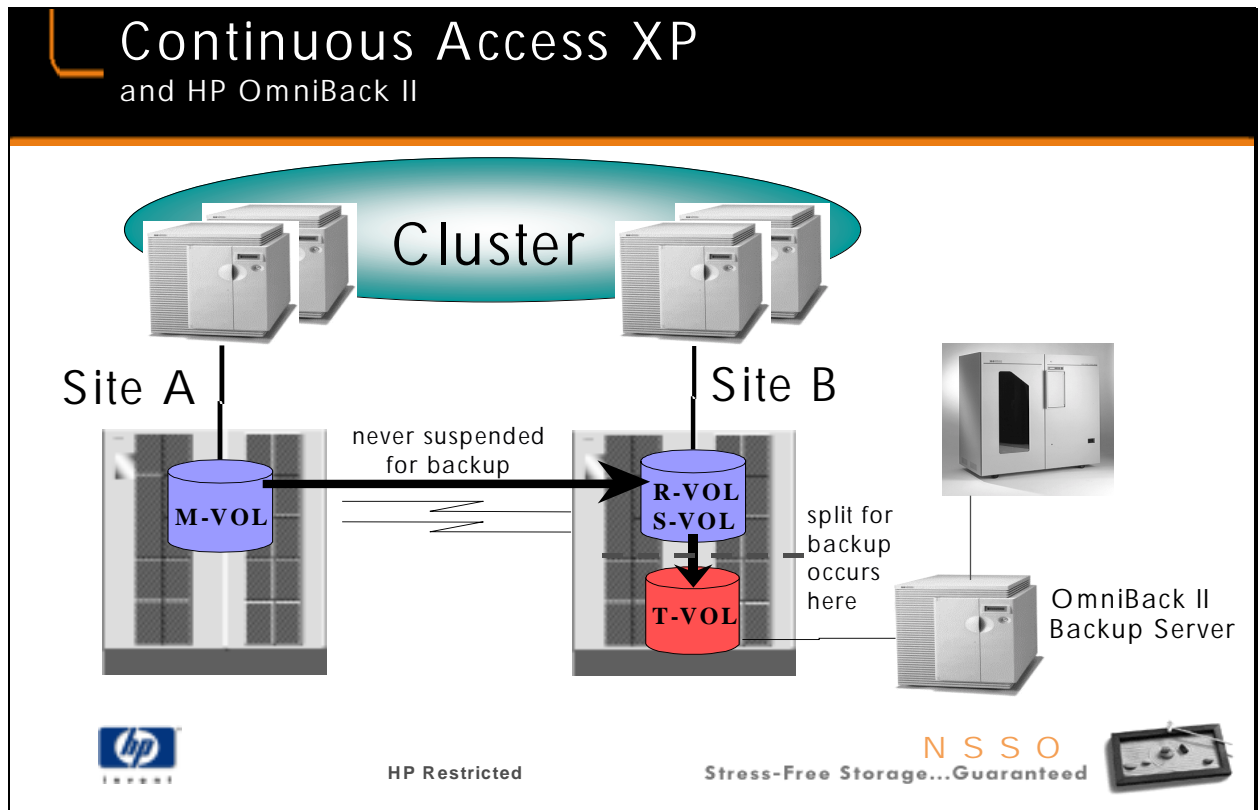
- continued split mirror backup after failure via business copy

- caveat is performance hit for remote synchronous updates.

Although this configuration eliminates performance issues associated with synchronous remote backup, care must be taken to not use Business Copy and Continuous Access at the same time. The business copy pair must be created first and then split before creating the continuous access pair. The continuous access pair must be suspended prior to each re-syncing of the business copy pair. By using Business Copy XP and Continuous Access XP together, you can provide on-site backup copies of M-VOLs and remote backup copies of R-VOLs.

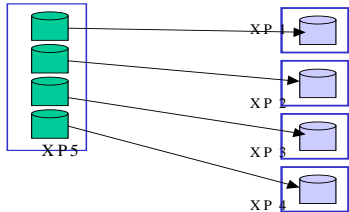
When BC sources a CA , these are all the rules:

1. The BC pair must be made before the CA pair
2. When you create the CA pair, the BC pair must be PSUS
3. If either the BC or CA is PSUS, the other is not allowed to be PSUE
4. If either the BC or CA is COPY, the other is not allowed to be COPY
5. If either the BC or CA is PAIR, the other is not allowed to be PAIR

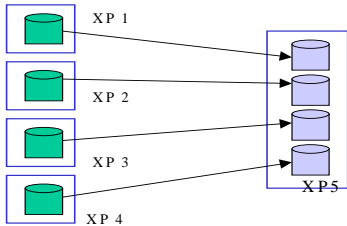


# NxM Array Configurations


NxM Configurations



NxM Fan **out**, up to 1:4  
OK for sync-CA & Async-CA  
as long as volumes, device groups or Consistency groups only see 1:1 (do not span arrays)




NxM Fan **in**, up to 4:1  
OK for sync-CA & Async-CA  
as long as volumes, device groups or Consistency groups only see 1:1 (do not span arrays)




HP Restricted

NSSO  
Stress-Free Storage...Guaranteed




Connecting 1 array to up to 4 arrays is possible, with restrictions. Connecting up to 4 arrays to 1 array is also Possible (with restrictions). The restrictions differ between sync-CA and Async-CA as described below.

## NxM Array Configurations (cont.)



sync-CA - **Allowed**  
async-CA - **Allowed**

- 1:1 array configuration for any given pair
- single volume pairs
- no XP knows about more than 4 other XPs



Sync-CA - **Allowed** (device group)  
async-CA - **Allowed** (Consistency Group)

- 1:1 array configuration for any given group
- async-CA consistency assured as a group. If pair #2 has an error, pair #1 will stop svol updates, so data consistency is maintained at the svol group level
- sync-CA consistency NOT assured as a group. If a write goes to pair #1 pvol first and #2 pvol second, there is no guarantee that the svols will be updated in that same order. The device group is just a convenience for multiple pair creates/suspends with one command. With **ever?** fence, a failure with pair #2 will not stop pair #1 from talking, thus the group svol data consistency is compromised. With **ata?** fence, a failure on pair #2 will cause all pairs in the group to refuse writes (i.e. to be fenced). Its up to the host to ensure ordering between the pairs in a group such as both pvals are the same files system so it won't try a second write until the first write completes.
- no XP knows about more than 4 other XPs




HP Restricted

**NSSO**  
Stress-Free Storage...Guaranteed




## NxM Array Configurations (cont.)



Aggregated device groups (sync-CA)  
that span arrays are **allowed**,

but the application must use Data fence, and the application (e.g. dB) must  
be prepared to deal with the data consistency issues and more possible  
points of failure



Consistency Groups (async-CA) that  
span XP arrays are **not** allowed

One Device or  
Consistency Group

One Device or  
Consistency Group



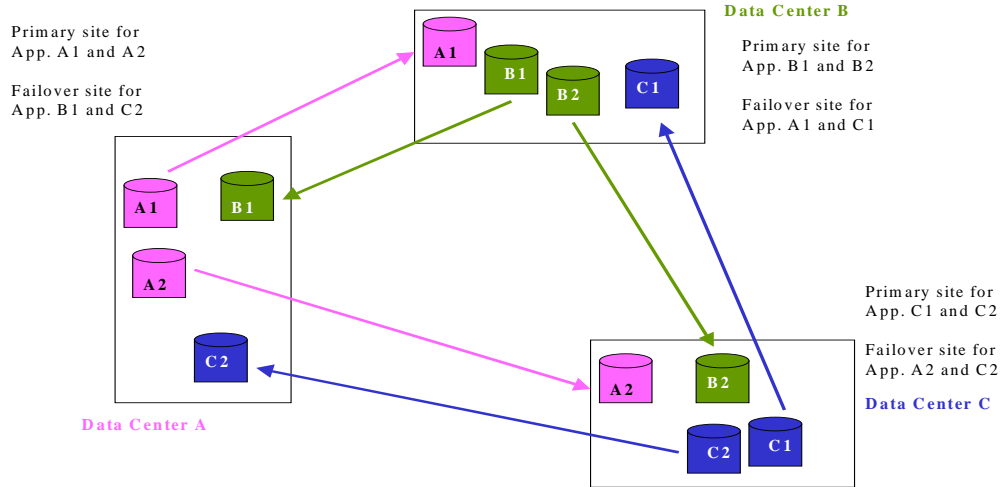
HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## NxM Array Configurations (cont.)

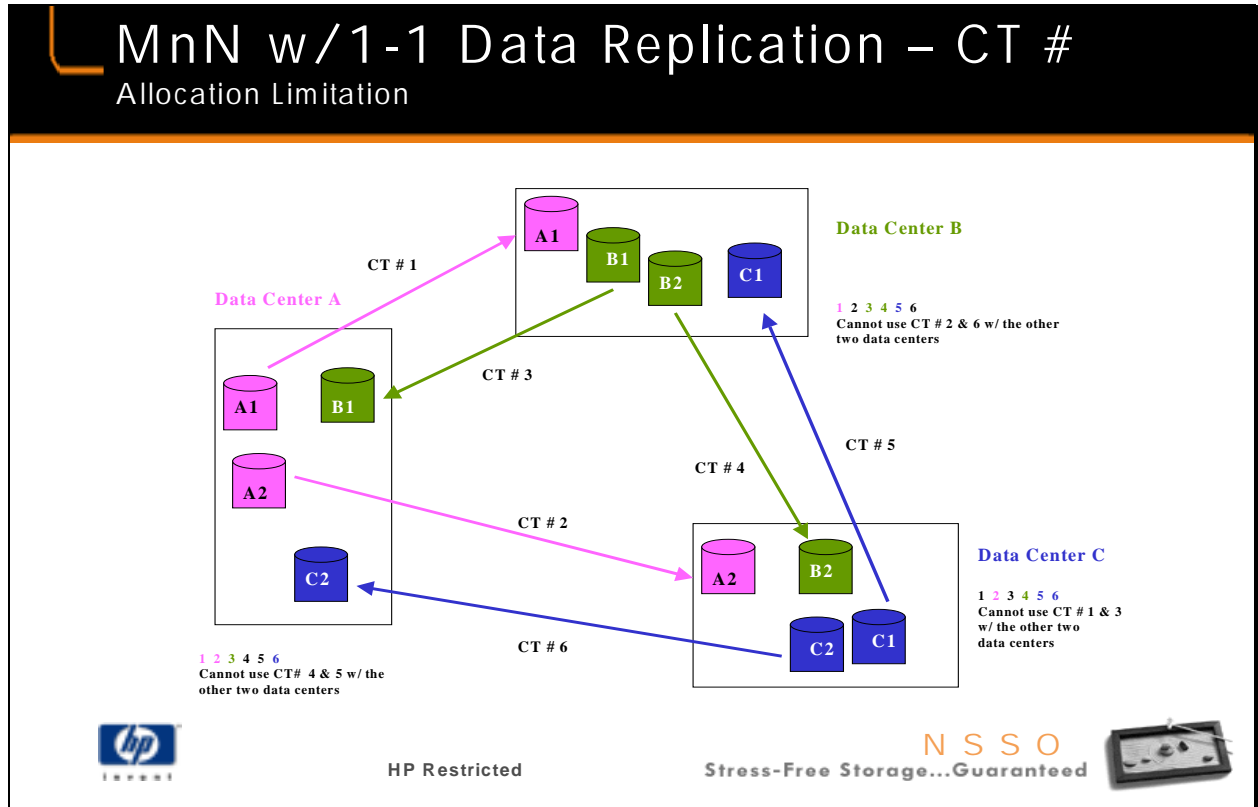
Multi Data Center MxN using 1-1 Data Replication- **allowed**



HP Restricted

NSSO  
Stress-Free Storage...Guaranteed







## Continuous Access Rules of Thumb

- One CA ESCON link can control a maximum of 1023 paired volumes for XP256 and 4095 for XP48/512.
- ESCON links should be configured in pairs.
- ESCON performance (Link Speed)
  - Link distances greater than 3km : 12-13MB/sec
  - Link using repeater / director : 5-10MB/sec
- Write performance will be limited to ESCON link speeds for sync-CA pairs.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



# Continuous Access

## Competitive Alternatives

- Competitive alternatives:
  - EMC Symmetrix SRDF
  - Amdahl XRC and AXRC
  - HDS 7700E Remote Copy
  - IBM PPRC and XRC
  - Sun A7000 Remote Dual Copy



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Cluster Extension XP

- The case for Cluster Extension XP
- Value Proposition
- Product Structure
- Competitive Differentiation
- Technical Overview



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



This is a new product and new product software that also functions as a new disaster recovery solution available from HP. And what we want to talk about briefly is specifically Cluster Extension XP which was launched in November of last year -- so, just a couple months ago. What we want to accomplish is giving an overview of what Cluster Extension XP is, to talk about the value proposition as well as the product structure, and then some of the competitor differentiation factors. We will then go into the technical details of the product.

## Integrated High Availability Cluster Solutions

- ✓ Cluster solutions, ensuring seamless continuity of business operations
- ✓ Extends Datacenter/Campus Cluster Solutions over greater distances
- ✓ Integrates continuous access xp fast failover and failback functionality with leading cluster software for reliable, automatic site recovery in the event of a fault, failure, or disaster
- ✓ delivers the potential for consolidated, heterogeneous cluster solutions




HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



"Integrated High Availability Cluster Solution," data availability overall starts with the XP disk array which in itself is 100 percent fully redundant with no single point of failure. From here, we can add higher levels of availability by incorporating software such as Business Copy XP or Continuous Access XP for remote data mirroring. And then to provide even further levels of availability for your customer, you can incorporate clustering as part of the XP disk array solution. So whether availability is just the XP itself which is fully redundant, the XP array with software such as Continuous Access XP, or XP disk array with Continuous Access XP and Clustering Solutions, all of these solutions are focused on providing data availability for the customer and ultimately, seamless business continuity. Specifically, XP array focused Clustering Solutions integrate the XP disk array, Continuous Access XP, and leading cluster software depending on the type of OS (I.e. MC/ServiceGuard for HP-UX for example). And so this truly provides a heterogeneous solution for the customer. When you look at clustering overall, basic clustering starts with clustering within a single data center or local campus environment to provide fail over fail back capability with stored subsystems. The clustering can also be extended over even wider and wider distances to provide a even more robust disaster recovery solution.

## HP-UX/XP High Availability Cluster Solutions



Delivering city-wide automated fail-over



HP Restricted

### ➤ What does it offer?

- Full and seamless integration XP remote mirroring with HP High Availability Cluster Solutions:
  - MC/Service Guard
  - Metro Cluster
  - Continental Cluster
- A complete, end-to-end HP solution, delivering 5x9's Availability
- Datacenter, Metropolitan, or Continental distance solutions
- the Most robust clustering solution for open systems available

Stress-Free Storage...Guaranteed



This slide shows the HP-UX specific clustering solutions that work with the XP disk arrays today without getting into too much detail. Therefore, clustering for HP-UX starts with MC/Service Guard which is the solution for the data center or campus environment which we had discussed earlier in this module.

And then if you extend that to even wider distances, we have metropolitan wide clustering called Metro Cluster which works with MC Service Guard and to geographically disperse XP disk arrays. And then for even wider multi-clustering needs, we have Continental Cluster, which, like Metro Cluster, works with MC Service Guard as well as multiple disk arrays over wider distances.

So this overall -- with all these solutions -- these are excellent solutions for the HP-UX specific environment using XP disk arrays. But the question is, "What about the other OS environments?"

## Non HP-UX/XP Cluster Solution

The diagram illustrates a cluster solution where two servers, labeled 'Cluster', are connected to two storage subsystems (represented by purple cylinders) via a network. The network is labeled 'Datacenter/Campus' and the distance between the storage subsystems is labeled 'Datacenter/Campus distance'. The HP logo is visible at the bottom left of the diagram area.

- **Existing Solution**
  - Datacenter/campus cluster solutions today can be implemented using more than one storage subsystem as shown
  - However I/O paths to storage subsystem limited by link distance (10km)
- **The case for Cluster Extension XP**
  - Customers are asking, "How can I extend my datacenter/campus cluster solution over longer distances without compromising the failover/failback, high data integrity performance of a datacenter/campus cluster solution?"

HP Restricted

N S S O  
Stress-Free Storage...Guaranteed

for the other OS environments, the XP disk array can be and has been integrated with leading cluster ware such as Veritas Cluster Server for Sun Solaris and (HACMP) for IBM AIX. And historically, in terms of how clustering generally works, for a single data center or local campus environment, the customer typically uses one, maybe two disk arrays. And as the customer attempts to extend this single cluster over wider distances to create a disaster recovery solution, the customer will require two disk arrays.

However, for this type of extension, it's typically limited by the distance link, which for fiber channel it's about ten kilometers. So the question is, "How is the customer able to extend the single campus clustering solution over wider distances without compromising the fail over fail back capability of the cluster software with the disk array and ultimately to be able to maintain data integrity and business continuity?"

## Cluster Extension XP Overview





The diagram illustrates the Cluster Extension XP architecture. At the top, a 'Cluster' (represented by two server icons) is connected to a 'Cluster Extension XP' layer (represented by a blue oval). This layer is connected to two 'Datacenter A' and 'Datacenter B' units. Each datacenter contains an 'XP Array' with a 'Data' disk and a 'Mirror' disk. The two datacenters are connected via 'Continuous Access XP' (represented by a double-headed arrow). The distance between the datacenters is labeled 'Campus/Metropolitan/Continental distances!'. The HP logo is visible at the bottom left, and the text 'HP Restricted' is at the bottom center. The N S S O logo is at the bottom right, along with the slogan 'Stress-Free Storage...Guaranteed' and a small image of a storage device.

- **What does it offer?**
  - Offers disaster-recovery, cluster-specific automatic failover/failback solutions for Veritas Cluster Server and HACMP, extending them over Metropolitan Distances for Sun Solaris and IBM AIX environments
  - Seamlessly integrates cluster operations with XP remote mirroring operations
- **How does it work?**
  - Provides XP-specific commands and interface to tightly integrate Cluster Software functions with Continuous Access XP and XP Extension
  - Automatic takeover and redirection of mirrored pairs in failover scenarios
  - checks over 400 rules for disk state scenarios based on data consistency and currency
  - provides auto-recover pair-resynchronization monitor

The solution is Cluster Extension XP. Now, Cluster Extension XP is a new middleware product that has been announced in November 00 that offers a disaster recovery cluster specific fail over fail back solution for the Sun Solaris environment using Veritas Cluster Server and IBM AIX environments using (HACMP). A simple way to think about this is that today we have Metro Cluster for HP-UX. Now we have Cluster Extension XP for the Solaris as well as the AIX environments using XP disk arrays. This slide gives a few bullet points as to what it does and what it offers which we will talk about in more detail in follow-on slides.

# Value Proposition

<i><b>features</b></i>	<i><b>benefits</b></i>
<b>Cluster Extension XP enables:</b>	<b>for the organization, this means:</b>
Integration with Veritas Cluster Server on Sun Solaris and HACMP on IBM AIX	Ensures automatic fail-over and fail-back for a truly heterogeneous environment
Automatic mirrored volume pair management	Monitors volume and link status to detect potential link failures and automatically manages recovery when links are restored
Extensive volume state rules database	Ensures fast and reliable fail-over recovery for all possible fail-over scenarios
Flexible programming interface	Allows customized integration of Cluster Extension XP and XP Disk Array remote mirroring solutions with other cluster and host compute environments







This slide shows you the value proposition for Cluster Extension XP along with some of the features and benefits of Cluster Extension XP. And overall, Cluster Extension XP provides XP-centric clustering solutions for Solaris, AIX, and in the future, Windows 2000 and Linux, which overall, we're able to provide a fail-over fail-back solutions with the XP disk array for truly heterogeneous environments. Now, Cluster Extension XP has many features including, as shown here, automatic mirrored volume peer management and extensive volume state rules database, which with all these features, it provides tremendous flexibility to the customer to meet their exact clustering needs with the XP disk array which will discussed in greater detail shortly.






## Product Structure

- **Productization:**
  - B9531A, Cluster Extension XP for VCS on Sun Solaris
  - B9532A, Cluster Extension XP for HACMP on IBM AIX
- **Licensing:** LTU per cluster
- **Availability:**
  - Cluster Extension XP for HACMP: Live CPL on Feb'01
  - Cluster Extension XP for VCS: Live CPL on Mar'01
  - Future: Cluster Extension XP for MS DataCenter/Advanced Server: TBD

  HP Restricted  Stress-Free Storage...Guaranteed 



"Product Structure," this gives you an overview of how the product is structured. As you can see from the slide there will be two separate product numbers for Cluster Extension XP for each version of the OS. For VCS on Sun Solaris the product number is B9531A, for (HACMP) on IBM AIX, it's B9532A as shown. Licensing will be on a per cluster basis. In terms of availability, Cluster Extension XP will be on CPL for (HACMP) on February 1. The Veritas Cluster Server on Sun Solaris support with Cluster Extension XP will be on CPL beginning March 1.

## cluster extension xp support



### Veritas

- Stress Free Central
- HP and Veritas jointly tested and certified VCS and VxVM to work seamlessly with XP Disk Arrays
- VCS offers an Agent Development Kit to allow custom agents; Cluster Extension XP is a custom agent
- Pass Through Support Model between HP and Veritas




### IBM

- Stress Free Central
- Avallant (formerly Clam Associates) designed HACMP for IBM
- HACMP designed to allow flexible customization scheme
- IBM support available through TSAnet membership

N S S O

Stress-Free Storage...Guaranteed



There is probably a concern of this type of heterogeneous solution in that “how is HP going to support this new product and solution given that Cluster Extension XP works with cluster software from IBM and Veritas?” HP has worked the support issues to make sure that Cluster Extension XP will truly be supportive. A first clarification of all, given that Cluster Extension XP works only with the XP Disk Array, we do have dedicated resources within Stress-Free Central to resolve support issues around XP and also for Cluster Extension XP. In addition, HP also has a few dedicated resources within Stress-Free Central who are knowledgeable of the Sun Solaris environment as well as the IBM AIX environment. And these people have been given sufficient training on Cluster Extension XP for Veritas Cluster Server and (HACMP). An important point to remember is that VCS has been jointly tested and certified by both Veritas and HP which overall provides assurance that VCS does work seamlessly with the XP Disk Array.

## Competitive Differentiation for Solaris

- EMC: No comparable solution!
- SUN: No comparable solution!
- Veritas: Volume Replicator with VCS, VxVM, Solaris
  - Host-based solution
  - "Write-twice" host-based mirroring means negative I/O performance impact
  - Extending cluster distance further negatively impacts I/O performance
- HP: Cluster Extension XP with VCS, VxVM, Solaris
  - Array-based solution
  - No I/O performance impact
  - No performance impact by extending cluster distances
  - Fully utilizes high performance capability of Continuous Access and XP Disk Arrays



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed




Now, Cluster Extension XP is a key differentiator for HP, given that neither Sun nor EMC have a comparable solution. The EMC Disk Array today does work with leading cluster software in a single data center or campus environment, but when you try to extend the solution over wider distances, EMC cannot offer this. Now, it's true that Metro Cluster works with EMC to provide this solution over (Metropolitan-wide) businesses, but this is just for HP-UX. EMC can't offer this for the Solaris environment. And, unlike Metro Cluster, EMC will not, nor will it ever work with Cluster Extension XP. Sun does have a cluster solution for the single data center or campus environment, but like EMC, they also do not have a comparable solution to be able to extend this cluster over wider distances using two metropolitan separated disk sub-systems. Looking at Veritas, they do have a similar solution if you combine their Volume Replicator plus their server and Volume Manager Solution which does provide migration of replication services. It is HP's concern that the Veritas solution may not be quite robust enough to provide a disaster recovery solution over wider distances. And the reason for this is because this type of solution is a host-based solution where the host has to write twice, once to the primary array as well as to the remote disk array. And as a result of this, the solution puts additional burden on the server by using up additional CPU cycles which negatively impacts I/O performance. In comparison to Cluster Extension XP, which is a very unique solution that is an array-based solution that offers no performance impact even when extending the cluster over

wider distances. So this slide gives you an overview of how Cluster Extension XP differs in the Sun Solaris environment.


## Competitive Differentiation for AIX

- EMC: No comparable solution!
- SUN: No comparable solution!
- IBM: HAGEO with HACMP, AIX
  - Host-based solution
  - "Write-twice" host-based mirroring means negative I/O performance impact
  - Extending cluster distance further negatively impacts I/O performance
- HP: Cluster Extension XP with HACMP, AIX
  - Array-based solution
  - No I/O performance impact
  - No performance impact by extending cluster distances
  - Fully utilizes high performance capability of Continuous Access and XP Disk Arrays




HP Restricted

N S S O  
Stress-Free Storage...Guaranteed




This slide gives you the differentiation of Cluster Extension XP in the AIX environment. Much like the Solaris environment, Cluster Extension XP is a key differentiator for HP given that neither Sun nor EMC have a comparable solution. When you look at IBM, they do have a similar solution, which on a functional basis provides something similar in functionality. And this is what their HAGEO solution on AIX which does provide migration replication and data replication features. But much like Veritas, HAGEO is probably not as robust enough to provide a complete disaster recovery solution over wide distances. And similar to Veritas, again, because it's a host-based solution where the host has to write twice, once at a primary array and then sending a copy to the remote disk array. And as a result, this solution uses up additional CPU cycles. Looking at the last bullet with Cluster Extension XP, un-like the HAGEO and Veritas solution, Cluster Extension XP is an array-based solution that offers no performance impact to the host and overall it's able to do this by integrating Cluster Extension XP with the features of the cluster software HACMP and VCS), while also providing the full mirroring capability of Continuous Access XP where it's able to use both synchronous and asynchronous copies along to both XP disk arrays.

## cluster extension xp - what is it?




- software, which helps your customer to survive any kind of disaster
- automated server-based metropolitan-wide solution
- needs cluster software to ensure application failover, availability and consistency
- controls the XP disk array family by making sure your data is safe, wherever



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



Cluster Extension XP is a software product which is a middle ware piece between the HP RAID Manager XP miniature component which gives you all the HP disk states and handles fail over between two disk areas from the host side and the cluster software which resides on top of that. We will see a diagram of this in several slides. In general, Cluster software is usually not aware of that there are two different disk arrays where the servers connect to. The biggest issue you face with Continuous Access XP solutions is that the primary disk area is in a “read-write” mode while the secondary disk array , or copy is in read-only mode. Therefore, if the cluster software fails over from one server connected to the first disk array to the second server connected to the second disk array, the cluster software isn’t necessarily aware that this is a read-only disk. And this is the problem that the middle-ware product such as Cluster Extension XP or Metro Cluster works to solve. The concept of what HP does in these kinds of environments utilizing products such as Cluster Extension XP or Metro Cluster is to automate these fail over process and to allow the cluster software to fail over to the secondary side. In order to do that, the cluster software (MC/ServiceGuard, VCS, HACMP), has to follow some rules, meaning that to make sure that the secondary copy is current and consistent. And this is the goal of Cluster Extension XP and Metro Cluster. What is required of Cluster Extension XP is definitely the fail over capability of the cluster software. Without that, Cluster Extension is not able to automatically fail over from one disk array or from one server to the other server. However, Cluster Extension can be used to

fail over from one server to another without any cluster software, but – what I mean is it could be used to actually check the disk array state in the same manner as we would do with any integration into cluster software. So Cluster Extension actually controls the disk array state and makes sure that the data is safe on every side and that we only activate the data or the disk set when the data is safe to do so.

## cluster extension xp- what does it do?

- automated takeover and redirection of mirrored pairs
- checks over 400 rules for disk states based on data consistency and concurrency considerations
- provides auto-recover pair-resynchronization monitor
- full and seamless integration in leading cluster software
- cli for easy integration into any today available cluster software



HP Restricted


N S S O  
Stress-Free Storage...Guaranteed



How do Cluster Extension XP do all this you might be asking yourself? First Cluster Extension XP uses an internal rule engine with along with it's object database stores all kind of past data information. For example, if we were to fail-over from one site to another, what do we do when I have a secondary disk and it's write protected? Can I easily activate this disk or can I easily do a takeover on this disk? This is one of the questions that would be contained in this database. Therefore, and all this information is stored in these databases and these databases are accessed each time you fail over your application with the class of software. Currently, there are actually over 550 rules in this database to make this all happen. These rules usually must be checked by the system administrator or the storage administrator when you fail over from one site to the other to be sure they are correct for this environment. With HP's recently available fast-failback capabilities with Continuous Access XP, the former primary copy is now a secondary copy. HP basically reverses the personalities of the disks in more or less any case so that the customer has the ability to keep data consistency and provide a quick recovery once the problem has been found and corrected. What Cluster Extension will also provide is an automatic error recovery monitor, which means that in most cases, Metro Cluster and other clustering extensions only works when you activate an application or when a class of software fails over an application to another server. But after the package failover there is no monitoring going on. If the link is still healthy or if the pair is still pair state for example. Therefore, if you have, for example, a very




unreliable network in between the two disk arrays you would never know when the pairs are suspended or get suspended. So, what we provide with this Cluster Extension XP is a monitoring facility and you can actually set that up in the cluster software or in the cluster extension part of the software. And this piece will monitor your health between, or the health of the pair state between the two disk areas. A very nice feature of this capability is that you can set up this pair monitor so that it automatically resynchronizes when the pair is getting suspended. That means it would provide the kind of self healing feature in here so that pairs between the two disk arrays are always in sync. This prevents fail over situations where you activate an application on the disk set, which is not current. Cluster Extension XP is designed and fully integrated in the Veritas or HACMP product by IBM meaning we HP doesn't have different products or any other software requirements. For easier integration into any other kind of clustering environment, HP provides a command line interface, which must be integrated by HP consultants or other consulting organizations. And the purpose of this command line interface is to provide fail over capabilities for continuous excess or various disk arrays on platforms where we haven't yet integrated with the Cluster Extension XP software yet.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed

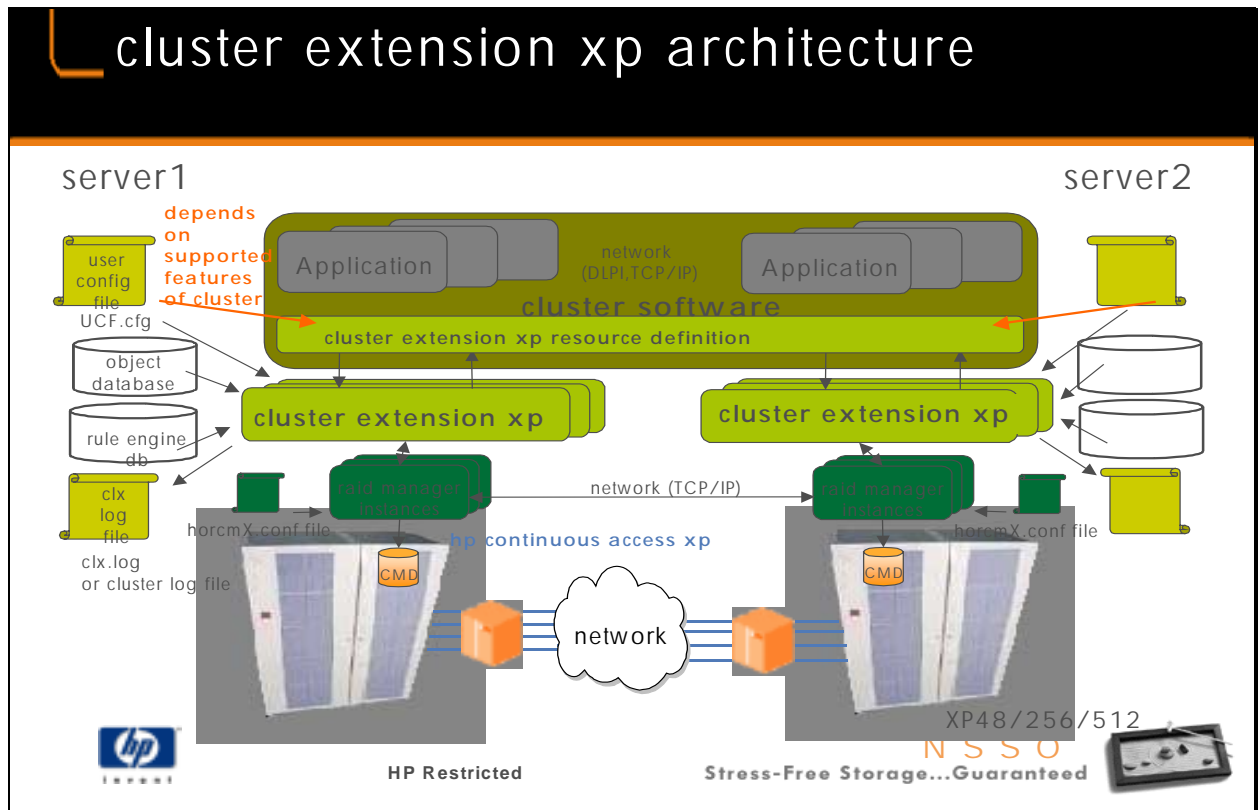


## cluster extension xp - unlimited control

- DataLoseMirror/  
DataLoseDataCenter
- AutoRecover
- ResyncMonitor
- ApplicationStartup
- Pre/PostExec
- PostExecCheck
- FastFailbackEnabled
- alternate raid manager instances
- automatically raid manager startup
- force startup
- ca xp fast failback support

This page summarizes some of the features and functionalities of Cluster Extension XP. I will start on the right side to explain it a little bit more. For people who know metro cluster, this is to be seen as an enhancement for that kind of functionality. What we do in cluster extension, for example, is we alternate between different RAID Manager instances. For example, if one RAID Manager instance were to go down, maybe it is mis-configured and cluster failover software depends on RAID Manager. So, if we cannot talk to RAID Manager we cannot provide fail over capabilities. So, what you can do as the user is you can set up several RAID Manager instances running in parallel on each server. Therefore, if one of the RAID Manager instances is not available it would not stop functionality or the program, it would just alternate to the next available RAID Manager. Also, if there is not a RAID Manager instance running at the time you were to start up cluster functionality, then Cluster Extension XP would automatically start up RAID Manager to use RAID Manager's functionality. Another feature is if a customer doesn't like or want to use Cluster Extension's rule engine you would have the ability to force the application start up without using Cluster Extension's internal logic. Cluster Extension XP supports the fast fail back capabilities of the most current CA HP firmware. With this version of firmware, HP supports Continuous Access XP fail over operations which doesn't lead into simplex or split pairs of drives. Meaning whenever you fail over from the primary side to the secondary side Continuous Access XP will fail over the disk area pairs in a way such that you would only need to do Delta

resynchronizations to get back to the pre-failover state. Meaning that you would not have to recreate pairs as is or has been in metro cluster over the last couple of years. You only have to do a pair resynchronization and this will only copy changed tracks between the two disk areas. So, there are some new pair states in the HP disk area, which have to be handled for supporting these kinds of features. And because Cluster Extension XP is tightly integrated with the disk array HP supports these and all future coming features. So how can you set up the front fold of the fail over process? For this HP has introduced some parameters, which have self explaining names. Cluster extension is the middleware that controls the HP disk array in case of a fail over situation. For this, HP has more or less the storage point of view in the cluster software. So, if your customer wants to activate data based on Cluster Extension's rule engine, Cluster Extension's internal logic is built on the single question is it safe to activate data on the disk area? So therefore, you can set up several different options and they each have names. The only thing you can do though is to set them to a yes or no state. So, for example, we have options where you can say I want to lose my remote mirror when I fail over or I don't want to lose my remote mirror when I fail over. So, if you say you want to loose a mirror when you fail over and you set that parameter to yes, Cluster Extension XP will start up the application even though it knows that we don't have data replication going on between the two disk areas any longer. If your customer chooses to say I want to keep my copy consistent between the two disk areas at any time, then you would set this parameter to no and the fail over would be stopped or the application start up process would be stopped. The other recover feature in Cluster Extension XP and is another differentiator to any of the current available products is the auto recover feature which fully utilizes the fast fail back feature of the Continuous Access XP software and associated firmware. What this means is Cluster Extension XP tries to automatically recreate the pairs from the most current volume. Therefore, after you have started up the application, these devices all ready in pair state so that you never lose your secondary copy. The ResyncMonitor can be started from the cluster extension software. Another feature is that it is possible to specify how you would like and where you would like to start up the application. Meaning the application start up feature is more than to specify if you want to start up on the most current copy or if you are not starting up on the most current copy and if you want to wait until the copy you start on is resynchronized from the most current copy. Also possible in Custer Extension XP is pre and post executables so that you can actually send e-mails before you fail over the application to notify somebody that there is an application fail over taking place. And the last option listed here is FastFailback enable and this feature is more or less implemented into Veritas cluster server. The reason is because only in Veritas cluster server HP has changed the fail over behavior of the cluster software. Therefore, if somebody doesn't like this level of control HP has done they can enable or disable this feature to say you want to have it in the way Veritas does it or you want to have it in the way HP does it.



This is more or less the architectural overview of how Cluster Extension XP and how it works internally. If you look at the cluster software piece you will see that it is the upper layer of the entire software level. This part controls the application fail over process. For the Veritas cluster server software, Cluster Extension XP is completely integrated into the clustering software. With (HACMP), Cluster Extension XP is a layer below. That means Cluster Extension XP is it's own middle ware part and Cluster Extension XP is not that tightly integrated into the processor. So, with Cluster Extension XP you have two ways to integrate into a cluster software product. One is to be fully integrated and using the functionality of the software completely or using part of the cluster software functionality as part of cluster extension. Or you are not that tightly integrated with the cluster software and then Cluster Extension XP shows up as another layer. The lower layer to control the disk area itself is the RAID Manager XP instance, which controls the disk area which we call upon to fail over this pair between the two disk arrays. So, cluster extension internally uses an object database to define parameters. And a rule engine, which is more or less the brains of the whole software. The only thing the user or support people have to do is to specify the user configuration files shown there, which defines the fail over and server behavioral cluster extension.

## cluster extension xp - what is supported?

- ibm hacmp 4.3.1\*  
on aix 4.3.3 with or w/o  
auto path
  - veritas vcs 1.1.2a\* (VxVM 3.0.2)  
on sun solaris 2.6 , 7
  - veritas vcs 1.3.0\* (VxVM 3.1)  
on sun solaris 2.6 , 7 , 8
- \*various restrictions apply
- continuous access xp and  
continuous access extension\*
  - any currently supported  
extension hardware for  
escon mirroring\*
  - fast failback with xp48,  
xp256 and xp512




HP Restricted

N S S O  
Stress-Free Storage...Guaranteed




This slide lets you know what cluster software we integrate with and this is currently HACMP version 4.3.1 and only on AIX or 4.3.3 at the moment. We prefer to integrate the whole solution with auto path, but Cluster Extension XP doesn't depend on auto path to work correctly. Cluster extension is fully integrated with VCS 1.1.2a and also 1.3.0. So, this is a new feature here, too. What HP does support is Solaris 2.6 and up to Solaris 8, which is the latest version of the Solaris operating system software. The tests state that various restrictions apply. These restrictions are mainly restrictions in the use of the cluster software and are more or less explained in the user configuration guide of cluster extension and also in the installation guide.


## cluster extension xp for veritas cluster server




- full integration in vcs
- based on veritas development framework
- online addition possible
- metropolitan and global solution




HP Restricted





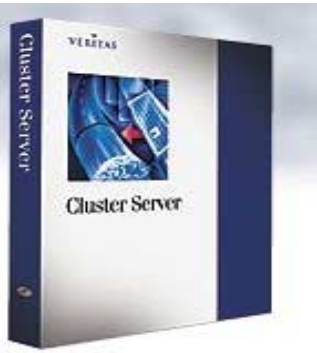
- escon link auto-recover feature build-in and configurable from the vcs gui
- full cluster manipulation to ensure data consistency at the xp array level



Stress-Free Storage...Guaranteed



This slide explains a little bit more about the features of Cluster Extension XP for the Veritas cluster software. So, what HP has done here, and what will be shown on some of the next slides with some screen shorts is that Cluster Extension XP is completely integrated into Veritas cluster server so it looks like this piece is actually coming from Veritas, but it isn't. HP developed Cluster Extension XP based on the very best development framework, which gives the customer the ability to create either some of their own scripts or even a full binary which would leverage or use the full feature set of the cluster software. Cluster Extension XP could be edited online if the customer would want to do that. Cluster Extension XP also supports metropolitan and global solutions using cluster extensions with Veritas cluster server.




## veritas cluster server

- heartbeat support/cluster reformation over DLPI protocol (private network, that is a single subnet; no distance limitation)
- supports 3 level of application dependencies (local and remote)
- up to 32 systems per cluster (os support includes sun solaris, hp-ux, windows nt)
- all clusters can be configured and administered from a single java gui
- vcs disables failover when less then two network connections are available
- vcs1.3.0 comes with a configuration wizard to copy complete clusters

HP Restricted

N S S O  
Stress-Free Storage...Guaranteed

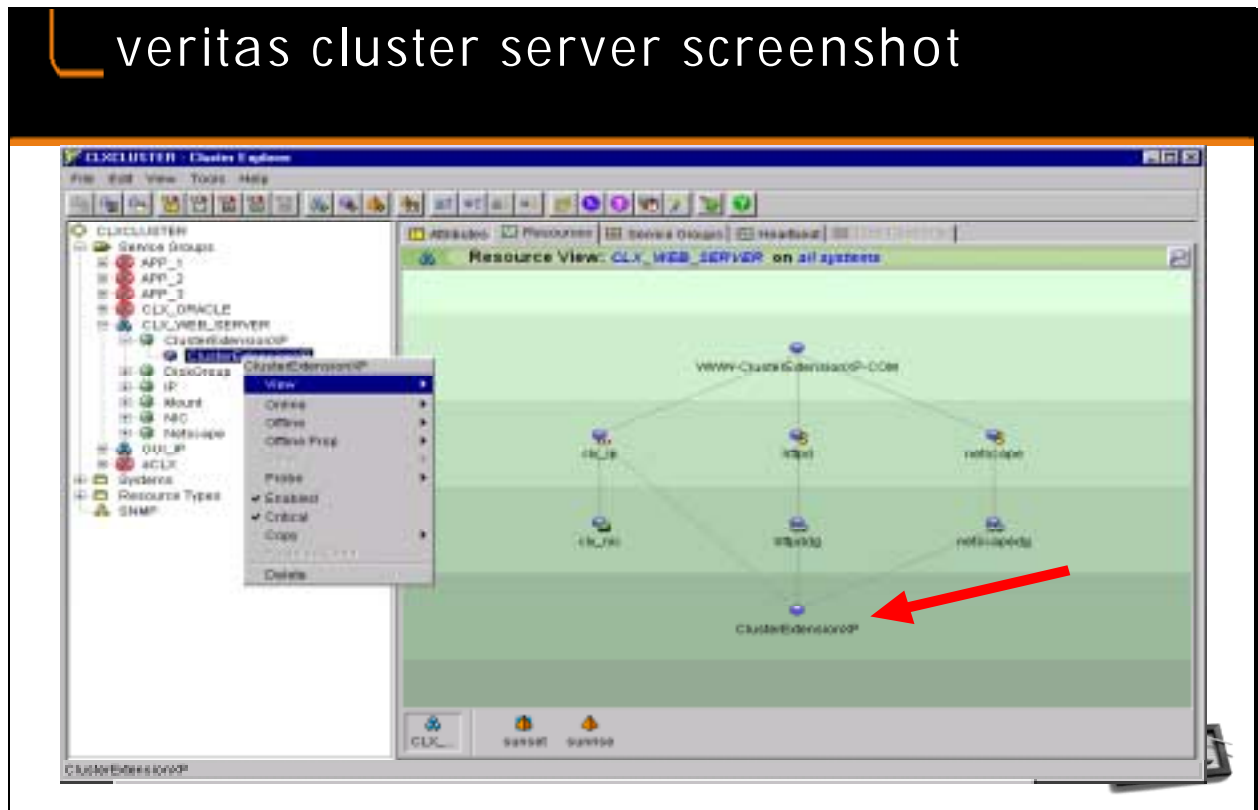


Veritas cluster software has the same kind of heartbeat mechanism as MC/ServiceGuard. Basically, their heartbeat and cluster Information protocol works in the same way as MC/ServiceGuard. So this means the heartbeat must be routed over a private network and can actually not be routed over a longer distance without some limitations. So, what that means is you have to follow the rules for private networking. A differentiation to MC/ServiceGuard is that Veritas cluster software actually supports application level dependencies. Veritas cluster software actually supports package dependencies, meaning one application package is depending or can depend on another application package. And this is done on a local base, meaning on the same server or between different servers. Veritas supports up to 32 nodes at the moment being of types Sun Solaris, HP-UX and Windows NT. HP's Cluster Extension XP software is currently built to run only on Sun Solaris. The nice feature of Veritas cluster server is that you have a nice Java GUI which you will see in a few slides. Veritas cluster server deals with split brain syndromes completely different than MC/ServiceGuard. How to protect against split brain syndrome so that one application can only run on one system and not on two at the same time is quite different from how MC/ServiceGuard deals with it. And it works in a way that you must have at least one network connection between all the servers in the cluster. Otherwise, application startup will be disabled. There are two different ways to integrate the Cluster Extension XP software into Veritas. HP has chosen to implement that in a way that nobody can re-engineer the software, which is

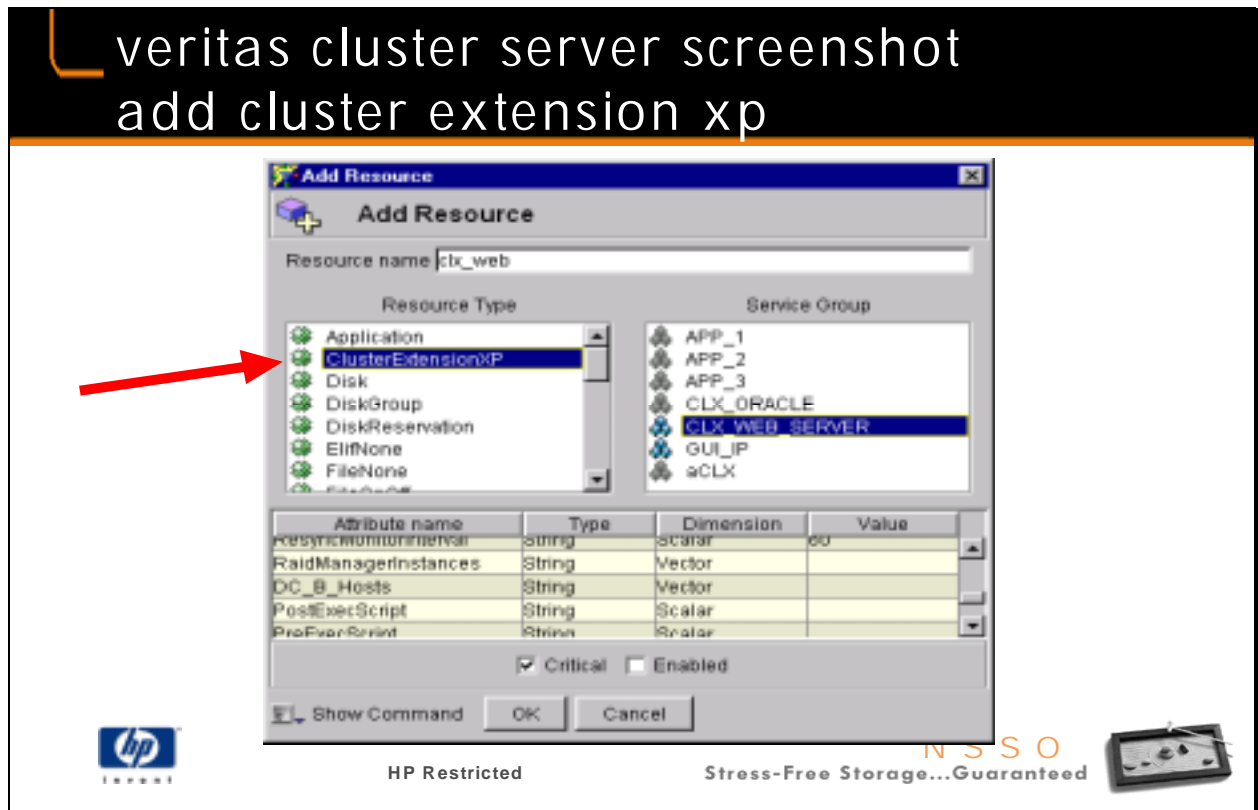
SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Continuous Access XP

binary format which is an executable which is linked and compiled with the VCS agent framework library to form the agent binary. The code is not viewable.

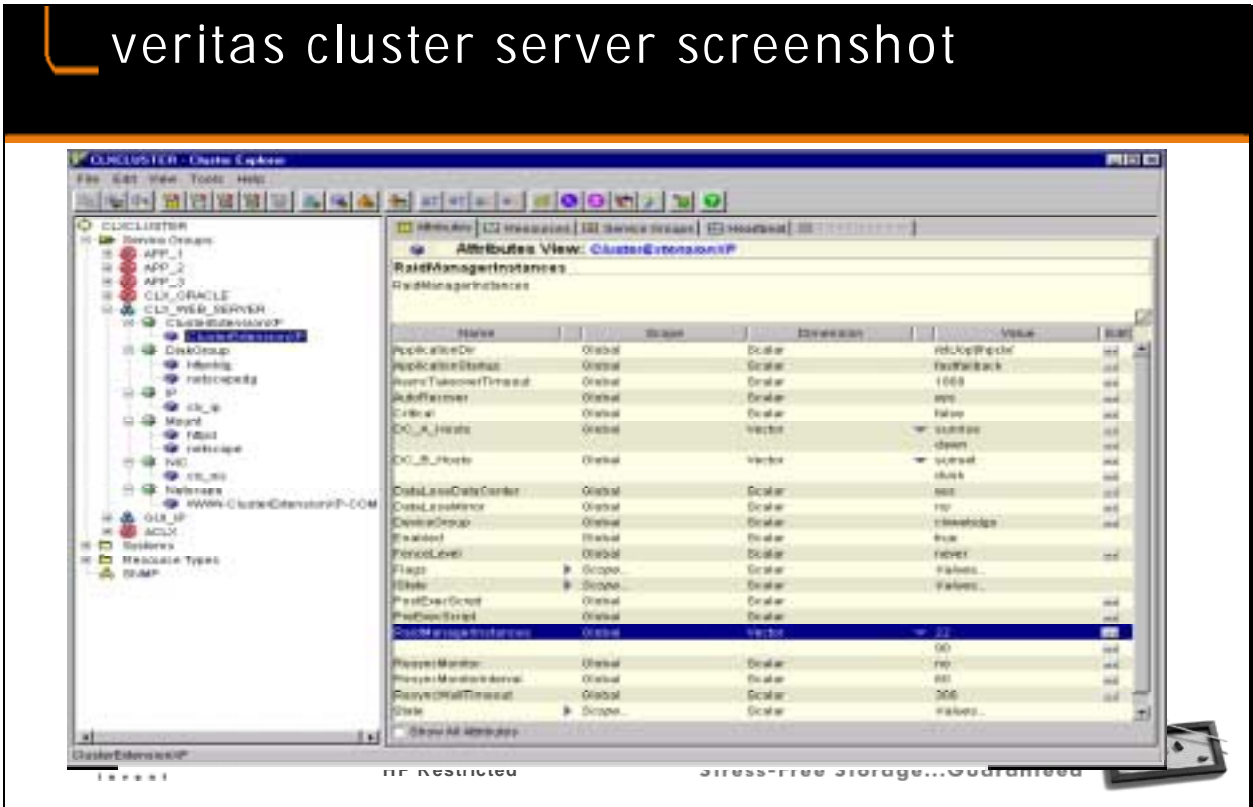




This slide actually shows you a complete picture of a high availability application in a Veritas cluster server configuration. Starting on top is the application which is described as a Web server. The next layer is the IP address and some mount points for the drives. Actually the important part is the lowest layer, which is HP's Cluster Extension XP product which checks the disk array before it starts up all the other resources of the application.



This slide shows you how easy it is to integrate Cluster Extension XP in Veritas Cluster Server. Here you can actually see how Cluster Extension XP is integrated in Veritas Cluster Server. If you want to add Cluster Extension XP to your application you would proceed to a menu which you choose from the resource types and you highlight one of the applications on the right side titled "Service Group". You can actually then set the behavior of cluster extensions fail over capabilities in the box below by specifying the RAID Manager XP instance you use along with some other options. You would then hit the OK button and your cluster extension is more or less set up for a Veritas cluster server.



On the right side of this screen shot you can see all the features or options that you can set up in Cluster Extension XP for each application.


# veritas cluster server screenshot

This slide shows a complete picture of an Oracle high availability application in Veritas cluster server configuration. Where we show by the red arrow HP's Cluster Extension XP is the first level which means it is the lowest level where the application is activated. Therefore the Cluster Extension XP software must be implemented first before any other resources of the software for the application.

## continuous access xp and multi-platform cluster solutions with cluster extension xp

cluster software ca configuration	hp mc sync	hp mc async	hp cc sync	hp cc async	vcs sync	vcs async	hacmp sync	hacmp async	ms sync	ms async
1	3	3	3	3	3	3	3	3	3	3
2	43	43	43	43	43	43	43	43	NSC43	NSC43
3	50	50	50	50	50	50	50	50	NSC50	NSC50
4	SP	SP	SP	SP	SP	SP	SP	SP	SP	SP
5	SP	SP	SP	SP	SP	SP	SP	SP	SP	SP
6	NS	NS	NS	UI	SP	UI	SP	UI	SP	UI
7	NS	NS	NS	UI	SP	UI	SP	UI	SP	UI
8	NS	NS	NS	UI	NS	UI	NS	UI	NS	UI


UI = unlimited SP = support pending NS = not supported ? = to be clarified NCS =not in single cluster



HP Restricted

**N S S O**

Stress-Free Storage...Guaranteed



## continuous access xp fence level\* and multi-platform cluster solutions with cluster extension xp

Max. Distance	Escon direct	Escon directors/ repeaters	Escon over DWDM <sup>1</sup>	Escon over T3 or ATM OC3 (static)	Escon over T3 or ATM OC3 (dynamic) <sup>2</sup>	Escon over IP (static) <sup>1</sup>
3 km	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>
43 km	n/a	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>
50 km	n/a	n/a	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>
> 50 km	n/a	n/a	n/a	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>	<b>NEVER, DATA, ASYNC</b>

\* bold is recommended fence level

1 Heartbeat network over DWDM or IP. The cluster heartbeat communication could be routed over the same DWDM or IP links as the Continuous Access XP data traffic. However, the networks should be separated because of link congestions and the higher potential of "Split-Brain" syndrome in case of network component failures or link failures.

2 Dynamic routing. Dynamic routing in the wide area network must be transparent to the Continuous Access XP links otherwise dynamic switched network connections could lead to suspended XP disk pairs.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## EMC Symmetrix Remote Data Facility

- SRDF has a master and slave site. If the master site goes down, the slave site continues. After restoring the master site, resynchronization requires system downtime.
- CA solves this problem. When the primary site goes down, primary switches to the secondary site and stays online. After restoring the site, the new primary is mirrored. Then the primary is switched to the original site.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Continuous Access XP Extended

- Asynchronous mode is compatible with IBM's DFSMS and XRC software
- Subsystems that honor the IBM compatible software commands may be from different vendors
- "Semi-Synchronous" on EMC Symmetrix
  - Only one outstanding write per volume before it blocks. This write must be acknowledged before another write is allowed
  - Guarantee of data consistency on the remote Symmetrix frame
- "Adaptive Copy" on EMC Symmetrix
  - No transaction time stamping
  - Allows from 1 to 65,535 outstanding I/Os per volume
  - No guarantee of data consistency

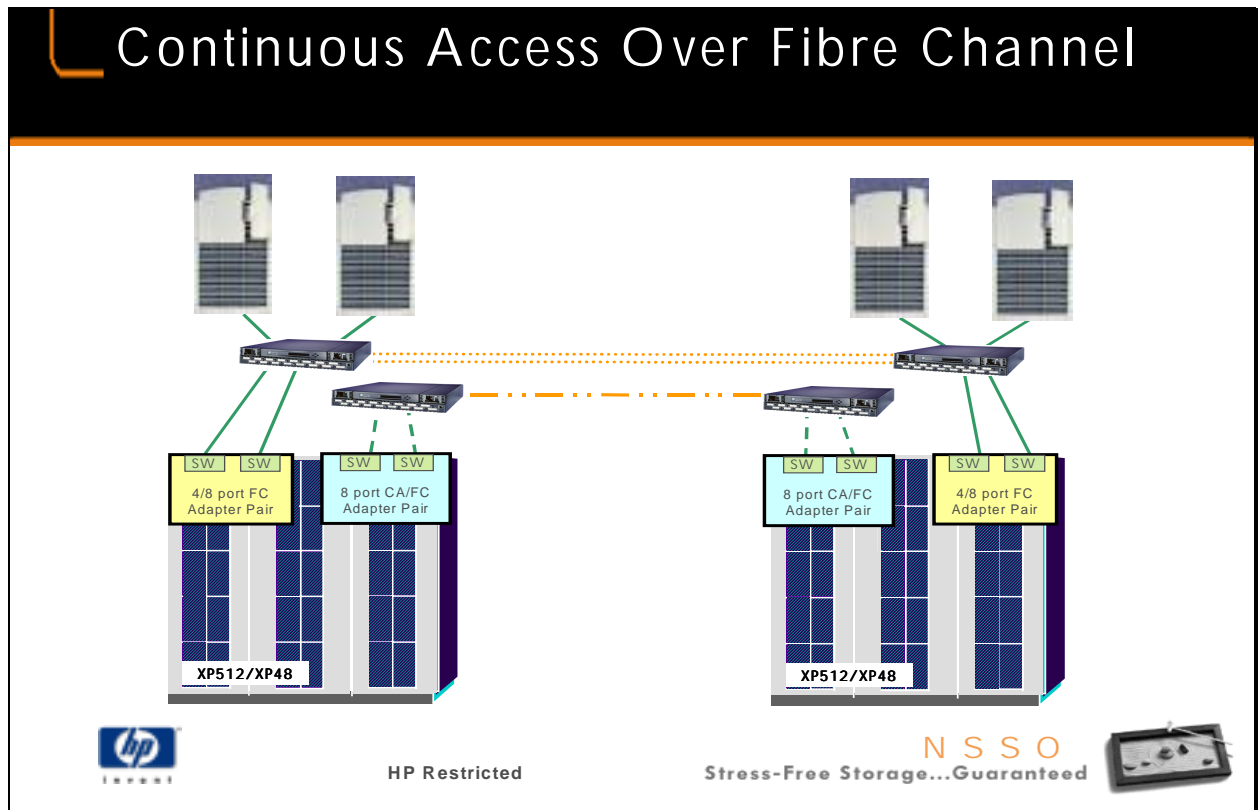


HP Restricted

N S S O  
Stress-Free Storage...Guaranteed







## New CA Over Fibre Channel Link

- CA over Fibre Channel now supported in addition to CA over ESCON
- New Fibre Channel CHIP supports CA as well as host connection as of (7/1/01)
- CA over Fibre Channel provided over all the historical CA link types except ESCON
  - DWDM
  - Switched Circuit (ATM, DS3, IP, T3)
- CA Over Fibre Channel supported on FC switches and N




HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Continuous Access XP Over FC: Overview

### High Performance Remote Mirroring Solutions!



### ➤ What does CA over FC offer?

- High Link Throughput (~50MB/s peak)\*
- Simplicity of Solution Management
- Power of Asynchronous Copy Fast Failover/Failback operations
- Seamless Integration with High Availability Clustering Solutions for HP-UX, Sun Solaris, IBM AIX, and Windows NT/2000
- Part of HP's end-to-end Five Nines/Five Minutes Program
- The most robust clustering solution for open systems available!

\*Fibre Channel standard provides for 100MB/sec. Initial release of CA over FC only provides 50MB/sec



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## CA Over FC: CHIP Product Structure/Pricing

### ➤ Product Structure

XP48: A5927A/U, XP512: A5957A/U (A suffix: initial array, U suffix: upgrade)

### ➤ Pricing

Once we roll to new 8-port adapter pair for host connect, position existing 4 port FC adapter pair as low cost entry point & New 8 port FC/CA adapter pair as having more functionality at a slight premium

8-port CA/FC adapter pair WW Reference = \$104,491 (existing 8-port FC adapter pair : \$94,992)



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## CA over FC: CHIP Board

### FC-CA: New FC CHIP Pair

- **Specifications (XP48: A5927A/S, XP512: A5957A/S)**
  - 8-port card pair only (short-wave GBICs only)
  - 1Gb/sec FC, 5X ESCON per link throughput (using 4 ports/pair)
  - Replaces existing 4/8 port ESCON CHIP pair for CA connect
  - Support for synch and asynch copy modes
  - XP512 and XP48 only (no support for XP256)
  - **CA and Host connects on the same CHIP pair as of 7/1/2001**
- **Dramatic throughput improvement**
- Will coexist with existing 4/8 port FC CHIP pairs which will continue to be used for host connect only
- Continue to make available existing ESCON CHIP for installed base (product is otherwise required for M/F and XP256 CA connect)
- Requires firm ware release 1.13.xx or later



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## CA over FC: Solution Requirements

### ★ Solution HW requirements:

XP512/XP48 (no XP256 support)

New CA/FC Adapter Pair (one per array)

Unique Remote link/infrastructure requirements:

- ◆ #1 Direct Connect Fibre Channel cables
- ◆ #2 Extended Connect via Dedicated FC Switches
- ◆ #3 Extended Connect via Shared FC Switches (hard zoning required)
- ◆ #4 Extended Connect via FC DWDM interface(s)
- ◆ #5 Extended Connect via FC WAN Converter interface(s)



HP Restricted

Stress-Free Storage...Guaranteed



## CA over FC: Solution Requirements

### ➤ Solution SW requirements:

Continuous Access XP/XP Extension (existing product): no new product requirements, license per TB

Dependency on new Array Firmware (> 1.13.xx)

All other SW requirements same as existing ESCON offering

### ➤ Intermixing ESCON and FC CA in the same array MCU only



HP Restricted


Stress-Free Storage...Guaranteed



## CA over FC: CHIP Board

### FC-CA: Configuration Considerations

- CA FC data flow is uni-directional (data only moves from the MCU -> RCU) for a given FC port even though standard Fibre Channel is bi-directional
- Adjacent ports on the CHIP must be configured the same (either both "Initiator" or both "Target")
- A minimum HA solution with a single CA FC CHIP pair would require at least one initiator and one target on each CHIP board to support bi-directional HA capabilities



HP Restricted

IN S S U  
Stress-Free Storage...Guaranteed





## CA over FC: CHIP Board

- Adjacent ports on the CHIP must be configured the same (either both “initiator” or both “target”)
- Two types of targets, “Target” (for host connectivity) and “RCU Target” (for CA connectivity)

The diagram shows a 4 port FC CHIP card with two Initiator ports (green) and two Target ports (blue). Each port has a T (Target) and an R (RCU Target) sub-port. Two orange cables connect the RCU Target ports of two cards. A green starburst contains the text: "Host Connectivity Not Available Until 7/1/2001".

**Port Change**

Port#	Cluster	Type	Mode
CH J/K	1	RCU Target	Standard
CH L/M	1	Initiator	Standard
CH J/K	2	RCU Target	Standard
CH L/M	2	Initiator	Standard

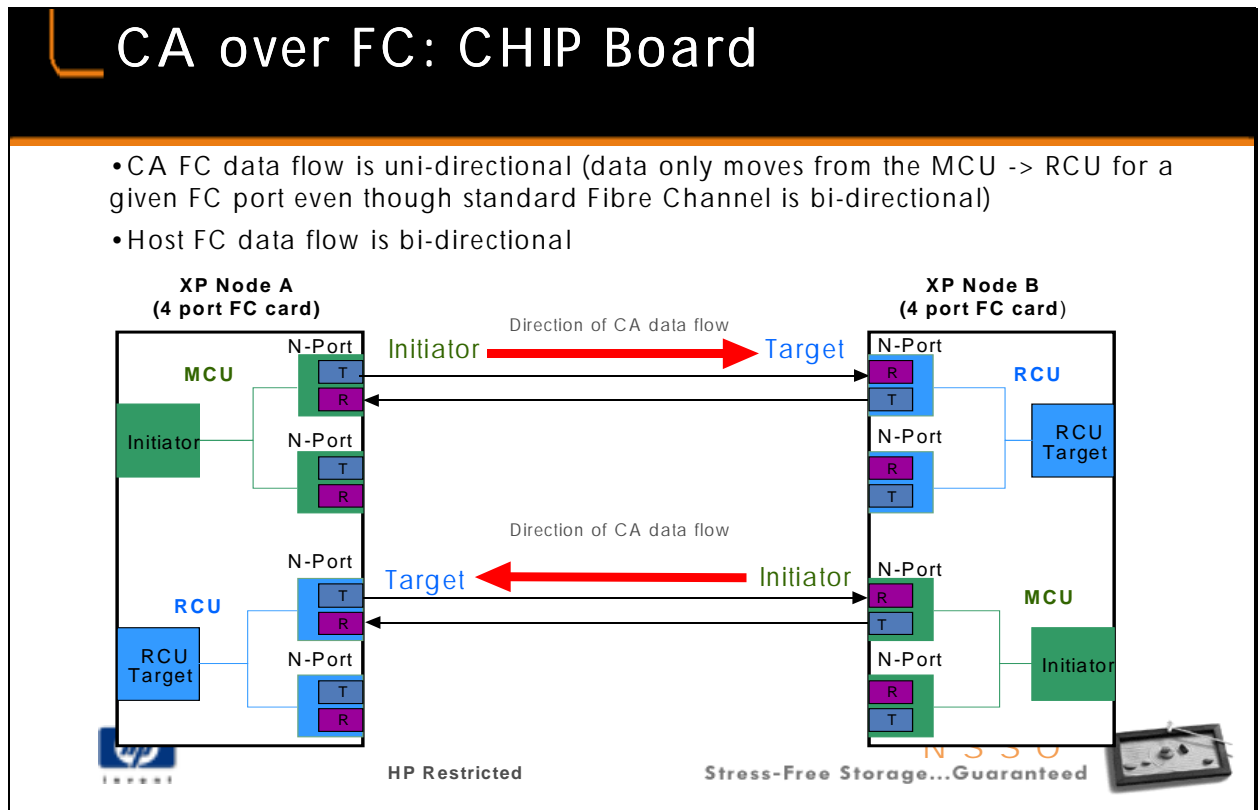
Buttons: >RCP, >LCP, >Initiator, >Target, >RCU Target, Close

Message :

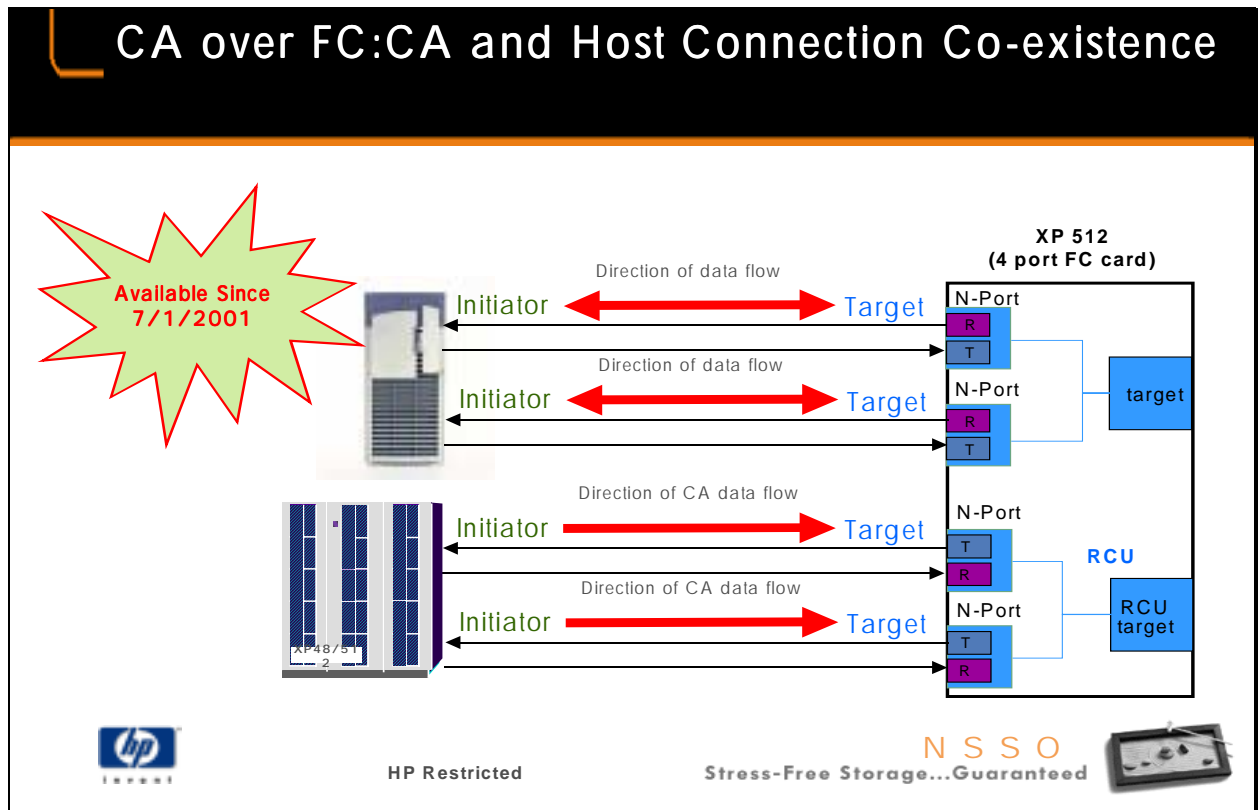
HP Restricted

N S S O  
Stress-Free Storage...Guaranteed

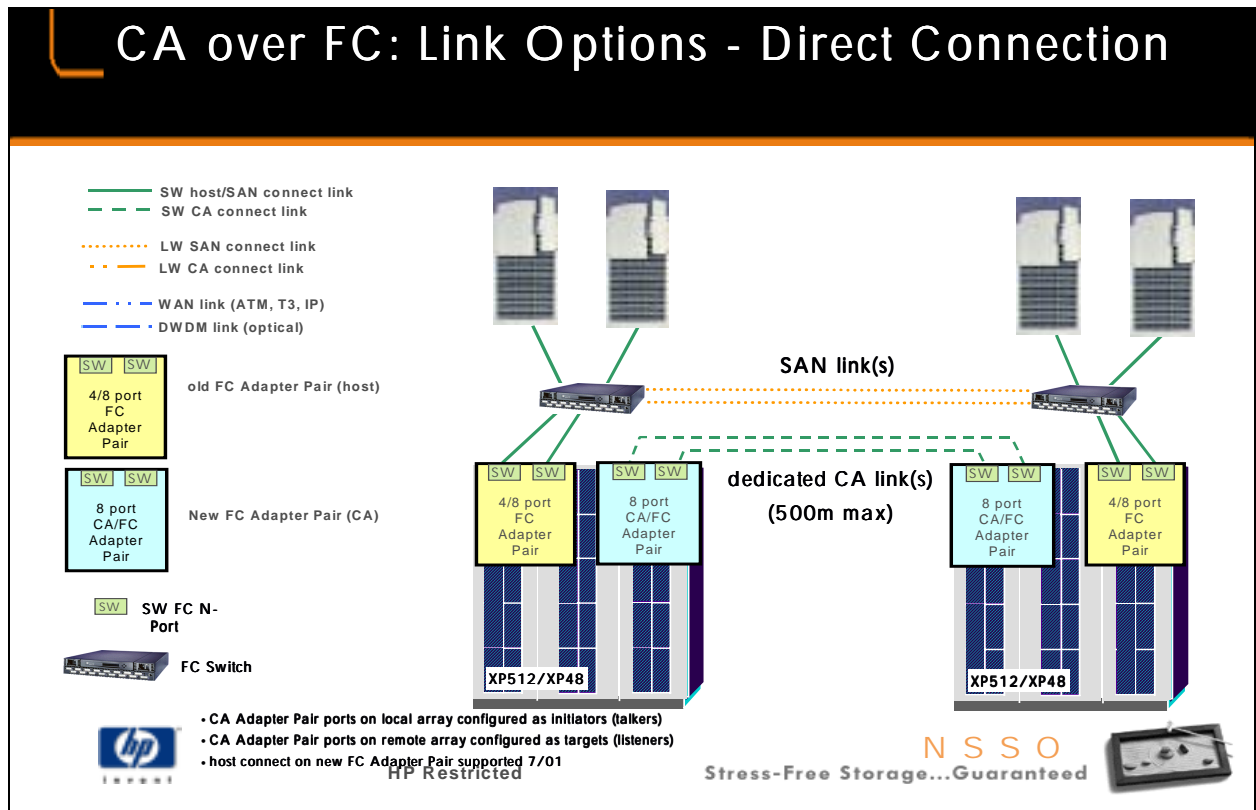
Notice that on a given card, 4 FC ports are serviced by 2 FC controller ICs. This dictates that 2 ports in a group (if both are used) must be both initiators (talkers) or targets (listeners). Since I/O cards always come in pairs, and each FC-CA card contains two 50MB/s ports (for high-speed use, or four 25MB/s ports for low speed), the recommended FC-CA configuration is one FC connection in each direction (one port as initiator and one port as target). This provides for more than adequate bandwidth across a FC CA CHIP pair (100MB/s each way) plus redundancy and bi-directional failover /failback ability.



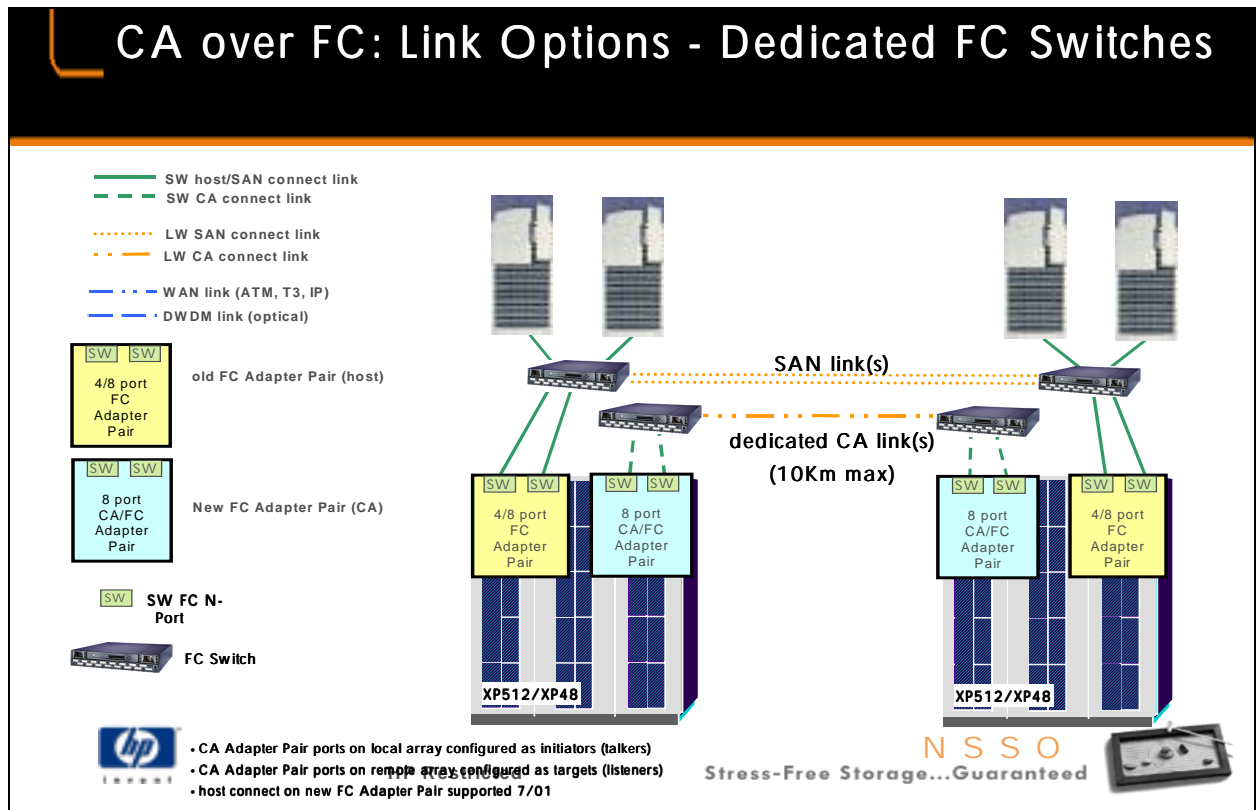
Notice that on a given card, 4 FC ports are serviced by 2 FC controller ICs. This dictates that 2 ports in a group (if both are used) must be both initiators (talkers) or targets (listeners). Even though a standard Fibre Channel connection is capable of bi-directional data transfer, CA over FC only moves data in one direction (MCU -> RCU) even though acknowledgments move from RCU->MCU.



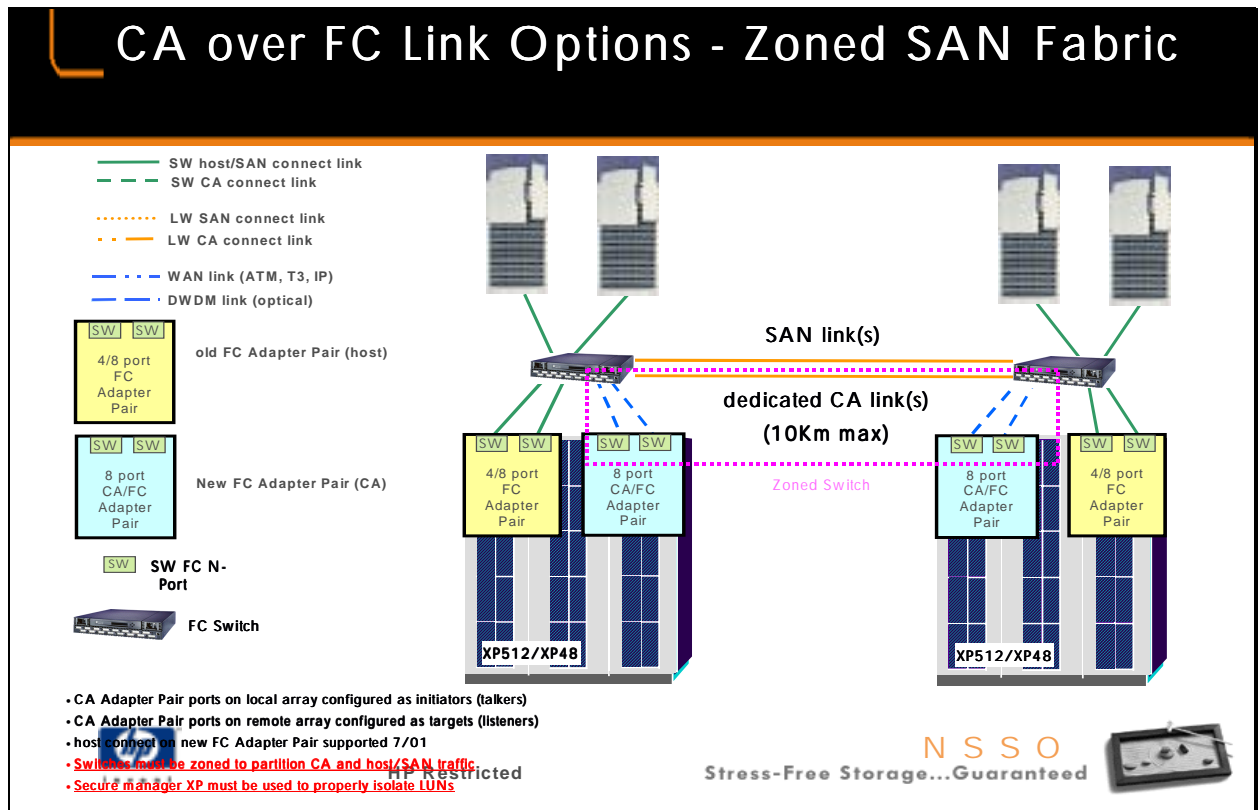
Notice that on a given card, 4 FC ports are serviced by 2 FC controller ICs. This dictates that 2 ports in a group (if both are used) must be both initiators (talkers) or targets (listeners). When connectivity is paramount over performance, use all 4 ports @25MB/s. When performance is paramount over connectivity, use just 1 port in each portpair as shown @50MB/s each. The configuration shown is capable of 100MB/s real world throughput (2x50MB/s per card). This linear scalability is expected to continue up to at least 150MB/s by also using the second card in the pair. A maximum configuration employing FC-CA is expected to provide a total throughput which is greater than 3 times what was possible using ESCON-CA, while using fewer slots. Since I/O cards always come in pairs, and each FC-CA card contains two 50MB/s ports (for high-speed use, or four 25MB/s ports for low speed), the recommended FC-CA configuration is one FC connection in each direction (one port as initiator and one port as target). This provides for more than adequate bandwidth across a FC CA CHIP pair (100MB/s each way) plus redundancy and bi-directional failover /failback ability.



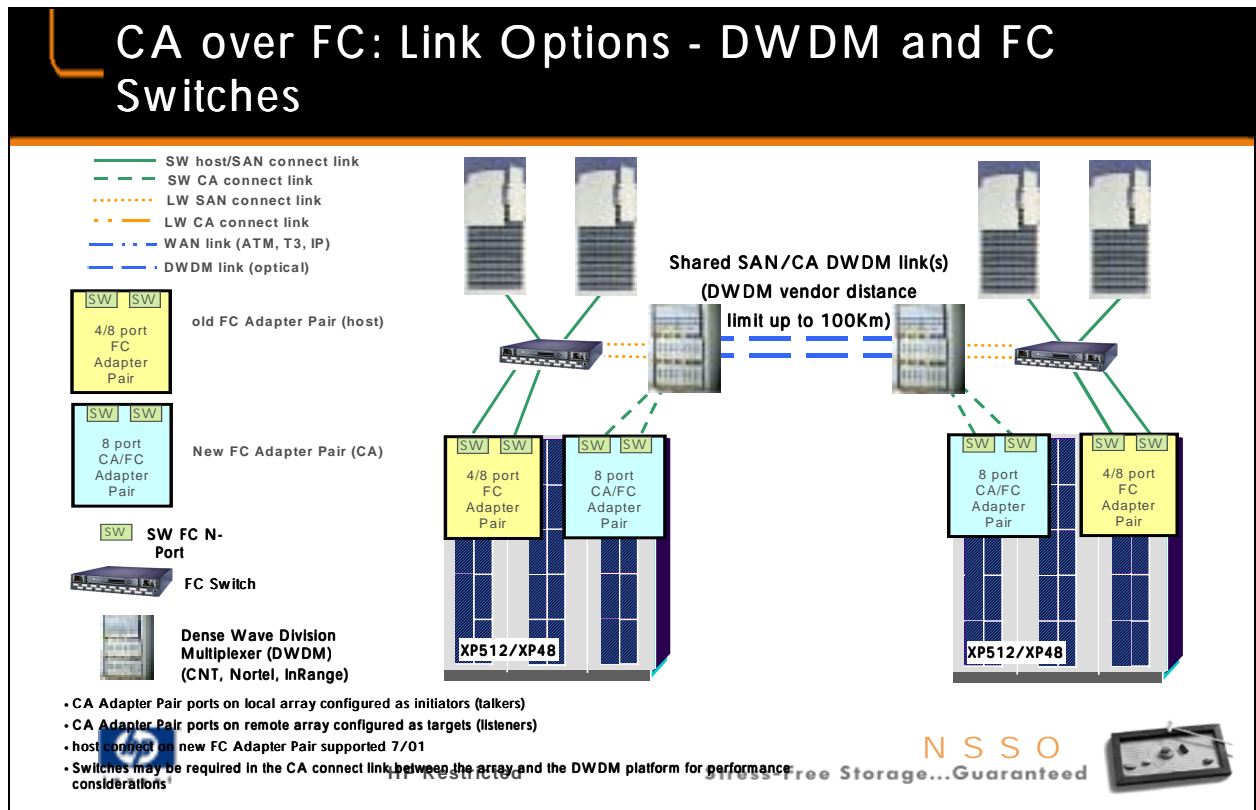
The first variation (below) is “point to point” or direct connect FC. The CA FC CHIP port must be configured as NL-Port. Standard arbitrated loopFC uses a short-wave optical connection between NL-ports at distances up to 500 meters. Notice that the CA link connection is dedicated and not mingled with the regular SAN. As in any SAN situation when LVM's are in use, the owner of the CA primary volume (e.g. on the left) can not be allowed to “see” the secondary volume of the pair (and Visa Versa). To do so would cause an LVM error. Therefore, Secure Manager/XP or other suitable mechanism (e.g. RAID Manager Security or a special option to paircreate that hides the s-vol) must be used to prevent this. Unless otherwise noted, this requirement also exists for all to configurations in the following pages. Each FC link consists of an optical fiber pair. Each of the two fibers in the pair is used in a one-way fashion. As a pair, the bulk of the traffic is intended to be in a single direction (e.g. talker-to-listener or visa-versa) with the second fiber typically used just for acknowledgements.



This configuration employs a dedicated SAN utilizing E-Ports and long-wave GBICS to achieve up to 10KM distances. CA FC ports on the CHIP can be configured as either N-Port or NL-Port. E-Ports ("E" stands for Extended) between the FC switches allow for up to 10km long-wave optical connections. If both arrays have a 500meter connection to the switch, the total distance can be up to 11km. The CA FC CHIP card has a 4 buffer credit limit per port, meaning that up to 4 writes can be outstanding on the CA link at one time. Those 4 buffer credits are more than enough to support the latency involved with 500 meters and keep the "pipe" full. The 16 buffer credit feature of the added FC switches allows a distance of up to 10km. Solutions longer than 10km(e.g. via DWDM) should ideally involve FC switches capable of extended fabric support (more than 16 buffer credits, e.g 64).



This interconnect option shows CA over FC and Fibre Channel connected hosts sharing the same FC fabric. CA FC traffic and FC connected host traffic are isolated from one another via switch zoning. As with the other interconnect options show, SecureManager must be used to ensure a host cannot see both the P-vol and the S-vol of a CA pair.



This interconnect option involves the use of a passive optical multiplexing device known as DWDM (Dense Wave Division Multiplexer). In this case, the regular FC SAN and CA FC traffic can share the DWDM link. However, Secure Manager/XP is required, because in the larger context it would still be possible for a CA primary volume owner to incorrectly also see the secondary volume via other (non-CA) ports. At these distances, it is necessary to insert a FC switch (i.e. Brocade Silksworm with the Expanded Fabric Option) between the arrays and DWDMs with special (64 buffer credit) firmware. Since the number of available buffer credits directly effects the number of IOs that can be outstanding, this enhanced ability becomes critical to performance beyond 10km and allows DWDM support up to 100km. In the case shown, the 100km max applies to the entire distance between the 2 switches because class-3 FC is being used in a buffer-to-buffer fashion.

## Module Wrap-up




HP Restricted

N S S O  
Stress-Free Storage...Guaranteed









# Module 10

## RAID Manager XP – Included in the purchase of Continuous Access XP and Business Copy XP



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



HP SureStore E RAID Manager XP (RM) utilities enable you to issue Business Copy XP (BC) and Continuous Access XP (CA) commands. These commands can be issued from the command line or built into a script (for example, a ksh or perl script or DOS batch file). By using RM scripting, you can set up and execute a large number of CA and BC commands within a short period of time.

## Module Agenda



- RAID Manager XP Overview & Specifications
- RAID Manager XP Specific OS Product Numbers



- RAID Manager XP Example Use & Configuration
- RAID Manager XP Software Files



- RAID Manager XP Commands
- RAID Manager XP Script Example
- Wrap-Up



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Overview of RAID Manager XP

- RAID Manager XP is a host-based software package initiated from a server via a command line interface.
- RAID Manager XP OS Specific (all versions currently ship on CD).
- Perform BC and CA commands through the command line or from a script.
- Main interest is that commands can be automated.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



The RM allows BC (and/or CA) pair management from the host. It exhibits the following characteristics:

- Supports the various BC Commands (via the Command Device): Pair Create, Pair Split, Pair Resync, Pair Display, etc.
- Requires at least One LDEV as Command Device in each XP (Min. size 36MB if using CVS)
- Notifies the OS (via syslog file entry) of any pair volume problems.

The HP SureStore E RAID Manager XP (RM) software product enables you to perform operations with HP SureStore E Continuous Access XP (CA) and HP SureStore E Business Copy XP (BC) on the disk array by issuing commands from the server host to the disk array. The RM software interfaces with the system software and high availability (HA) software on the server host, as well as with the BC and CA software on the disk array.

UNIX - RM runs on a UNIX host in the form of a daemon named HORCM.



Windows NT - RM runs on a Windows NT host as a service. RAID Manager XP enables the user to run individual or scripted Business Copy and Continuous Access operations from a host system. Business Copy and Continuous Access software functions as a component of RAID Manager displaying remote copy information and allowing continuous access operations through either the command line, scripts, or a GUI. Commands and communication from the host are sent to the command device on the HP SureStore E Disk Array XP Family of Arrays. The command device is a logical device (LDEV) which has been

selected to function as the command device. RAID Manager software consists of:

1. RAID manager instances
2. Configuration file
3. Command device
4. CA/BC volume pairs
5. CA/BC commands and shell scripts

# RAID Manager XP

- RAID Manager is OS-specific
  - RAID Manager XP for HP-UX
  - RAID Manager XP for Windows NT
  - RAID Manager XP for Solaris
  - RAID Manager XP for AIX
  - RAID Manager XP for Digital Unix
  - RAID Manager XP for Sequent DYNIX host systems
  - Linux

 HP Restricted N S S O  
Stress-Free Storage...Guaranteed 

RAID Manager is included with purchase of Continuous Access XP and Business Copy XP. RAID Manager provides OS-specific, host based software initiated from a server via a command line interface that allows full user control of Continuous Access XP (and Business Copy XP) from the host server. RAID Manager also provides functionality for synchronized operation with cache and OS, etc. If launching from a host using scripts, a consulting engagement may be recommended. RAID Manager must be installed on each host acting as failover controller/arbitrator. The customer can control this using Command View XP software. OS-specific, host based software initiated from a server via a command line interface that allows full user control of Continuous Access XP and Business Copy XP from the host server. RAID Manager XP must be purchased together with these two SW packages for full functionality and synchronized operation with cache and OS, etc. If launching from a host using scripts, a consulting engagement may be recommended. EXAMPLE: Scripting can be used to implement Metro Cluster ServiceGuard fail-over. The control lines are LAN-based and the data lines are ESCON-based. If a customer wants to manage from the host server using RAID Manager, and it is scripted at both the local and remote sites, then one RAID Manager will talk to a daemon at the copy of RAID Manager at other site through the LAN connect. In the event of a failure the remote site is notified, the mirrors are broken and the remote site

takes over. RAID Manager XP provides a number of commands to perform maintenance and control functions of the volume pairs. Among the functions RAID Manager XP can perform, it can generate a volume pair, split a pair, resynchronize a pair, take over a pair, and monitor and display information about the pair. Commands can be automated in scripts.

## RAID Manager XP Example

- Scripting can implement a MetroCluster / ServiceGuard fail-over.
- RAID Manager scripted at both local and remote sites
- One RAID Manager talks to a daemon at the other site through the LAN interconnect
- If the local site fails, the remote site is automatically notified
- Mirrors are broken and the remote site takes over



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



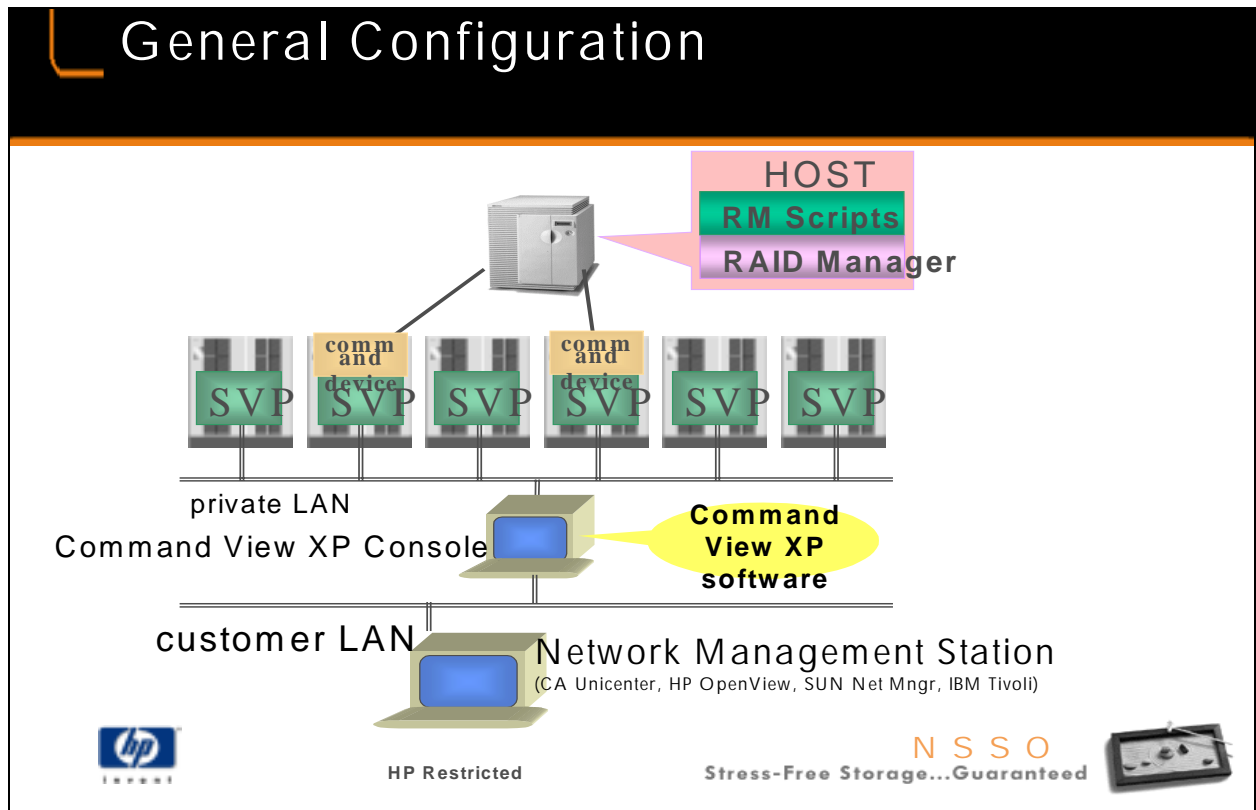
Each copy of RM is known as a RM instance. Instances are local or remote and can run on the same host. A local instance is the instance currently being configured or the instance that commands are currently being issued to. A remote instance is the instance that the local instance communicates to, as configured in the HORCM\_INST section of the RM instance configuration file. Local instances link to remote instances by using UDP socket services. RM provides fail-over and operation commands to allow mutual hot standby, in cooperation with fail-over software (for example MC/Service Guard or Metro Cluster).

### Using RAID Manager in HA Environments

When using HA software (such as MC/ServiceGuard), application packages can be transferred to the takeover node at any time. If the application package transfer operation is performed in an environment where CA is used, you may need to switch the CA secondary volumes to primary volumes. The horctakeover command provides this function. The horctakeover command provides macro functions to determine the data consistency of the S-VOL and to perform takeover functions:

- takeover-switch
- swap-takeover
- SVOL-takeover
- PVOL-takeover

The horctakeover command is not available for BC.



RAID manager communicates in-band with the command device of the HP SureStore E Disk Array XP Family of Arrays via the SCSI or FC interface. A user-selected logical volume on the disk array functions as the command device. The command device accepts CA or BC control operations, seen as read and write commands, which are then executed by the disk array. The command device returns read requests to the open-system host. The volume designated as the command device is used only by the disk array and is blocked from the user. The command device can be any OPEN-3/8/9/E device that is accessible by the host. The command device uses 16MB of space. RM issues disk read/write commands to the command device through either a UNIX or Windows NT interface. If the command device fails for any reason, all CA/BC commands terminate abnormally and the host cannot issue commands to the disk array. To avoid data loss and system downtime, it is possible to set an alternate command device. When RM receives an error notification in reply to a read/write request, it switches to the alternate command device. Commands and communication from the host are sent to the command device on the HP SureStore E Disk Array XP Family of Arrays. The command device is a logical device (LDEV) which has been selected to function as the command device.

Each copy of RAID Manager is called an instance.

RM commands support fail-over mechanism (MC/ServiceGuard or MetroCluster).



SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
RAID Manager XP

A RAID Manager script is a list of instructions or macros in a file that execute CA/BC operations.

Scripting operations for UNIX uses a shell script file.

Scripting operations for NT are defined in a text file.


The host executes the script file and sends the commands to the command device on the disk array for automatic execution.

## HORCM File

- HORCM stands for **H**itachi **O**pen **R**emote **C**opy **M**anager  
(also known as RAID Manager)
- There are four major sections in the horcm file.
- These sections contain ...

Information about:

1. Local Host  
(HORCM\_MON)
2. Command (CMD)  
Devices  
(HORCM\_CMD)
3. Pair Members  
(HORCM\_DEV)
4. Remote Host  
(HORM\_INST)



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Components of RAID Manager Software

- RAID Manager instances
- Configuration file
- Command device
- CA/BC volume pairs
- CA/BC commands and shell scripts



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



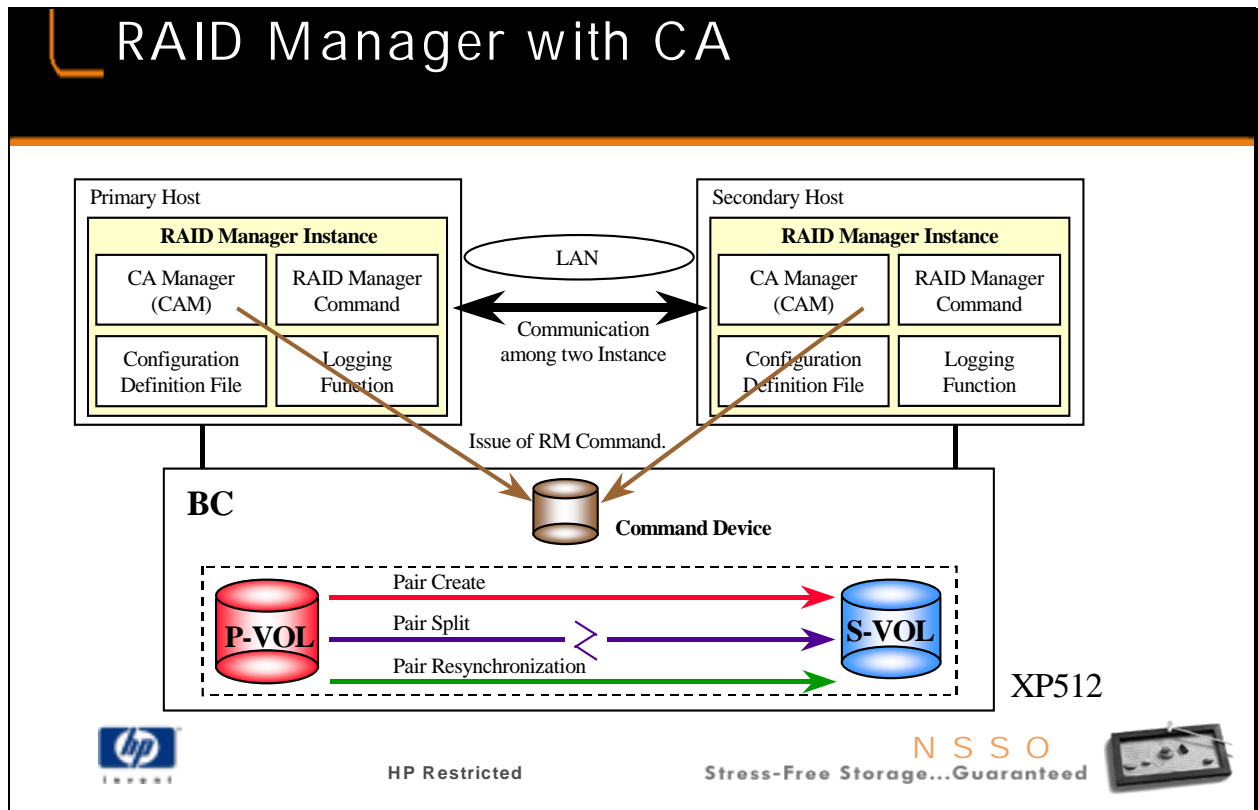
RAID Manager software consists of:

1. RAID Manager instances - Each copy of RAID manager is known as a RAID Manager instance. It provides fail-over and operation commands in cooperation with host fail-over software (MC/ServiceGuard) (Instance per volume involved in a Business Copy or Continuous Access pair.) The INST section of the configuration file contains network information about the other volume of the pair.
2. Configuration file - a text file that provides a definition of hosts, volumes, and groups known to a RAID Manager instance. It includes a) monitoring and communication parameters, b) command device definition, c) volume pair member device definition and identification, and d) physical device mapping and network information (i.e. IP address of the host)
3. Command device - a designated LDEV on the HP SureStore E Disk Array XP Family of Arrays which accepts the Business Copy and Continuous Access read and write commands from the host. This LDEV is blocked from any other use. It is not accessible by the user and cannot be part of an extended LDEV.
4. Volume pair - pairing of source and target logical units for Business Copy; pairing of main and remote logical units for Continuous Access. The DEV section of the configuration file contains the address and pairing information of the volume pairs.

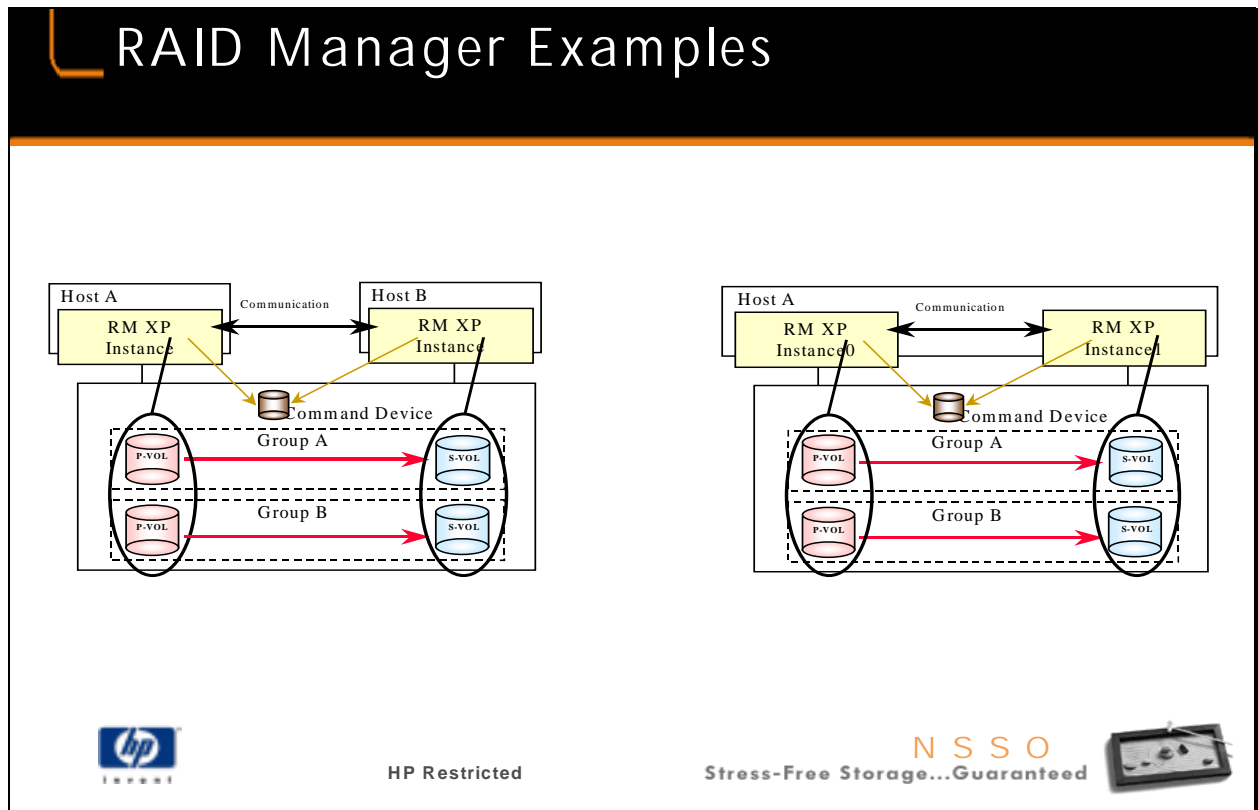
SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
RAID Manager XP

5. Scripts - a shell program containing RAID Manager commands which initiate Business Copy and Continuous Access operations. Data duplication and recovery procedures are thus automated by the host. The major HP application of for this feature is MetroCluster.

RAID manager operations require root access.



A RM instance manages/controls BC (and/or CA) paired volumes by communicating with another RM instance via LAN. The two sides of a volume pair must be registered in different RM config (configuration) files. The pair of RMs required to manipulate a pair may or may not be on the same host. Volume pairs are created by using CA or BC commands. The pair consists of a primary volume and a secondary volume. The minimum size of a volume pair using VSC is 36MB. The maximum size of a volume pair using LUSE is 267GB. The maximum number of pairs for CA is 128. The maximum number of pairs for BC is 512.



**RM XP Instance0**

**Host A**

**RM XP Instance1**

**P-VOL**

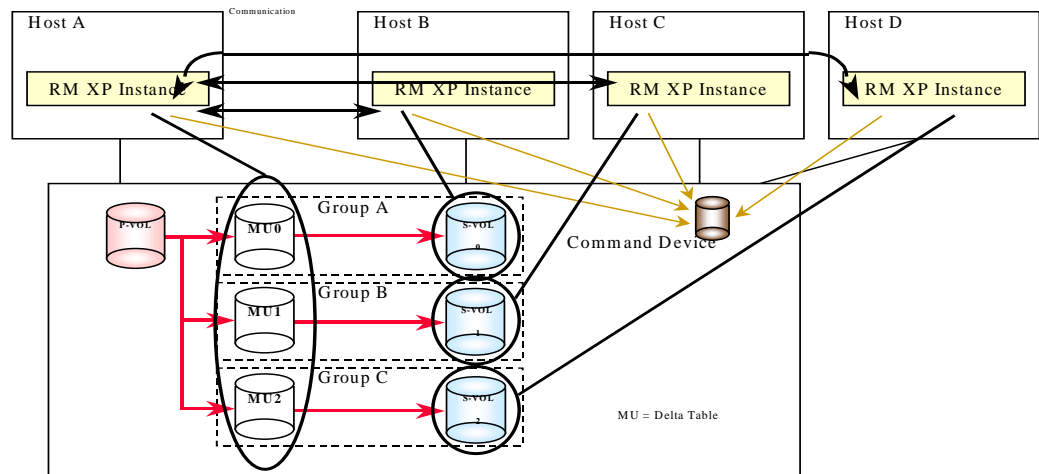
**Command Device**

**S-VOL**

**Group A**

Here are some examples of using pairs of RMs (on one or more hosts) to manage BC (and/or CA) pairs:

## RAID Manager Examples (cont.)

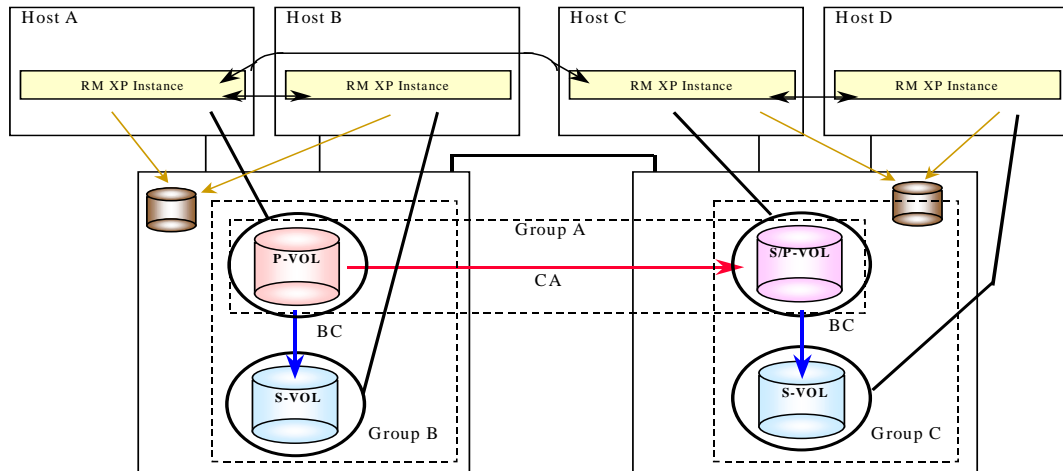


HP Restricted

NSSO  
Stress-Free Storage...Guaranteed



## RAID Manager Examples (cont.)

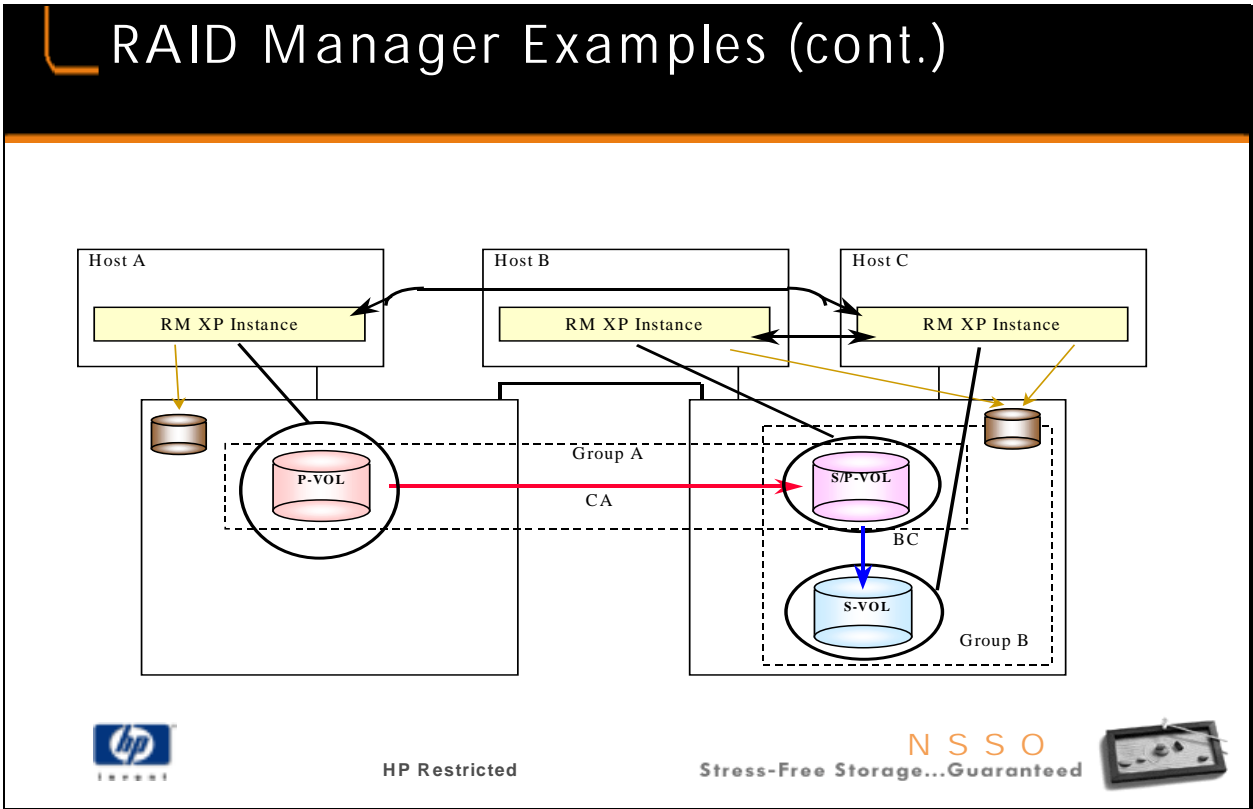


HP Restricted

NSSO  
Stress-Free Storage...Guaranteed










## Configuration File Definition

- **HORCM\_MON** This is where the local instance communication parameters are set.
- **HORCM\_CMD** Here the command device, used for in-band communication between the instance and the array is defined.
- **HORCM\_DEV** This section defines groups of devices. Requires Device Group, Device Name, Port, TID, LUN
- **HORCM\_INST** Here the groups defined in the DEV section are mapped to hostnames of the instances that are physically connected to them.
- \* **HORCM** is Hitachi Open Remote Copy Manager



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



The RM instance configuration file defines the link between a volume and an RM instance. It also defines the relationships between RM instances. You must create a configuration file for each RM instance. The Raid Manager configuration file has four sections consisting of: HORCM\_MON, HORCM\_CMD, HORCM\_DEV, HORCM\_INST. When using HA software (such as MC/ServiceGuard), packages can be transferred to the standby node at any time. If the package transfer operation is performed in an environment in which CA is used, you may need to switch the CA secondary volumes to primary volumes. The MON section is where the communication parameters are defined for RAID Manager instance on the local host.

**ip\_address** - This is a host name corresponding to the interface on which RAID Manager should communicate.

**service** - The name of the service (or port number) defined in /etc/services (or %SystemRoot%\system32\drivers\etc\services for NT)

**poll** - This is the interval in milliseconds in which the pair volumes are monitored for errors.

**timeout** - This sets the communications timeout period with any remote hosts, in a multiple host topology.

Although HA software packages are classified into primary (active) and secondary

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
RAID Manager XP

(standby) packages, they are not strictly distinct. The active package is switched to the standby package, and vice versa, making it necessary to be able to switch the primary

volume to the secondary volume depending on the package state. The takeover command (horctakeover) provides this function. The CMD section is where the command device on the disk array is defined. Multiple command devices can be defined to provide failover functionality in case the original command device fails.

dev\_name - A raw device filename

UNIX: /dev/rdisk/c8t15d0

NT: \\.\PHYSICALDRIVEX

(where X is the number of the unassigned disk device as assigned by the NT Disk Administrator.)

The DEV section specifies the mapping between groups of device names and the port, SCSI ID and LUN mapping of a LDEV, LUSE or VSC volume defined on the array.

dev\_group - An unique arbitrary string tag used to name a group of devices, this group is then associated with the host name of an instance in the INST section.

dev\_name - An unique arbitrary string tag describing a single device on the array, which is specified via the port, target ID and LUN mapping which has already been defined on the array via the Remote Console or the SVP.

**\*\*Note\*\*** - dev\_group and dev\_name must match in the config file on two instances to define a valid volume pair.

port # - The physical port description on the DKC (eg CL1-D)

TargetID - The SCSI target ID assigned to the volume.

LU# - The Logical Unit Number assigned.

MU# - The mirrored unit number assigned for BC. 0 is the original copy, 1 the first mirror, 2 the second mirror. A maximum of 9 mirrors are allowed.

The INST section defines the host names for the arrays which hold the volumes defined in the DEV section.


dev\_group - The device group name specified in the DEV section.

ip\_address - The host name (or IP address) of the instance where the device group has been defined.

service - The name of the /etc/services entry (or port number) for the instance specified. (or SystemRoot%\system32\drivers\etc\services for NT)

## HORCM File – Local Host (HORCM\_MON)

- Local Host  
Information includes



1. Local Host Name(or IP Address)
2. Local Service Name entry (or Number)
3. Poll & Timeout values




HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## HORCM File – CMD Devices (HORCM\_CMD)

- Command (CMD) Devices are device files that allow RAID Manager to access the XP Array in-band through Fibre Channel.
- Command Devices are identified by a "-CM" in the ProductID field of the SCSI standard inquiry data.



- Example:  
/dev/rdisk/c0t0d0




HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## HORCM File – Pair Members (HORCM\_DEV)

- Pair Member information includes...



1. One or more groups (containing pair members).
2. A unique dev\_name for the pair.
3. Path information (host perspective):
  - XP CHIP Port,
  - SCSI TID,
  - SCSI LUN



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Pair Member Information

- The Local Host horcm.conf file contains information about only one member of the CA pair.
- The Remote Host horcm.conf file contains information about the "other" CA pair member.
- The "**dev\_name**" links the pair members together – and MUST be unique within a horcm file for the pair.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Pair Information Example

Host 1

#Grp_Name	Dev_Name	Port	TID	LUN
Group1	dev_1	CL1-N	1	0
Group1	dev_2	CL1-N	1	1

Host 2

#Grp_Name	Dev_Name	Port	TID	LUN
Group2	dev_1	CL1-N	3	0
Group2	dev_2	CL1-N	3	1

CA Pair Members

HP logo


HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## HORCM File – Remote Host (HORCM\_INST)

- Remote Host Information includes



1. Remote Host Name (or IP Address)
2. Remote Service Name entry (or Number)



HP Restricted

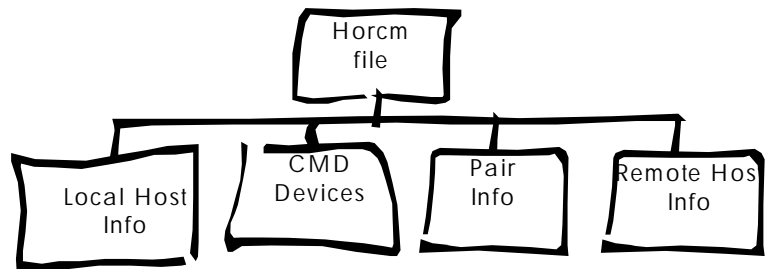
N S S O  
Stress-Free Storage...Guaranteed



## HORCM File – Demystified

- The HORCM file is 95% comments. You should remove the comments from your horcm file copy. The file will be much easier to use & it is advised by RAID Manager (see the last line of the horcm file)!

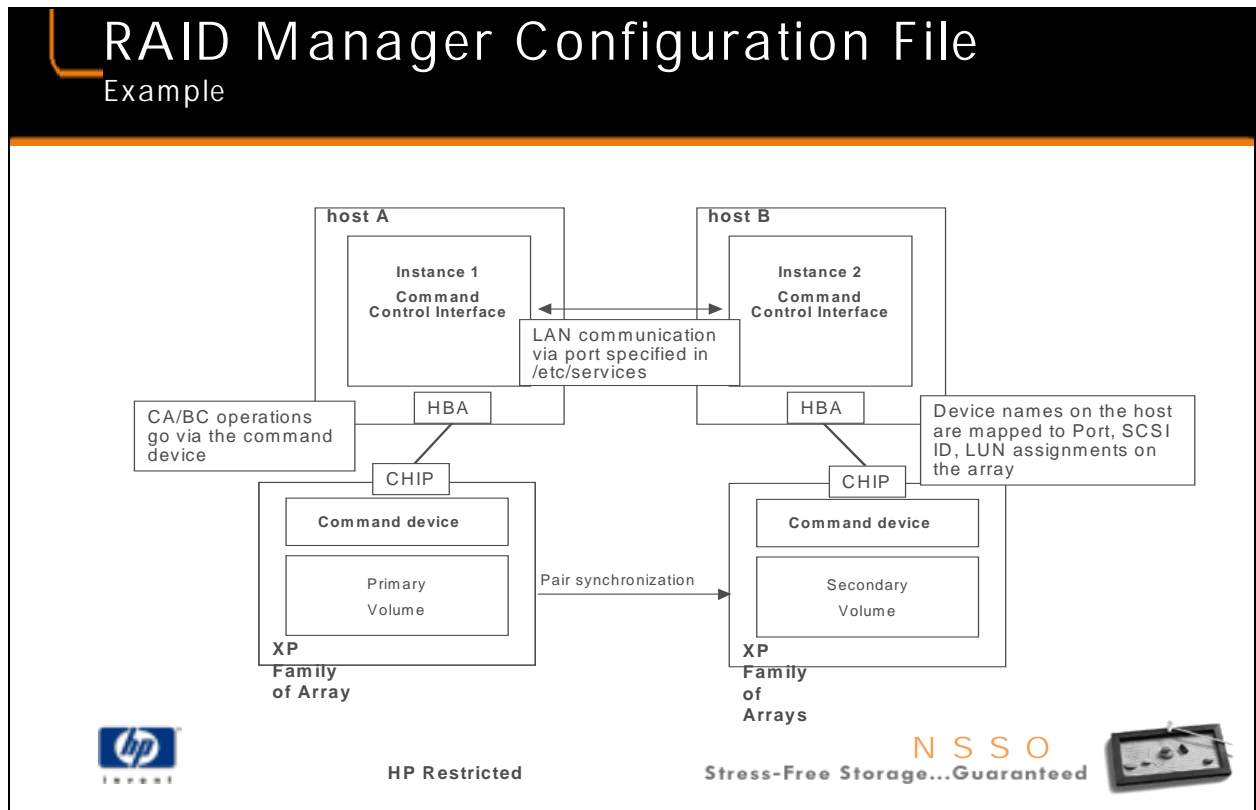
- Have fun! 😊



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed





The primary volume controls the pair status.

The pair status is changed via CA or BC commands.

I/O requests from the host to the secondary volume are allowed or denied on the basis of the status.

## RAID Manager Commands

RAID Manager Command	Description
horctakover (CA only)	Host executing horctakover takes ownership of the pair.
paircurchk (CA only)	Checks the consistency of the data on the secondary volume.
paircreate	Creates a pair.
pairsplit	Splits a pair.
pairsync	Resynchronizes a pair.
pairevwait	Event waiting command.
pairmon	Monitors a pair and reports changes in the pairs status.
pairvolchk	Checks the attributes of a volume connected to the local or remote hosts.
pairdisplay	Confirms the configuration of a specified pair.
raidscan	Lists the SCSI/fibre port, target ID, LUN number, and LDEV status.
raidar	Reports the I/O activity of a specified LDEV.
raidqry	Confirms the connection of the disk array and the open system host.
horcctl	Displays the internal trace control parameters.
horcmstart.sh	A shell script that start RAID Manager.
horcmshutdown	Stops the RAID Manager shell script.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



# RAID Manager Commands

## NT Specific Commands

RAID Manager Command	Description
Findcmddev	Searches for the command device.
Drivescan	Displays relationship between hard-disk number and physical drive
Portscan	Displays the physical device on a designated port.
Sync	Flashes the remaining unwritten data to the physical drive.
Mount	Mounts the specified device.
umount	Unmounts the specified device.
Setenv	Sets the environmental variables
Usetenv	Deletes the environmental variables
Env	Displays the environmental variables
Sleepetenv	Sets the sleep time for a specified environmental variable.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## RAID Manager Startup

- Startup script for RAID Manager - start the RAID Manager
- daemon and read the config file
  - UNIX: `horcmstart.sh [Instance]`
  - NT: `horcmstart` from DOS
- The instance number is passed as an argument to the startup script for RAID Manager:
- Command to stop the instance
  - `horcmshutdown [Instance] [Instance]`
- Shut down the given instances.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## RAID Manager Script Example

- You can use the RAID Manager example script handout as a reference, `rm_script.txt`
- It provides a sample of how some of the commands can be used in a "real world" situation to manipulate Business Copy volume pairs via the command line.
- The shell script contains code to:
  - create BC pairs
  - split BC pairs
  - delete BC pairs
  - resync of BC pairs
  - display BC pairs, and more . . .



Text Document



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



RAID Manager offers a bundle of commands for the customer to use to leverage the XP Family of Arrays's capabilities. Unfortunately, the manual doesn't offer much in the way of examples in their use, and so each individual must investigate and experiment on their own. You can use the RAID Manager example script handout as a reference. It provides a sample of how some of the commands can be used in a "real world" situation where Business Copy (possibly CA also) is used. In addition to the actual commands, it is an example of how waiting for an event or checking for status/state can be done. It has several switches:

create - creating pairs from a simplex state

split - splitting pairs after being copied

resync - resyncing previously split pairs

delete - deleting pairs back to a simplex state

fullrestore - copying the remote data on top of the local data (a reverse create)

diffrestore - reverse differential resync (needs microcode 52-42-51)

status - check the status (including the percent copied) of the pairs

Two useful switches not covered in the manual are worth mentioning here regarding the "pairedisplay" command. If you read the manual closely, you see these tips used in the examples, but not defined in the online help or manual text. The "-fx" switch displays the LDEVs in HEX rather than decimal (can be a real pain). Normally, CU 0, device 10 is displayed as "16" decimal, and CU 1, device 10 is displayed as "272"

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
RAID Manager XP

(256+16) decimal. Seeing that we configure the array with the RMC in HEX, this can be inconvenient to say the least. Using this switch forces a display of CU 0, device 10 as "010" ("0:10" without the colon), and CU 0, device 10 as "110" respectively. The "-fx" switch works also for the "raidscan" command. The "-fc" command displays the "percent copied" for a given pair. It is essentially a "percent" indication of the number of invalid (dirty) tracks for the pair. Since writes are asynchronous for the copies, its quite possible to see this display as less than 100% for a pair that is synced (PAIR state). If there is considerable write activity on the pair, the number may even go down for a time. If you start a pairing operation, then intensively write to the volumes (like a "dd") while they are copying, this behavior can be observed. If you wish to use both switches, put them together: "-fxc".

```
#!/usr/bin/ksh
```

```
# UNISRC_ID @(#)pairs_control: $Revision: 1.0 $ $Date: 99/10/07  
16:38:00 $
```

```
# MWR991007
```

```
# Sample script to illustrate how to manipulate Hewlett-Packard
```

```
# Business Copy volume pairs via command line
```

```
# Alter PATH if needed
```

```
PATH=$PATH:/HORCM/usr/bin/
```

```
# Set Environment for HORCM instance
```

```
export HORCMINST=0
```

```
# Set Env to select Business Copy operations vs CA operations
```

```
export HORCC_MRCF=1
```

```
# Set timeout variables in seconds
```

```
let MASTERTIMEOUT=10000
```

```
let SHORTTIMEOUT=600
```

```
# Set return or completion codes
```

```
RC_COMPLETE=0
```

```
RC_SIMPLEX=1
```

```
RC_PAired=3
```

```
RC_SPLIT=4
```



```
# Minimum rev of RAID Mgr for differential restore
DIFF_RESTORE_REV="010203"
# Minimum rev of microcode for differential restore
DIFF_RESTORE_MICROCODE="524251"
```

```
#####
# Functions Declarations #
#####
```

```
check_raidmgr ()
# Check if RAIDMgr exists and set revision
{
whence raidqry > /dev/null 2>&1
if [ "$?" -ne 0 ]
then
    echo "RAID Manager Software not detected on system!"
    exit 1
fi
```

```
# Get Rev of RAIDMgr
RMGR_REV=`raidqry -? 2>&1 | grep "^Ver&Rev" | awk '{print $2}' | tr -d
" "`
if [ $RMGR_REV -ge $DIFF_RESTORE_REV ]
then
    DIFF_RESTORE_RMGR=1
else
    DIFF_RESTORE_RMGR=0
fi
```

```
# Verify HORCM instance connection
raidqry -l > /dev/null
if [ "$?" -ne 0 ]
then
    echo "Cannot make connection to horcm instance!"
    exit 1
```

fi

# Get array microcode version

MICROCODE\_REV=`raidqry -l | tail +2 | awk '{print \$7}' | cut -c 1-8 | tr -d  
"\_"`

if [ \$MICROCODE\_REV -ge \$DIFF\_RESTORE\_MICROCODE ]  
then

DIFF\_RESTORE\_MICROCODE=1

else

DIFF\_RESTORE\_MICROCODE=0

fi

if [ \$DIFF\_RESTORE\_MICROCODE -eq 1 -a \$DIFF\_RESTORE\_RMGR -eq  
1 ]

then

CAN\_DIFF\_RESTORE=1

else

CAN\_DIFF\_RESTORE=0

fi

#echo "DIFF\_RESTORE\_MICROCODE=\${DIFF\_RESTORE\_MICROCODE}  
\012

#DIFF\_RESTORE\_RMGR=\${DIFF\_RESTORE\_RMGR} \012

#CAN\_DIFF\_RESTORE=\${CAN\_DIFF\_RESTORE}"

}

create\_pairs ()

{

# Create Business Copy pairs

date

echo "Creating Business Copy pairs for group \$GROUP"

# Check status of volumes to be paired. Must be in SIMPLEX state.

pairevtwait -g \$GROUP \$MEMBER -nowait -nomsg

```
RETURN="$?"
if [ $RETURN -ne $RC_SIMPLEX ]
then
    echo "Error: Pairs $GROUP are not in SIMPLEX state. Cannot
initiate copy."
    pairevtwait -g $GROUP $MEMBER -nowait
    echo "Return code was: ${RETURN}."
    exit 3
fi

# Create pairs with maximum track rate (15)
paircreate -g $GROUP $MEMBER -vl -c 15

# Monitor Status
echo "Waiting for pairs to transition to PAIRED state."
let COUNTER=0
RESULT=-1
while [ "$RESULT" -ne "$RC_COMPLETE" ]
do
    pairevtwait -g $GROUP $MEMBER -s pair -nomsg -t
$SHORTTIMEOUT
    RESULT=$?
    let COUNTER=$COUNTER+$SHORTTIMEOUT
    if [ "$COUNTER" -ge "$MASTERTIMEOUT" ]
    then
        echo "Error: Timeout occurred while waiting for pair
$GROUP to transition to PAIRED state".
        pairevtwait -g $GROUP $MEMBER -nowait
        echo "Return code was: ${RESULT}."
        exit 4
    fi
done
echo "Volumes in group $GROUP synced. Copy completed"
date
}
```

```
split_pairs ()
{
# Split Business Copy pairs
date
echo "Splitting paired volumes for group $GROUP. "

# Checking for required state of PAIR
pairevtwait -g $GROUP $MEMBER -nowait -nomsg
RETURN="$?"
if [ $RETURN -ne $RC_PAIED ]
then
    echo "Error: Pairs $GROUP are not synchronized."
    pairevtwait -g $GROUP $MEMBER -nowait
    exit 5
fi

# Split pairs
pairsplit -g $GROUP $MEMBER

echo "Waiting for split to complete."
let COUNTER=0
RESULT=-1
while [ "$RESULT" -ne "$RC_COMPLETE" ]
do
    pairevtwait -g $GROUP $MEMBER -s psus -nomsg -t
    $SHORTTIMEOUT
    RESULT=$?
    let COUNTER=$COUNTER+$SHORTTIMEOUT
    if [ "$COUNTER" -ge "$MASTERTIMEOUT" ]
    then
        echo "Error: Timeout occurred while waiting for pair
$GROUP to transition to PAIED state".
        pairevtwait -g $GROUP $MEMBER -nowait
        echo "Return code was: ${RESULT}."
        exit 6
    fi
done
}
```

```
        fi
done
echo
echo "Split has completed. We may now resume processing on source
volumes."
echo
}

delete_pairs ()
{
# Split Business Copy pairs
date
echo "Deleting pairing of volumes for group $GROUP. "

# Checking for required state of SPLIT (guarentees syncing)
pairevtwait -g $GROUP $MEMBER -nowait -nomsg
RETURN="$?"
if [ $RETURN -ne $RC_SPLIT ]
then
    echo "Error: Pairs $GROUP are not in PSUS state."
    pairevtwait -g $GROUP $MEMBER -nowait
    exit 7
fi

# Split pairs to Simplex state (delete)
pairsplit -g $GROUP $MEMBER -S

echo "Waiting for delete to complete."
pairevtwait -nomsg -s smpl -g $GROUP $MEMBER -t $SHORTTIMEOUT
RETURN="$?"
if [ "$RETURN" -ne $RC_COMPLETE ]
then
    echo "Error: Pairs did not transition into SMPL state in
$SHORTTIMEOUT seconds."
```

```
    pairevtwait -g $GROUP $MEMBER -nowait
    echo "Return code was: ${RETURN}."
    exit 8
fi

echo
echo "Delete has completed."
echo

}

resync_pairs ()
{
# Resync previously split Business Copy pairs
date
echo "Resyncing ${GROUP}."

# Check if in required Split state
pairevtwait -g $GROUP $MEMBER -nowait > /dev/null 2>&1
RETURN=$?
if [ "$RETURN" -ne "$RC_SPLIT" ]
then
    echo "Error: Volumes $GROUP are not in a PSUS state!"
    pairevtwait -g $GROUP $MEMBER -nowait
    echo "Return code was: ${RETURN}."
    exit 9
fi

# Resync Split volumes back to paired state
pairresync -g $GROUP $MEMBER -c 15
echo
echo "Waiting for differential synchronization to complete."
let COUNTER=0
RESULT=-1
```

```
while [ "$RESULT" -ne "$RC_COMPLETE" ]
do
    pairevtwait -g $GROUP $MEMBER -s pair -nomsg -t
$SHORTTIMEOUT
    RESULT=$?
    let COUNTER=$COUNTER+$SHORTTIMEOUT
    if [ "$COUNTER" -ge "$MASTERTIMEOUT" ]
    then
        echo "Error: Timeout occurred while waiting for pair
$GROUP to transition to PAIRED state".
        pairevtwait -g $GROUP $MEMBER -nowait
        echo "Return code was: ${RESULT}."
        exit 10
    fi
done
echo
echo "Synchronization completed."
}

restore_pairs_diff ()
{
# Restore volumes by making remote volume the source or P-VOL
if [ $CAN_DIFF_RESTORE -ne 1 ]
then
    echo "Insufficient microcode or RAIDMgr version for differential
restore!"
    exit 15
fi
date
echo "Creating Business Copy (restore) pairs for group $GROUP"

# Check status of volumes to be paired. Must be in SUSP/SPLIT state.
pairevtwait -g $GROUP $MEMBER -nowait -nomsg
RETURN="$?"
```

```
if [ $RETURN -ne $RC_SPLIT ]
then
    echo "Error: Pairs $GROUP are not in SPLIT state. Cannot initiate
differential restore."
    pairevtwait -g $GROUP $MEMBER -nowait
    echo "Return code was: ${RETURN}."
    exit 13
fi

# Restore Business Copy Volumes with differential mode
pairresync -g $GROUP $MEMBER -c 15 -restore

# Monitor Status
echo
let COUNTER=0
echo "Waiting for differential re-synchronization to complete."
RESULT=-1
while [ "$RESULT" -ne "$RC_COMPLETE" ]
do
    pairevtwait -g $GROUP $MEMBER -s pair -nomsg -t
$SHORTTIMEOUT
    RESULT=$?
    let COUNTER=$COUNTER+$SHORTTIMEOUT
    if [ "$COUNTER" -ge "$MASTERTIMEOUT" ]
    then
        echo "Error: Timeout occurred while waiting for pair
$GROUP to transition to PAIRED state".
        pairevtwait -g $GROUP $MEMBER -nowait
        exit 14
    fi
done
echo "Volumes in group $GROUP synced. Restore completed"
date
}
```



```
restore_pairs_full ()
{
# Restore volumes by making remote volume the source or P-VOL
date
echo "Creating Business Copy (restore) pairs for group $GROUP"

# Check status of volumes to be paired. Must be in SIMPLEX state.
pairevtwait -g $GROUP $MEMBER -nowait -nomsg
RETURN="$?"
if [ $RETURN -ne $RC_SIMPLEX ]
then
    echo "Error: Pairs $GROUP are not in SIMPLEX state. Cannot
initiate copy."
    pairevtwait -g $GROUP $MEMBER -nowait
    echo "Return code was: ${RETURN}."
    exit 11
fi

# Create Business Copy volumes in reverse (restore)
paircreate -g $GROUP $MEMBER -vr -c 15

# Monitor Status
echo
let COUNTER=0
echo "Waiting for pairs to transition to PAIRED state."
RESULT=-1
while [ "$RESULT" -ne "$RC_COMPLETE" ]
do
    pairevtwait -g $GROUP $MEMBER -s pair -nomsg -t
$SHORTTIMEOUT
    RESULT=$?
    let COUNTER=$COUNTER+$SHORTTIMEOUT
    if [ "$COUNTER" -ge "$MASTERTIMEOUT" ]
    then
```

```
        echo "Error: Timeout occurred while waiting for pair
$GROUP to transition to PAIRED state".
        pairevtwait -g $GROUP $MEMBER -nowait
        exit 12
    fi
done
echo "Volumes in group $GROUP synced. Copy completed"
date
}

display_pairs ()
{
# Display config and status of Business Copy pairs
pairdisplay -g $GROUP $MEMBER -fxc
}

#####
# Main Body of Code    #
#####

# Check if required number of parms
if [ "$#" -lt 2 ]
then
    echo "Usage: $0 GROUP ACTION [MEMBER]"
    echo "Must supply the HORCM Group name and action (member
optional) to be performed."
    exit 1
fi

check RAIDmgr

# Get BC group name and action to perform
GROUP=$1
ACTION=$2
DEVICE=$3
```

```
if [ -z "$DEVICE" ]
then
    MEMBER=""
else
    MEMBER="-d $DEVICE"
fi

# Select and execute action
case "$ACTION"
in
'create')    create_pairs
              break;;
'split')     split_pairs
              break;;
'delete')    delete_pairs
              break;;
'resync')    resync_pairs
              break;;
'diffrestore') restore_pairs_diff
              break;;
'fullrestore') restore_pairs_full
              break;;
'display')   display_pairs
              break;;
*)           echo "Error: Unknown Action $ACTION requested!"
              exit 2;;
esac

exit 0
```

## Module Wrap-up



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Module 11

# Configuring and Mapping for Performance on the HP SureStore E XP Family of Disk Arrays



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Module Agenda



- RAID Levels / Effects on Performance
- Cache / Shared Memory Effects on Performance
- ACPs & CHiPs / Effects on Performance
- Drive Types / Effects on Performance
- Business Copy / Configuring for Performance
- ACP & CHiP Mapping
- Cluster Configurations
- Oracle Layout for Performance
- Wrap-Up



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Cache Affects I/O

- XP Family of arrays are cache-centric system, that is, no data flows directly between the disks and the hosts
- All data must flow through the cache
- Reads
  - A record to be read which is not in cache already is put into cache by an ACP, then the data is read from cache by a CHIP and sent to a host
- Writes
  - No record is written directly to disk, but is written by a CHIP to cache, then read by an ACP from cache and written to disk



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



Reads, if random reads were truly random, adding more cache would do nothing for random read I/O, since the data would not likely be in cache. That segment would be de-staged before another access occurs. In many cases, the same random record is re-accessed within a short time. That record is promoted and will remain in cache. Writes, when random writes are held in cache and another read or update request occurs for the same block, the first write need not be de-staged from cache. Some random disk writes will be avoided by this means, and that segment is promoted in cache to avoid de-staging. HP-UX host cache buffers read data, so the XP256 cache services write data. Oracle stalls all transactions while flushing its logs to disk. The cache write area must be big enough to contain an entire "flush." Databases (especially Oracle) tend to issue a burst of writes to disk when synchronizing logs. It is very valuable to process this burst of writes at cache speed, since user applications stall until the burst completes. This means that there must be enough cache to absorb the entire burst, as well as enough ACO and disk resources in the back end to de-stage the burst before the next arrives. In write-intensive workloads, such as database logging volumes, what could you recommend to your customer to improve their performance on the XP family of arrays? Spread write-intensive data across several volumes to minimize queuing. Also consider increasing the duplex write line (DWL) of the disk array using the Performance Manager XP software product, or using a Cache LUN XP (cached) LDEV for this operation.

# XP512 Performance Overview

**Front-end Performance:**  
100% Cache Hits

- **165,000 IO/sec**
- **1,560 MB/sec**

**CHIP Pair Limits:**

- 41,000 IO/sec
- 464 MB/sec

**Single FC PORT:**

Standard Mode:

- 10,500 IO/sec
- 90 MB/sec

High Performance Mode:

- 20,000 IO/sec

**Back-end Performance:**  
Cache Avoidance

- **31,000 IO/sec**
- **840 MB/sec**

**ACP Pair Limits:**



- 7,300 IO/sec
- 318 MB/sec

**Crossbar Bandwidth:**  
6.4GB/sec

**1 Array Group:**




- 450 IO/sec
- 65 MB/sec

**Configuration:** RAID 0/1, 84 RAID Groups, 47GB Disks, 32 GB Cache, 4 ACP, 4 CHPP, 32 FC PORTS, FW rev. 05-02-00/00, 4 N-Class Hosts, 8 Tachlite ports each; Random Performance: reads, 2k block size; Sequential Performance: reads, 64k block size. Note: Single FC-port and CHPP-Limit tests are 100% cache hit. ACP-Limit and 1-Array Group tests are cache avoidance. All values in this overview are sustained.

 **HP Restricted** **NSSO** **Stress-Free Storage...Guaranteed** 



XP vs. Symmetrix Raw Performance Analysis				
	<b>XP256</b>	<b>Sym4.8</b> <small>(3630, 3830, 3930) (5630, 5830, 5930)</small>	<b>Sym5.0</b> <small>(8430, 8730)</small>	<b>XP512</b>
<b>Disks</b> (big box)	240	256	384	512
<b>Host FC Ports</b>	16	16	16	32
<b>Internal bus</b>	0.7GB/s	0.7GB/s	1.4GB/s	6.4GB/s
<b>MB/s sustained</b>	170 MB/s	170MB/s	340 MB/s	842MB/s
<b>IO/s read cache</b>	51,000	40,247	75,000est	165,011
<b>IO/s read cache avoidance</b>	10,900	9,000	15,000est	30,924
<b>IO/s read/write cache avoidance</b>	6,500	6,400	10,000est	17,403

HP Restricted

## XP512: What's New

- **Crossbar replaces internal buses: faster**
  - *None of the old boards fits in the new XP arrays*
- **FibreChannel back end: faster, more disks**
  - *None of the old disks fit in the new XP arrays*
- **RAID 0/1 groups are now 4 disks**
  - *Sequential read of a single LDEV is striped across all 4 disks, with a stripe size of 8\*48KB per disk*
- **Cache expander now affects performance**
  - *4 memory controllers instead of 2*
- **CHIP and ACP boards both 3x as fast**
  - *CHIP board ~20,000 IO/s compared to ~6,000*
  - *ACP board ~3,850 IO/s compared to ~1,350*



ACP board ~3,850 IO/s compared to ~1,350

HP Restricted

N S S O

Stress-Free Storage...Guaranteed



## Should I Configure the Cache Expander?

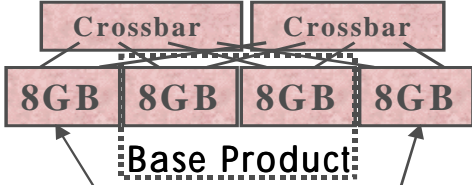


Diagram illustrating memory configuration options:

- Two Crossbar units are shown at the top.
- Below each crossbar are two 8GB memory modules.
- The first two 8GB modules are grouped under a dashed box labeled "Base Product".
- The next two 8GB modules are grouped under a dashed box labeled "Additional Cache Platform Board".

*Internal memory bandwidth is about 2GB/s with additional cache platform boards, about half that as the base product. About 80% is usable. Memory gets used 2 times to read a disk, and 5 times to write in RAID1. In general, bandwidth sensitive apps will require the option and "enough" memory DIMMs.*

HP Restricted

N S S O  
Stress-Free Storage...Guaranteed

Notice that both cache memory expanders must be installed in order to achieve full sequential performance. We haven't characterized how much cache must be loaded in each expander to achieve performance, yet. You can see that random I/O performance scales 2x by the simple fact that there are twice as many i960 ACP processors. Sequential performance is architecturally limited to about  $0.7 \times (3.2\text{GB/s} / 2)$  or 1.1-1.2GB/s, but will be below that at first release.

## Shared Memory Contains

Table of what is stored in  
the cache, by 16KB block

**Grows with  
size of cache**

Table of information on  
each LDEV

**Grows with  
# of LDEVs**

*Information on each CHIP and ACP in the array, and  
other service information: relatively small, and is  
included in the values shown in the table.*



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## 18GB-15K FC Disk Array

- Overview
  - customer needs & hp's solution
  - high performance
  - positioning
  - competition
  - pricing
  - compatibility
  - services & support
  - key take aways

### What does 18GB-15K offer?

- Among the first to market with high performance 15 K rpm FC disk array
- The industry's fastest array becomes even faster
- Exceptional random I/O performance
- Up to 25% faster than today's 18GB 10K rpm FC disk in I/O'sec
- Up to 25% more performance at 5% price increase
- Scalability to match customer's performance and capacity needs



HP Restricted

Stress-Free Storage...Guaranteed



## High Performance

➤ **need:**

➤ Random I/O's are a bottleneck

➤ **solution:**

➤ The 18GB-15K FC disk increase up to 25% the number of random I/O's per spindle, by increasing the data transfer rate and average seek time



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## positioning within the hp product line/family

New/Current	18GB-15K	18GB-10K
Target Customers	XP 48/512	XP 48/512
Key Needs Met	25% more random I/O's	100 I/O's per disk
Capacity	18GB	18GB
Speed	15K RPM	10K RPM
I/O Performance	125 I/O/disk	100 I/O/disk




HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



competition


Features, Benefits & Strengths	HP SureStore	EMC	HDS
	18GB	18GB	18GB
RPM	15K	10K	15K
Capacity	18GB	18GB	18GB
Support	XP48/512	Sym 8000	??
Availability	May 2001	Today	May 2001
Price	NTE \$7,400	\$9,000	??



HP Restricted

NSSO

Stress-Free Storage...Guaranteed



#### Product Structure

XP48 – Product #A5940A/U/S/SU

XP 512 – Product #A5970A/D/U/S/SU

Up to 25% more performance at a 5% increase over today's 18GB-10K rpm FC disk

Firmware/software

01.13xx FW required

(5/15/01)



## 18GB-15K FC Disk Array : Key Takeaways

- Up to 25% faster than today's 18GB-10K disk:
- Fastest RPM FC disk for Enterprise Storage Arrays:
- Exceptional random I/O performance:
- Scalability to match customer's performance and capacity needs:




HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## 18GB-15K RPM FC Disk Overview




*18GB-15K rpm disk replaces the existing 18GB-10K rpm High Performance disk*


➤ **What does it do?**

- More random I/O's per second ~ 120 I/O's/sec
- More I/O per GB

➤ **How does it work?**

- The 50% increase in rotational speed allows random access time to improve up to 25%
- Smaller capacity per spindle allows more I/O's per GB - 120 I/O's per 18GB of storage

 HP Restricted



**N S S O**  
Stress-Free Storage...Guaranteed 

### target customers

High Performance random access based applications like OLTP, e-Commerce, Data Warehousing. Particularly suited for applications with High I/O demand on storage

## XP512 2k Random Performance Per-Spindle Performance

	RAID 5	RAID 1
2k random reads	100 IO/sec	100 IO/sec
2k random writes	30 IO/sec	60 IO/sec
2k random 60/40 mix	50 IO/sec	85 IO/sec

 HP Restricted N S S O  
Stress-Free Storage...Guaranteed 

There are 2 processor bottlenecks

- CHIP or frontend processor
- ACP or backend processor

When you issue a write, it's guaranteed you are going to burn time in the CHIP, and at some point you are going to burn time in the ACP to de-stage it out of the cache. If you have to de-stage something to make room for the write of the cache, you're going to burn time in the ACP right up front when the write happens. The sustained writes, writes that actually go to the backend, are actually performance bottlenecks in the ACP processor in the back end doing the de-staging. That back end limits to 1/5, maybe even 1/10, the rate at which the front end can operate. Even though it appears that the cache gets rid of the bottleneck, when the I/O rates are high enough the cache fills up and you are in a steady state. You have to send one to the back for every one that comes in the front. Therefore getting the data to the disks is in fact processor bound.

## XP512 Peak Performance Measurements

Test	RATE	Response Time
2k Reads From CACHE	165,012 IO/sec	0.8 msec
2k Reads From DISKS (RANDOM)	31,166 IO/sec	24.6 msec
64k Reads From CACHE	1,570 MB/sec	2.7 msec
64k Reads From DISKS (SEQUENTIAL)	842 MB/sec	21.2 msec

All measurements were sustained for over 5 minutes. Although these numbers do show some of the "highest rates", they are obtained with a moderate to heavy load. Hence, reducing the load can show substantial Response Time improvements. e.g. 64k reads from Disk: 835 MB/sec @ 11ms. All tests were measured on XP512 in RAID 0/1.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## RAID Levels and Random I/O

- RAID-1 and RAID-5 each have advantages depending on the circumstances
- Single random read cache miss, considered in isolation
  - RAID-1 has 2 locations on 2 disks with 4 paths to ACP pair
  - RAID-5 has 1 location on 1 disk with 2 paths to ACP pair
  - With more paths and more disk choices to read the record, RAID-1 has less interference from other I/Os and less delay
- At moderate I/O rates, where the XP arrays cache destage limit is not exceeded, all random writes are cache hits and have the same response time, regardless of RAID level.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



RAID-1 always reads the primary copy today. Balancing load between primary and secondary copy is a future, targeted for late fall 2000.

## RAID Levels and Sequential I/O

- Some database read operations are sequential and very I/O intensive, at least in bursts
- In a best case scenario, it depends on front end connection speed (SCSI, Ultra-SCSI, Fibre Channel, server speed)
  - for front end I/O rate slower than the sequential stage rate, sequential reads are almost all cache hits
  - for front end I/O rate faster than the sequential stage rate, sequential reads are nearly all cache misses
  - RAID-1 can use both copies of data to double data rate
  - RAID-5 can triple the rate by using 3 disks at once
- Sequential writes are all cache hits unless the XP256 is overloaded with writes
  - normally no difference in RAID levels for sequential writes



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Impact of RAID Levels on I/O Rates

- RAID-5 allows a higher I/O rate, so the limits take longer to be exceeded
  - RAID-5 writes a parity track for every 3 data tracks, 4 I/Os per 3 tracks
  - RAID-1 writes to both disks, requiring 50% more disk activity, 6 I/Os per 3 tracks



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Drive Types

- For random cache-unfriendly I/O with short blocks, the choice of disk will overwhelm RAID level for performance.
- Cache-unfriendly random I/O - use 15GB drives
  - twice as fast, twice as expensive in cost per megabyte
- Cache friendly I/O - use 36GB drives
  - save cost, increase capacity
- New 47GB drives have performance and capacity



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



New line up of disk drives will be

18GB

47GB

73GB

They have identical performance at different capacity points.



## ACP Pair Considerations

- The number of ACP pairs controls the back end I/O power
- Each ACP pair provides 4 back-end SCSI strings or FC-AL loops.
- Calculate the number of front end CHIPs and their aggregate I/O power, multiply by the back end I/O ratio expected
- Consider cache hit rates and I/O types, random / sequential / read / write / heavy or moderate I/O
- Determine the number of ACP pairs for balance
- Alternatively, get all 4 ACP pairs and balance the load
- Granularity is at the ACP processor pair level. Evenly spread workload across ACP processor pairs and array groups to optimize backend performance



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Distribute Load

- Spread load across as many physical devices as possible
- Balance data transferred across ACP pairs
- Spread "hot spots" across multiple LDEVs, to make use of the I/O capability of more disks and more paths
- If you cannot isolate high-activity I/O subsets, spread I/O as much as possible across ACP pairs
- LDEVs derive from physical volumes. Prevent spreading I/O across LDEVs that map to the same physical volume where feasible
- Spreading I/O is important when using HP-UX LVM "striping" or equivalent (Microsoft stripe sets or Veritas volume groups)
- Plan assignments to CHIPS to evenly distribute I/O activity across the ports. Consider primary and alternate paths.
  - Keep load balance at the CHIP level



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Parity Group Recommendation

- Within each parity group, combine LDEVs with high I/O rates and LDEVs with low I/O rates to ensure an overall moderate I/O rate for that parity group.
  - excessive head movement occurs when 2 or more high I/O rate volumes are in the same RAID group
  - head movement drastically increases I/O response time (when data is not in cache)



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Consider Future BC

- With BC, I/O is posted to at least 2 volumes
  - assure additional writes do not impact performance
  - additional volume should be placed in a different RAID group or even a different ACP pair when feasible to limit the load on the ACP and minimize head movement
    - ♦ when the volumes are split during normal production and re-synced prior to starting backups, the performance impact will be small and no accommodation is needed



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



# ACP Mapping: Maximum Performance

Add L1 to distribute workload to free SCSI controllers

R1 is default

Guidelines for maximum backend/ACP performance, add ACP pairs

1. Expand by adding ACP pairs & disk groups
2. Add second disk group per ACP pair
3. Add third & fourth ACP pairs & disk groups
4. Stop at two DKU's for maximum performance

HP Restricted

N S S O

Stress-Free Storage...Guaranteed


When it comes time to expand a customer's XP Family Array, the question always becomes "Which way do we go?" The default configuration is the DKC and R1. Do we add R2 or L1? Does it matter? The answer will depend on the ultimate goal of the expansion - immediate cost efficiency or better performance. For best performance, we should add L1. Besides producing aesthetically pleasing symmetry to the system, sound reasoning dictates this solution. If you were adding a disk to a system running JBODs, would you add the disk to the SCSI controller with 4 devices already attached or would you attach it to another controller if available? You'd choose the free controller to distribute the workload. The same reasoning applies in this example. The subsystem will perform better if we have additional ACP pairs to manage the traffic. More controllers means more throughput and improved performance.

## ACP Mapping: Cost/Density Configuration

Add R2 first (no additional ACPs, SCSI cables, or power supplies)


Guidelines for maximum Cost/Density configuration:

1. Expand by fully configuring each ACP pair using dense disks
2. Configure DKU 1, then DKU 2 for ACP pair 1
3. Add and fully configure ACP pair 2
4. Repeat procedure for DKU 3 and 4 / ACP pairs 3 and 4



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



If cost is the most compelling factor in the expansion decision, another option would be to add R2. This strategy will cost less since it does not require additional ACP pairs, SCSI cable sets, or power supplies. The subsystem's performance will degrade slightly due to the higher load being placed on the processors, particularly the ACPs. This can be offset by the addition of extra cache memory. If four DKUs are required, the best performance solution is to have one DKC per DKU pair, cost permitting.

## CHIP Mapping

- For a “first pass” approximation, distribute LUNS evenly across CHIP ports (round robin LUN assignment) starting with most heavily used LUNs.
- Write queue depth is assigned per CHIP port (not per LUN). Watch this when configuring large numbers of LUNs.
- Use supplied cache & shared memory tables for basic configuration, but...
  - Expand Cache size when using HP CacheLUN XP
  - Expand Shared memory when using larger ACP counts
- NT does not support LVM, Use LUN expansion to allow larger mass storage sizes using XP Family Arrays



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed




Record LUN to port mappings carefully for host connection. Assume that CHIP/Cache “front end” is seldom a bottleneck, ACP “backend” performance is critical.

## CHIP/Cache Mapping

### 1. Determine the number of ports

- For every Fibre Channel CHIP pair, install 1 ACP pair
- For every SCSI CHIP pair, install 1 ACP pair



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



Point here is to focus on XP256 config NOT host:

- Use the following basic assumptions but...
- don't discuss them at length!

FWD SCSI Rules of thumb

- 17-18 MB/s Sequential IO performance per FWD-SCSI port
- ~3000 IOPS Random IO performance per FWD-SCSI port
- ~6000 TPM-C supportable per FWD-SCSI port


TPM-C Rule of thumb:

- Double IOPS for approximate TPM-C
- Will vary based on DB type



## CHIP/Cache Mapping

### 4. Map LUN's in a 'round robin' across ports based on LUN access patterns.



- Optimize LUN mapping as needed for Performance & Isolation
- Watch queue depth and maximum LUN per port rules
- Carefully document LUN/Port Mapping (for Host/FC Bridge installation and configuration)

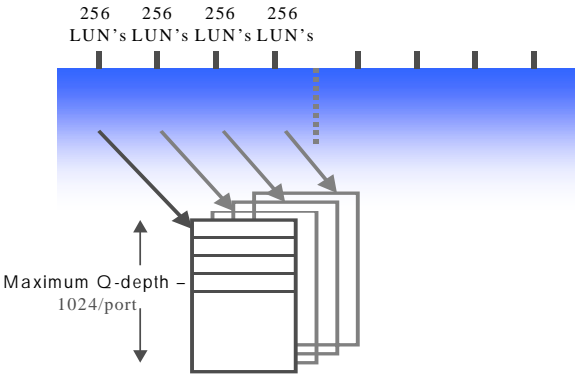


HP Restricted

N S S O  
Stress-Free Storage...Guaranteed






## CHIP/Cache Mapping (XP512)



- Write queue depth is per port, dynamically allocated to LUNs based on when I/Os are received
- XP512 write queue depth per port is 1024, HP-UX Write queue depth is 8! \*\*
- Maximum LUN count per port is 256 or 480 LUNs per CHIP.
- Maximum LUN count for the entire XP256 array is 2048 at launch

\*\* This yields ~8 queue entries per LUN when 120 LUN's are active on a port!

HP Restricted

HDS sales representatives have recently been claiming that all differentiation between the HDS 9900 and the XP512 will disappear with the advent of HP-UX 11.11 for HP-UX systems. In fact, HDS has been extremely aggressive in some sales situations with HP-UX customers, where they are trying to “level the playing field” by saying that the two products will become functionally equivalent.

This statement, by HDS sales teams, is emphatically false, and here are the facts to help you to refute claims made by HDS. The following differentiators will still exist even after the release of HP-UX 11.11.

### HP-UX Queue Depth

HP retains the same Q-depth performance advantage here, on HP-UX, with the XP512 that we have enjoyed with the XP256.

#### Maximum Queue Depth / Port

#### Queue depth / LUN

HP XP512

1024

8 per LUN

HDS 9900

256

8 per LUN

This limited HDS 9900 queue depth per port may create serious performance problems for the HDS 9900 once they have reached the end of the command queue, since this could slow down the HP-UX application by as much as 10X in some workloads.

### LUN Security

HDS's Zone Allocation Manager (their version of Secure Manager XP) does not allow for LUN security using the HDS 9900 when connected to HP servers.

When the HDS 9900 is connected to an HP 9000 server or to an NT server with an HP HBA, there is no LUN security for these servers through the HDS Zone Allocation Manager. As you can imagine, this gives HP a major LUN security benefit for the XP products in multivendor SAN environments that also include HP-UX or HP NT servers as well as HP only server environments.

### LUN Limits

Our lab has confirmed with the HP-UX kernel lab that no LUN address changes will be implemented in HP-UX 11.11. Nothing has changed! This means that the XP512 can have up to 256 LUNs per port while the HDS 9900 still has only 8 LUN's per port.

This still is a very serious LUN addressing advantage for the XP512 on HP-UX systems. HP enjoys a clear-cut advantage here with major differentiation on the number of LUN's. Databases are easily starved of LUNs on an 7700E platform connected to HP-UX, reducing performance. The result is increasingly slow application response time for business users

XP512

HDS9900


Max. LUN's per port

256

8

## Cluster Considerations

- Map existing LUN's to second set of Hardware Paths using the SVP or Remote Console.
- Spread primary and alternate LUN's across all CHIP ports.



- Import LVM configuration on new host



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



*HP firmware tuning gives HP a huge advantage over HDS in LUNs per fibre channel port. An XP512 has 256 LUNs per fibre channel port in an HP-UX environment. A comparable HDS Lightning 9900 still has only 8.*

In a non HP-UX environment, both the HP products and the comparable HDS products have an identical number of LUNs per port. For example, in a Sun environment an XP256 and a 7700E would both have 120 LUNs per port. An XP512 and a 9900 would both have 256.

### Maximum LUNs per port

HP-UX

Non HP-UX

HP XP256

120

120

HDS 7700E

8

120

HP XP512

256

256

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Configuring for Performance

HDS 9900  
8  
256

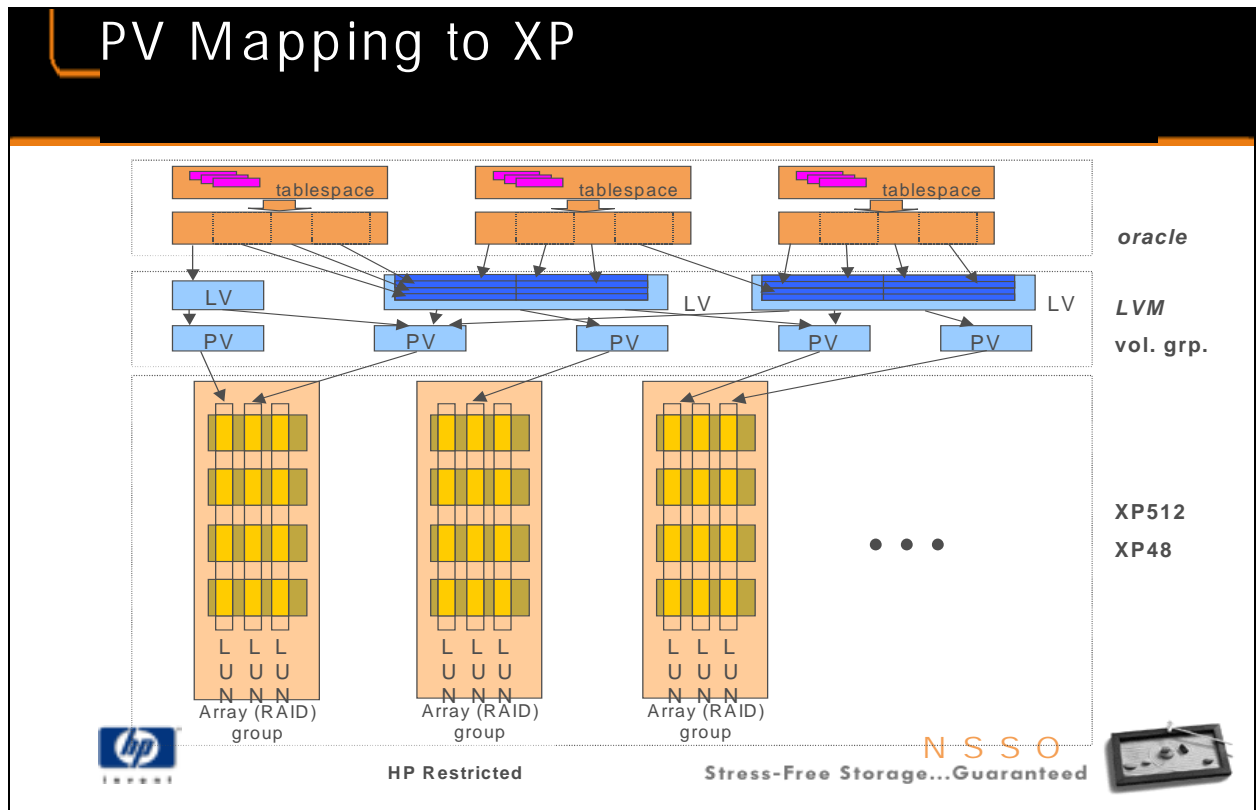
An XP256 has 1024 LDEVs and a maximum of 1920 LUN instances in all environments (16 fibre channel ports x 120 LUNs per port).

In an HP-UX environment a comparable HDS 7700E would have a maximum of 128 LUN instances (16 FC ports x 8 LUNs per port).

For an XP512 vs. a 9900, the HP-UX number would be 8192 LUNs (32 ports x 256 LUNs) vs. 256 (32 ports x 8 LUNs).

**Maximum LUNs per Array**

HP-UX	
Non HP-UX	
HP XP256	
1920	
1920	
HDS 7700E	
128	
1920	
HP XP512	
8192	
8192	
HDS 9900	
256	
8192	




Now we want to introduce the XP into the overall equation. Now the first thing to notice here is the tan colored rectangular boxes at the bottom of the slide, they represent the array groups located within the XP array. As you have learned over the past several days, the array group is the fundamental building block of an XP disk array. Whenever you add space to a XP array you add an array group and as we have explained in previous modules, an array group consists of four physical discs. All the data that goes in the array group gets striped across discs configured in an array group. The key thing to notice is that we carve out LUNS out of those array groups which are shown in the rectangular boxes that span all four physical disks in an array group. It is the LUNS that map to the physical volumes and logical volume manager. So this idea of these independent spindles as we described in the previous slide, if you take a look at the individual PVs -- the physical volumes -- they could very well map to LUNS that are in the same array group and therefore are no longer independent of one another but instead residing on the same physical volumes. As we stated in the first slide talking about Oracle, you are trying to put data and indexes on separate devices. With an XP disk array you wind up putting the data on one PV that maps to a LUN that happens to be in the same array group as a different LUN where you put the index. And then you basically just put everything on the same set of spindles (i.e. physical disks). So you don't have that independence that you used to have utilizing the smaller spindles. So you need to be a little more concerned if

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Configuring for Performance

that's what your goal is. Or put it this way, it is a little harder to implement that strategy that we used before.

## Array Group Configuration



- Each AG must be same RAID
  - can now mix RAIDs on ACP
- Use LUSE to coalesce LUNs
  - fewer LUNs: shorter IOSCANs
  - limit volgroups to ~16 devices
- Leave space for growth



### Array (RAID) group

HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



This slide details a little more about the array group configuration itself. In review, four disc mechanisms creates an array group and your data gets striped across all four disks and you have to determine whether that array group is going to be RAID 0/1 or RAID 5. Notice that if you choose RAID 0/1 and that we have 73 GB discs, that your array group is roughly about 146 GBs in size. If it's RAID 5 you are talking about 214 GBs in size. That is your minimum unit of independence with the XP array. This was done to illustrate that this is quite a bit different from the old days where we had discs that are one or two gigabytes that gave you a lot more flexibility. So again if you are looking for ways to deploy your data to make sure that all the different components are located on different independent spindles, the term spindle now refers to a chunk of like 140 or 214 Gbs, essentially much larger in size.



## RAID 0/1 or RAID 5?

- Economics, competition, and customer preference will probably dictate mix
- Place random-write-intensive objects on RAID 0/1
  - Data/indexes can be random
  - Redo/rollback heavily written but sequential
- Lots of cache helps overcome RAID 5
  - Tradeoff: cash versus cache!
  - Use high-end formula for determining cache




HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



This slide is placed here to address the issue of discussing RAID levels. The RAID level is going to probably be determined more by your customer, the economics in a competitive situation, et cetera and so forth. Or maybe just a customer wants to have the fastest available so they are going to specify RAID 0/1. But obviously you'll have cases where you'll have a mix. But there are also certain cases where for economic reasons it's going to be all RAID 5 or for performance reasons it's going to be RAID 0/1. If you are being asked to make recommendations for the customer, I guess it's obvious that you want to put your write intensive objects on RAID 0/1. You would start by picking the objects that you believe will be most write intensive in a random fashion, such as the data itself. The indexes and temp spaces will be a higher potential for being random as well. The redo logs and rollback segments are definitely heavily written as well, but are more sequential in nature. So there is less of an issue with sequential writes than there is with random writes with the RAID 5 penalty. If you only have a very limited amount of RAID 0/1 to deal with then hopefully that is a rule of thumb that will guide you in that direction.



This is a graph basically to show that with the XP256 we talked about how good the RAID 5 performance was. To back this up, we actually had some Hitachi data and some tests where we had performed Oracle benchmarks and were able to demonstrate that the RAID 5 performance was in some cases actually better than RAID 1. And what you see from this graph is that yes, that was very true with the XP256 but with the 512 it's a little different. There is a much larger disparity between the RAID 1/I/O performance and RAID 5 performance for the XP512. So what we have in this chart is a series of tests that basically each test has nine different data points. The tests were divided up into three different areas where there was a write intensive data point, which consists of 20% reads and 80% writes, an even read/write data point and a read intensive data point. Now within each one of these data points, we've buried the random/sequential ratio from 20% to 80%. So by twisting all those parameters we end up with nine different things. And you're going to kind of see an evolution of how things happen between the 256 and the XP512. So the very first test, which is the very first line on the left in each one of those data points, is basically the initial XP256, which was SCSI RAID 5. And then after that you are going to see what SCSI RAID 1 performance numbers are. What you'll see is that in many of the data points there is a very small, especially when you get up in the read intensive step -- difference between what you get with RAID 5 and RAID 1. Now there is one additional data point in here. And that is the blue line which is for the native fibre connect on the XP256. So we put in a RAID 1 fibre test there to show some of

the increases that we got in fibre vs. SCSI. What is interesting here is that you'll notice that fibre helped us out much more on the tests that were sequential where we had lots of bigger transfers versus SCSI which actually did better when we had lots and lots of smaller transfers and lots of I/Os. The two lines on the right of each one of the tests show the XP512. And you'll see that we have a significant performance increase between the XP512 and the XP256. But the other thing that you'll also notice is there is a much larger delta now between RAID 1/0 tests and RAID 5 tests than there were in the XP256. So the whole reason for this graph is just - it's not to let you know what maximum performance or I/O rates or anything like that is but just give you an idea of how that evolved between the 256 and the 512 and the fact that, you know, the RAID 1/0 on the XP512 is much, much better than the RAID 1 is on the XP256. And therefore there is a bigger delta in performance now between RAID 5 performance and RAID 1/0 performance.

## Cache LUN for Oracle?

- Cache LUN: carve out part of XP cache and use as "solid-state disk"
- Basically, ensure 100% cache hits for all I/O to the cache LUN
- Tradeoff: same cache could be helping performance of ALL I/Os to the XP!
- Not recommended
- Consider solid-state FC10 instead



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



We are next going to address the issue of creating volume groups and actually carving up what is on the XP to make it useable. You want to make sure that I carve the XP up into more than one group. We do want to implement some striping, but the intent here is to group things together on the XP so that we have maybe three, four, or five different volume groups that we put Oracle objects into. And then if we happen to find that we have a performance issue with one of the objects and one of those volume groups we can see that. For example, maybe we are doing 50% more I/O in this section than the other section. We want to have the ability to move things around. If you stripe everything over everything you have lost the ability to move things around and you lose that flexibility. So basically what we are going to do is we're going to take the XP, we are going to take the number of array groups that we have in the XP and we are going to divide them up into maybe four to eight sets. A good rule of thumb is if you have an XP array that is fully configured with four ACP pairs, take an array group from each ACP pair and put it into a volume group. And then take the next four from each ACP pair and put it into another volume group. In most cases four arrays in a volume group have been sufficient. In some cases we have gone six to eight arrays per volume group. But it is difficult to say for sure whether that is going to give you a significant advantage. In almost all the cases it has been sufficient to have four arrays in a volume group. And then what we will end up doing is actually striping across those arrays. We'll use a stripe size of typically 128 to 256. And what we want to do is we want to use a stripe size that is small enough

so that we end up spreading the objects -- when talking about the objects it's the actual tables that go into table spaces -- across multiple arrays so that we get the advantage of using multiple spindles. But we want it also to be large enough so that we don't have a lot of objects that over span the stripe size. So when we go out there to write something we don't - we want to be able to write to one array group and not have to wait and write to two different array groups. This is the LVM level striping size. So what we end up having is our array split into multiple volume groups. We never take an array group and span it across volume groups because then it is harder to control. And what we end up with then is maybe a volume group has got all RAID 5 in it or two or three volume groups that have RAID 5, maybe a couple of volume groups that have four RAID 1 arrays in there. And we can use those for our write intensive objects.

## Create Volume Groups

- Separate Array Groups into 4-8 sets
  - span ACP pairs (for XP512)
- Create one Volume Group per set of AG's
  - Use LUSE to limit # PVs (LUNs) per VG
- Specify alternating paths to PVs (LUNs) to spread I/O across paths

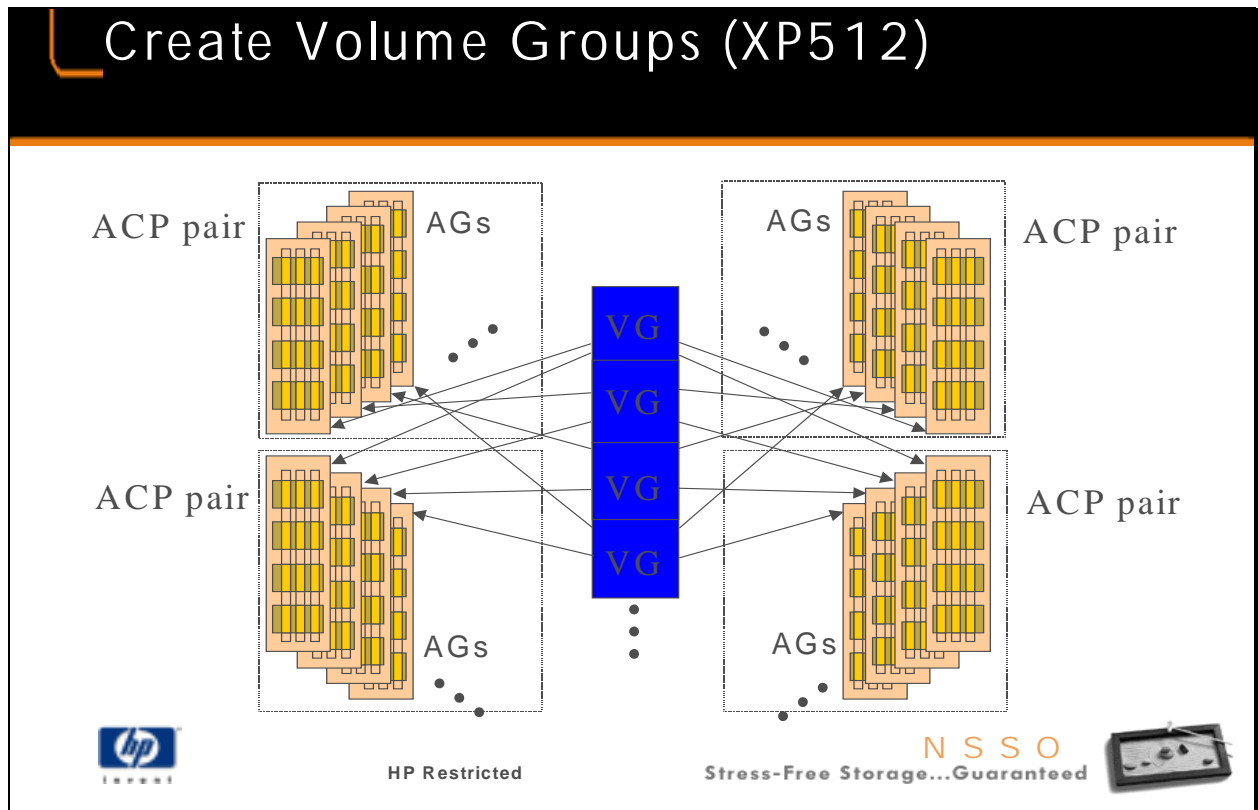


HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



An important point is that when you set up your physical volumes and your volume groups, make sure to alternate the paths so that as you do your I/O you are going to spread across multiple paths as well. I think that's probably pretty standard and understood by now but it was worth pointing out.



This page illustrates basically what was discussed on the previous page. It shows a bunch of volume groups, each one having as a constituent of LUNs in separate array groups spread across multiple ACP pairs. Of course this depends on the configuration of the XP512 as well because you don't always necessarily have all of the ACP pairs. But you basically do the best you can do. Take each volume group and make sure you have components within all the different array groups, within the array groups from different ACP domains. And this example shows that there is one array group for each volume group. But you could have multiple array groups in each ACP domain in the same volume group. For performance reasons, the idea is that when you wind up creating the logical volumes and striping them across the PVs and your volume group that each of the various PVs, the result will be that you will stripe your I/O across all the ACP pairs if that is possible.

## Ready to Create Logical Volumes

- Create LVs using LVM striping
  - Use 128k or 256k stripes
- Divide objects evenly among VGs
  - Keep each object wholly contained in a VG
- For raw I/O, create standard-sized LVs
- Use symbolic links in Oracle to point to file locations



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed





## Monitor Performance & Balance I/O

- Look for high I/O wait states on shadow processes (gpm)
- monitor VG I/O rates (gpm)
- Performance Adviser/XP
  - backend disk & bus utilization
- Move Oracle files (if nec) to balance I/O
  - Work with Oracle DBA to identify & correct problems

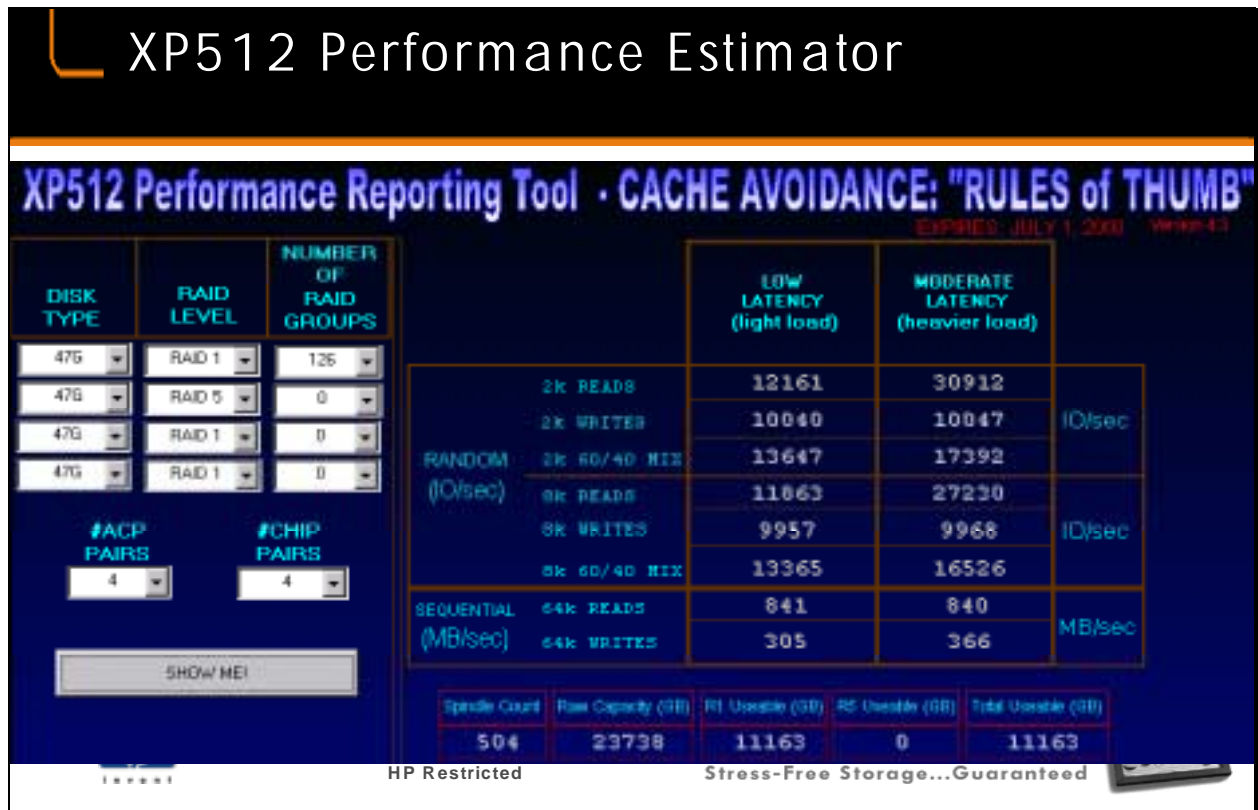


HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



And as far as actually using the tools to identify this we use GPM a lot. Look at weight states on channel processes. There are actually many tools that you can get from Oracle. From a basic level is take a look at the weight states to make sure that you are not doing lots and lots of wait on I/O, take a look at the I/O by volume group with the volume groups the way you've got them spread out on your XP. And make sure that you don't have one or two of them that are doing 80% of the I/O. And if that is the case you always have the opportunity to move the objects between the volume groups.



HP has developed an Excel based performance estimator tool that lets you pick the different types of discs and RAID levels and gives you some idea of what the high level performance capabilities of the system will be. Please check with your local Mass Storage TC to obtain or use this tool. This tool is called pet47p3\_XP512.xls For those teaching the class, please demonstrate the use of this tool to the class.

## Document

- Good documentation is crucial
- Provide the customer with a storage map defining
  - how each LDEV maps to hosts and addresses
  - LDEV number, RAID Group, Emulation, Size, RG Members, Port Assignments, host modes, device files/volume identifiers, Volume Groups (if used), Business Copy / Continuous Access associations
  - where each volume (canister) is physically located
  - physical and logical diagrams, cabling charts (Visio diagram) (diagram samples at [//visio-cafe.hp.com/](http://visio-cafe.hp.com/))



HP Restricted


N S S O  
Stress-Free Storage...Guaranteed



Before you try to resolve performance issues, this data should be available from the technical account manager. Tracking down a hot CHIP or ACP processor, or a hot SCSI bus or spindle, is much easier when you know what data is on that resource.

## Sample Storage Map


Sample XP512 Storage Map								Date	14-Jul-00	Total Space (MB)		
								Cache GB	32	Space assigned to		
								Shm MB	512	Space not assigned		
Enter values only in VG fields. All others referenced from other sheets and protected.										Space assigned to		
Ser# CU:LDEV	RAID	Disk	RAID	Emul	LVM MB	Raid Group	Host1	Host1	Host1	Host2	Host2	Host2
	GRP	Size	Lvl			Members		Prim Pvol	Alt Pvol		Prim Pvol	Alt Pvol
35358 0:00	1-1	47	5	OPEN-9	7040	R100,R110,R120,R130	HOSTA	c0t0d0	c6t8d0	HOSTB	c2t0d0	c8t8d0
35358 0:01	1-1	47	5	OPEN-9	7040	R100,R110,R120,R130	HOSTA	c0t0d1	c6t8d1	HOSTB	c2t0d1	c8t8d1
35358 0:02	1-1	47	5	OPEN-9	7040	R100,R110,R120,R130	HOSTA	c0t0d2	c6t8d2	HOSTB	c2t0d2	c8t8d2
35358 0:03	1-1	47	5	OPEN-9	7040	R100,R110,R120,R130	HOSTA	c0t0d3	c6t8d3	HOSTB	c2t0d3	c8t8d3
35358 0:04	1-1	47	5	OPEN-9	7040	R100,R110,R120,R130	HOSTA	c1t0d0	c7t8d0	HOSTB	c5t0d0	c9t8d0
35358 0:05	1-1	47	5	OPEN-9	7040	R100,R110,R120,R130	HOSTA	c1t0d1	c7t8d1	HOSTB	c5t0d1	c9t8d1
35358 0:06	1-1	47	5	OPEN-9	7040	R100,R110,R120,R130	HOSTA	c1t0d2	c7t8d2	HOSTB	c5t0d2	c9t8d2
35358 0:07	1-1	47	5	OPEN-9	7040	R100,R110,R120,R130	HOSTA	c1t0d3	c7t8d3	HOSTB	c5t0d3	c9t8d3



HP Restricted

N S S O

Stress-Free Storage...Guaranteed



Please contact your local ASE about obtaining the “sample” storage maps for implementation needs.

## Summary of Performance Recommendations

- Larger cache improves performance
- Number of spindles = performance
- Spread load across as many physical devices as possible
- Evenly spread workload across ACP processor pairs and array groups
- Plan assignments to CHiPs to evenly distribute I/O activity across the ports. Consider primary and alternate paths.
- Correct an imbalance of I/Os going to either even or odd LUNs behind a port
- For random cache-unfriendly I/O with short blocks, the choice of disk will overwhelm RAID level
- Within each parity group, combine LDEVs with high I/O rates and LDEVs with low I/O rates
- Additional BC volumes should be placed in different RAID groups or different ACP processor pairs



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



For best performance when hitting the “backend”

- use more ACP pairs
  - 2 ACP pairs per SCSI CHiP pair
  - 4 ACP pairs per FC CHiP pair
- evenly populate disk groups across ACP pairs
  - Minimum 2 disk groups per ACP for max. bandwidth
  - Do not use DKU 2 & 4 (the outer DKU's)
- use fast 12,000 RPM drives

For best cost or storage density

- build unit using DKU's R1 & R2 first, share ACPs
- use high density drives

## Module Wrap-up



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



# Module 12

## XP Volume Concepts



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Module Agenda



➤ RAID concepts

➤ XP Free Space

➤ The XP Array "Volume"



➤ XP Array Custom Volume Size (CVS) concepts.

➤ Wrap-Up



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed





## RAID Concept – a Solution

- **RAID** (Redundant Arrays of Independent Disk drives)
  - Provides availability of user data despite a disk failure.
  - Provides for the on-line reconstruction of the contents of a failed disk to a replacement disk.
  - Offers robustness (via redundancy) at the cost of lost storage capacity.
  - Defines standard redundancy algorithms, such as RAID 1 and RAID 5.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## RAID Concept - Parity

- Storage Capacity = (Disk Capacity) - (Parity Data)
- Parity Data is redundant information that can be used for data reconstruction.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## RAID Concept – the Cost

- For RAID 1, the Parity Data uses 50% of the total disk space; data is either mirrored, or there are two copies of the data.
- For RAID 5, the Parity Data uses  $1/(\# \text{ of disks})$  of total disk space and parity data is interleaved among all disks
- For more information:  
<http://www.raid-advisory.com/rabguide.html>  
<http://www.raid-advisory.com/hp.html>



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## RAID Concept: XP Array RAID Group

- Think of RAID protected space as available space – NOT as individual disks.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## RAID Concept: XP Array RAID Group

➤ Think of RAID protected XP Array storage space in a logical way:

- RAID data is protected usable storage space
- users DO NOT know exactly where the data is
- vendors implement RAID types in different ways.




HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## General Concept: From Disks to RAID Group

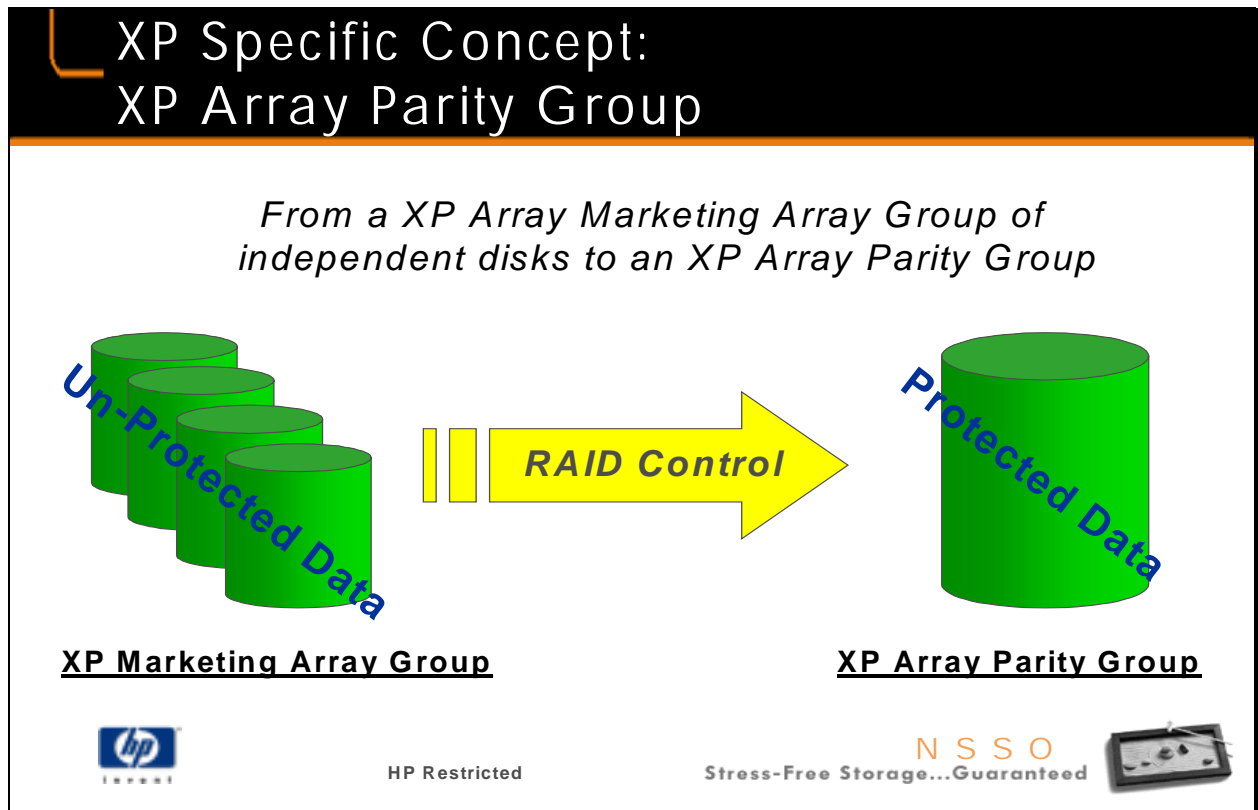
*From a group of independent disks to a RAID group*



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed





## XP Array Free Space

- Usable space in the XP Array is found within a Parity Group and is called Free Space.
- Free Space is available to the customer, but NOT in its raw form.
- Free Space is used AFTER the Array has been configured.
  - An array may have 200 Gbytes of raw Free Space per a Parity Group that must be configured for customer usage.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed





## Creating a Parity Group

- Q. How is Free Space created?
  - A. By creating a Parity Group.
- Process:
  1. A RAID level must be defined for an XP Array Domain – this should already have been performed.
  2. A Marketing Array Group of disks must be installed into the array Domain.

Now, we have a Parity Group with Free Space available!



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Using Free Space

- To use Free Space for data storage, four tasks must be performed.
  1. Define an Emulation Type for the Parity Group, thereby creating Volumes - virtual storage devices/disks.
  2. Assign Volume information to the newly created Volumes to uniquely identify the devices.
  3. Map the Volumes to a CHIP\* port.
  4. Connect a server to the CHIP port.

\*CHIP means Client Host Interface Processor (analogous to an HBA on a server)



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Emulation Types

- Normal disk drives come in different sizes; XP Arrays offer customers a set of predefined "disk" sizes.
- These predefined "disk" sizes are known as different Emulation Types.
- An Emulation Type must be defined for a Parity Group, thereby partitioning the Free Space into predefined virtual disk (Volume) sizes.
- Only one Emulation Type may be assigned to a Parity Group.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Emulation Types & Normal Volumes

- Normal Volumes are of a fixed size.  
Example: OPEN-3 Volumes are 2,347 MBytes in size.
- When a Parity Group's Free Space is partitioned during the Emulation Type Configuration task, any remaining space left over is still Free Space.
- Remember, Free Space is not usable unless configured.



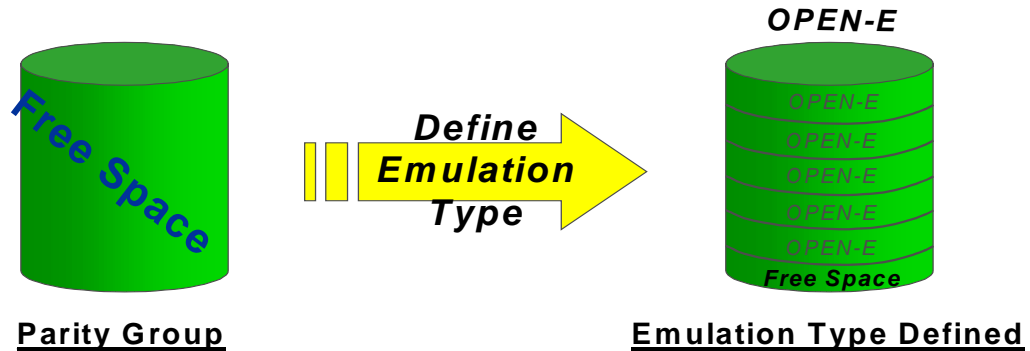
HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Using Available Space

- First, assign an Emulation Type to the Parity Group, such as OPEN – E.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Normal Volumes & Free Space

### Example:

Let's say I have the following –

- An Emulation Type with a defined size of 3.0 Gbytes
- A Parity Group with 32 Gbytes of Free Space before setting the Emulation Type

After setting the Emulation Type –

- How many Normal Volumes will be created?
- How much Free Space will be left? Normal Volumes & Free Space



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Normal Volumes & Free Space

### Solution:

How many Normal Volumes will be created?

Initial Free Space / Emulation Size per a Volume = #  
Volumes

(32 Gbytes of Free Space ) / (3.0 Gbytes per a Volume) =  
???

**32/3.0 = 10 Volumes Created**



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Normal Volumes & Free Space (cont.)

### Solution:

How much Free Space will be left?

Initial Free Space – Used Free Space = Unused Free Space

(32 Gbytes of Free Space) –

(10 Volumes \* 3.0 Gbytes per a Volume) = Unused Free Space

32 – 30 = 2.0 Gbytes of Free Space left



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed





## XP Array Volumes

- Regarding the XP array, a **Volume** refers to the CU:LDEV number.
  - The CU number is defined as a Control Unit and is used by the XP's internal data structures to manage its virtual disks. One CU number can contain 256 LDEV's
  - The LDEV (Logical Device) number is an arbitrary number from 0 – 255.
  - Together, the CU and the LDEV number uniquely identify a virtual XP Array disk – we call this a Volume.
- Example: for CU=01 & LDEV=2a, the Volume = 01:2a



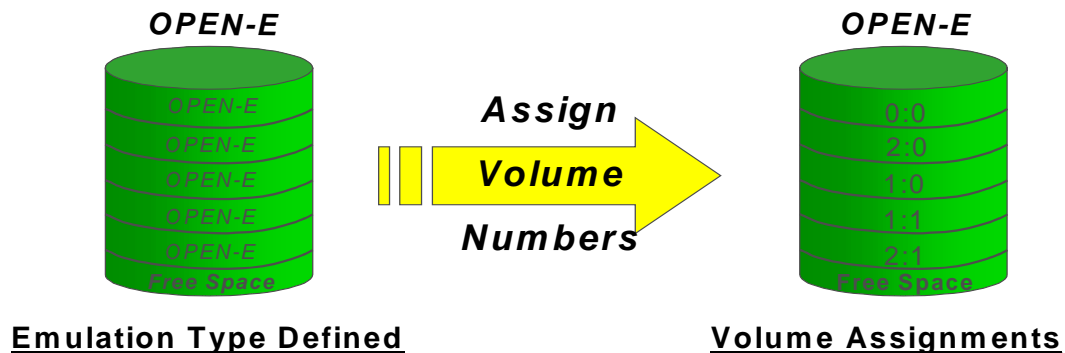
HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Assigning Volume Information

- After specifying an Emulation Type for a Parity Group, next assign Volume Numbers to the OPEN-E's, thereby uniquely identifying the new virtual disks.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## LUN Mapping?

- “LUN Mapping” should really be called “Volume Mapping”, since that’s what it is ☺
- LUN Mapping is the process of mapping a **Volume** to a CHIP (Client Host Interface Processor) port for the purpose of allowing an external host to use the volumes for storage.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Mapping a Volume to a Port

- An XP **Volume** (CU:LDEV) is visible to a server as a logical storage device (a LUN or Disk), only after being mapped to an array port that is connected to the host.
- Alternate Paths:
  - A **Volume** mapped to more than one array port is said to have an alternate path.
  - A **Volume** with multiple paths to a host, will appear to a host as multiple and separate storage devices.
  - Two devices on a host with the same XP Array **Volume** (CU:LDEV) number are really alternate paths to the same Volume.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Volume Types

- Normal Volumes are sized according to the type of device being emulated (Emulation Type).
- Open Systems have normal volume "OPEN-\*" Emulation Types, such as OPEN-3 or OPEN-E.
  - Normal Volumes are of a fixed size.  
Example, OPEN-3 Volumes are 2,347 MBytes in size.
  - Only one Emulation Type may be assigned to a Parity Group.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Module Wrap-up



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Module 13

### XP Family of Arrays SPI Training

#### HP OpenView Operations Center SMART Plug-In for SureStore E Disk Arrays B9367A



N S S O  
Stress-Free Storage...Guaranteed



One of the fundamental components of the 5-nines infrastructure is the HP SureStore E Disk Array XP Family of Arrays, which serves as highly-available storage facility in the architecture. In order to provide full end-to-end management of the infrastructure in HP OpenView IT/Operations (ITO), the XP Family of Arrays must be managed from the central console. ITO intercepts XP Family of Arrays fault events and provides them to the 5nines error analyzer component as well as to operator consoles. An ITO Smart Plug-In (SPI) will be developed that configures ITO in order to fulfill these tasks.

HP OpenView Smart Plug-In For XP Family of Arrays – B9367A

The OpenView Smart Plug-In (SPI) for XP Family of Arrays connects HP OpenView Operations Center Operations (Operations Center) to the HP SureStore E Disk Array XP Family of Arrays enabling proactive monitoring and management from a single point – the Operations Center management server.

Features of the OpenView SPI for XP Family of Arrays include:

- intercept, filter and prioritize events from the XP Family of Arrays

- automatically detect XP Family of Arrays configuration and upload to Operations Center Navigator

- service view of XP Family of Arrays low-level components and automatic mapping of Operations Center messages to the appropriate service objects

- root-cause analysis of component failures

- propagation scenarios to visualize the wider impact of component failures

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
HP OV Operation Center SPI

Supported OS: HP-UX 11.0, 10.20; Sun Solaris version 7.8




Operations Center Server: Operations Center Operations 6.00; Operations Center  
Navigator 6.00 XP25 Firmware: 52-47-06-00+

Reference Price: \$36,000




# Module Agenda

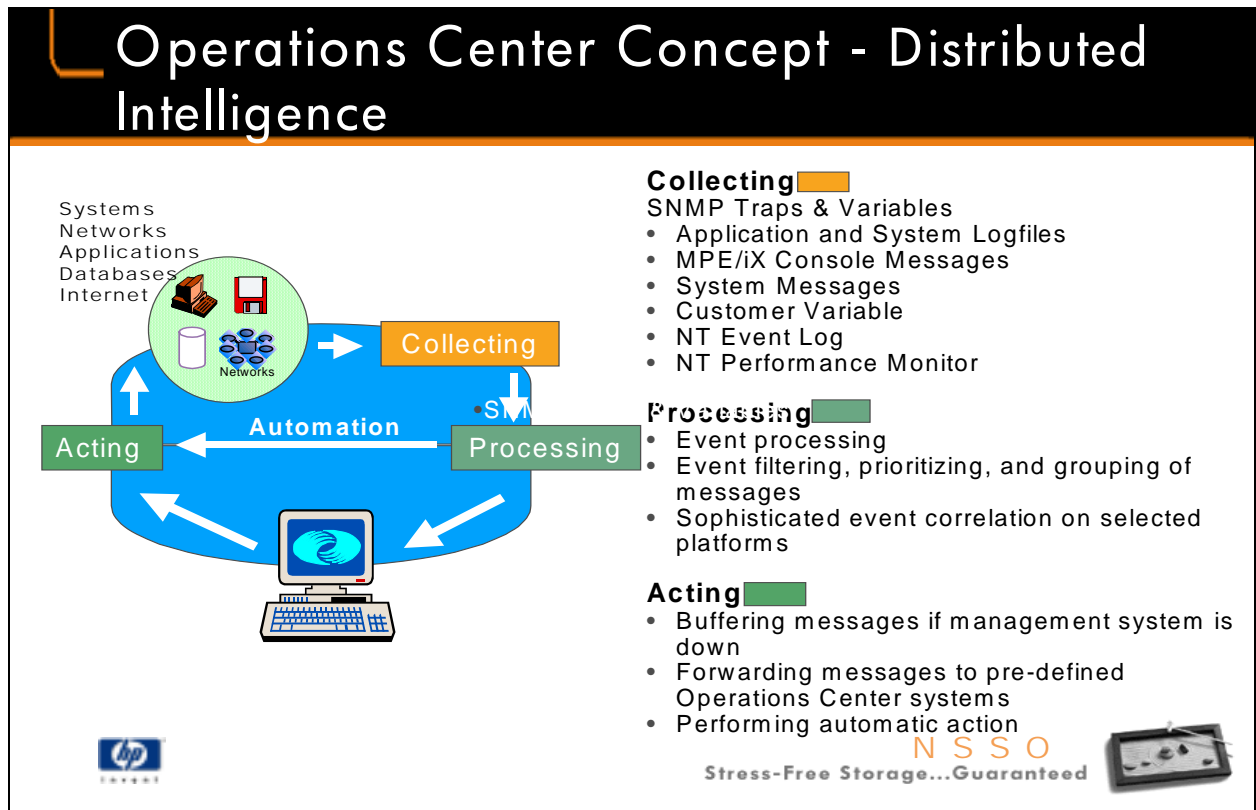
- Operations Center Concept
- Architecture
- Instrumentation
- Installation
- Configurations
- Discovery
- Service Model
- Technical Details
- Wrap Up



Stress-Free Storage...Guaranteed



HP OpenView SMART Plug-Ins are fully integrated, ready-to-go solutions which "plug into" the HP OpenView platform, extending the managed domain to industry leading business, Internet, middleware, database and management applications. A natural extension of the HP OpenView platform, these tightly integrated modules leverage the familiarity, security and scalability of HP OpenView solutions, while eliminating redundant infrastructure and processes. SMART Plug-Ins enable the rapid incorporation of new business and management management applications into the managed IT environment, allowing IT organizations to provide the highest level of support to their business customers. Designed as preconfigured, ready-to-go modules, SMART Plug-Ins are flexible and scalable, requiring minimal configuration effort.



One of the fundamental elements of any network infrastructure is a high-availability storage facility such as the HP SureStore E Disk Array XP Family of Arrays (XP Family of Arrays). The HP SureStore E Disk Array XP Family of Arrays is designed as a complete, high-availability, storage solution. The HP OpenView Operations Center SMART Plug-In for SureStore E Disk Array 256 (XP Family of Arrays-SPI) connects Vantage Point Operations Center to the XP Family of Arrays environment and allows events and error conditions in the XP Family of Arrays environment relating, for example, to disks, processors, or shared memory, to be proactively monitored and managed from one, central point: the Operations Center management server. Once configured, the XP Family of Arrays-SPI allows Operations Center operators to see and react to problems in the XP Family of Arrays environment before the problems become serious enough to affect applications, and ultimately the end-users, who are using the affected applications. The OpenView SPI for XP Family of Arrays is the intelligent choice for managing your storage environment from a central, best-in-class console. The OpenView SPI for XP Family of Arrays will increase availability and performance, let you visualize root-cause analysis and lower the cost of maintaining your storage environment. The HP OpenView Operations Center Smart Plug-In for the SureStore E Disk Array XP Family of Arrays (XP Family of Arrays) connects HP OpenView Operations Center to the XP Family of Arrays environment enabling management from a central point: the Operations Center management server. The

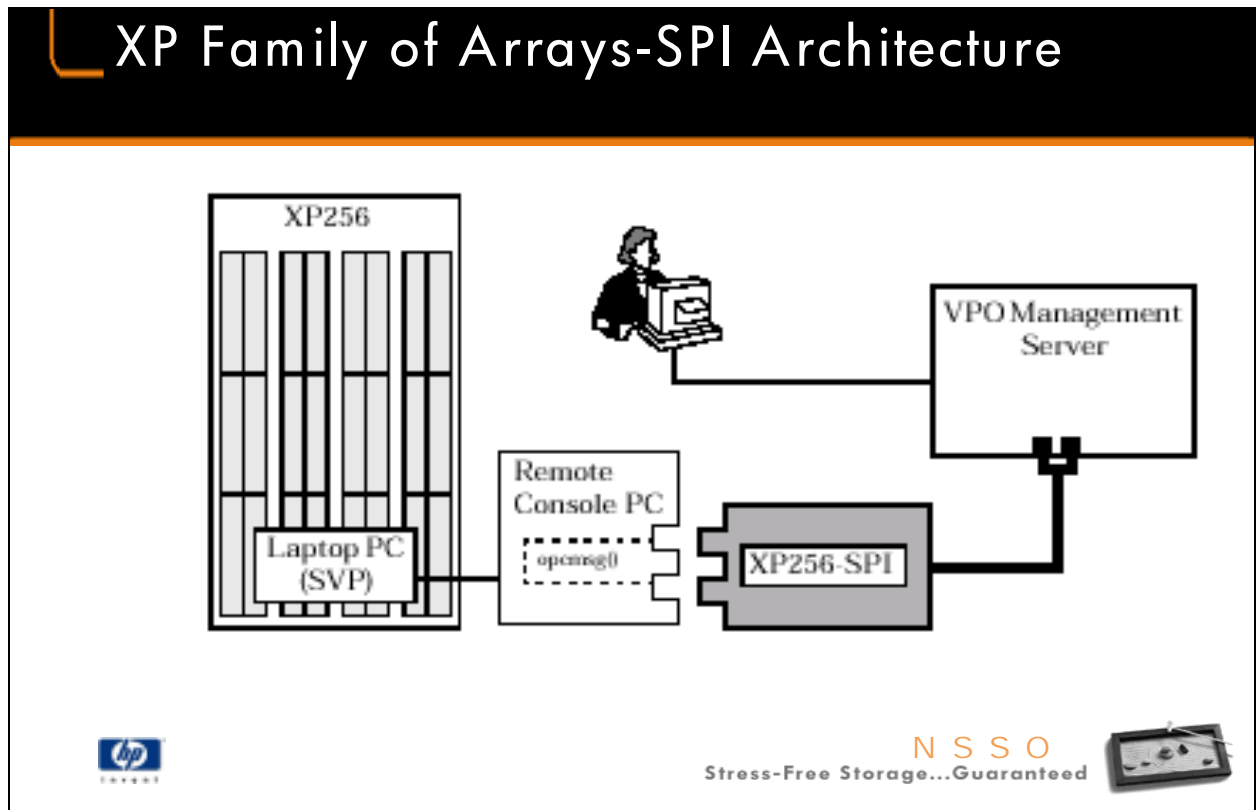
XP Family of Arrays-SPI allows operators to see and react to problems in the XP Family of Arrays environment before the problems become serious enough to affect applications and ultimately end-users.

HP OpenView Operations Center SMART Plug-Ins Defined:

HP OpenView Operations Center SMART Plug-Ins are fully integrated, ready-to-go solutions which "plug into" HP OpenView, extending the managed domain to industry leading business, Internet, middleware, database and service desk applications. A natural extension of the HP OpenView Operations Center platform, these tightly integrated modules leverage the familiarity, scalability, reliability and security of HP OpenView solutions, while eliminating redundant infrastructure and processes.

How Does it Work?

The XP Family of Arrays generates messages related to everyday operations and service requirements in the form of warnings, alarms, events and component failures. These messages are stored in a log on the XP Family of Arrays Remote Console, which is a PC linked to the XP Family of Arrays Disk Array. During the installation process, Operations Center agents are installed and configured on the Remote Console. These agents analyze, filter and prioritize the messages generated by the XP Family of Arrays before forwarding them to the Operations Center management server, enabling the appropriate operator to see and react to the messages before end-users are affected.



### The Features of the **XP Family of Arrays-SPI**

The following is a list of the features provided by the HP OpenView Operations Center SMART Plug-In for SureStore E Disk Array 256: - intercept, filter, and prioritize events received from the **XP Family of Arrays** Remote Console - monitor the Remote Console (RC) components - automatically detect the configuration of the **XP Family of Arrays** and upload the data into HP OpenView Operations Center Navigator - tight integration with Operations Center Navigator to provide service view of the **XP Family of Arrays** components and, in addition, automatically map Operations Center messages to the appropriate service objects - root-cause analysis of component failures - propagation scenarios to visualize the wider impact of component failures The **XP Family of Arrays** generates Service Information Messages (SIMs) that are related to everyday operations and service requirements such as warnings, alarms, events and component failures. SIMs are stored on the XP Family of Arrays internal Service Processor (SVP) and log a Remote Service Information Messages (R-SIMS) on the XP Family of Arrays Remote Console, a PC linked to the XP Family of Arrays disk array. Operations Center intelligent agents are installed, configured and running on the Remote Console, which is a managed node in the Operations Center management domain. The Remote Console uses Operations Center's opcmgr interface to pass on the R-SIMS to the Operations Center agent. The Operations Center agent analyzes, filters, and prioritizes the R-SIMS before packaging this information into a Operations Center message, which it forwards to the Operations Center management server.

The appropriate operators can see and react to the messages as they appear in the message browser window.

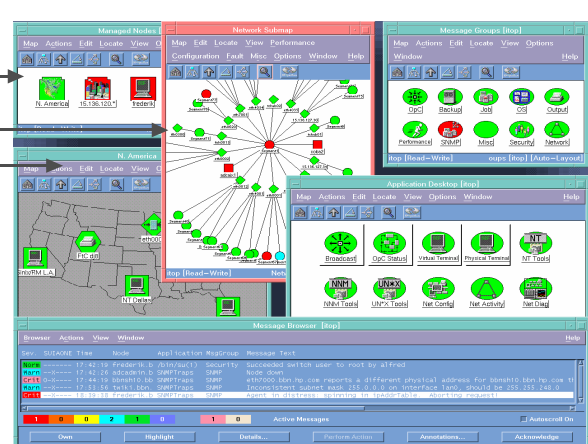
The XP Family of Arrays-SPI monitors and reports on the following areas of the XP Family of Arrays:

- System Faults
- Environmental Alarms
- SVP Communication Problems
- Disk Error Problems
- Hot Spare Disk Invoked
- Business Copy Abnormal Copy Problems
- Continuous Access links not operational
- Asynchronous Access links not operational
- Auto LUN Volume Migration failed


# Operations Center – Motif GUI

Fully-Customizable Task Oriented Working Environments

- **Managed Node Window/IP Map**
  - End-to-end view of operator's management domain
  - Problems identified with severity colors
  - Automatic discovery
  - On-Demand submaps
- **Message Browser**
  - Display of all incoming messages
  - In-depth message details
  - User guidance
  - Event escalation
  - Message ownership
  - Consistent presentation




The screenshot displays the HP Operations Center Motif GUI. It features a central network map with nodes and connections, a message browser window showing incoming messages, and an application dashboard with various icons for system components like CPU, Memory, and Network. The interface is designed for task-oriented operations with multiple panes and toolbars.



N S S O

Stress-Free Storage...Guaranteed



HP OpenView Training Schedules are located at URL:

<http://openview.hp.com/train/index.asp>

The tight integration with the HP OpenView Operations Center Navigator allows the XP Family of Arrays-SPI to illustrate clearly and graphically, by means of a service hierarchy, both physical and logical views of the hardware and software components as well as the relationship between components and the SIMS intercepted by the Operations Center agent. Once configured, this service view provides a quick and efficient way to isolate the root-cause of failures in low-level components and, in addition, graphically propagate the effect of such a failure to the higher levels of the hierarchy in order to predict the practical consequences.

hp openview Operations Center overview

HP OpenView Operations Center is a distributed client/server software product. Clients, called Intelligent Agents, and the Central Management Console--the server--work together to provide a mission control center for the entire distributed environment. From the central management console, IT staff can take full control of distributed IT resources across the enterprise, identifying potential problems before they occur and resolving them before end-users are affected. IT staff can also configure the powerful Intelligent Agents to solve problems immediately, without having to interact with the central management console--a major step toward a self-

managing system. Operations Center offers management add-ons to extend its functionality:

- hp openview Operations Center advanced security
- hp openview Operations Center developer's toolkit
- hp openview Operations Center JAVA-based user interface
- hp openview Operations Center navigator

features

- manages system and network complexity
- provides stronger, more competitive service
- implements flexible solutions
- delivers reliable, cost effective management

- Tight integration with Operations Center Navigator to provide a service view of the XP Family of Arrays components and automatically map Operations Center messages to the appropriate service objects
- Root-cause analysis of component failures
- Propagation scenarios to visualize the wider impact of component failures
- Integrates storage into network & system management solution for single end-to-end management

The OpenView SPI for XP Family of Arrays is the intelligent choice for managing your storage environment from a central, best-in-class console. The OpenView SPI for XP Family of Arrays will increase availability and performance, let you visualize root-cause analysis and lower the cost of maintaining your storage environment.

# HP OpenView Operations Center Navigator *Features and Benefits*

Manage services, not nodes any longer.  
The XP Family of Arrays-SPI monitors and reports on the following components of the HP SureStore E Disk Array XP Family of Arrays:

- System Faults
- Environmental Alarms
- SVP Communication Problems
- Disk Error Problems
- Hot Spare Disk Invoked
- Business Copy Abnormal Copy Problems
- Continuous Access links not operational
- Asynchronous Access links not operational
- Auto LUN Volume Migration failed

During installation, the XP Family of Arrays-SPI creates an application group that contains a number of applications designed to help the Operations Center operator manage and monitor the machines in the XP Family of Arrays environment.

Supported OS Platforms Supported Platforms for Operations Center Management  
Server 6.0



SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
HP OV Operation Center SPI

Platform	OS Supported OS Versions
HP 9000/700/800	HP-UX 10.20, 11.0 SUN Solaris 7,8

Supported Platforms for the Remote Console (Operations Center Managed Node 6.0)

Platform	OS Supported	OS Versions
Intel WinNT	4.0 (SP 5+)	

## Value of the XP Family of Arrays-SPI

- **Integrate the XP Family of Arrays into a Enterprise Management Solution**
- **Proactively monitored and managed from one, central point: the HP OpenView Operations Center management**
- **Display XP Failure Events in the central Message Browser**
- **Illustrate graphically the view of the hardware components in the Operations Center Service Navigator**

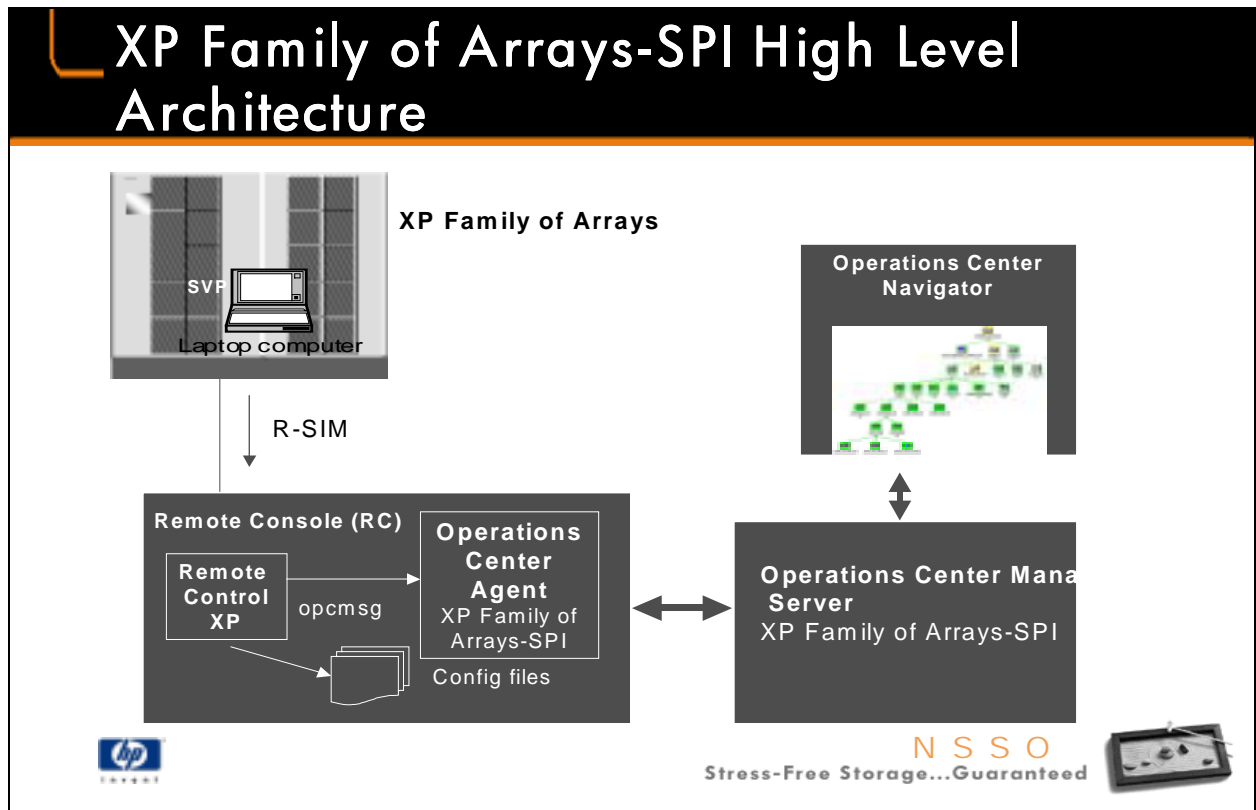


N S S O  
Stress-Free Storage...Guaranteed



### Introduction

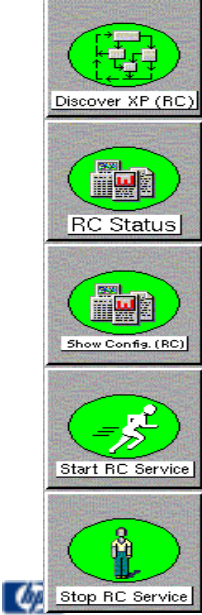
One of the fundamental elements of any network infrastructure is a high-availability storage facility such as the HP SureStore E Disk Array XP Family of Arrays (XP Family of Arrays). The HP SureStore E Disk Array XP Family of Arrays is designed as a complete, high-availability, storage solution. The HP OpenView Operations Center SMART Plug-In for SureStore E Disk Array 256 (XP Family of Arrays-SPI) connects Vantage Point Operations to the XP Family of Arrays environment and allows events and error conditions in the XP Family of Arrays environment relating, for example, to disks, processors, or shared memory, to be proactively monitored and managed from one, central point: the Vantage Point Operations (Operations Center) management server. Once configured, the XP Family of Arrays-SPI allows Operations Center operators to see and react to problems in the XP Family of Arrays environment before the problems become serious enough to affect applications, and ultimately the end-users, who are using the affected applications.





The XP Family of Arrays generates Service Information Messages (SIMs) relating to normal operations and service requirements as well as warnings, alarms, events, and component failures that occur in the XP Family of Arrays environment. SIMs are stored on the XP Family of Arrays internal Service Processor (SVP) and logged as Remote Service Information Messages (R-SIMS) on the XP Family of Arrays Remote Console, a PC linked to the XP Family of Arrays disk array. Operations Center intelligent agents are installed, configured, and running on the Remote Console, which is a managed node in the Operations Center management domain. The Remote Console uses Operations Center's opcmgs() interface to pass on the R-SIMS to the Operations Center agent. The Operations Center agent analyzes, filters, and prioritizes the R-SIMS before packaging vital information into a Operations Center message, which it forwards to the Operations Center management server. There, the appropriate operators can see and react to the messages as they appear in the Message Browser window. The tight integration with the HP OpenView Operations Center Navigator allows the XP Family of Arrays-SPI to illustrate clearly and graphically, by means of a service hierarchy, the view of the hardware and software components as well as the relationship between components and the SIMS intercepted by the Operations Center agent. Once configured, this service view provides a quick and efficient way to isolate the root-cause of failures in low-level components and, in addition, graphically propagate the effect of such a failure to the

higher levels of the hierarchy in order to predict the practical consequences. The service discovery feature generates the XP Family of Arrays Service model automatically, by using the configuration files, which are provided by the Remote Control software.

# Applications



- **Discovery XP**  
Starts the discovery of all XP Family of Arrays on the selected Remote Console node
- **Show Config**  
Displays the XP Family of Arrays configuration in textual form.
- **Start, Stop, Status of RMT\_CNSL service**  
Applications can be used for troubleshooting, because the Remote Console service should always run.

 Stress-Free Storage...Guaranteed 

The installation of the XP Family of Arrays-SPI creates an application group, called XP Family of Arrays SPI, which contains a number of applications designed to help the Operations Center operator to manage and monitor the machines in the XP Family of Arrays environment. Following applications are available:

Discover XP (RC)  
RC Status  
Show Config. (RC)  
Start RC Service  
Stop RC Service

The Discover XP application contacts the Remote Console and interrogates it about the contents of the configuration files. The configuration files reside on the Remote Console and contain information relating amongst other things to; controllers, adapters, physical and logical devices, parity groups, logical units, and program products. If the files are present, the information is used to build a service tree which is uploaded to the Operations Center database.


The RC Status application checks the current status of the Remote Console XP services and displays the results in an Output of Application window on the Operations Center operator's system.

The Show Config. (RC) application allows you to find out quickly if the configuration files required for the Discover XP application are present on the Remote Console and whether or not the Remote Console is correctly configured.


The Stop RC Service and Start RC Service applications provide a convenient way to stop and start the Remote Console XP services from the operator's system.

## Default External Node

The XP Family of Arrays Node Group



The XP Family of Arrays External Node Group



- Node Group XP Family of Arrays
- Node Pattern: xp.<\*>
- Uploaded during installation.



### The XP Family of Arrays-SPI Node Group

The installation of the XP Family of Arrays-SPI creates a node group, XP Family of Arrays, which is intended to be used for the systems in the XP Family of Arrays environment. You use this node group for the XP Family of Arrays Remote Console and the XP Family of Arrays disk arrays. By assigning the XP Family of Arrays node group and the Storage message group to the Operations Center operator who is responsible for the XP Family of Arrays environment, the Operations Center administrator can ensure that messages relating to errors and potential problems with the XP Family of Arrays appear in the appropriate operator's message browser. The installation of the XP Family of Arrays-SPI software creates an external node, called ALL XP Family of Arrays and adds it to the XP Family of Arrays node group. The Node Pattern field for the external node ALL XP Family of Arrays contains an asterisk (\*) in place of the nickname for a particular XP Family of Arrays machine. In this way, messages from all XP Family of Arrays machines that are added to the node bank appear in the Message Browser window of the operators to whom the Storage message group and XP Family of Arrays node group are assigned. If you want individual external-node icons in the XP Family of Arrays node group to represent particular XP Family of Arrays machines, you can make copies of the ALL XP Family of Arrays icon and replace the asterisk (\*) with the nickname of each XP Family of Arrays you want to define.

#### XP Family of Arrays-SPI Service Model

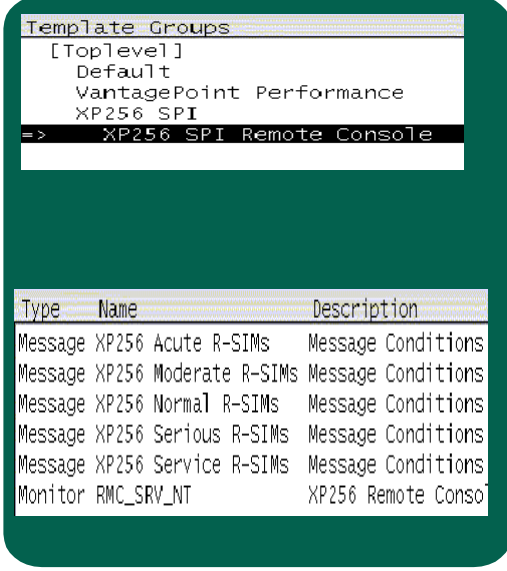
#### XP Family of Arrays-SPI Applications

During installation, the XP Family of Arrays SPI will create a series of groups: a node group, external node, and application group. Each of these groups is described below.

The XP Family of Arrays-SPI creates a node group, called XP Family of Arrays, which defines the group of XP Family of Arrays devices to be monitored and assigned to Operations Center. Assigning the XP Family of Arrays node group to the Operations Center operator ensures that messages relating to errors and potential problems with the XP Family of Arrays appear in the operator's message browser. The XP Family of Arrays-SPI software creates an external node, called ALL XP Family of Arrays, and adds it to the XP Family of Arrays node group. This enables messages from all XP Family of Arrays machines in the environment to appear in the message browser window of the storage administrator/operator.



# Templates & Template Groups



**Template Groups**


```
[Toplevel]
  Default
  VantagePoint Performance
  XP256 SPI
=> XP256 SPI Remote Console
```

Type	Name	Description
Message	XP256 Acute R-SIMs	Message Conditions
Message	XP256 Moderate R-SIMs	Message Conditions
Message	XP256 Normal R-SIMs	Message Conditions
Message	XP256 Serious R-SIMs	Message Conditions
Message	XP256 Service R-SIMs	Message Conditions
Monitor	RMC_SRV_NT	XP256 Remote Console

## Template Groups


XP Family of Arrays Error Severity Level	XP Family of Arrays Impact	Operations Center Message Severity
Serious	Subsystem stops	Critical
Moderate	Error section stops	Major
Service	Partial error occurs	Minor
	Minor error occurs	Warning

## Templates



N S S O

Stress-Free Storage...Guaranteed



### Default XP Family of Arrays-SPI Templates

The XP Family of Arrays-SPI message-source templates intercept SIMs, which are generated by error conditions on the XP Family of Arrays and logged on the Remote Console. The SIMs are forwarded to Operations Center using the opcmmsg() interface. Messages generated by the XP Family of Arrays-SPI templates pass through a filtration process and are then prioritized and assigned a default Operations Center severity level according to the rules described in the slide.

The only exception to this rule concerns messages generated by the XP Family of Arrays Miscellaneous R-SIMs template and relating to configuration changes in the XP Family of Arrays. The messages generated by all the XP Family of Arrays-SPI templates belong by default to the Storage message group. An additional monitor template verifies the status of the XP software services running on the XP Family of Arrays Remote Console and generate messages to indicate any changes in status.

### XP Family of Arrays-SPI Message Severity

The messages generated by the XP Family of Arrays-SPI templates appear in the Message Browser window and contain information that is vital for your understanding of the problem which is being reported. The severity of messages generated by the XP Family of Arrays-SPI templates is determined by the severity of




the original R-SIM error codes. As a general rule, XP Family of Arrays errors are mapped to message severities according to the values listed in this slide.

## opcmmsg Message Format

Input Message:  
**Appl:** XP Family of Arrays   **Obj:** <Failure Unit>   **Msg Grp:** Storage  
**Msg Text:** Ref.-Code <R-SIM> in <nickname> (Ser # <SN>)

Formatted Message:  
**Appl:** XP Family of Arrays   **Obj:** <Failure Unit>   **Msg Grp:** Storage  
**Nodename:** xp.<nickname> (!)  
**Msg Text:** <R-SIM>: <Error type>: <Error> in <nickname> (# <SN>)

Example Message Text:  
307361: DKA processor error: Processor blocking in DKA 2K  
(MP1 in PCB6) in 35390 (# 35390)



Input Message:

Appl: XP Family of Arrays   Obj: <Failure Unit>   Msg Grp: Storage  
Msg Text: Ref.-Code <R-SIM> in <nickname> (Ser # <SN>)

Formatted Message:

Appl: XP Family of Arrays   Obj: <Failure Unit>   Msg Grp: Storage  
Nodename: xp.<nickname> (!)  
Msg Text: <R-SIM>: <Error type>: <Error> in <nickname> (# <SN>)

Example Message Text:

307361: DKA processor error: Processor blocking in DKA 2K  
(MP1 in PCB6) in 35390 (# 35390)

Message Attributes:

**Node Name**   The name of the XP Family of Arrays machine on which the SIM error originated

**Application**   The application reporting the error. XP Family of Arrays-SPI messages always display the application name XP Family of Arrays

**Object**   In the event of an error, the name of the unit which failed. For example, DKC-Processor, DKU-Power Supply, or DKC-Drive. In the event of a change to the configuration of the Remote Console, the string: Configuration Change

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
HP OV Operation Center SPI

**Message Text**        The text of the message consisting of the following elements separated by a colon (:):

**SIM-RC**        The original SIM error code

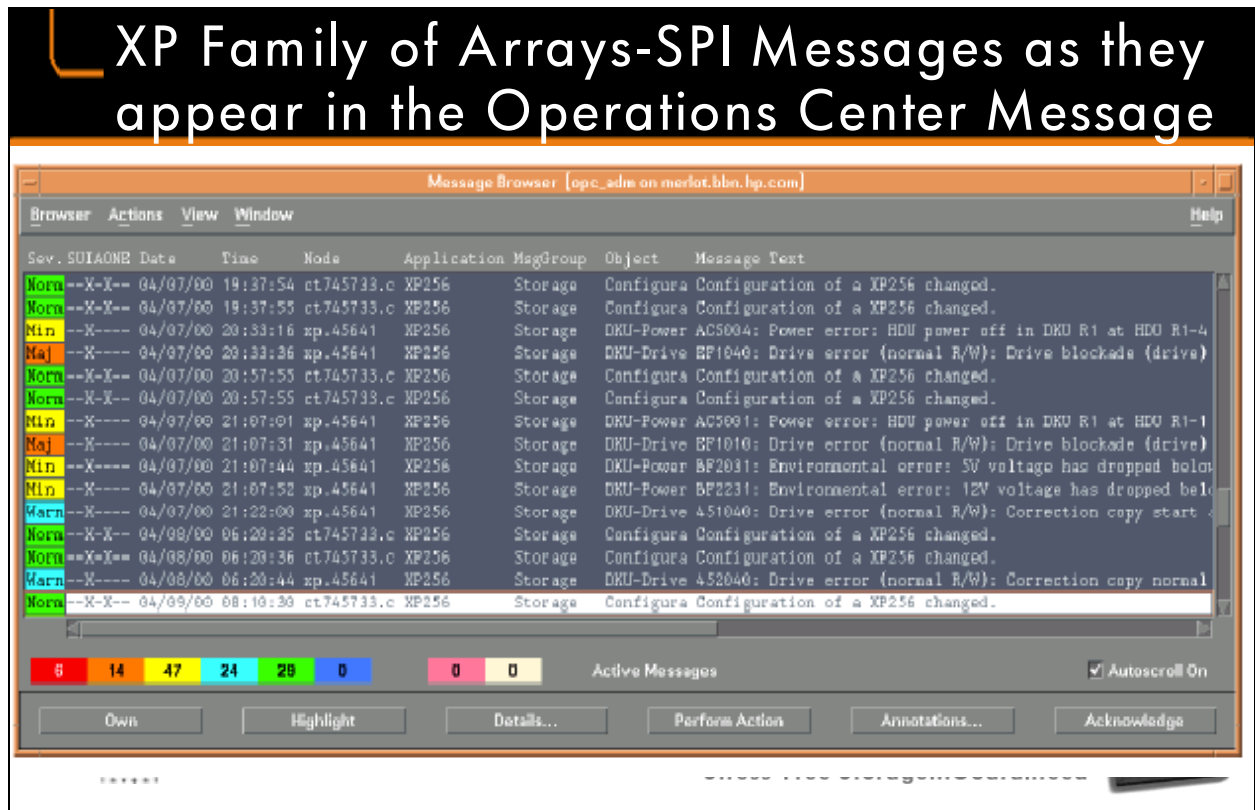
**Error type**    The type of error which occurred

**Description**   A short description of the error including the location

**Nickname**    The nickname assigned to the XP Family of Arrays machine. This is typically the serial number of the XP Family of Arrays, for example; (#35399).

However, it may    also be a name you choose, for example; XP Family of Arrays-1.

**(#S/N)**        The serial number of the XP Family of Arrays. For example,  
**(#35399)**



In addition to the message templates, the XP Family of Arrays-SPI also provides a miscellaneous R-SIMs template that generates messages regarding configuration changes in to the XP Family of Arrays. Messages generated by the XP Family of Arrays-SPI templates appear in the message browser window of Operations Center. The message headline provides detailed information of the problem which, in the context of the XP Family of Arrays-SPI, can be broken down in the following way:

**Node** Name Name of the XP Family of Arrays machine, external node, or the Remote Console on which the R-SIM error originated

**Application** The application reporting the error. XP Family of Arrays-SPI messages always display the application name XP Family of Arrays

**Object** In the event of an error, the name of the component that failed. For example, DKC-Processor, DKU-Power Supply or DKC-Drive.

In the event of a change to the configuration of the Remote Console, the string: Configuration Change will appear.

**Message Text** Text of the message consisting of the following elements separated by a colon (:)

- SIM-RC – the original SIM error code
- Error Type – type of error occurred

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
HP OV Operation Center SPI

Description- brief description

Nickname – typically the serial number of the XP Family of Arrays or a name chosen by the operator

(# S/N) – serial number of the XP Family of Arrays.

An example of an error generated by an XP Family of Arrays-SPI message template is as follows:

BF1220: Environmental error: Abnormal temperature (over 45 degrees C) detected in front of Logic 2 in XP Family of Arrays-1 (#35390).

## R-SIM Condition - Example

<b>Description</b>			
140100 SIM transfer failure to SVP in CHA 1P			
<b>Condition</b>			
<b>Severity</b>	<b>Node</b>	<b>Application</b>	<b>Message Group</b>
normal warning minor major critical			
<b>Message Text</b>			
ode<S><[140100].sim><S>in<S><@.Nickname><S><Ser#[<S>]><@.SerialNo>>			
8	+	140160 SIM transfer failure to SVP in CHA 2X	
9	+	140170 SIM transfer failure to SVP in CHA 2X	
10	+	150100 SIM transfer failure to SVP in CHA 1P	
11	+	150110 SIM transfer failure to SVP in CHA 1P	
12	+	150120 SIM transfer failure to SVP in CHA 1P	
13	+	150130 SIM transfer failure to SVP in CHA 1P	
14	+	150140 SIM transfer failure to SVP in CHA 1P	
15	+	150150 SIM transfer failure to SVP in CHA 1P	
16	+	150160 SIM transfer failure to SVP in CHA 1P	
17	+	150170 SIM transfer failure to SVP in CHA 1P	
18	+	150180 SIM transfer failure to SVP in CHA 1P	
19	+	150190 SIM transfer failure to SVP in CHA 1P	
20	+	150200 SIM transfer failure to SVP in CHA 1P	
21	+	150210 SIM transfer failure to SVP in CHA 1P	
22	+	150220 SIM transfer failure to SVP in CHA 1P	
23	+	150230 SIM transfer failure to SVP in CHA 1P	
24	+	30906Y LDEV blockade in effect of processor blockade	
25	+	30907Y LDEV blockade in effect of processor blockade	
<b>Node</b>			
xp.<Nickname>			
<b>Application</b>			
Storage			
<b>Message Group</b>			
Storage			
<b>Object</b>			
<b>Message Text</b>			
<sim>: SVP interface error (CHA side): SIM transfer failure to SVP in CHA 1P in <Nickname> (#<SerialNo>)			
<b>Service Name</b>			
XP256.<SerialNo>.CHA_1P.SVP_IF			

Pattern Matching

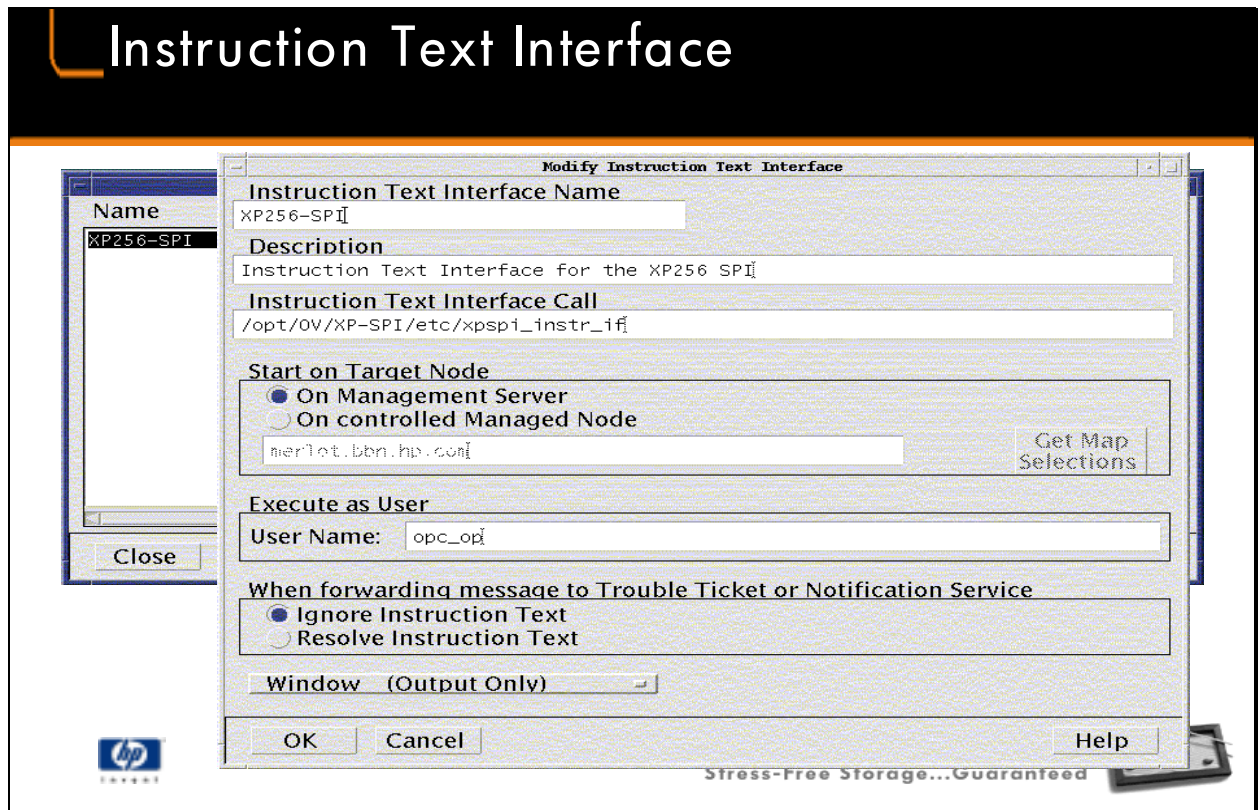
Set Attributes

Detailed Message Text



N S S O  
Stress-Free Storage...Guaranteed





### Instruction Text Interface

All R-SIM conditions use the XP Family of Arrays-SPI Instruction Text Interface.

Here is an example of a serious R-SIM.

A serious R-SIM arrived. This means, the error section (see Object) stopped!

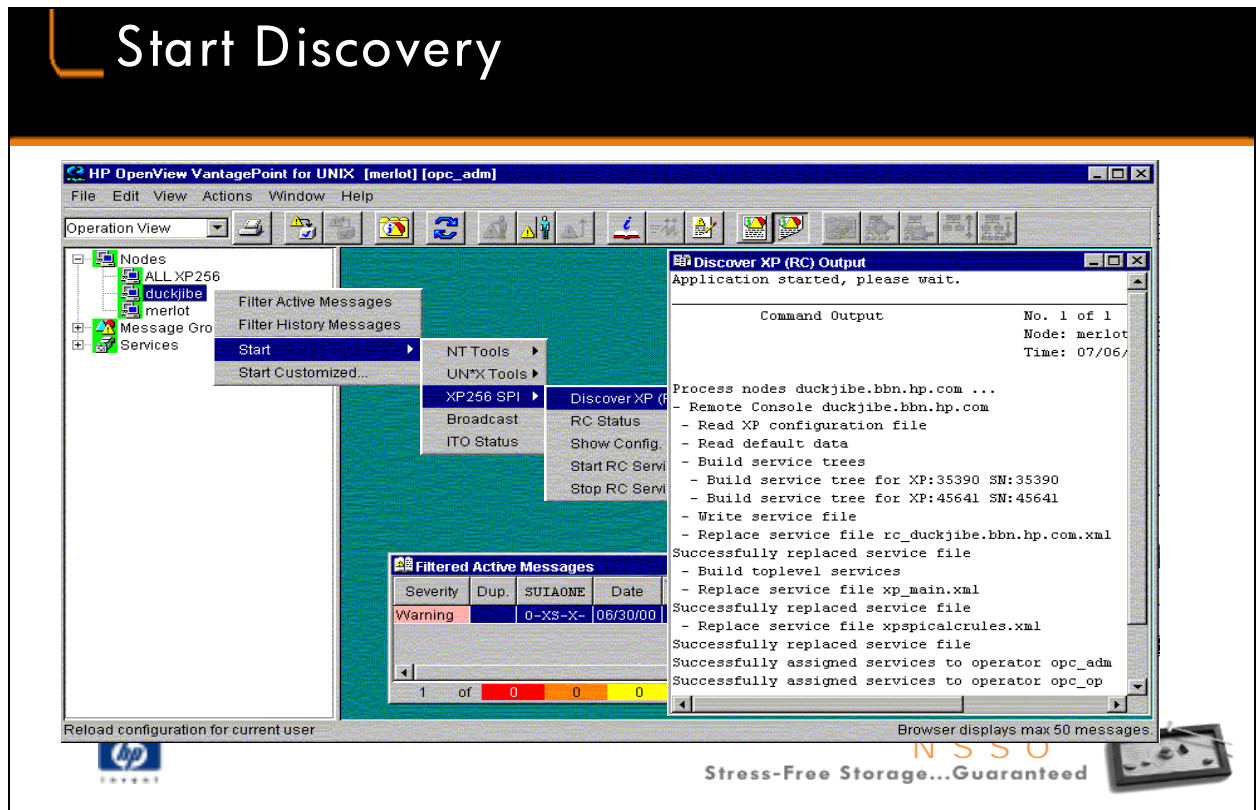
The XP Family of Arrays SIM reference codes are classified into four levels:

XP Family of Arrays Level of Error    XP Family of Arrays Impact    ITO  
Severity

Acute	Subsystem stopped entirely	Critical
Serious	Error section stopped	Major
Moderate	Partial error occurred	Minor
Service	Minor error occurred	Warning

Please contact your XP Family of Arrays Administrator or call the HP Response Center!





The XP Family of Arrays-SPI Service Model, the tight integration with the HP OpenView Operations Center Navigator allows the XP Family of Arrays-SPI to illustrate graphically a service hierarchy tree detailing the associations and dependencies that exist between components in the XP Family of Arrays environment. The service hierarchy tree is generated by running the Discover XP application on the Remote Console. The operator can use the Java-based GUI to view the hardware and software components as well as the relationship between components and SIMS intercepted by the Operations Center agent. Once configured, this service view provides a quick and efficient way to isolate the root-cause of failures in low-level components and graphically reproduce the effect of such a failure to the higher levels of the hierarchy in order to predict the consequences.

Configuration of XP:35390 SN:35390

DKCMAIN Microcode Version: 52-47-02-00/00  
SVP Microcode Version: 51-47-01/00

CHA:  
CHA 1P: CHS(SCSI)  
CHA 1Q: CHT

....

-----  
DKA:

DKA 1B installed

DKA 1C installed

....

-----  
DKU R1:

HDU-R10

R100: Model=DK30A, Capacity=32G, Serial=81050536, PDEV=00/01,  
Type=Data Drive

R101: Model=DK30A, Capacity=32G, Serial=81050860, PDEV=00/03,  
Type=Data Drive

....

-----  
LUN:

CHA-Port=1A, SCSI-ID=0, LUN=0, Host-Mode=08, LDEV=0:00

CHA-Port=1A, SCSI-ID=0, LUN=1, Host-Mode=08, LDEV=0:01

CHA-Port=1A, SCSI-ID=0, LUN=2, Host-Mode=08, LDEV=0:02

....

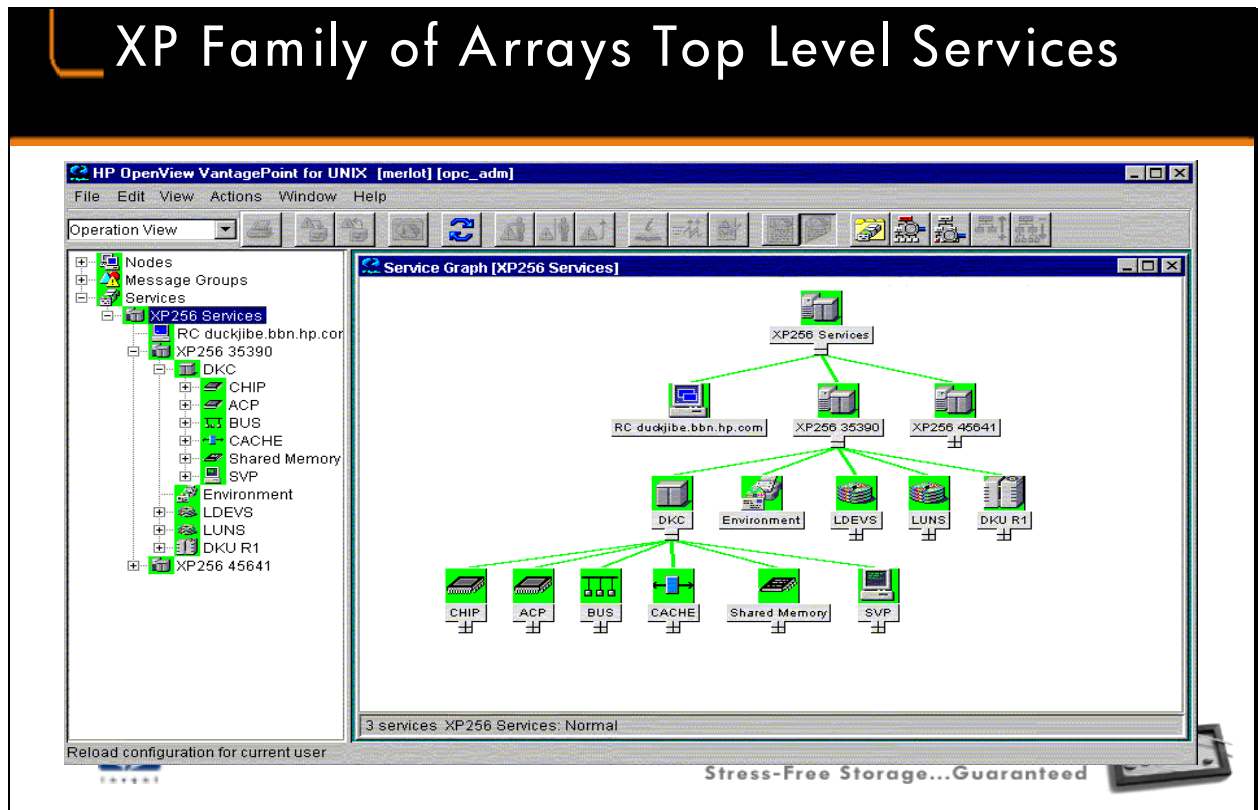
-----  
LDEV:

0:00 Type=OPEN-9, Capacity=7211520 Bytes, Not Configured CVS, Not  
Command Device

0:01 Type=OPEN-9, Capacity=7211520 Bytes, Not Configured CVS, Not  
Command Device

0:02 Type=OPEN-9, Capacity=7211520 Bytes, Not Configured CVS, Not  
Command Device

0:03 Type=OPEN-9, Capacity=7211520 Bytes, Not Configured CVS, Not  
Command Device



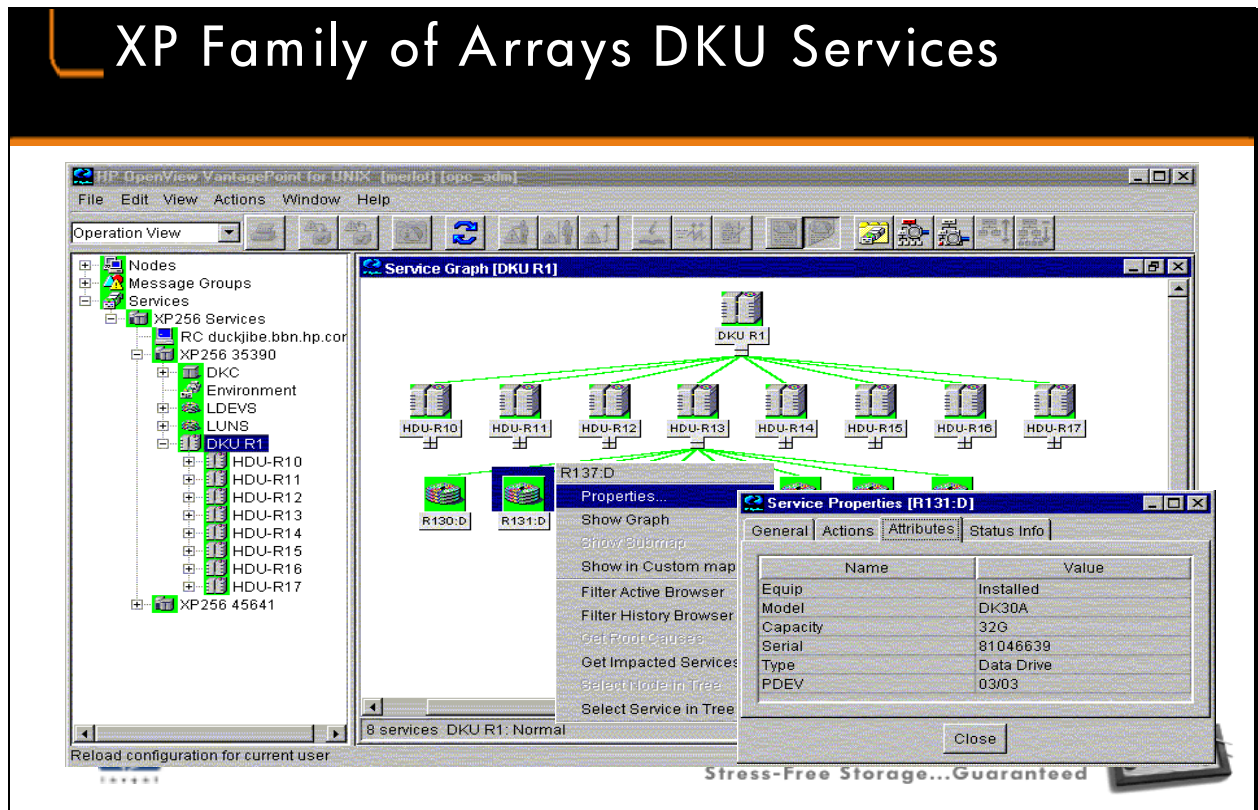
The Service Tree Graph, above, details the elements present in the top-level service hierarchy of an XP Family of Arrays. The top-level service in the XP Family of Arrays-SPI is XP Family of Arrays Services. XP Family of Arrays Services contains icons representing the Remote Console PC and each XP Family of Arrays in the environment. The icons representing the Remote Console and the XP Family of Arrays systems have the following naming scheme: RC.<hostname> and XP Family of Arrays.<nickname>. Each of the XP Family of Arrays devices has five sub-elements structured hierarchically. These include:

**DKC** Disk Control Frame which includes CHIP, ACP, Bus, Cache, Shared Memory, SVP

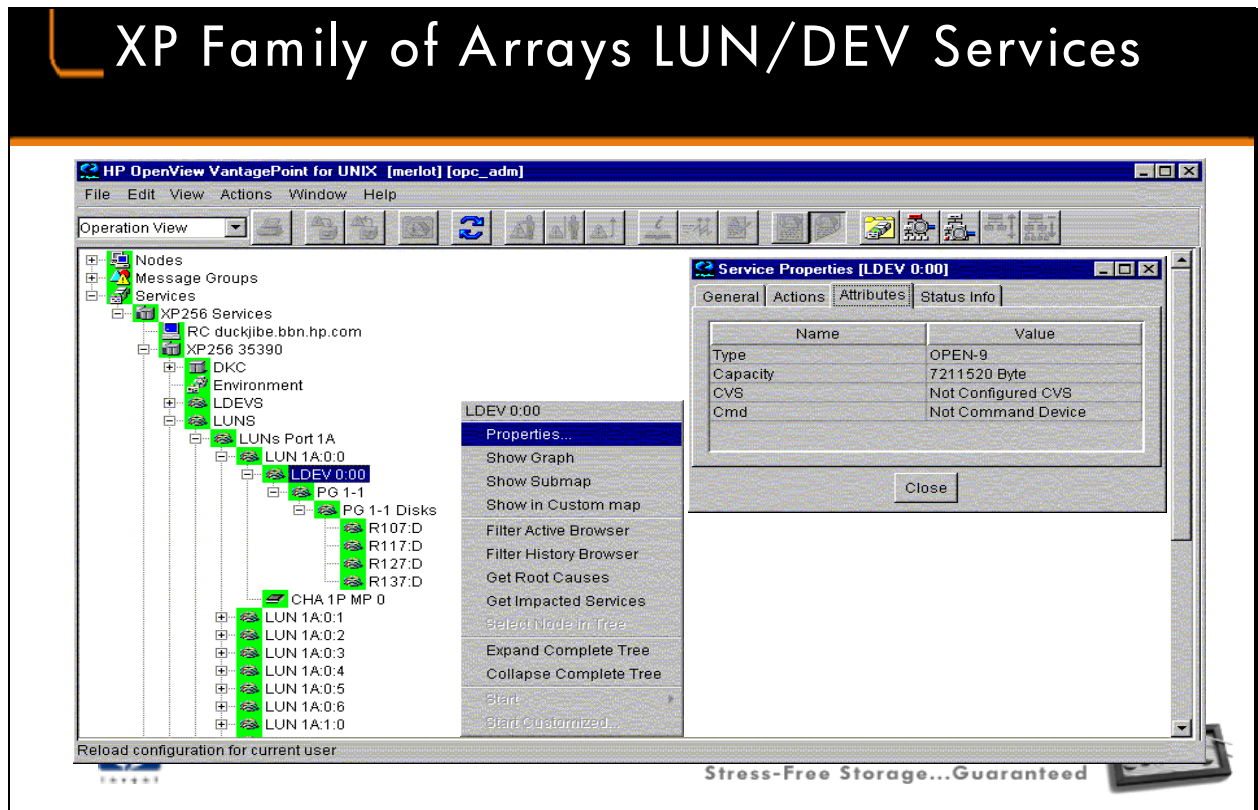
**Environment** General environment conditions such as temperature, voltage, power supply

**LDEVs** Logical Devices including parity groups (parity levels), disks, parity-group disks, DKA SCSI Ports

**LUNs** Logical Units such as SCSI and fibre-channel paths, LDEVs (emulation type and capacity)

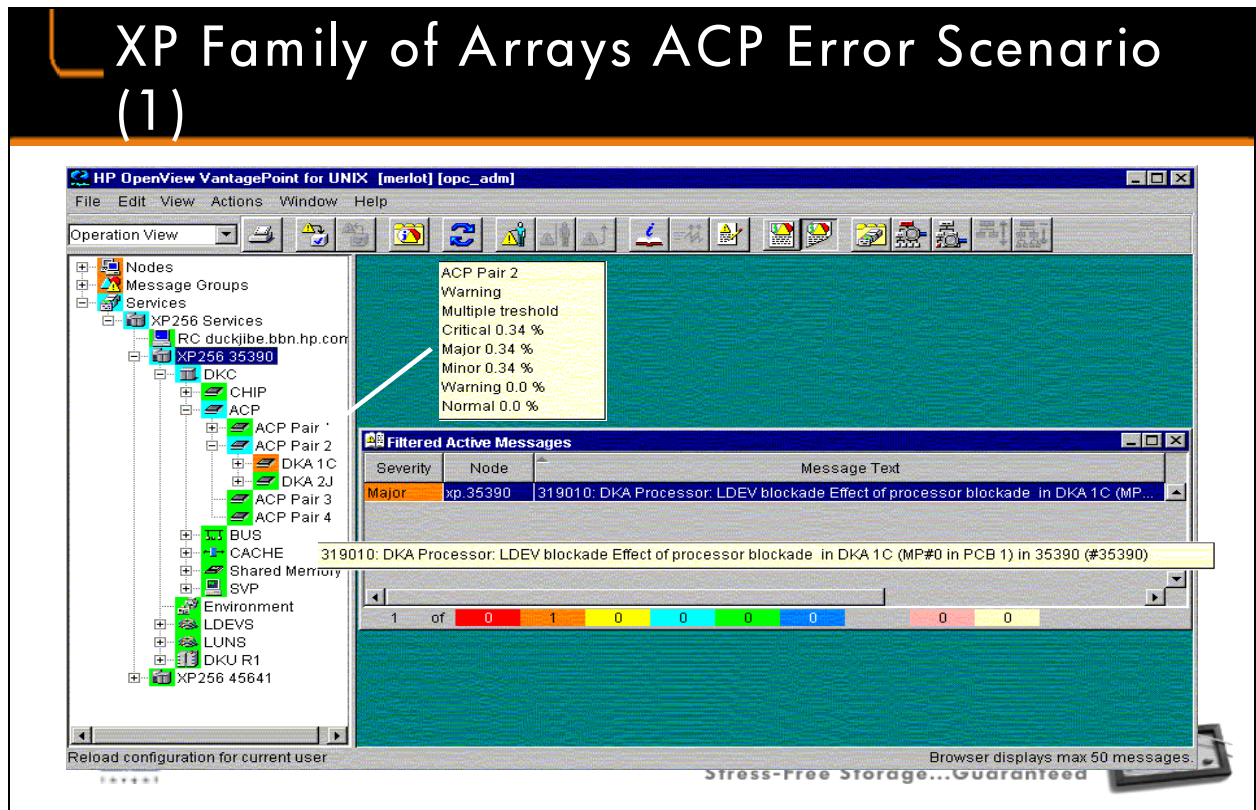


Clicking the (+) sign located beneath each of the icons expands the icon graphically displaying the lower-level components. This graphical hierarchy enables the operator to determine the root cause of a problem by viewing the dependencies and associations of the components that will speed-up the problem resolution process and increase the efficiency of the operator. A right-mouse click on any of the icons will result in a pop-up menu with a number of actions available.



Shows emulation type, capacity, if the LDEV is a Command or CVS volume.





### Monitor Template

The monitor checks if the remote console service is running. If not, it sends a ITO message. The operator can investigate the problem with the start/stop/status application for this service.

### Opcmsg Templates

The XP Family of Arrays-SPI message-source templates intercept SIMs, which are generated by error conditions on the XP Family of Arrays and logged on the Remote Console. The SIMs are forwarded to Operations Center using the `opcmsg()` interface. Messages generated by the XP Family of Arrays-SPI templates pass through a filtration process and are then prioritized and assigned a default Operations Center severity level.

The only exception to this rule concerns messages generated by the XP Family of Arrays Miscellaneous R-SIMs template and relating to configuration changes in the XP Family of Arrays. The messages generated by all the XP Family of Arrays-SPI templates belong by default to the Storage message group.

### XP Family of Arrays Configuration and Reference Tables

The XP Family of Arrays Configuration and Reference tables are generated by the Remote Control XP software. They are used to build the service model. The Operations Center service files (XML) are generated on the Management Server.

The script xpspi\_catcpnf.bat on the Agent writes the table to STDOUT and the output is captured by the Management Server.

## XP Family of Arrays Service Discovery Mechanism

- 1) Select the **Remote Console node** in the Operator UI and start the **Operations Center application "Discovery XP"**
- 2) This application **launches** the script xpspidiscovery on the **Operations Center Management Server**
- 3) The script xpspidiscovery starts **xpspi\_catconf.bat** on the **Remote Console** via opctrnm and writes the output (config files) of opctrnm into a **temporary file**.
- 4) xpspidiscovery starts **xpspi\_disc.pl** on Management Server.  
This perl script generates the service files (XML) for **each XP Family of Arrays** and also for the **main XP Family of Arrays Service**
- 5) Finally, it **upload** all service files and the file with the calculation rules.

### Remote-Console Configuration Files

Configuration	File Description
Controllers	Information relating to controllers such as; BOX Serial #
Adapters	Information relating to adapters such as; CHA, DKA, DKU, etc. Physical Devices Information relating to hardware such as; the model name of the hard disks
Parity Groups	Information relating to parity groups such as; RAID level, etc.
Logical Devices	Information relating to logical devices such as; emulation type and capacity
LUNs	Information relating to Logical Units such as; SCSI/Fibre-channel path, LUN, and LDEVs
PPs	Information relating to software options

### ITO Management Server Default Files

xpspicalcrules.xml All calculation- and propagation rules that are used  
by the service model. Could be changed by the  
customer to implement a different behavior of the status  
propagation.

## Module Wrap



N S S O  
Stress-Free Storage...Guaranteed





## Module 14

# Supported XP Disk Array Family Windows NT 4.0 & 2000 Configurations



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Module Agenda



➤ hp surestore xp fast recovery  
for Exchange 2000



➤ Supported Windows NT & 2000 Configurations

➤ Supported Microsoft Cluster Server (MSCS)

➤ Microsoft Certification

➤ Assured Availability, and the XP Family of  
Arrays



➤ Cluster Extension XP for Microsoft

➤ Wrap-up




HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



Module Introduction, MS Exchange deployment provides real opportunity to sell storage. Exchange administrators are now considering consolidation of Exchange servers and users in general have realized the criticality of their messaging solutions. The XP Family of Disk Arrays offer solutions to meet Exchange consolidation and availability needs. This section is not intended to teach the student MS Exchange or MS NT. Rather, it is intended to provide an overview of the Exchange environment and to provide background information valuable in pre-sales venues. The HP Microsoft Services Organization (MSO) can provide MS Exchange expertise for deal support where that expertise is critical. The second part of this section provides information relating to Exchange deployment on the XP Family of Disk Arrays. The student will gain an understanding of the issues and factors to consider when deploying Exchange on the XP XP Family of Arrays.



HP Restricted

hp surestore xp fast recovery  
for Exchange 2000

recover Microsoft  
Exchange stores  
in minutes, not hours:

hp surestore xp  
fast recovery  
for Exchange 2000:

manageable


fast

Reliable

available

efficient

Stress-Free Storage...Guaranteed



No organization wants to take precious time to recover data. In fact, many enterprises have a 24x7 availability goal. Fast Recovery for Exchange 2000 is manageable because it eliminates backup device access conflict delays for multiple consolidated Exchange servers. It enables Exchange Service Providers to meet or exceed higher levels of Service Level Agreement (SLA) requirements. And it can be integrated with nightly backup to attain higher levels of confidence in the integrity of Exchange 2000 data stores.

It's fast and reliable. It can recover a corrupted Exchange 2000 database in minutes versus hours.

And Fast Recovery for Exchange 2000 provides for larger stores without increased recovery windows, reducing cost per MB.

HP's Surestore XP Fast Recovery for Exchange 2000 is one tool that can help Exchange 2000 meet these requirements. In the event of an Information Store corruption (either Private, Public or Directory), it can use a secondary volume created by Business Copy XP to recover the Information Store instead of having to use a tape backup.

This method of recovery is significantly faster than recovery from tape backup – more than an order of magnitude! With other tools like zero-downtime backup, Exchange can approach the goal of an always-on infrastructure. Fast Recovery for Exchange 2000 can increase the availability of Exchange to allow enterprise IT managers, application service providers and their customers more confidence that Service Level Agreement (SLA) requirements meet. They can be confident that in the event of Exchange database corruption, the downtime window for recovery are as short as possible.

As you can see, today, there are two methods of backup.

- 1) An online backup, which occurs while Exchange services are still available, requires no special software or solution other than backup software that supports interaction with the Exchange Backup API, or
- 2) A snapshot backup. Here a secondary volume is created by Business Copy XP. At backup time, the Exchange services are halted for a short time to establish logical consistency of the database, the secondary volume is “split” from the primary volume, Exchange services are re-started, and the backup application uses the now-independent secondary volume to do the backup. Once the backup is completed, the secondary volume is resynchronized to the primary volume. Either method can backup to either tape or disk.

In method #2, an additional copy volume of the secondary volume is split from the secondary volume, and reserved for Fast Recovery for Exchange 2000. In the event of an Exchange database corruption, it would use this additional secondary volume for Exchange recovery instead of having to use the tape backup. This recovery is over 10 times faster than recovery from a tape backup.

In method #1, if the online backup was to disk, this solution would restore the online backup onto a secondary volume, and then use that secondary volume for Exchange recovery.

Today's solution is focused on Exchange 2000. This capability will be extended to other application environments in the future.

## hp surestore xp fast recovery for Exchange 2000: target customers.



The image contains two side-by-side photographs. The left photograph shows a woman with dark hair, wearing a red shirt, smiling and resting her chin on her hand at a desk. The right photograph shows a row of five tape drives in a rack, with a greenish-blue textured background. Below the left photo is the text 'HP Restricted'. Below the right photo is the text 'Stress-Free Storage...Guaranteed' and a small icon of a tape drive.

HP Restricted

Stress-Free Storage...Guaranteed

Even for small Exchange deployments, recovery is an activity that can take several hours. For larger environments, recovery could literally take days.

Fast Recovery for Exchange 2000 is targeted at customers who consider Exchange a business-critical application. They may have more than 2,000 mailboxes, or a high cost of downtime. It's very possible that these customers have already implemented other practices to minimize downtime, such as a zero-downtime backup snapshot backup solution, or limited the size of Information Stores based on a predicted recovery window.

For example, if a measured recovery pace from backup is 100Gb/hour, an administrator may limit Information Stores to 100Gb. This can ensure that once the recovery process has started, the duration of downtime will not be longer than one hour.

Best of all, with it, recovery can begin immediately, with information coming directly from disk rather than from tape.

Messaging and collaboration is a business-critical application. The cost and loss of productivity of downtime is significant. Even under ideal conditions with no server or storage load, recovery of a 100 GB Exchange database would take three to four hours.

GartnerGroup estimates that the recovery time to restore a 30 GB information store would be four to six hours, and for larger systems eight to 30 hours.

By using the secondary volume technology of Business Copy XP, Fast Recovery for Exchange 2000 can recover a database in minutes.

Exchange Service Providers (xSPs) often have to offer a guaranteed level of availability – say, 24x7 at a 99.9% uptime. But even 99.9% means up to 8 hours of downtime per year. If the Exchange database is large, a single failure/recovery could take several days.

That's why xSPs are anxious to eliminate risk of catastrophic downtime. For a given history or forecast of Exchange corruptions, Fast Recovery for Exchange 2000 can directly forecast improved SLA.

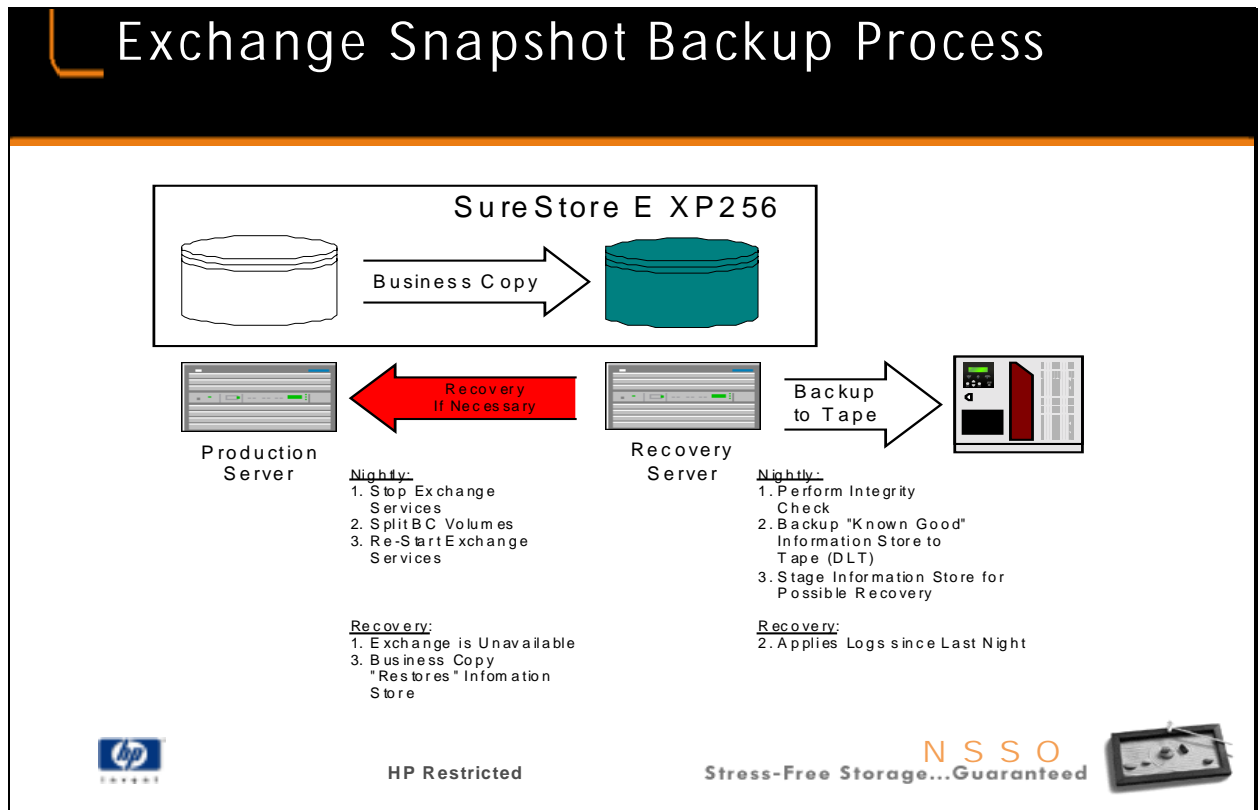
By using Fast Recovery for Exchange 2000 in conjunction with Business Copy, Exchange 2000 can be considered a much more robust solution.

It can significantly reduce or eliminate long Exchange service outages during the recovery process. This duration could be measured in a test environment without affecting in-production Exchange services. Exchange configurations or XP disk configurations could be “tuned” to regulate the anticipated duration.

For example, some Exchange administrators limit individual Information Store sizes to meet a target recovery duration. Because Fast Recovery for Exchange 2000 is so much faster than recovery from a tape backup, the Exchange administrator could increase Information Store sizes to reduce management resource overhead while still achieving target recovery durations.


Also, Information Store corruptions can be identified sooner, perhaps before any users have been negatively affected.

Whichever solution you choose, you'll benefit from HP's innovative 5nines:5minutes Program. It's an end-to-end solution with leading industry partners that delivers 99.999% availability or less than five minutes of downtime per year using HP Fast Recovery for Exchange 2000 in conjunction with XP Disk Arrays—perfect for every business need.



High Availability Exchange XP Family of Arrays Solution means:  
NO downtime during backup  
Minimal downtime during recovery

# Architecture Overview of Exchange XP Family of Arrays solution



- Production Exchange server holds the Information Store that is backed up nightly.
- Recovery Exchange Server is used to integrity check the database nightly, and be recovery-ready.
- XP Family of Arrays Business Copy is used to hold current and previous night's "recovery-ready" databases.
- In the event of a recovery, business copy is used to "restore" from Recovery server to Production server.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



Here's how Fast Recovery for Exchange 2000 works: By maintaining a secondary copy of the Exchange database as a Secondary Volume created with Business Copy XP, it can overlay the corrupted Exchange database with this secondary copy.

Here's how Fast Recovery for Exchange 2000 works: By maintaining a secondary copy of the Exchange database as a Secondary Volume created with Business Copy XP, it can overlay the corrupted Exchange database with this secondary copy.

The production server is up and running during the entire backup process. Exchange Online Backup APIs are used during the backup process thus allowing users access to the information store and also verifying checksums of pages. Nightly integrity checks guarantee you have a consistent database. The information store is in place for a fast recovery on the Recovery server.



# **Supported XP Disk Array Family Windows NT 4.0 & 2000 Configurations**

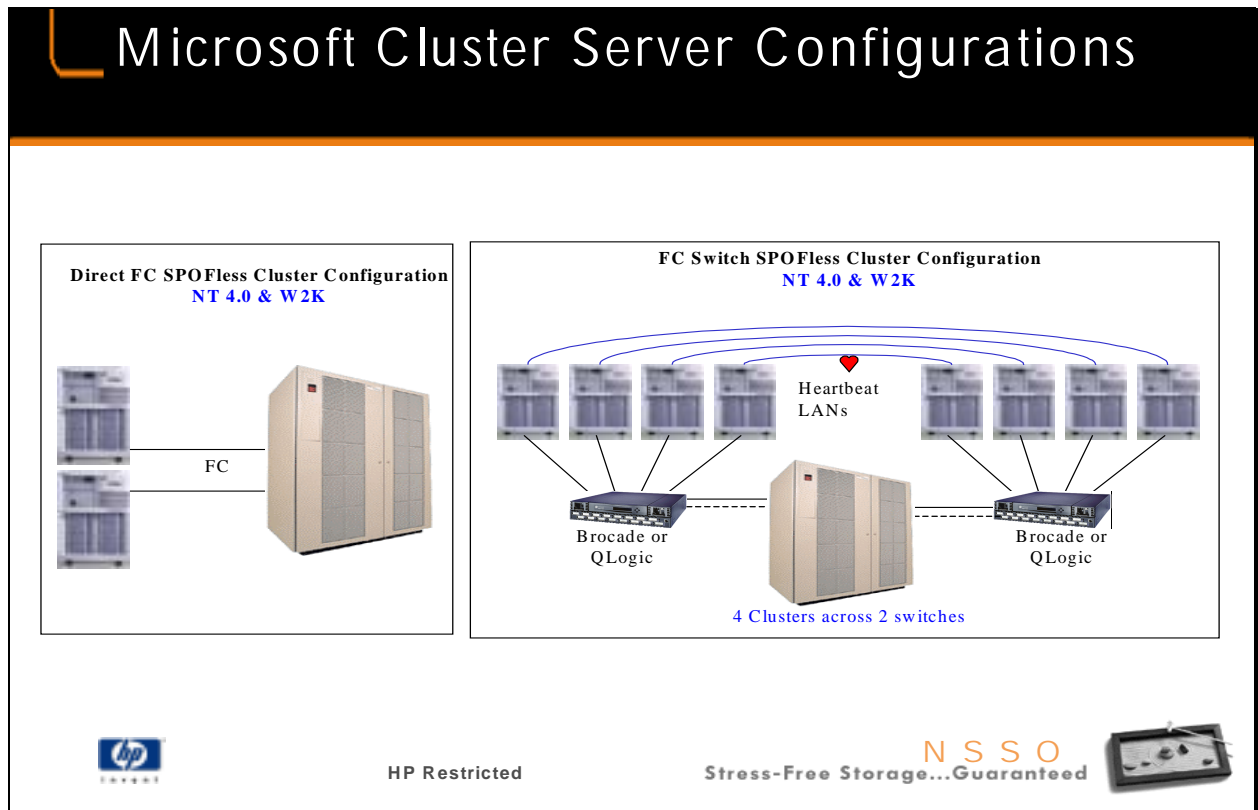
**(Please see Configuration Guide in Appendix)**



HP Restricted

**N S S O**  
Stress-Free Storage...Guaranteed





The XP256 and XP512 have been certified as Cluster RAID devices for MSCS for both NT 4.0 and W2K Advance Server MSCS configurations. In order for a cluster configuration to be fully supported by Microsoft, the entire cluster (specific servers with specific firmware, HBAs, storage devices etc.) must also be tested and certified. Because of the wide range of servers, HPSO can certify only select configurations. Refer to <http://www.microsoft.com/hcl/default.asp> for Microsoft certified configurations.

Up to 8 shared LUNs per cluster are supported. Maximum LUN size for the XP256 is 262.8GB using Open 9's and LUSE (7.3GB x 36). Maximum LUN size for the XP512 is 525.6GB using Open E's and LUSE (14.6GB x 36). Microsoft does not support stripe or volume sets for shared storage in MSCS.

Refer to the HPSO matrix, sections [HBA Support for the XP512 in MSCS Configurations](#) and [HBA Support for the XP256 in MSCS Configurations](#) for the updated list of supported servers:  
[http://hpsu.rose.hp.com/config\\_matrix/config\\_matrix\\_table.htm](http://hpsu.rose.hp.com/config_matrix/config_matrix_table.htm)

Refer to the NetServer Tech Forum site for more details on MSCS, including NSD storage, futures, and config guides:


SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Microsoft Windows Configurations

<http://gss.cup.hp.com/netserver/products/ha/mscs/mscs.htm>

There are 2 types of Fibre Channel MSCS configurations: SPOF (Single Point of Failure) and SPOFless (No SPOF).

Always recommend SPOFless configurations where possible.

## W2K clusters, Auto Path XP can provide HBA redundancy

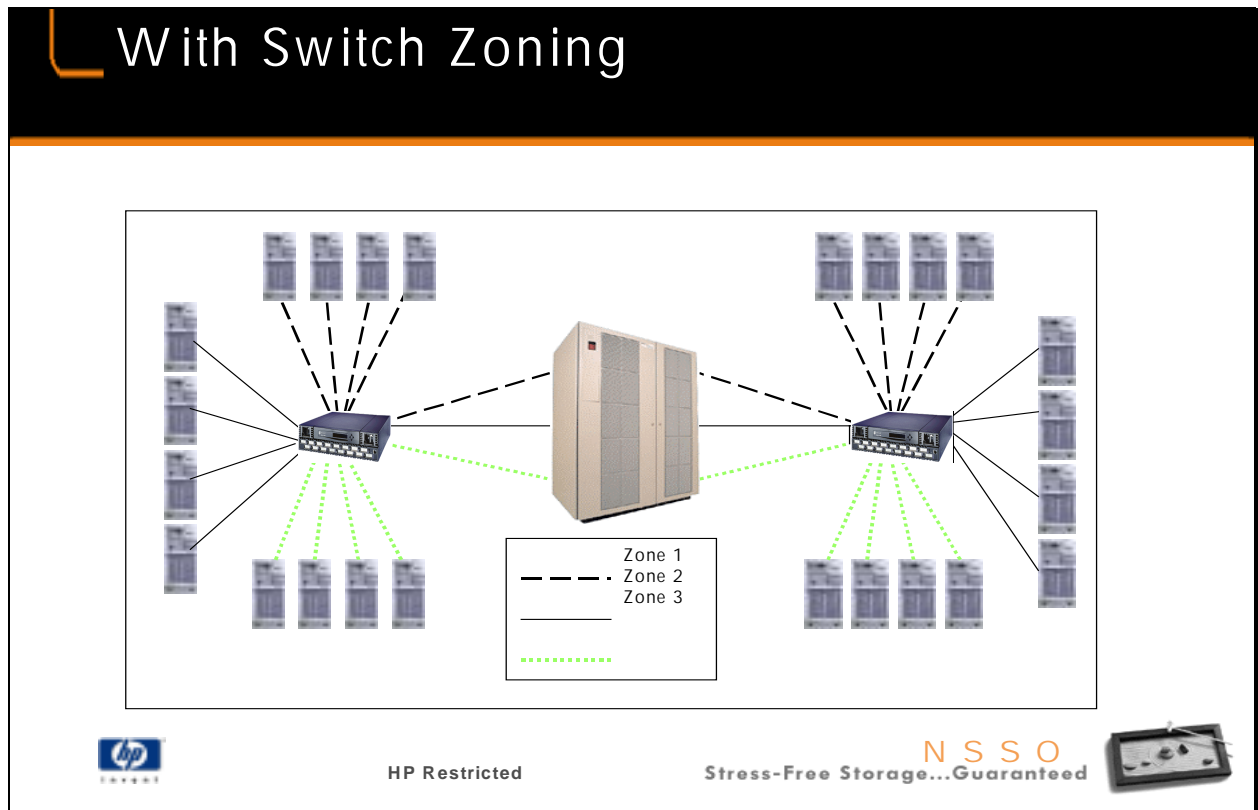


HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



In this example, if a switch or XP port fails, all 4 servers will access the XP through the remaining switch and XP port. If additional servers or cluster nodes are added to the switches, additional connections to the XP are necessary to ensure that no more than 4 servers access an XP port.



With switch zoning, a total of 12 cluster nodes could connect to a single 16-port switch with up to 3 zones and 4 nodes per zone. Alternatively 4 zones and 3 nodes per are supported as well.

Zone 1  
Zone 2  
Zone 3



The Microsoft "Designed for Windows" logo identifies hardware and software products that have been designed and tested to meet Microsoft standards for compatibility with Windows® 98 and Windows NT® Workstation 4.0. Below are key benefits you get when you look for the logo when purchasing products for your computer. Tested

Products displaying the logo have passed rigorous testing that assures they are easy to use, stable, and take advantage of the new features of Windows 98 and Windows NT Workstation 4.0. Software is tested by VeriTest, an independent testing lab. Hardware--both PCs and peripheral devices--are tested by Windows Hardware Quality Labs (WHQL).

Works well with other logo'd software and hardware-when you mix and match products with the Designed for Windows logo, you can be assured they will work well with other products that bear the logo.

Installs properly and uninstalls completely

Logo'd products have been checked to be sure they don't overwrite key system components when they install--a phenomenon known as "DLL Hell" that is the main reason new installations sometimes trash your system.

WHAT DO LOGO'D PRODUCTS OFFER BUSINESSES?

The Designed for Windows logo can play a key role in reducing total cost of ownership for organizations. The Windows logo criteria is one of Microsoft's key

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Microsoft Windows Configurations

vehicles in communicating to software developers how to incorporate Zero Administration initiative for Windows (ZAW) features into their applications in preparation for the next releases of Windows. Businesses that use products meeting the Designed for Windows logo criteria stand to gain the following benefits:

Lower support costs: Logo'd applications follow standard Windows look-and-feel guidelines. That helps users get up to speed without phoning your helpdesk.  
Help manage "DLL Hell": Logo'd applications do not overwrite applications or uninstall key components.

Support mixed Windows environments: Logo'd products work on both Windows 9x and Windows NT to help you manage the mix of Windows desktops in your organization

Proper use of operating system: Logo'd products make proper use of the Windows registry and other key operating system files.

Compliance with the Americans with Disabilities Act and other equal rights legislation: Logo'd applications meet usability standards for a wide range of people.

L

# XP512

The screenshot shows the Microsoft Windows Hardware Compatibility List website. The search bar contains 'XP512'. The results table lists compatible hardware configurations. The taskbar at the bottom shows the Start button, several open applications, and the system clock.

Type	Company Name	Product Name	Windows 98	Windows Me	Windows NT 4.0	Windows 2000
Cluster	Heckel- Dahl and Company	Heckel- Dahl and Company 3500 / Compact StorageWorks 44- 8x16-MHz Fibre Channel Host Adapter [cluster] [Cluster]				
Cluster	Heckel- Dahl and Company	Heckel- Dahl and Company 3500 / Compact StorageWorks 44- 8x16-MHz Fibre Channel Host Adapter [cluster] [Cluster]				
	Heckel- Dahl and Company	Heckel- Dahl and Company 3500 / Compact StorageWorks 44- 8x16-MHz Fibre Channel Host Adapter [cluster] [Cluster]				

HP Restricted      Stress-Free Storage...Guaranteed

Test and certification is in progress for Data Center. Right now the only supported Data Center cluster is using the XP256 and point to point connections, single HBA per server. AutoPath XP is not yet tested or supported with Data Center. The XP512 is planned to be supported very shortly as well, though HP does not have the exact date.

Data Center is very different. It can only be purchased and supported from OEM vendors, HP included. The only server HP currently offers which supports DC is the LXR8500dc. HP announced a 32 way system which will ship early 2001. Below is a website with more W2K DC info. Look for more updates on this site in the future.

<http://gss.cup.hp.com/Netserver/Products/W2K/Default.htm>

<http://www.microsoft.com/hcl/default.asp>







## MS Certification Nomenclature

**Single Initiator Device**  
A single server may be directly connected to the storage via an interconnect type. Certification results in device inclusion on MS's "Hardware Compatibility List" (HCL). XP Family of Arrays submittals made for SCSI and FC.

**Cluster Device**  
This defines that a MSCS cluster may be connected to the storage via an interconnect type. Certification results in device inclusion on MS's HCL. XP Family of Arrays submittals made for SCSI and FC.

**Multi-Cluster Device**  
A new classification for the HCL, this defines that multiple MSCS clusters can be connected to the storage. Expected availability early 2000. ESBUS is working with MS to certify the XP Family of Arrays. Only FC will be submitted.

**Cluster System**  
A specific cluster with named components (server, HBA, storage). All components must be on the HCL. Certification results in inclusion on MS's "Microsoft Cluster Server" (MSCS) list.



### Windows Certification, XP 256 Status

Microsoft sets the bar, passes out the tests, and reviews the results. Vendors (such as HP) participate as the price of admission for sales. Customers use the certification lists to determine which solutions to deploy.

MS certifications for NT 4.0 and Windows 2000 are found at

<http://www.microsoft.com/hcl/default.asp>.

Choose your hardware option and drill down. Choose 'clusters' for clusters.

A good overview of MSCS can be found on the MS web site at


<http://www.microsoft.com/ntserver/ntserverenterprise/exec/overview/clustering.asp?RLD=36>.

HP's solutions are described at:

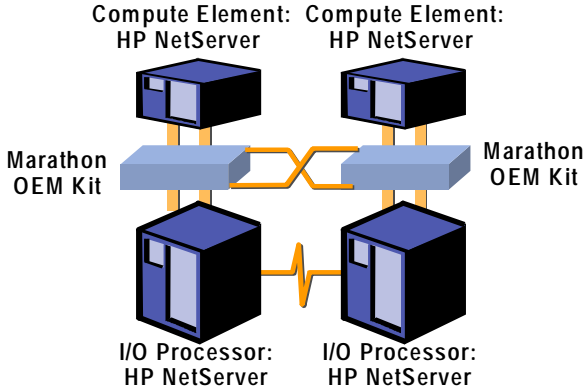
[http://netserver.hp.com/netserver/products/highlights\\_ms-clusters.htm](http://netserver.hp.com/netserver/products/highlights_ms-clusters.htm)

# HP NetServer Assured Availability

## Solution Components



**The OEM Kit**



**The Array**

HP Restricted

N S S O  
Stress-Free Storage...Guaranteed

The XP256 is support on the AA4000 with SCSI connections. Please review the configurations page for information on the AA6200 and XP256 (with FC support).

### HP NetServer Assured Availability

#### Solution Components

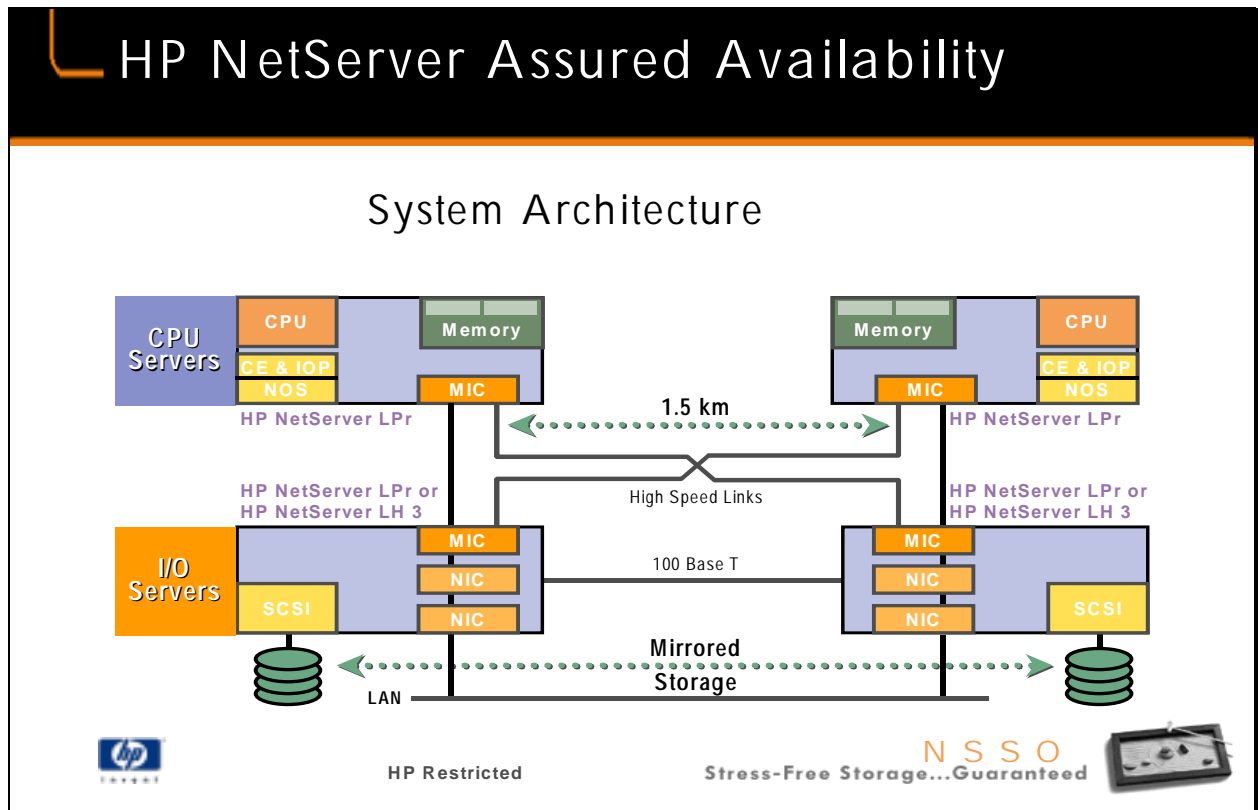
Marathon provides the enabling technology for the HP NetServer Assured Availability Solution. HP has been testing, evaluating and planning for a range of Marathon Technologies products for the last nine months. Through its current multi-product, multi-year original equipment manufacturer (OEM) agreement with Marathon Technologies, HP is purchasing the Marathon Endurance 4000 kit, consisting of four (4) PCI cards, cables, two (2) switches and software. This kit is combined with award winning HP NetServers to create a single logical server array from four different, low cost servers. The array works as one virtual server with no single point of failure (SPOF). The entire array is stored in a rack to optimize space utilization. HP NetServer Assured Availability Requires Industry standard components to optimize investment protection and lower the total cost of ownership: General Purpose HP NetServers bundled with Marathon Endurance 4000 kit Industry-Standard Microsoft Windows NT Server Unmodified, "Off the shelf" Applications such as Microsoft Exchange

Overview of the Marathon Product- The Marathon E4000 is a "kit" containing Hardware and Software. This kit, together with four (4) standard IA32 Servers,

## SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training Microsoft Windows Configurations

creates a solution that acts like a single, Fault Tolerant NT server. HP is offering 3 general configurations of NetServers, based on this solution, to cover a wide range of customer needs. These solutions will be delivered as pre-configured systems to the customer. Assured Availability Description-Introduction-The Assured Availability product is OEM'd from a company called Marathon Technologies. This is their Endurance 4000 product with no changes. NSD has determined that this product can provide many customers with MUCH greater availability of their applications than by any other means, including clustering. Since this product is complex to install, HP has decided to deliver this as a "Solution". This means that the customer selects a configuration and has a limited number of options to select from. At this stage, the options are the amount of memory in the IOPs and CEs, the amount of CE data storage, and the number of network connections. These configurations are built and installed only by organizations that have been trained and are approved by HP.

Overview-As was described above briefly the Assured Availability kit is combined with 4 servers to implement a high availability solution. The picture above shows a logical view of how these pieces are put together. It is useful for understanding how it works at a higher level. There are two servers shown on the left and two on the right. This left/right implementation allows a server on one side to take over for its "mate" in the case of a failure. These left/right pairs are called "Tuples", with Tuple 1 on the left.



### HP NetServer Assured Availability

#### System Architecture Takes the Logical Server Approach

An array of four servers acts as one logical server (i.e. it has a single IP address on a LAN). The application runs simultaneously on two symmetrical halves (or tuples), which, combined together, do not have a single point of failure. In other words even if one or both servers in a tuple fail the system will keep operating with no impact to users. Even during failure there is no performance degradation.

One of the biggest causes of failure ("the blue screen of death") is the complex interaction of computing and I/O processing within a system. The timing issues between the two functions can often bring down a system. In order to address this critical cause of failure, HP's Assured Availability solution distinctly separates computing from the input/output (I/O) processing. So within each of the symmetrical halves (or tuples) of the array one server performs the computing function while the other server performs the I/O processing.

I/O processors run asynchronously to further minimize the chances of failure which would occur if both I/O elements failed at the same time. As a result of asynchronous operation, if one I/O processor fails the second one will just keep on

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Microsoft Windows Configurations

processing the I/O operations. When the failed I/O processor has been restored and brought back on line the data in mirrored storage will be synchronized.

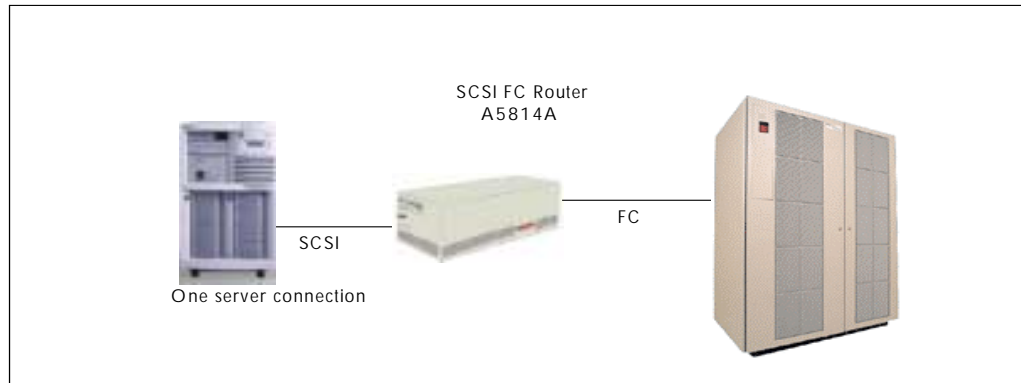
The compute elements run synchronously. If one fails the users don't experience any interruptions as the other compute element (server) continues the computing functions uninterrupted.

The array uses pairs of network cards (NICs) which are treated as one card when interfacing with the operating system; each card is utilized for each network in each of the I/O Processors. Consequently, if a single point of failure occurs on the network, users never experience an interruption of services. The Windows NT Registry of each I/O Processor has the same MAC address on file for both cards, comprising one Ethernet rail. If the alternate Ethernet card takes over sending data, clients on the network see the same address as the failed Ethernet card. This insures that only one address for the array is ever broadcast over the network at any time eliminating any potential network or database conflict problems.

MIC (Marathon Interface cards) provide communication between servers, as well as fault detection and correction.

The only way this array can fail is if both of compute elements or both of I/O elements failed but the system is built in away that insures that such a simultaneous failure is virtually impossible.

## Router Connection for XP512 and XP48



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



The A5814A SCSI FC Router can be used to connect SCSI-based servers to the Fibre based XP512 and XP48. Hubs and switches are not supported with the router.

## Hub Configurations - Standalone Servers

The diagram illustrates a 'Single FC Hub Connection' setup. On the left, three server racks are shown with lines indicating 'Up to 9 server connections' to a central 'S10 (A3724A) or L10 (A4839A) FC hub'. The hub is labeled 'FC'. A single line labeled 'FC' connects the hub to an 'XP Disk Array' on the right. Text next to the array states 'Up to 9 connections from one hub to one or more XP Disk Arrays'. A red note at the bottom of the diagram reads: 'Note: The D8602A can only access two XP ports'. The HP logo is in the bottom left, 'HP Restricted' is in the bottom center, and 'Stress-Free Storage...Guaranteed' with 'N S S O' is in the bottom right.

The 6 port hub from NSD (D6976A) is not supported with the XP Disk Arrays. The short-wave ports on the S10 and L10 hub support a distance of up to 500 meters. The long-wave port on the L10 hub supports up to 10 km and can only be used to connect to another long-wave port on a second L10 hub or switch. When multiple servers are connected to hubs, some type of LUN security must be used to prevent the disks from being accessed by all servers causing data corruption. A single XP256 port can address 120 LUNs; the XP512 supports 256 LUNs. Only recommend multiple connections to a single XP Disk Array if more than maximum addressable LUNs per port are needed. Since hubs are a single loop with a shared 100 MB/s bandwidth, multiple XP Array connections from a hub will not offer any performance advantage.

## Cascaded FC Hub Connection

The diagram illustrates a cascaded FC hub connection. Two HP FC hubs, labeled 'S10 (A3724A) or L10 (A4839A) FC hub', are connected to each other via a single FC connection. Each hub is connected to multiple servers (labeled 'FC') and can provide 'Up to 17 connections across both hubs to one or more XP Arrays'. A note specifies 'Only one FC connection between hubs'. Another possible connection is shown between the hubs and the XP Arrays. A red note states: 'Note: The D8602A can only access two XP ports'.

HP Restricted

N S S O  
Stress-Free Storage...Guaranteed

Two hubs can be cascaded to offer more port connections. Cascading hubs for greater distance is not recommended because of severe performance degradation over long distances. If a distance greater than 500 meters is required between servers and the XP Disk Array, strongly recommend switches instead of hubs.



## Switch Configurations - Standalone Servers

Single FC Switch Configuration

Brocade or Ancor

FC

FC

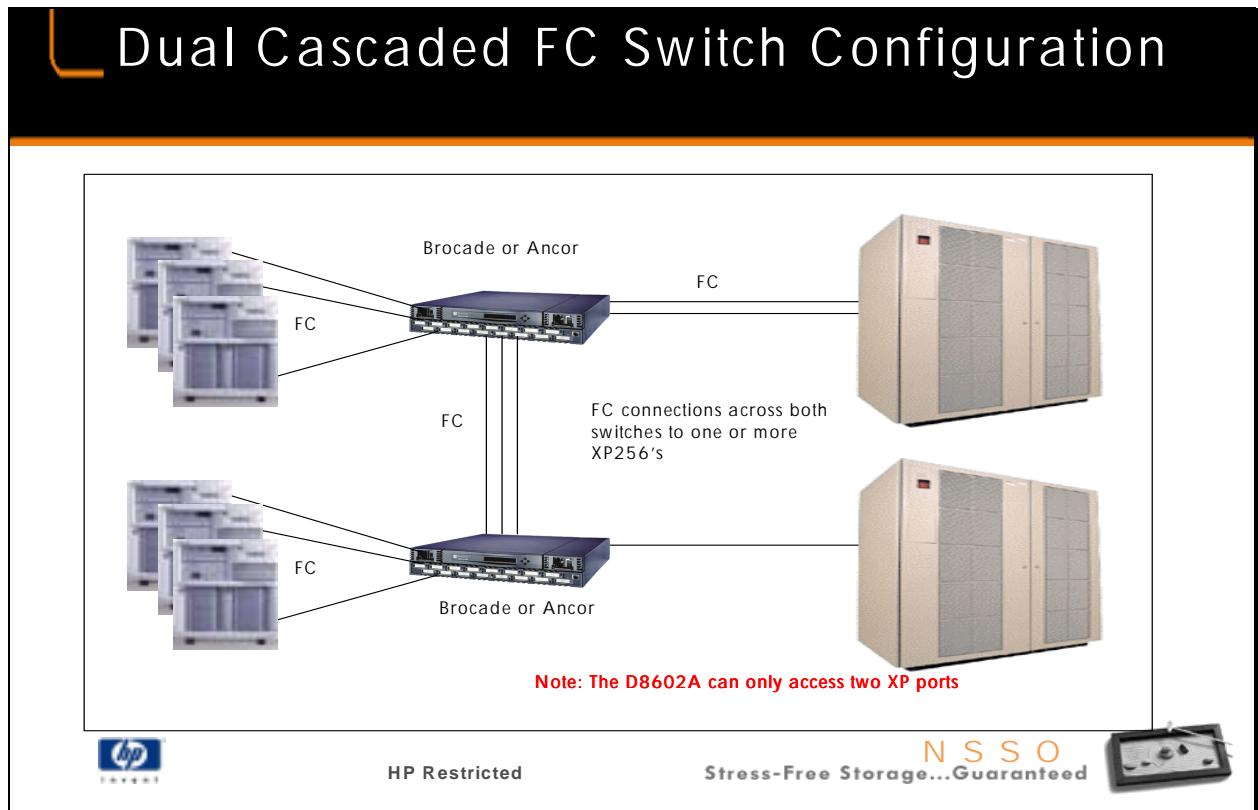
Note: The D8602A can only access two XP ports

HP Restricted

N S S O  
Stress-Free Storage...Guaranteed

The XP family of disk arrays supports the 8-port Ancor (A5665A), 16-port Ancor (A5666A), 16-port Brocade 2800 (A5624A), and 8-port Brocade 2400 (A5625A for SF21 & A5667A for SF88). Short-wave (A5225A) and long-wave (A5226A) GBICs must be ordered for each port. The short-wave port can support a distance of 500 meters. The long-wave port can support up to 10 km. The long-wave port can only be used to connect to another long-wave port on a second switch or L10 hub. All short-wave ports can be used for server or XP port connections.

When multiple servers are connected to switches, some type of LUN security must be used to prevent the disks from being accessed by all servers causing data corruption. Multiple connections to an XP256 can improve performance. A single connection from a switch to the XP256 does not offer any significant performance advantage over a hub. Switches offer advantages over hubs by overcoming any issues or limitations associated with LIPs.



The XP family supports up to 2 cascaded switches. The 16-port switches support up to 8 interconnects; the 8-port support up to 4 interconnects between switches. The Brocade and Ancor switches cannot be cascaded with each other.

## Mixed Switch-Hub Connection

The diagram illustrates a Mixed Switch-Hub Connection. On the left, multiple server racks are shown, each labeled 'FC'. These servers are connected to two FC Hubs, labeled 'S10 (A3724A) or L10 (A4839A) FC Hub'. Below the servers, it states 'Up to 9 server connections per hub'. These two FC Hubs are connected to a central switch, labeled 'Brocade or Ancor'. The switch is then connected to a large storage array on the right, labeled 'FC'. Below the storage array, a red note states: 'Note: The D8602A can only access two XP ports'. Another red note below the switch states: 'Note: Only one server using the D8602A can connect to a hub which is connected to a switch.' A blue note at the bottom of the diagram states: 'Maximum number of server connections depends on LUN security'.

HP Restricted

Stress-Free Storage...Guaranteed

Hubs can attach to single or cascaded switches to add additional FC connections at a low cost. Multiple hubs can be connected to the switch but only a single hub can be connected to a switch port, cascading hubs is not supported. A hub can only connect to a single switch, it cannot be connected between 2 switches.

# Meshed FC Switch Configuration

(support under investigation)

Brocade or Ancor

Note: The D8602A can only access two XP ports

Maximum number of server connections depends on LUN security

HP Restricted

N S S O Stress-Free Storage...Guaranteed

More than 2 cascaded switches is called a mesh. Up to 39 switches can be connected in a mesh with up to 8 connections between individual switches. Up to 7 hops between switches is allowed. There is no SAN imposed limit to the number of servers that can connect to the SAN. Hubs may also be connected to switches in the mesh. The XP Disk Arrays do not currently support meshed switches, date TBD.

# Switch Zoning

The diagram consists of two side-by-side illustrations. The left illustration, titled 'Zoning without Secure Manager', shows three server icons on the left connected by dashed lines to a central switch, which is then connected by a solid line to a storage rack on the right. A red text box at the bottom states 'No advantage over direct connections'. The right illustration, titled 'Zoning with Secure Manager', shows the same setup but with an additional server icon below the first three, and a storage icon below the storage rack. Dashed lines connect the servers to the switch, and the switch to the storage rack, indicating a more complex zoning configuration.

HP Restricted

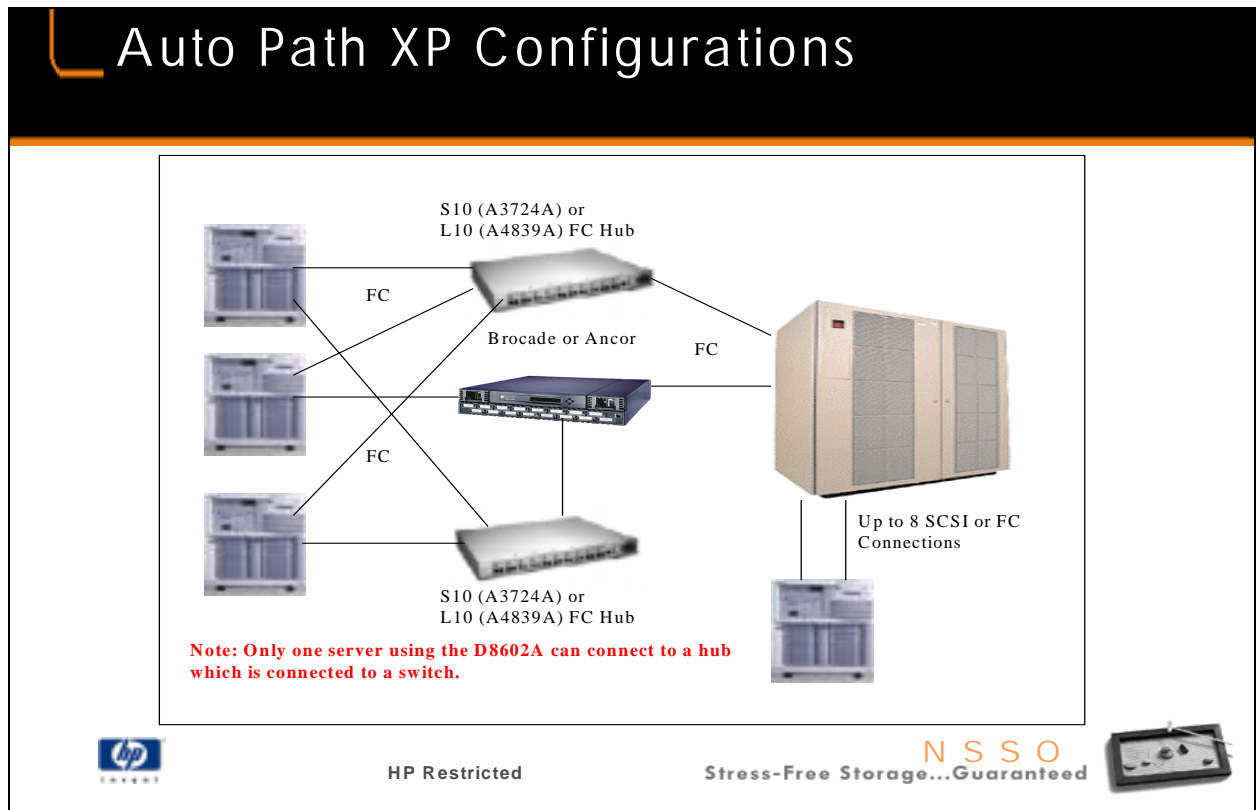
N S S O  
Stress-Free Storage...Guaranteed

Switch Zoning can also be used to provide LUN security by creating virtual SANs within the entire fabric. Zoning is only supported on the current Brocade (firmware 2.1.9) and Ancor switches, it is not supported on the F16 switch. Zoning is OS independent and currently only supports F and FL ports on the switch. Zoning by itself is not sufficient to provide LUN security, it can only provide security for one server to one XP port. For example, if there are 3 servers on a switch with 3 XP connections, a zone could be made for each server and XP connection. This does not provide any advantage over point to point connections, the switch is basically just added hardware. Zoning is usefull when combined with Secure Manager XP and can be used to create multiple zones of servers and XP ports. Zoning can also be used to isolate mulitple OS's on switch or switch fabric.

SAN Manager LM (SMLM) is the third option for LUN security. SMLM is host based software that currently supports Windows NT only; W2K support is planned for Q1 2001. SMLM must be installed on every server in the SAN. It currently supports only the Qlogic 2100 or 2200 HBA. SMLM will support the next version of the D8602A HBA and driver, target tade is late December 2000. One final restriction is that Auto Path XP cannot be used with SMLM and there is no timeframe for when the two products can be used together.

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Microsoft Windows Configurations

SMLM has an advantage of providing security for hundreds of hosts and LUNs. It allows the SAN administrator to add storage to the SAN or easily move storage from one node to another on-the-fly without rebooting the server.



Auto Path XP is host-based software that provides IO path load balancing and fail-over. Up to eight FC or SCSI (XP256 and NT only) I/O paths from one server to a LUN are supported. For the NT version (B7934A - OEMed from Hitachi), the software must be purchased for each server that requires multiple IO paths. For the W2K version (B9500A - designed and developed by HP and covered in a module 3 of this training), there are multiple bundled LTU packages. Refer to the HP Auto Path XP for Windows 2000 presentation in ESP for more details (keyword: AUTOPATH). Auto Path XP for NT works with any of the supported HBAs. However, it supports only stand-alone servers; MSCS will be supported once the HP developed version is ported to NT. The target date for NT 4.0 support is December 2000. The HP W2K version currently only supports the D8602A; other HBAs TBD. Although this version does support MSCS, XPSO currently only supports a single HBA per cluster node; support for multiple HBAs per cluster node TBD. Only W2K Server and Advanced Server are supported. W2K Data Center support requires certification; support TBD. The above diagram is just one example of a valid configuration with Auto Path XP. Mixing SCSI with FC connections in the XP256 is supported. Fibre Channel connections are supported with any combination of the hub, switch, and point-to-point connections described in the previous sections. HPSO has not tested and will not support Auto Path XP with the Bridge FC 4/2.

## Cluster Extension XP for Microsoft

The disaster recovery solution over metropolitan-wide distances for Solaris, AIX, and now for Windows 2000. Cluster Extension XP extends a single Windows 2000 Microsoft cluster over metropolitan-wide distances, while ensuring that the automatic cluster failover capability works seamless with Continuous Access XP remote mirroring.



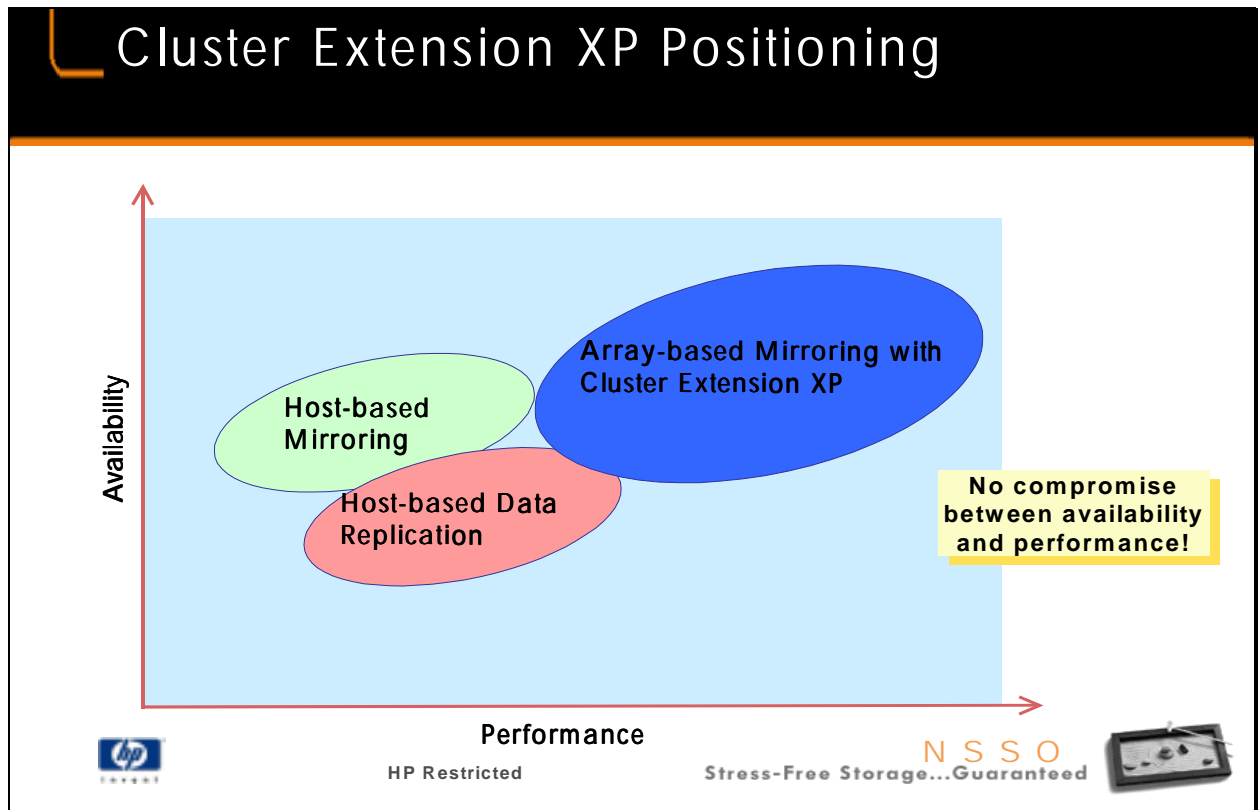
HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



Cluster Extension XP offers protection against system downtime – from fault, failure, or disaster – by extending a single cluster over metropolitan-wide distances. Cluster Extension XP ensures that your cluster software - VERITAS Cluster Server for Sun Solaris, HACMP for IBM AIX, and Microsoft Cluster Service for Windows 2000 - works seamlessly with the XP mass storage resources to ensure data availability, data integrity, and ultimately business continuity.





When we think of business interruption, each of us probably have a few things in mind - earthquake, hurricane, flood, fire. But have you considered all the possibilities of system downtime? Sabotage? Rodents? Sprinkler discharge? Business interruption can be caused by more than just disasters that we think of in catastrophic terms, or that we hear about in the news, such as a building fire, natural flood, earthquake, or hurricane. In fact, natural disasters account for only a small percentage of all data loss.

An interruption can be caused by something as routine as system upgrades or routine system maintenance. All in all, there are numerous causes of downtime, whether planned, such as for maintenance and upgrades, or for unplanned, such as catastrophes.

So what's so terrible about this? It's terrible because of all the consequences of downtime.

Think through the consequences downtime. Lost revenue, customers, reputation, image, lost business

So how bad can it be? According to Forrester Research, the average data loss per only one hour of downtime is \$6.45 million for the average retail brokerage business and \$2.6 million for a credit card authorization corporation. So the financial consequences can be severe. (Cost factors include lost revenue, lost

productivity, and replacing lost data. Also included are damages such as loss of good will and image, which may result in loss of clients and future revenues)

A carefully planned local disaster recovery solution such as local clustering will help to minimize downtime. (A local cluster offers failover to secondary node within the same local environment, such as within the same datacenter) However, a local disaster recovery solution will not offer protection for all disaster scenarios, as shown in the slide.

But with extended Cluster solutions, such as Cluster Extension XP, disaster recovery can be extended beyond just a local datacenter, thus offering customers the protection from even the worst disasters imaginable

## cluster extension xp -product structure overview requirements

### •Requirements:

- synch copy: Continuous Access XP
- asynch copy: Continuous Access XP + Continuous Access XP Extension
- Raid Manager XP (bundled w/ CA)
- Cluster Software:
  - VERITAS Cluster Server for Sun Solaris
  - HACMP for IBM AIX
  - Microsoft Cluster Service for Windows 2000 Advanced Server or Datacenter Server
- Secure Manager XP (for Microsoft Cluster Service for Windows 2000 only)
- Remote Control XP (included with Command View) on management PC
- XP Disk Arrays
- Cluster Extension XP for Windows2000 Advanced Server/  
Datacenter Server: Live CPL on July'01**



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Cluster View Products – B9533A

### ►Productization:

- B9531A: Cluster Extension XP for VCS on Sun Solaris
- B9532A: Cluster Extension XP for HACMP on IBM AIX
- B9533A: Cluster Extension XP for Microsoft Cluster Service on Windows 2000
- B9320A: Continuous Access XP
- B9325A: Continuous Access XP Extension

### •CLX Licensing: LTU per cluster



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## cluster extension xp for microsoft cluster service - what is required?

### Prerequisites:

- Microsoft cluster-certified Intel-based servers with FC HBA running Windows 2000 (see Microsoft HCL and HP connectivity streams document)
- Windows 2000 Advanced Server or Datacenter Server (Service Pack 1 or later)
- Microsoft Cluster Service



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



### Solution Requirements:

B9533A: Cluster Extension XP for Microsoft Cluster Service on Windows 2000 (Live CPL July'01)

B9320A: Continuous Access XP (for synchronous copy)

B9325A: Continuous Access XP Extension (for asynchronous copy mode; also requires B9320A)

Certified DWDM, converters to extend Continuous Access XP to wider distances (CNT, Inrange, Nortel Networks)\*

Raid Manager XP (included w/ CA)

B9351A: Secure Manager XP (for Microsoft Cluster Service for Windows 2000 only)

Remote Control XP (included with Command View B9357A) on management PC

Two XP Disk Arrays

### Implementation & Support Requirements:

Continuous Access XP Implementation Service

Software Enablement Service (Option OSY) for Cluster Extension XP

SW Foundation Services (Options OS6, 3Y6)

## cluster extension xp for microsoft cluster server

- full integration in mscs and cluster admin gui
- online addition possible
- full online help
- metropolitan solution (tested up to 20 km )
- escon link auto-recover feature build-in and configurable from the gui
- full cluster manipulation to ensure data consistency at the xp array level



Windows 2000 Advanced Server (Service Pack 1 or later)

## microsoft cluster server

- heartbeat support over IP  
(private network, that is a single subnet; no distance limitation)
- up to 4 systems per cluster (windows nt (2), windows 2000 advanced server (2), windows 2000 datacenter server)
- all clusters can be configured and administered from a single gui
- mscs tries to online resource at any cost
- mscs does not support dynamic disks, raid controllers and many other windows2000 features




HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## clx for mscs - configuration example (first release)

- first release requires shared quorum disk connected to all cluster systems
- fiber channel infrastructure must be in place
- secure LUN manager xp must be configured
- manual intervention required if primary data center fails
- manual cluster service startup
- clxsplitbrain program to protect against false app startup

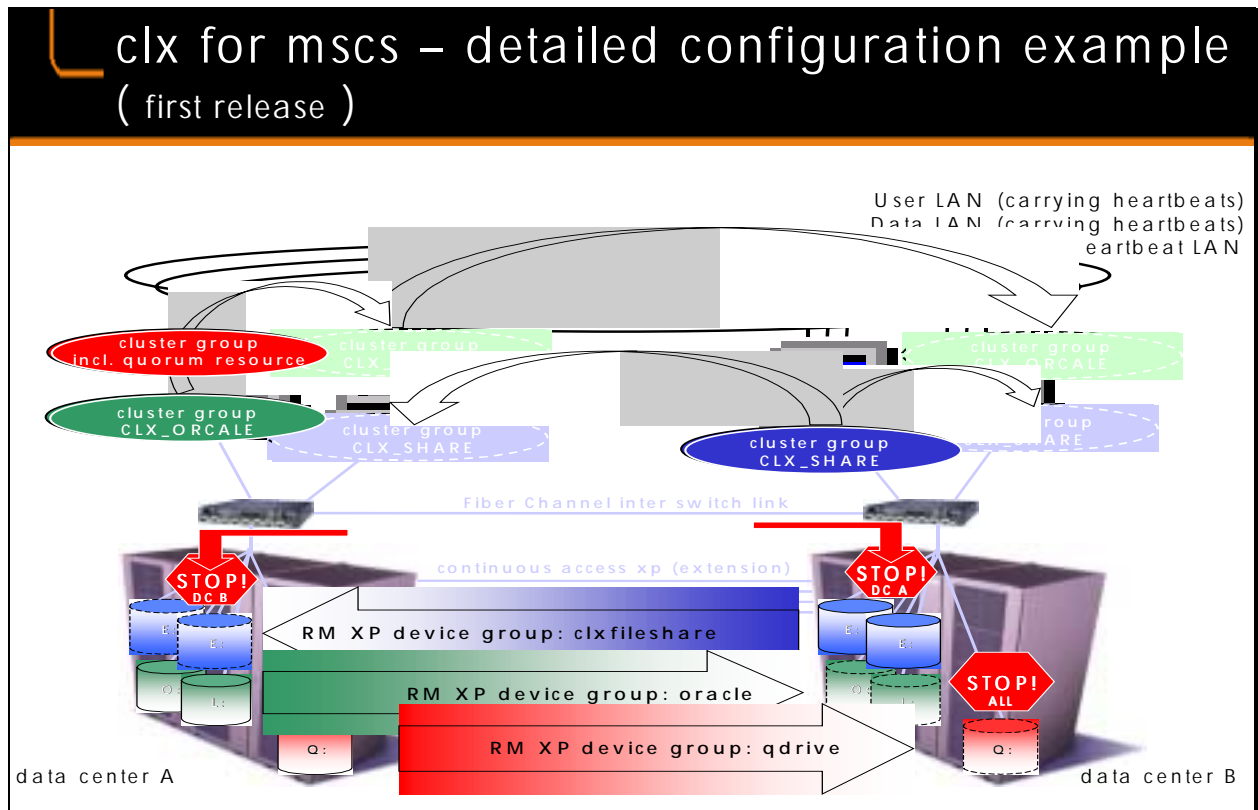


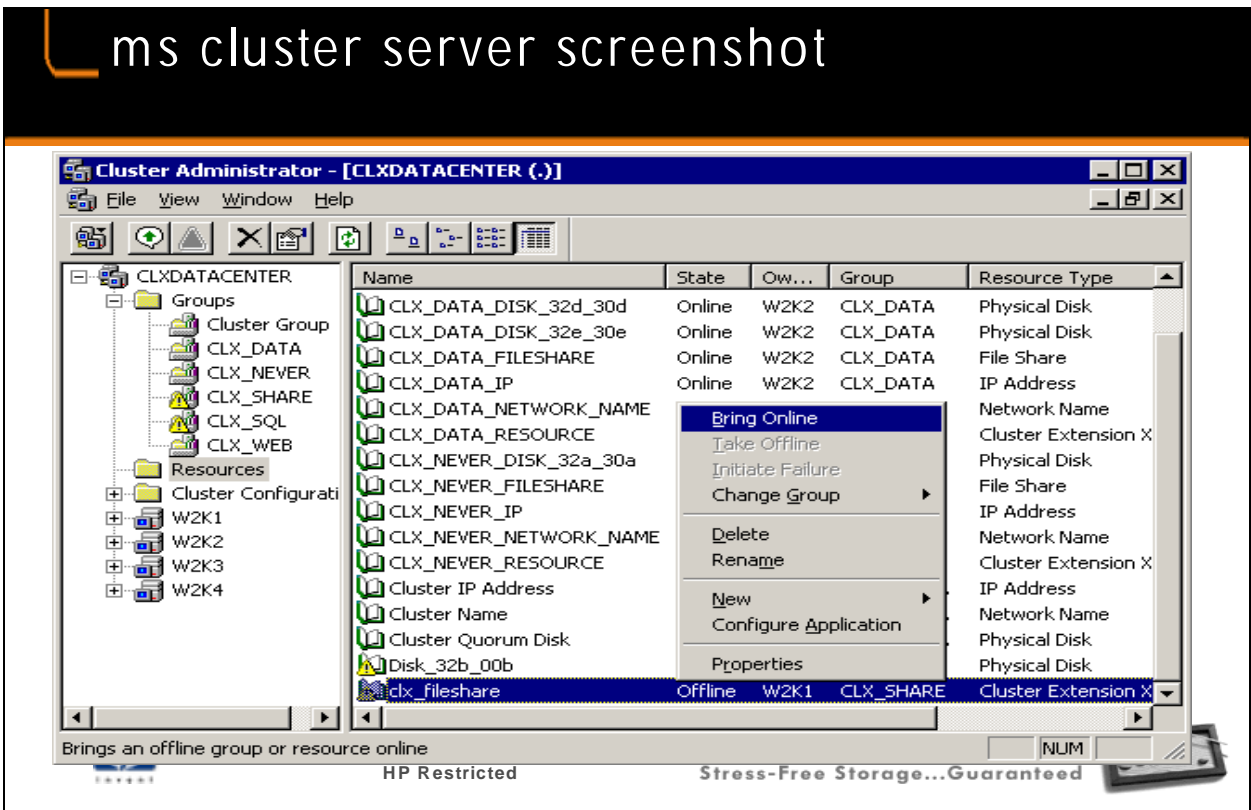
HP Restricted

N S S O  
Stress-Free Storage...Guaranteed

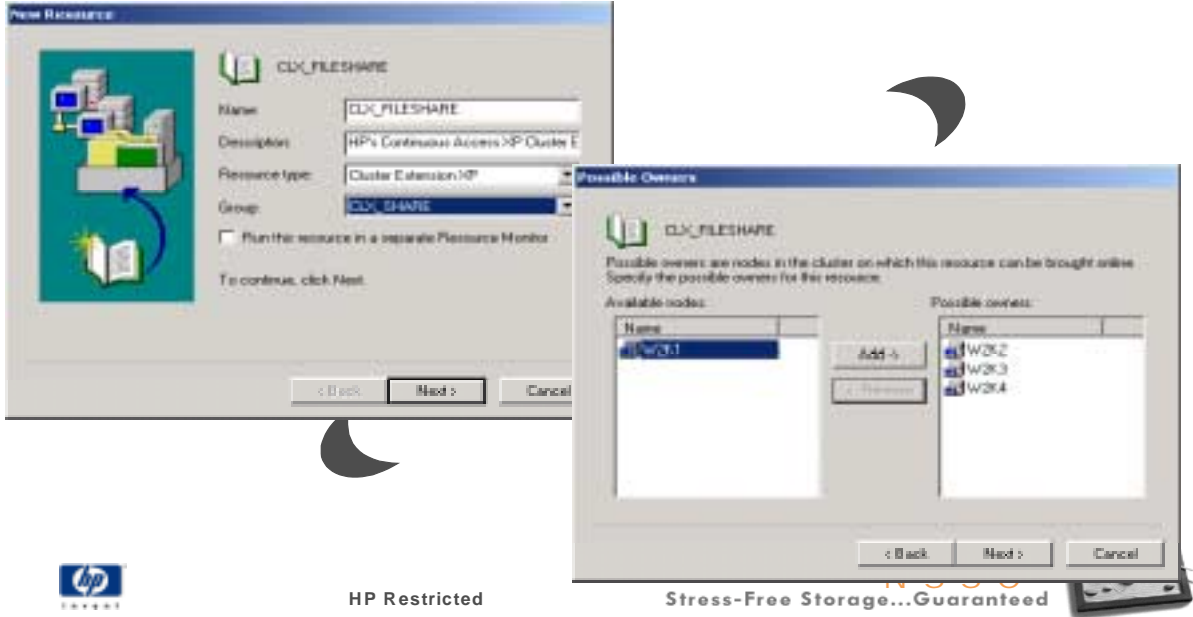








# add cluster extension xp resource specify possible owner



The first screenshot shows the 'New Resource' dialog box for 'CLX\_FILESHARE'. The fields are: Name: CLX\_FILESHARE, Description: HP's Continuous Access XP Cluster E, Resource type: Cluster Extension XP, Group: CLX\_SHARE. There is a checkbox for 'Run this resource in a separate Resource Monitor' and a 'Next >' button.


The second screenshot shows the 'Possible Owners' dialog box. It lists 'Available nodes' and 'Possible owners'. The 'Available nodes' list contains 'W2K1'. The 'Possible owners' list contains 'W2K2', 'W2K3', and 'W2K4'. There are 'Add >' and 'Remove <' buttons between the lists, and 'Back <', 'Next >', and 'Cancel' buttons at the bottom.

HP Restricted

Stress-Free Storage...Guaranteed

## clx for msccs - configuration example (second release Q4'01)

- second release provides quorum disk driver for **split quorum disk control**
- fiber channel infrastructure **not** necessary between data centers
- secure LUN manager xp **not** necessary
- manual intervention required if primary data center fails (to prevent split-brain syndrome)




## Module Wrap-up



HP Restricted


N S S O  
Stress-Free Storage...Guaranteed






# Module 15

## Support and Services Value Chain Guidelines



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



HP has had some significant competition, namely EMC, to live up to in more ways than one. They've got a strong product and they've got a strong support organization. They also have a strong pre-sales organization. And it has really tested a lot of the organizations within H-P, to some degree, like we've never been tested before. And it's also really forced HP's diverse organizations to work much closer together with the support organizations, and in many cases we haven't had to do that so much in the past couple years. We have realized that we could do a lot better job at delivering the whole value chain of support services surrounding primarily the XP Family Array.

## Module Agenda



- Understand the Value Chain of a XP Family Array Sale, including implementation and support



- Articulate Roles & Responsibilities between teams.



- Identify the different HP organizations and their responsibilities

- Introduce the SFIP (Stress Free Implementation Plan)

- HP Channel Partner's Role & Responsibilities

- Wrap-UP  
HP Restricted



N S S O  
Stress-Free Storage...Guaranteed



# Sales, Support and Services Value Chain

	Sales Process	Initial Configuration	Proposal Process	Support and Services Proposal	ASO Notified of Sale	SFTP (Stress Free Installation Plan)	Customer Meeting	Installation Project Mgmt.	Site Prep	System Integration (Optional)	LUN Map Design	Validate Raid Level	Operating System Readiness / Patches	Installation	Hardware / Firmware Readiness	Software Enablement	Data Migration	Implementation and Consulting Services	Account Management	Ongoing support
ESBU					X															
GSS Pre-Sales						X	X													
ASO												X								
Factory										X										
HSSO									X					X	X					
HPC																	X	X		
TAC																				
RC																				

X - Primary Responsibility

Blue and Green column headers indicate bundled services

HP Restricted

NSSO  
Stress-Free Storage...Guaranteed

This slide which is basically a matrix of the sales and services value chain for an XP Family Array sale. It's got a number of significant process steps across the top and then the different HP organizations across the left hand side. It's obvious when you see how many organizations are involved with this, and how many process steps there are, that it's easy to understand now, in hindsight, why we've had some troublesome installations, because there is so many handoffs in the process and there's so many people involved. What we really tried to do with the task force that developed the SFIP, and this task force was made up of people from ASO, HSO, and GLO. The focus of the task group was to really nail down, again, the roles and responsibilities in the process of selling and implementing a XP Family Array array. So this slide pitch does a pretty good job of that. The other collateral that we'll talk later in this module is the stress free-installation plan, and what that's used for. We will also talk briefly about the mass storage design team. This has cross-functional endorsement on a world-wide basis. So if you start living through this process with local support people you may need to help get them on the same page. So looking at the matrix, obviously TCO has a key role in the sales process as far as, you know, all the obvious things that appear in TCO. So the first couple slides within the slide pitch, really just represent what TCO is responsible for. Just to make a few key points. One of the big areas of confusion during the early days, which has been clarified, is the fact what is our responsibility with respect to doing detailed (LUN) maps for the



## SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training Services Value Chain

customer? There was some documentation that said that TCO was responsible for doing it. And what we have resolved is that, to the extent that TCO in a pre-sales role is responsible for helping the customer conceptually understand what (LUNs) are and what, you know, mapping of data throughout the array, what are the implications there, the functional implications. But it is truly the responsibility of the installing ASE to do the detailed storage LUN map. There's one key element in the pre-sales phase, that was one that I want to point out. The other one that I'll say is still a work in progress, is we are driving towards being able to deliver pre-configured arrays of various RAID levels. So TCO does have a role in the pre-sales phase of getting that document initiated.

## The Site Prep Team

The site prep planning team is responsible for determining site location and location size, ensuring that construction requirements and local codes are met, and scheduling all events related to site completion to prepare for the successful installation and maintenance of the HP SureStore E Disk Array XP512. The site prep team consists of the following personnel:

- sales representative
- customer engineer
- application software engineer
- customer



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



From the HP Site Prep Guide

## XP Family Array Warranty Offering

A 2 year, 24x7x365 onsite warranty that includes:

- Site Prep
- Installation
- Availability Support
  - Phone-Home Support
  - Remote Repair
  - Onsite Part Replacement
  - Redundant Parts in the box
  - Account Management



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Availability Support - Service Features

**Phone-Home Support** - system calls home once daily to report system status; will proactively call when there is a potential problem. The system also provides self-diagnosis to Stress Free Central to assist with the determination of the appropriate response.

**Remote Repair** - performance problems and microcode updates are managed remotely by Stress Free Central without affecting system availability.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Availability Support – Service Features

**Onsite Part Replacement** - if a redundant component requires replacement, it will be dispatched and a CE will install it the next day. Stress Free Central may dispatch a CE and part (if required) with 2 hour\* response to prevent or address a hard system down or customer visible problem.

**Account Management** - Proactive partnership services tailored to meet customer needs. Provided by the ASO.

*It is estimated that 95% of the onsite repairs will be dispatched with next day response.*

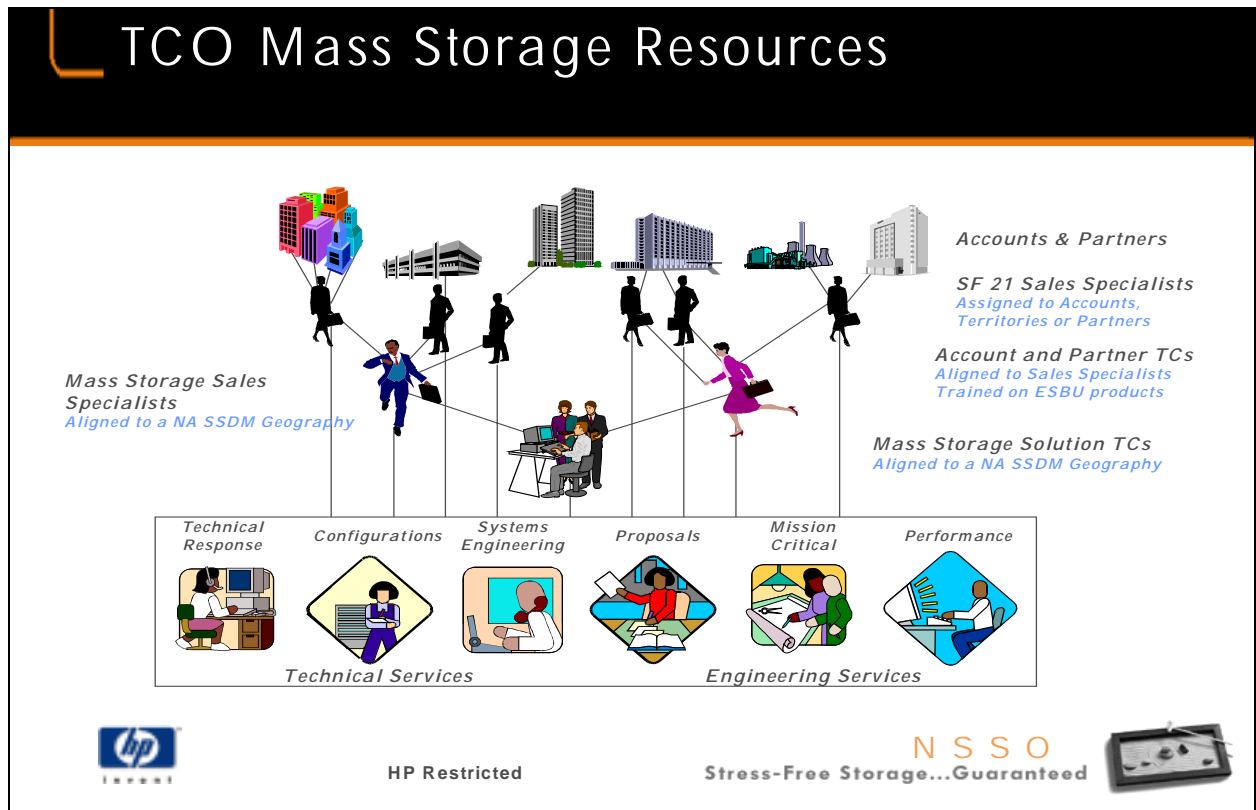
\* if the customer is within 50 miles of the support responsible office



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed









Enable Sales by having an immediate impact on the perception and value of HP's technical capabilities, products and services in the presence of the Customer. Understand current and emerging IT Infrastructure in developing winning strategies, technical ownership and technical selling for a deal. Assemble the technical team, lead TCO Sales Services Team members (Front/Back) Office. Utilize Partners wherever and when-ever possible. Serve as technical resource for Generalist TCs. Own the Pre-sales development of installation document. Align and Focus Consultants to win deals. Align to deals vs. accounts with sales specialists. Focus on Mass Storage deals, Online and Offline. Focus on technologies, (SAN, High-end Arrays, Storage Management, and High Availability Storage solutions). Match Heterogeneous skills to customer environment. Communicate conceptual understanding of LUNs and Mapping to customer and the functional implications. Establish desired RAID levels in order to propose enough RAW storage to meet customer requirements. Communicate RAID level to the factory for configuration. Develop Business Copy, Continuous Access, and other SW product needs which will effect design and storage space requirements. Engage any Back Office resources to design overall Server/Storage architecture. Develop a graphic to help capture high level design capturing network interfaces, SAN devices, and high availability architectures. Own creation of the SFIP (Stress Free Installation Plan). Ensure that additional ASE implementation time is included on the order for complex implementations. This

## SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training Services Value Chain

would be in addition to the bundled installation time, (i.e. ServiceGuard, high availability, etc.). Complete other details in TCO Configuration Tools which affect size and cost of array. Resolve any outstanding Pre-sales issues with customer which are preventing an order to complete pre-sales portion of SFIP. Contact assigned ASE to initiate SFIP review meeting with SR, TCO TC, CE and assigned ASE. Review SFIP and resolve any discrepancies with team and customer. TCO TC leads the customer hand off meeting with the customer, SR and CE to ensure a good hand-off and installation. Respond to ASE/CE requests for clarification of TCO design intentions.

## Storage Service Benefits

- Provide fixed price solutions intended to match majority of scenarios and customer needs – Ordering Simplicity
- Wrap individual products into Solutions (“walls” instead of “bricks”)
- Better control of the “customer experience”
- Increase revenue for HP
- Provide quota and commission credit for the sales team
- Products and Services are complementary business enablers
- Ongoing support relationship (TAS) above and beyond ‘break/fix’ that can be leveraged for future sales assistance

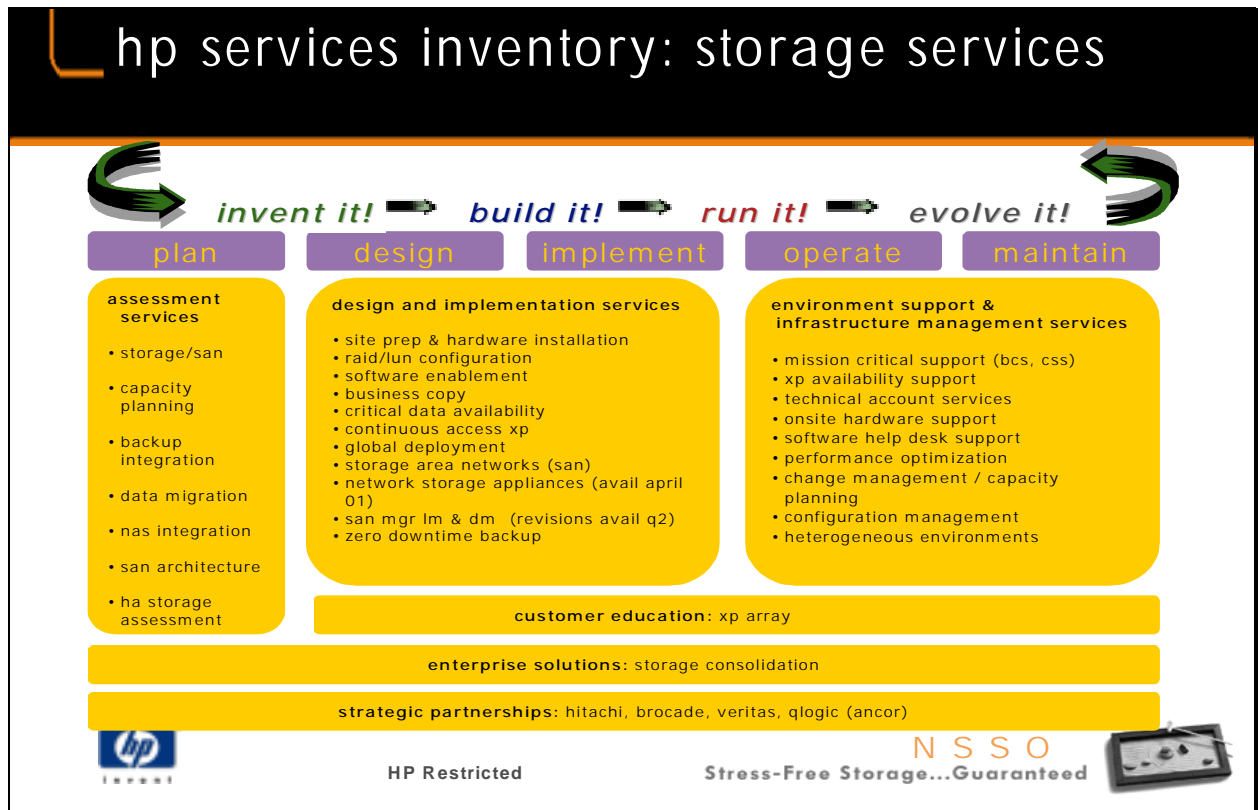
HP Restricted

The next several slides should be a nice, and very solid overview helping us to understand the storage support and services that the 3SO organization has to offer.

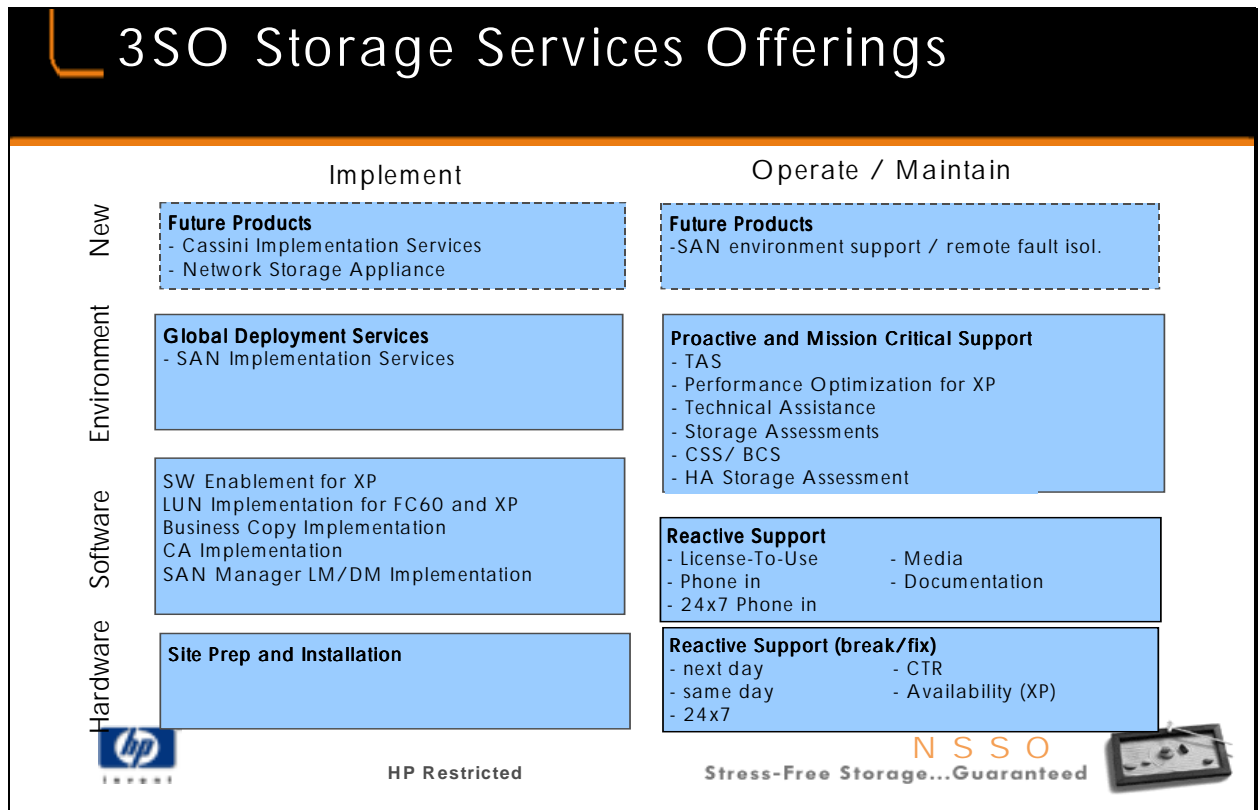
What are the benefits of storage services? Well, benefits are many. One is we offer a number of fixed price solutions that match what we hope is the majority of the customer's situation. Now there's no way that these fixed price offerings are going to match every scenario, but we're hoping they're going to get a large percentage to make it easy to order, easy to quote, easy to turn around the quotations. We're hoping to offer our customers more of the solutions, more a wall then instead of a brick.

And when Hewlett-Packard is involved in both sales of product and the actual execution, implementation of the solution, we have the greatest opportunity to influence, you know, what the customer's perception is going to be with their customer experience. And obviously every service we sell is more money for HP. Our sales force gets more quota and commission for every service sold up front. And as many of you are already aware, our products and services are self-perpetuating in complementary products. Good services make it easier to sell add-ons and easier to build reference accounts. And of course when we sell more product, these products have the potential to drag more services along. They're self-perpetuating.





This is HP's general inventory of services for the XP Family Arrays. We will be talking about some of these in several slides to follow. HP's services fall right into the plan of invent it, building, run it, evolve it strategy, and within those we have plan design, implement, operate, and maintain. Now these services that you're seeing here is all HP services including those provided by HPC and those provided by 3SO and also those provided by customer education and any other strategic partnerships that we have with Brocade, Hitachi, etc.



This slide is a drill down on the 3SO storage services.

## Foundation Services (OS6,3Y6)

- Upgrades standard warranty for HW and SW products
- Base on which to build higher level or comprehensive solutions
- Every option may not apply to every storage device or software product
- A key option: OS6/3Y6
  - Provides software update materials and ability to call the response center (SFC) software and implementation questions
  - Only get warranty HW break-fix without these options



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



Every house has a foundation, and every service solution has a foundation, and the foundation services we have for storage services is the (OS6) and (3Y6) options. These are basically warranty upgrades for the standard hardware and software products. HP builds on top of that to get the kind of level of responsiveness and services experience that we need to have for that given customer. It's kind of a modular approach. We build upon the base. Every option may not certainly apply to every storage device or product, but fortunately our order processing systems help coach us when there is an applicable, services product. One of the key options here is the (OS6) and (3Y6), and (3Y) being the three-year option. And (OS6) is the one-year upgrade option. And what this does is provide software update materials, the ability to get later versions of these software products, now available electronically on the web, which is good news for a lot of us and the ability to call stress-free central, for software implementation questions, break fix and implementation questions. And without this key service we end up having just hardware break fix capability. And the customer does not receive software updates. So it's important to get at least this foundation service on the order at the very, very least. We recommend a great deal more than that to kind of match what the competitor's provide for customer experience. But this one is very critical because the customer doesn't have access to the HP infrastructure without it.

## So what is Enablement?

- Intended to deposit, activate, test, and confirm that a given product is functioning normally
- Doesn't implement the product or customize to customer's environment
- Needs implementation to bring the product to a "solution ready" state
- Allows customer or reseller the option of doing their own implementation but have a stable base to build on
- Available for most XP software products



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



The question is, what is enablement? This is a very common question. And enablement is very simple. It's intended for you to deposit, activate, test, and you confirm that the product is functioning normally. If you were buying a lawn tractor, it's the equivalent of taking it off the truck, putting oil and gas in it, starting it up, and driving it up and down the driveway. But if HP were your gardener we would not be doing your lawn work and your gardening work and landscaping work for you. But we'd be making sure that that garden tractor you bought is ready to do what it's supposed to do and is working properly. It doesn't actually implement the product, doesn't do your yard work for you, and we don't customize the product to the customer's environment. We don't examine what the customer needs and tailor it. Basically it allows you to bring the product to a solution-ready stage. If we turn the keys over the customer, they would immediately be able to start using the product and start customizing it if they wanted to do it themselves or if they had a channel partner that was doing part of that work for us. And enablement is available for most XP software products and on the initial purchase of an array, enablement is bundled in to most of those products. It doesn't show up as a line item. So what's the difference between that and implementation, we will discuss that on the next slide.

## ... and what is Implementation?

- Intended to customize the product to the customers environment
- Provides discover -> plan -> implement - > confirm process (much larger time investment)
- Customization based upon customer dialog and study of the application and environment
- Provides customer with "ready to use" solution without investing resources in education and construction
- Allows customer to leverage HP's experience
- Stops short of host volume management and executing advanced integrations such as implementing OBII split mirror, SAP, Oracle, etc.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



Implementation does take it to the next step. The process involves going through a discovery process, we plan, we implement, we engage the customer, learn what it is they exactly want to do, and what they want to achieve. We select the options of the product. We learn how to maximize the use based on their criteria, whether it's performance or availability of the products. And we bring it all the way up to the ready-to-play stage. When we do this we take it up to the point where it's just short of advanced integration, such as where we're doing an enablement - not an enablement service - but one implementation service. We wouldn't do all their logical volume manager work for them. We would bring it up to the point where the system can see everything and is ready to roll, and the customer can immediately start building volumes. And we wouldn't integrate advanced solutions such as SAP or OmniBack for the customer. But the solution would be ready to go ahead and do those things. So it takes it much farther than enablement would go, and it really is like a mini-engagement all on its own.

And the next slide gives an example of an implementation service for the XP Arrays. The initial one, implementation service is bundled with all XP arrays. It includes enablement services for most of the software products. And what it does is it provides one configuration event, which includes all that planning and determining exactly how we're going to use the storage. And again we collect input from the customer on what their goals are and try to match what they've purchased to customize what they purchased to meet those kinds of goals. And we create a one design, one plan, execute that one plan, and then make sure -

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Services Value Chain

we close with the customer and make sure that we have executed to their needs, and that means whether it's connectivity performance or performance or availability needs.

## XP Initial LUN Implementation Service

- Included with all new XP arrays
- Provides one configuration effort including:
  - Planning and coordination of the array installation
  - Enablement of applicable SW products
  - LUN plan, design, and execution
  - Closure with customer
- Array implemented consistent with customer's connectivity, performance, and availability needs.
- Added services needed when SAN is part of solution



HP Restricted





N S S O  
Stress-Free Storage...Guaranteed



## XP LUN Implementation Service

- Provides LUN Planning, Design, and execution after the initial XP implementation
- Sell when:
  - existing array is being re-deployed
  - significant add-ons have been purchased (i.e.: > 20 RAID Groups)
  - the frame will move from development to production within a year
- \*Was\* used to capture SAN or complex installs and when initial LUN plan wasn't included with XP (no longer the case)

Product# H9273A-301  
\$14,400 Discountable


HP Restricted

HP has another service called LUN implementation service. And this is very to the one just described on the previous slide, the **initial** LUN implementation service. But this is really a best match when the existing array is being re-deployed if they're going to scrap all the work they did on it and re-do the LUN plan for a different purpose. Say they were running a SAN testing environment, and now they're deciding they want to move it into production. Or it's also appropriate if the customer has purchased a large amount of add-ons, say more than 20 array groups. That's going to be a significant planning and configuration exercise, and the service would be a fit under those conditions as all the work could not be accomplished in the LUN implementation services bundled with the array.




## XP Business Copy Implementation Service

- Implement Business Copy Solution on 1 frame and up to 4 hosts
- RAID Manager Implementation Included
- Includes:
  - Planning
  - Design
    - Considers performance/capacity
  - Execution
    - Volume placement, RM config, "sample" RM script
  - Verification
    - sync/split/resync ops, test vol/file access, test RM configs
- Added services needed for application integration such as: OBII, SAP, Oracle, etc



Product# H9273A-302  
\$14,400 Discountable

N S S O  
Stress-Free Storage...Guaranteed



The Business Copy XP implementation service. This is basically an implementation of HP's Business Copy XP replication product for the XP. This particular price point is for one frame and up to four hosts. And of course the real power of this particular product is the fact that it can be host controlled through a Command Line interface or APIs. And so the implementation of RAID Manager XP, which is key to that functionality and actually controls this product, is included with the service.

## Other XP Services

- ✓ XP Critical Data Implementation Service
- ✓ XP Performance Optimization Implementation Service
- ✓ SAN Implementation Service
- ✓ Continuous Access Implementation Service



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



The SAN Implementation Service takes us beyond the box and moves us into the storage area network. And what this does is implements that solution end-to-end in between all the storage appliances that the customer might have out there, the switches, the hubs, and attaches it to the host. HP works out all the details on who's going to be able to talk to what under what conditions, the LUN security, the connectivity, the high-availability multipathing.

## Account Support Organization

- Implement SFIP - Take over lead from TCO
- Installation project ownership
- Determine final configuration
- Collaboration with factory on changes
- Create LUN plan and Map with customer
- Host server verification
- Deliver the Basic Installation Bundle
  - LUN implementation
  - Software Enablement
- Additional consulting services
- On-going account support






HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



# ASO Installation

- Implement SFIP
- Installation project ownership
- Determine final configuration
- Collaboration with factory and ASO
- Design the LUN Map
- Verify Server readiness - OS, Patches, etc.
- XP Family Array System Verification
- Software signoff
- Enablement
- Customer acceptance

 HP Restricted  Stress-Free Storage...Guaranteed 

ASO is the primary delivery organization for Account Support and Implementation Services

H9273A opt 301	SSE LUN/SAN Implementation	
H9273A opt 302	SSE Business Copy Implementation	
H9273A opt 303	SSE Critical Data Availability	
H9273A opt 304	SSE Performance Optimization	
H9273A opt 305	Add'l ASE Technical Consulting (per day)	
H9273A opt 306	Add'l ASE Technical Consulting	- SAP (per day)

## HSSO Roles and Responsibilities

- Account Team Participation
  - Contact the ASE to make sure that the RAID level and LUN design are complete prior to scheduling installation
  - Ensure phone home modem line is available prior to installation
- Site-prep
- Install system
- Verify operation of hardware including firmware revision level
- Test Phone Home and callback
- Call the TAC Center to register the site for Phone Home activation
- On-going hardware support including microcode



HP Restricted




N S S O  
Stress-Free Storage...Guaranteed



Work with sales team to ensure correct HSSO services are sold. CEs Provide customers an excellent experience during the installation and implementation of their XP Family Array. Contact ASO manager for ASE assignment if necessary. Schedule site prep and on-site installation. Ensure phone capability is functional and working properly.

## HPC Roles and Responsibilities

- Executive Storage Vision Workshop
  - Analyze and set future strategy
- SAN Architecture
  - Detailed study and design resulting in Implementation Roadmap
- Storage Consolidation / Data Migration
  - Movement of data from old to new storage
- Backup
  - Design and Implementation of Complex Omniback and Veritas B/R Solutions
- Business Continuity
  - Study, Design, and Execution of complex Disaster Recovery Solutions
- Performance and Capacity Planning
  - Study of applications, servers, and storage resulting in status, tuning, and capacity recommendations

 HP Restricted  Stress-Free Storage...Guaranteed 

Now HPC has their own suite of services that complement and wrap around services from the 3SO that we have briefly looked at up to this point. They consist of the six that you see on this page. Either the executive storage vision workshop which is an exploratory service that helps the executive understand where he's at and where he could be but does not necessarily get into the detailed deployment or execution phases. There's a SAN architecture, which takes that activity a little farther, and actually builds a roadmap and a plan on how the customer is going to get there. HPC is in charge of storage consolidation and data migration. They are your first point of contact when a customer needs to move their data. That's part of the solution from old storage to HP storage. Same thing when we're talking about advanced backup solutions. HPC has advanced Veritas and OmniBack II implementation. HPC is your point of contact for building that solution. It is very possible that 3SO's Continuous Access and Business Copy services will be a component of that as well. But HPC may be building and providing the larger service which include some of the execution and the project planning as well. Same kind of thing with business continuity services. The large disaster recovery or, you know, multi-city, multi-continent kinds of disaster recovery services. And of course performance capacity planning and customer wants and kind of modeling assistance, do some what if scenarios based upon servers and storage at the different servers and different options. It is HPC that can provide that service.



## Service Brief

### Modem Connect for the HP Surestore E Disk Array XP Family Array

#### HP Services and Support System Support

##### Overview

HP Availability Support for HP Surestore E Disk Array XP Family Array combines leading-edge remote diagnostic technology with proactive hardware maintenance services that increase the availability of your storage systems, applications, and data.

The HPSurestore E Disk Array XP Family Array is designed to achieve the highest level of continuous access to your data. To accomplish this, modem access to the storage system and use of Continuous Track is a must.

HP Response Center engineers utilizes Continuous Track to monitor your storage systems and proactively investigate and resolve potential problems before your business is affected. The diagnostic "phone-home" technology, which identifies potential problems before they occur, and hardware redundancy embedded in your storage system allow HP to virtually eliminate downtime. If a problem cannot be solved remotely, a HP engineer will be dispatched to proactively go onsite before the storage system experiences a hard-down failure. This

preventive approach to support allows users continuous access to data and applications and eliminates stress for system managers.

#### Continuous Track

Continuous Track is used by the Response Center to proactively address potential problems. This tool is also loaded on the Service processor (SVP) located inside of your Storage System.

Continuous Track gathers maintenance information only. Customer data is never accessed.

In the Site Initiated mode the XP Family Array Storage System calls the Response Center to provide a daily status.

In the Center Initiated mode the Response Center contacts the XP Family Array Storage System, to upload and download micro-code files.

In the Remote Console Session mode the Response Center calls the XP Family Array Storage System identifies itself as an authorized caller, requests a remote session, disconnects and waits for a call back.

Both Center Initiated and Remote Session modes can be disabled using the Remote Enable switch. When this switch disabled micro-code can not be uploaded or downloaded to the SVP.

#### Benefits

Continuous data availability due to prevention of storage system interruptions.

Higher efficiency / productivity as a direct result of having your data continuously available.

Resolve problems faster due to the early warning system provided by Continuous Track.

Lower support costs as a result of remote analysis and problem resolution.

Part failures reduced to routine maintenance calls. Continuous Track early warning systems allows HP to schedule the repair of failed parts.

Uninterrupted 24x7 system monitoring Continuous Track monitors your storage system reporting any abnormality which may threaten your data availability

#### Security Features



SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Services Value Chain

- Continuous Track gathers maintenance information only, customer data is not accessible. Continuous track reads the System Information Messages loaded in the storage systems register to provide an early warning to HP on potential
- No physical or logical connection between Continuous Track and customer data on the disc
- Unique Protocol: Continuous Track uses a unique protocol and file formats, which is virtually unhackable.
- Outbound & Call Back: Continuous Track from the SVP calls the Response Center. When a remote console session is requested Continuous Track disconnects and the RC awaits a call back from the unit.
- Point to Point Communication: Continuous Track calls a pre-programmed phone number, the Response Center. Continuous Track does not use Internet.
- Multi Layered Passwords: Several layers of passwords must be received before any communication with Continuous Track on the SVP can be received.
- Time Limitation on Connection: If Continuous Track does not receive a valid command within a specified period time the connection is dropped and the modem is disabled until next day when Continuous Track calls the Response Center.

## Stress Free Central

- Located in Atlanta
- End point for all worldwide 'Phone Home' calls (95% of all calls)
- Support Provided 24x7x365 by trained HP Engineers
- Remote Repair can ONLY be done from Stress Free Central
- Determine priority of support call; ie Next day or Immediate
- Determine part(s) needed by On-site CE
- Notify CE of ETA
- Provide all worldwide CE XP Family Array technical assistance
- Provide all worldwide Escalation Support



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



## Stress Free Installation Plan

- ♦ **Background** - The SFIP was developed to detail the "Roles & Responsibilities" of the primary contributors who sell, consult, install and support XP Family Array disc arrays from HP. The SFIP is also used to document the support value chain and the high level processes associated with the XP Family Array. This plan was developed with the direct input from the multiple HP support organizations that are responsible for owning and delivering specific services. The SFIP is considered a "working document" and will be used by HP's various support teams throughout the process.
- ♦ **SR will initiate** - planning meeting with internal team (and customer?). The SFIP can be used to describe to the customer HP's comprehensive installation process.
- ♦ **Signoffs:** A key to Operational Excellence is OWNERSHIP. This section captures the completion of tasks by specific owners.



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



### SFIP Process Message

In order to achieve Carly's "Total Customer Experience" metric for 100% customer satisfaction, the Stress Free Installation Plan (SFIP) must be completed for every XP sale. The SFIP's usage will be tracked with respect to TCE results. This process will continue to be improved and will be integrated with the Superdome TCE process.

The SFIP process is a valuable part of the TCE because it:  
communicates the solution design and customer expectations to the support organization

results in a unified HP presence at the customer site. includes the XP512 information. This form should be completed in conjunction with booking the purchase order.

The North America, SFIP process is:

The Sales Rep is responsible for:

- initiating the SFIP process
- ensuring that HP delivers on its commitments.

The Front Office TC (or TC who designed the solution) is responsible for:

completion of Sections I & II of the SFIP. (The E-Storage TCs are available to assist in completion of the form if needed. They can be reached at t/n 447-7070, opt. 2,1

communicating the SFIP information to the installation ASE. Completed forms

should be sent to the assigned ASE and to Katrina Towle (katrina\_towle@hp.com).

The ASO requires the SFIP form before starting an installation. Therefore, please schedule the installation planning (hand-off) meeting with the assigned ASE, to review the SFIP, as early as possible.

An option is now available to PRECONFIGURE both XP256s and XP512s, at no additional charge. The following options can be preconfigured:

1. ACP RAID Level (0/1, 5)
2. Emulation type for the array
3. Array Grouping (location and size)

In NA, if you would like to implement a preconfigured system you must contact the E-Storage team early in the sales cycle.

If you are outside of NA, please contact your local Mass Storage Region Manager for the SFIP and preconfiguration process.

The local region managers are:

AP - Michael Young - t/n 839-8525

LA - Rafael Diez - t/n 258-4205

Europe - Alfred Haaga - t/n 778-2014

Please give me a call if you have any questions. Comments for improvement to this program are always welcome.

Warm Regards,

Marcy Pack

Project Manager Mass Storage Solutions

Global Sales Services Direct dial: (703) 204-2487 Telnet: 234-2487

## HP SurePartner's Role & Responsibility

- Complete Sections I & II of the SFIP for each pending deal.
- Contact Mass Storage TC for your area to get them engaged.
- After sale is made, arrange the customer meeting with the assigned ASE & Mass Storage TC.
- Deliver a well planned and detailed kickoff meeting with the same parties.
- A complete copy of the SFIP is located in the appendix



HP Restricted

### XP256 Stress Free Installation Plan (SFIP) Rev 404

<u>Process Sections</u>	<u>Responsible Owner</u>	<u>Participants</u>
I. Customer Information	Front Office TC (FOTO)	SR Channel Partner (CP)
II. Pre-Sales Phase	Front Office TC	BO/TC/CP/TC
III. Installation Planning Meeting	SR	FO/TC/ASE/CE/CP
IV. Storage Map	ASE	
V. Installation Phase	CE	
VI. Configuration Phase	ASE	
VII. TAM - Technical Account Manager	ASE-TAM	

Note: Upon future CP certification testing, if the Channel Partner is selling solution they will own the completion of Sections I, II, and arranging the customer meeting in Section III.

Suggested Attachments: (Check if attached)  
 I. Interactive Disk Capacity Spreadsheet  
 II. Connections Diagram (Hosts, SAN...)  
 III. Meeting Notes  
 IV. Additional documentation - Storage Map

<u>Signoffs</u>	<u>Completed By</u>	<u>Date</u>
Customer Information		
Pre-Sales Phase		
Installation Planning Meeting		
Storage Map		
Installation Phase		
Configuration Phase		
TAM- Introductions		

N S S O  
Stress-Free Storage...Guaranteed




## Module Wrap-up



HP Restricted


N S S O  
Stress-Free Storage...Guaranteed






# Module 16

## Cache LUN XP



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



For Your HP SureStore E Disk Array XP256 Manuals:

<http://www.hp.com/cposupport/prodhome/hpsurestor21443.html>

## Module Agenda



➤ Cache LUN Overview



➤ Cache LUN Screen Shots in  
Command View XP



➤ Wrap-Up



HP Restricted

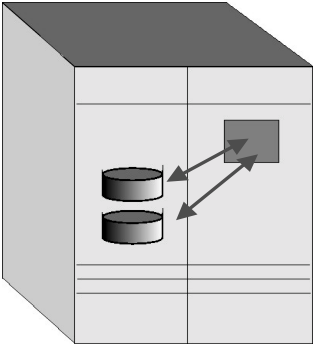
N S S O  
Stress-Free Storage...Guaranteed





## Cache LUN

- Cache LUN lets you reserve areas of the disk array's cache memory for storing high-use data
- Use Cache LUN to achieve higher data access rates
  - Use Cache LUN with VSC for even higher performance
- Only two Cache LUN operations
  - Placing data in cache
  - Removing data from cache




The diagram shows a 3D perspective of a storage array. On the left side, there are two disk icons. On the right side, there is a square icon labeled 'Cache memory'. Two double-headed arrows connect the disks to the cache memory, indicating bidirectional data flow.

Cache memory

**HP**  
INVENT

HP Restricted

**N S S O**  
Stress-Free Storage...Guaranteed







### Student Notes

The Cache LUN function enables you to reserve areas of cache memory on the disk array for storing specific high-usage data. This capability enhances data accessibility because the user-defined cache-resident data is available at host data transfer speeds for both read and write I/O operations. When VSC and Cache LUN XP are used together, you can achieve higher data access performance than when either of these features is used individually. Prior to using Cache LUN, space in the cache must be reserved. This can be done on install or the cache can be reconfigured later. This is an SVP activity, performed by the HP CE.

## Cache LUN modes

- Data is not read into cache until referenced
- There are two Cache LUN modes:
  - Priority mode
    - Generates traffic to disk
    - Cache areas can become free
    - Writes cause destaging to disk
    - Destaging to disk causes time delays
  - Bind mode
    - Keeps data in cache and doesn't generate traffic to disk
    - Keeps data until extent is removed
    - During a loss of AC power, battery can preserve data for up to 48 hours
- Mode is configurable for each cache extent



### Student Notes

A cache extent is a range of logical blocks within a LUN or as an entire LUN. A logical block is 512 bytes.

Data assigned to a Cache LUN extent is not read into cache until referenced.

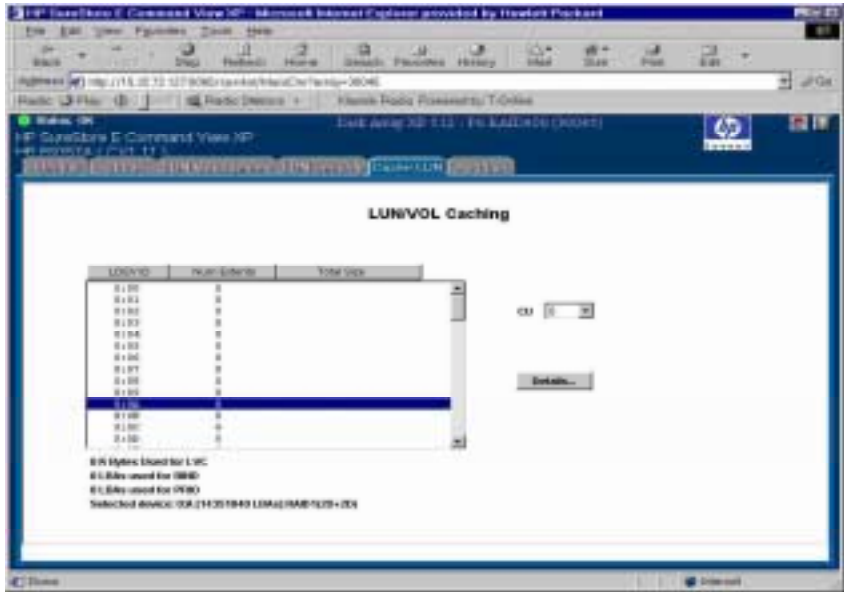
Priority mode requires fewer cache resources but has performance limitations over bind mode. Priority mode keeps all read referenced data in cache until you delete the extent. All writes in priority mode result in destaging the data to disk using the same algorithms as for regular, non-cache LUN data.

- All reads following the initial read to a block will result in cache hits.
- After destaging of the write block to disk, the duplicated write areas in cache become free.
- While destaging to disk, read requests to the same block are delayed until the disk operation is completed.
- Priority mode generates write traffic to disk, whereas bind mode does not. Bind Mode keeps data in cache, eliminating disk traffic. All data stays in cache until the extent is removed. Upon loss of AC power, data is maintained in cache for up to 48 hours, using battery power.
- All reads following the initial read to a block will be cache hits.
- Duplicated write blocks are stored in cache and are never "evicted."
- All read or write referenced blocks stay resident in cache until the extent is removed.

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Cache LUN XP

- Bind mode requires more cache because duplicated write blocks are not destaged to disk.

# LUN/VOL Caching screen



➤ Shows caching information for LDEVs in a CU

HP Restricted Stress-Free Storage...Guaranteed

## Student Notes

This screen shows Cache configuration information for all the LDEVs that are present in a particular CU. To view the configuration of LDEVs in another CU, select the desired CU from the CU dropdown list box.

For the selected CU, all the LDEVs, along with information about the number of cache extents and the size in LBAs, are displayed.

To view LDEVs in another CU, select the desired CU from the CU dropdown list box.

The screen shows:

LDEV ID The current LDEV selected

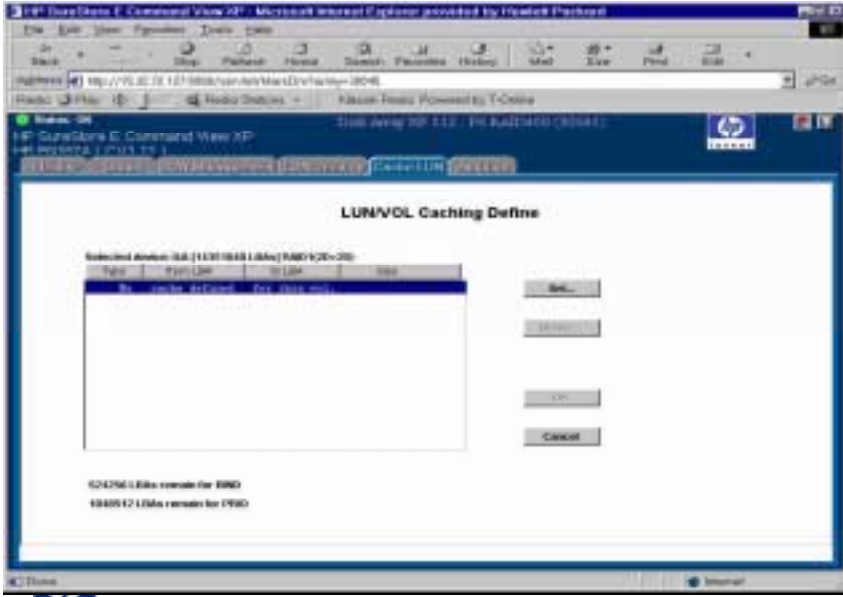
Num The number of cache extents

Size The size in LBAs

CU dropdown Selected CU

Select an LDEV and click on Details to set or delete cache extents for an LDEV.

## LUN/VOL Caching Define screen



- Displays when you select an LDEV
- Use this screen to choose to set or delete cache extents

HP Restricted Stress-Free Storage...Guaranteed

### Student Notes

The LUN/VOL Caching Define screen displays when you select a particular LDEV from the list of LDEVs for the selected CU. In this screen all the cache extents that have been defined for this LDEV are displayed with details.

#### Fields

Type Mode type (Bind or Priority)

FromLBA Starting Logical Block Address

ToLBA Ending Logical Block Address

Size The total size of LBAs (in MB) and RAID type for LDEVs is displayed about the multi-column list box.

A '+' sign at the end of a row in the multi-column list box means that the cache extent is not yet committed.

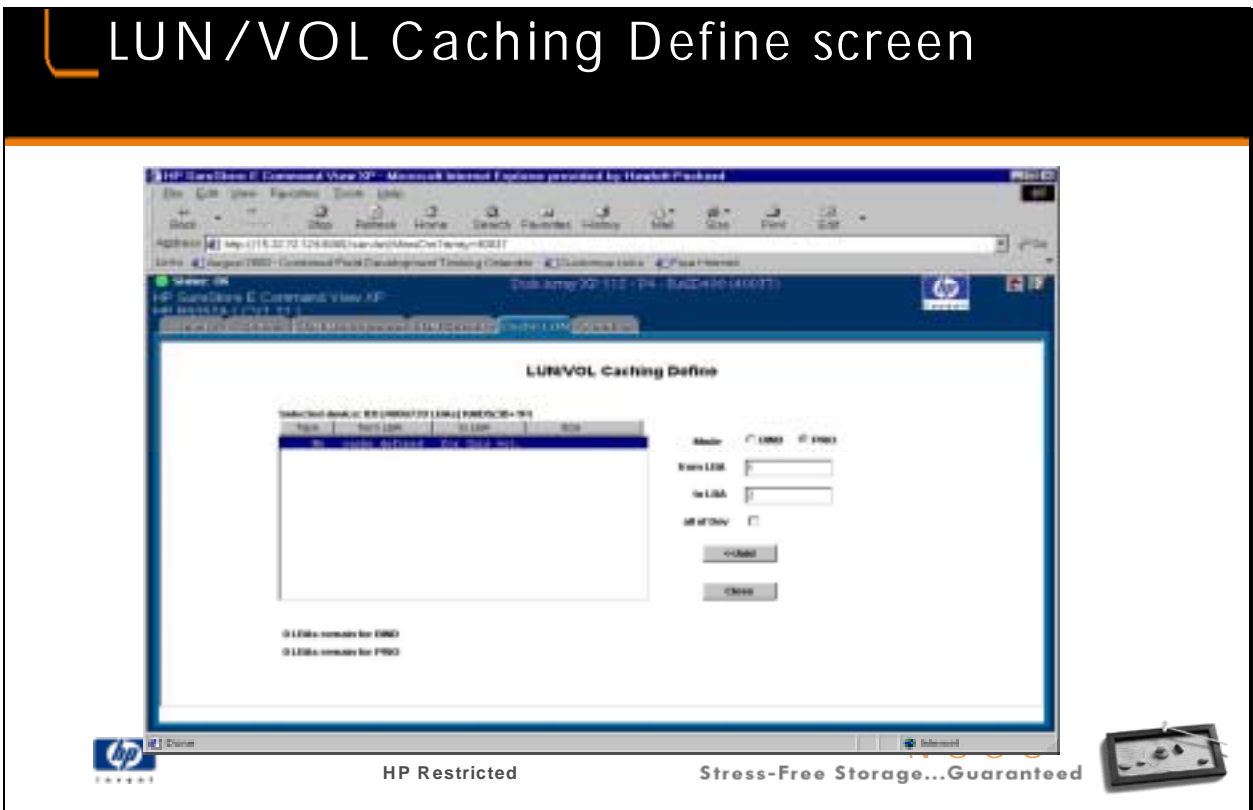
#### Buttons

Set Is used to add cache extents for the selected LDEV.

Delete Is used to remove cache extents for the selected LDEV.

OK Takes control back to the main screen after the requested cache extents are either set or deleted.

Cancel Cancels any selected items.

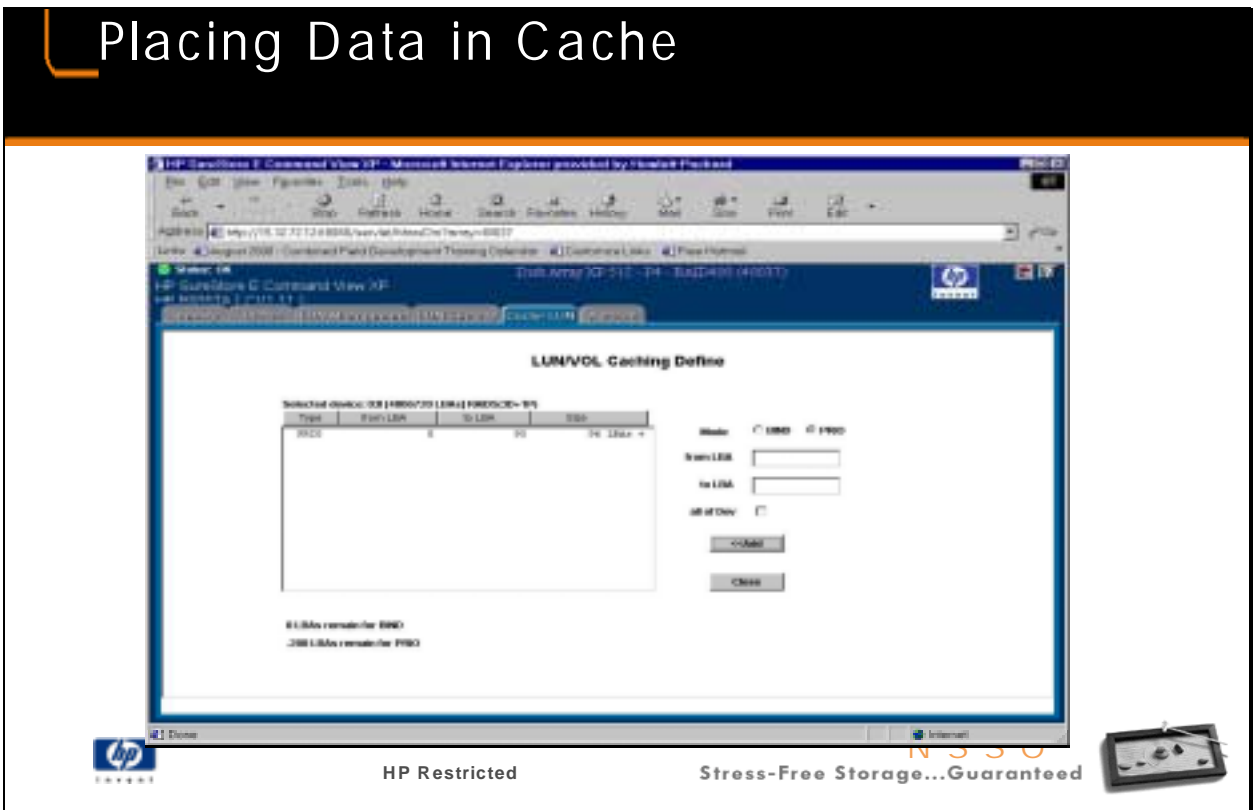


### Student Notes

Clicking the Set button on the previous screen, will replace the four buttons with a panel of text fields and buttons which are used to add cache extents from this LDEV to the cache area.

#### Buttons

**Add** - Used to add a cache extent from this LDEV to the cache area. You can use this panel repeatedly to add multiple cache extents in one operation  
**Close** - Removes this panel and replaces it with the original set of four buttons.

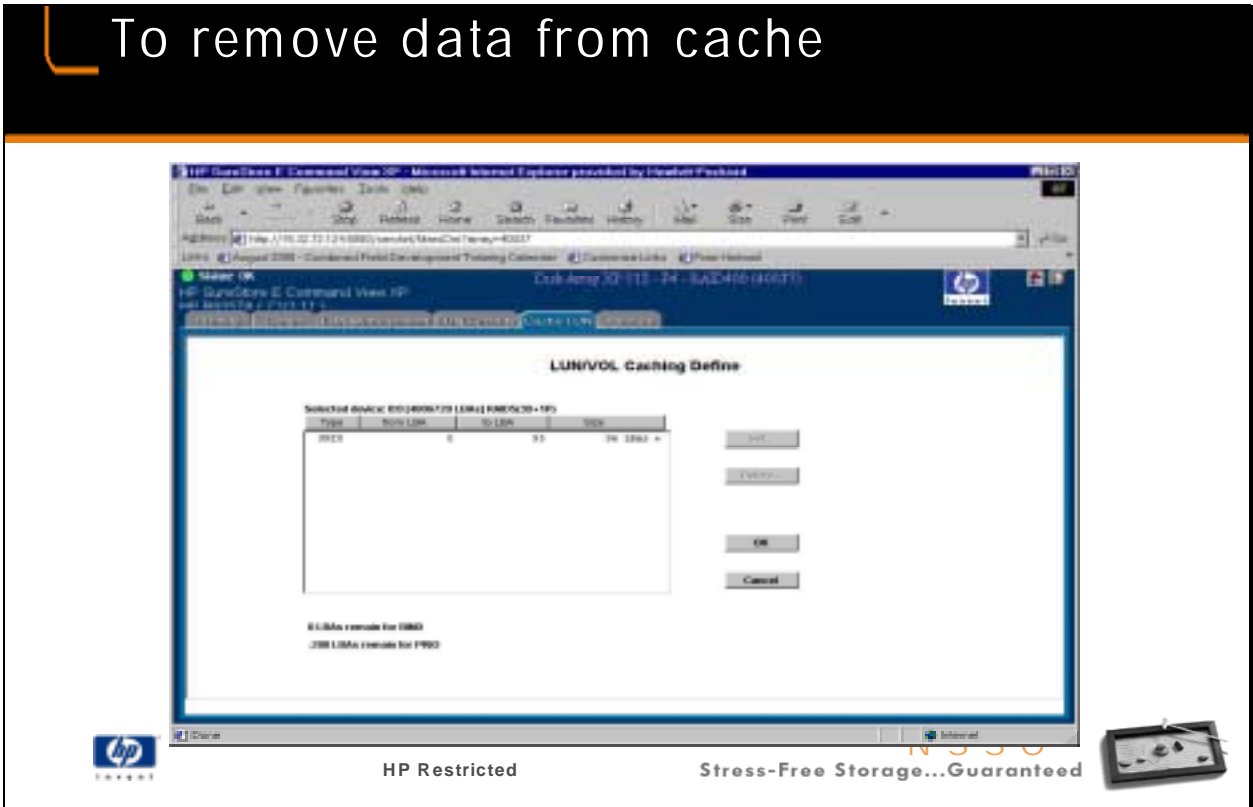


## Student Notes

The Cache LUN function allows you to place data in cache areas.

To place data in cache:

1. Select the CU in the CU pulldown box on the LUN/VOL Caching page. The configuration of LDEVs is displayed.
2. Select an LDEV. The caching information for the selected LDEV is displayed.
3. Click Define. The LUN/VOL Caching Define screen is displayed.
4. Click Set. The buttons are replaced with a new panel of buttons and text fields.
5. Add one cache extent from this LDEV to the cache area.
6. Repeat step 5 to add multiple cache extents from the same LDEV.
7. Click Close. This screen is replaced with another containing the original four buttons.
8. Click OK. All requested cache extents will be placed in the cache area.



## Student Notes

The Cache LUN function allows you to remove data from cache areas using the main screen.

To remove data from cache:

1. Select the CU from the CU dropdown list box.
2. Select the LDEV from the list box.
3. Click Define. The LUN/VOL Caching Define screen is displayed.
4. Select the cache extent to be removed.
5. Click Delete. The cache extent is removed from the list box.
6. Click OK. The cache extent will be removed from the cache area.

When you return to the previous screen:

8. Click OK.



## Module Wrap-up



HP Restricted

N S S O  
Stress-Free Storage...Guaranteed



# Module 17

## SHAC SCSI Host Array Connect

SCSI-Fibre Channel Router, Bridge/Extender



N S S O  
Stress-Free Storage...Guaranteed



## Module Agenda



➤ What is SHAC?

➤ Offerings?

➤ SHAC Benefits?



➤ Connectivity Options

➤ Documentation

➤ Wrap up



N S S O  
Stress-Free Storage...Guaranteed



## What is SHAC?

- The SHAC Project is SCSI Host Array Connect
  - Provides SCSI interface capability to the XP512 disk array through a 3<sup>rd</sup> party SCSI to FC bridge
  - Allows customers to connect their SCSI hosts to the Fibre Channel array
  - Crucial for allowing MPE/iX to attach to a Fibre Channel port
  - SHAC uses SCSI-Fibre Channel Router A5814A



N S S O  
Stress-Free Storage...Guaranteed



MPE/iX is Multiprogramming Executive with integrated POSIX for HP e3000 servers. The SCSI-Fibre Channel router/extender allows MPE/iX, to support current and future SCSI disks and tapes technologies at Fibre Channel distances (distance solution), or MPE/iX, HP-UX and NT to support Fibre Channel devices (bridge configuration). The SCSI-Fibre Channel router supports JBOD disk (SCSI), HP AutoRAID Model 12H disk arrays and XP256 (SCSI CHIPS) and one fibre channel device, namely the XP512. When two SCSI-Fibre Channel routers are connected back to back, the HP 3000 can extend the SCSI differential distance to 500 meters (distance solution). The distance solution is not supported on HP-UX and NT. Single-initiator testing took place in HP-Roseville on MPE (6.0, 6.5), HP-UX |(11.0) and NT 4.0 Server connected to the XP512.

## SCSI-Fibre Channel Router A5814A

- SCSI-Fibre Channel Router A5814A offers
  - Connectivity Solution, used as a bridge
    - SCSI-FC Bridge configuration for MPE/iX, HP-UX, and NT
  - Distance Solution, used as an extender
    - For the MPE/iX platform, supporting SCSI disks and tape at Fibre Channel distances
  - Public release for SHAC solution is *11 July 00*



N S S O  
Stress-Free Storage...Guaranteed



The HVD-FC SCSI-Fibre Channel Router is a simple to maintain stand-alone device that can be used as either a SCSI distance solution (extender) or a SCSI-Fibre Channel connectivity device that allows SCSI hosts to interoperate with FC storage loops. By converting SCSI to Fibre Channel, it can provide lengths of 30m (98 ft) to 500 m (1640 ft) and connect to a second router (distance configuration) or a Fibre Channel Connectivity device (bridge configuration). The distance configuration is specific to the HPe3000 platform whereas the connectivity option (bridge configuration) is available with the HPe3000, HP9000 and Windows NT platforms.

## SCSI-Fibre Channel Router Benefits

- Bridge for MPE/iX, HP-UX, NT
  - Provide access to native Fibre Channel disk arrays from SCSI hosts
- Extender for MPE/iX
  - Protect customer investments in SCSI storage devices
  - Improve use of computer room floor space
  - Move peripherals with sensitive data to vaults or secure rooms
  - Move critical data to secure areas for disaster tolerance



N S S O  
Stress-Free Storage...Guaranteed



## SCSI-Fibre Channel Router



31x11x7.6 cm  
11x4.25x3 in

A5814A SCSI-Fibre Channel Router


- OEM product from Vicom
- \$6400 on CPL
- Provides FC connectivity for MPE/iX, HP-UX and NT SCSI servers
- Can be used to with MPE/iX to extend SCSI distance



N S S O  
Stress-Free Storage...Guaranteed



## SCSI-FC Router • Rear View



- 1 parallel Fast Wide SCSI HVD 68-pin High Density connection
- 1 SCSI termination connection
- 1 Fibre Channel connection per box
- 2<sup>nd</sup> FC connection for future expansion



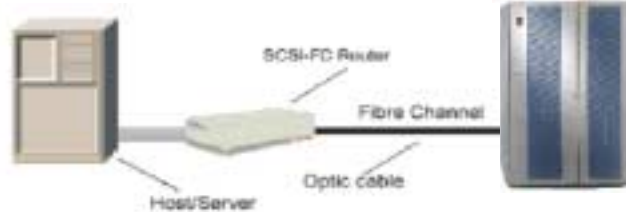
N S S O  
Stress-Free Storage...Guaranteed








## Channel Connectivity

### SCSI-FC Router as a Connectivity Device



XP512  
Standard  
Mode

- For HP e3000, HP9000 and Windows NT
- SCSI hosts connect to FC storage loops
- Hubs and switches are not tested and not supported
- 500m optic cable length for 50 micron, 175m for 62 micron
- 25m cable length for Fast Wide SCSI HVD

10km long wave single mode Fibre Channel cable not tested and not supported

Mode 08 for HP

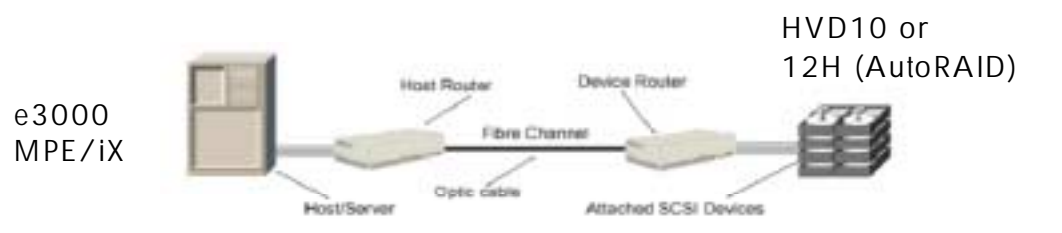
SCSI-Fibre Channel Connectivity (bridge configuration): In this configuration a single SCSI-FC router is connected to both the host/server and to the Fibre Channel storage loop. This allows the host to communicate with Fibre Channel arrays. At this time, the only FC supported array is the HP SureStore E Disk Array XP512. As newer FC arrays tapes and JBOD become available, the SCSI-FC connectivity mode will be extended to encompass them. The router is then referred to as the 'device router'. B. Configuration The SCSI-Fibre Channel Router is supported on MPE/iX, HP-UX, and NT using the Fibre Channel GBIC (Gigabit Inter-link Connection) only. SCSI-Fibre Channel router configured for distancing has minimum requirement of: Fibre optic cable: Short Wavelength 50 micron multimode fiber -Max Distance 500 meters -Connector Dual SC SCSI-Fibre Channel Router microcode revision 7.62 or later (revision 7.60 or later is supported on MPE/iX in the distance configuration) SCSI-2 Fast/Wide SCSI-2 differential host adapter HP28696A, firmware level 3944 HP-UX supports HSC (A2969A) but not the HP28696A adapter. NT also supports the FWD Adaptec AHA2944-UW with driver aic78xx.sys SCSI-2 compliant K-Class SCSI-2 compliant differential cabling

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
SHAC

Attachment \* Compatible with MPE/iX Operating Systems Versions 5.5 and 6.0 or later  
\* Compatible with HP-UX Operating System 11.0 (Patch Level: SP48) \* Compatible  
with Windows NT 4.0 Server (Service Pack 6) \* HP XLR1200 system




## Distance Solution

### SCSI-FC Extender as a Distance Solution



The diagram illustrates the SCSI-FC Extender setup. On the left, an 'e3000 MPE/iX' server is connected to a 'Host/Server' block. This block is connected to a 'Host Router' via an 'Optic cable'. The 'Host Router' is connected to a 'Device Router' via a 'Fibre Channel' link. The 'Device Router' is connected to 'Attached SCSI Devices' via another 'Optic cable'. The 'Attached SCSI Devices' are labeled as 'HVD10 or 12H (AutoRAID)'.

- Specific to the HP e3000 platform (MPE/iX)
- 500m optic cable length for 50 micron, 175m for 62 micron
- Hubs and switches are not supported
- 25m length for each Fast Wide SCSI HVD cable

10km long wave single mode Fibre Channel cable not tested and not supported

Distance Solution: Support of the SCSI-Fibre Channel router allows MPE/ix to support current and future disk and tape technologies. Today, the SCSI-Fibre Channel router supports JBOD disk and HP AutoRAID Model 12H disk arrays. When two SCSI-Fibre Channel routers are connected back to back, the user can extend the SCSI differential distance from 25 meters to 500 meters. SCSI-Fibre Channel Router setup in distance configuration: There are two SCSI-Fibre Channel Routers that make up the SCSI extender. One of the SCSI-Fibre Channel Routers directly attached to the host or the server is the 'host router'. The second SCSI-Fibre Channel directly connected to the device (RAID box or drives etc.) is the 'device router'. Both SCSI-Fibre Channel Routers are connected together via Fibre Channel point-to-point topology. The benefits of the SCSI-Fibre Channel router as an extender are: - Customer's investments in SCSI peripherals on NIO system are protected. - Customers enjoy improved I/O throughput for JBOD and disk arrays. - Customer has the ability to better manager computer room floor space. - Customers can move peripherals with sensitive data into vaults or secure rooms. - Customers can move critical data to secure campus areas for disaster tolerance. - Customers with NIO system get access to future Fibre Channel peripherals as as they become available. | In the distance configuration, the SCSI-FC router supports: | - HP SureStore E Disk Array XP256 & 12H | - HP SureStore E Disk Systems HVD 10 | - High Availability Storage System (A3311/12/A/AZ) | - HP SureStore E Tape Systems 2/20, 4/40, 6/60 | - DLT4000(FWD only), DLT7000 and DLT8000

## SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training SHAC

A minimum of 2 routers must be ordered for the distance solution. Two 1063MB shortwave GBICs must be ordered for each router. The router requires a fibre cable matched to the GBIC being used between the two routers. 1063MB Short Wave GBICs typically require 50 or 62.5 micrometer multimode fibre optic cable.

## Connecting e3000 to XP512

The diagram illustrates the connection between an MPE/iX e3000 host server and a Disk Array XP512 (Standard Mode). The host server is connected to four rack-mounted, HP-branded SCSI-FC Bridges via SCSI HBA cables. Each bridge is then connected to the disk array via FC optic cables.

- For MPE/iX limit, to support more than 8 LUNs per TID on a FC disk array, multiple SCSI-FC Bridges can be connected in parallel to spread the I/O over multiple SCSI HBAs in the e3000 host server
- Hubs and switches are not supported

HP logo: **HP**  
Stress-Free Storage...Guaranteed **N S S O**

Port Host Mode should be set to 08  
SHAC is not supported for the XP256

## SCSI-FC Router • Data Transfer Rates

### Data Transfer Rates

- SCSI-2 Fast Wide HVD (20MByte/sec) for host or target protocol
- Ultra Wide SCSI-2 (40MByte/sec) not tested or supported
- FC-AL 2-node loop topology supported
  - Other topologies not tested or supported







N S S O  
Stress-Free Storage...Guaranteed



# SCSI-FC Router • Documentation

[//docs.hp.com/mpeix/#doc](http://docs.hp.com/mpeix/#doc)

- *SCSI-Fibre Channel Router Installation and Users Guide*
- *SCSI-Fibre Channel Router Support Handbook*


URL for training docs (for MPE/iX, HP-UX, and Windows NT): | URL:  
<http://docs.hp.com/mpeix/#doc> Two docs are at this site:  
SCSI-Fibre Channel Router Support Handbook  
SCSI-Fibre Channel Router Installation and User Guide

SCSI-FC Router A5814A

Supported HBAs, *no HA Support*


MPE/iX 6.5	Patch MPE LX84	HPPB (NIO) 28696A
MPE/iX 6.0	Power Path 2	HPPB (NIO) 28696A
MPE/iX 5.5		not supported
HP-UX 11.0, K series	Patch Level SP48	HSC A2969A
HP-UX 10.20, K series		HSC A2969A
Windows NT 4.0	Service Pack 6	Adaptec AHA-2944UW
Windows 2000		HP equivalent A5252B

- Minimum router firmware requirement is rev 7.60. Firmware is field upgradeable through serial port connection
- Only direct connect FC configurations are supported (no hubs or switches)



N S S O

Stress-Free Storage...Guaranteed




SCSI-Fibre Channel Router microcode revision 7.62 or later (revision 7.60 or later is supported on MPE/iX in the distance configuration)

SCSI-2 Fast/Wide SCSI-2 differential host adapter HP28696A, firmware level 3944

HP-UX supports HSC (A2969A) but not the HP28696A adapter.

NT also supports the FWD Adaptec AHA2944-UW with driver aic78xx.sys SCSI-2 compliant K-Class SCSI-2 compliant differential cabling






## SCSI-FC Router A5814A

Supported HP9000 Servers

### SCSI to FC Bridge Support for XP512


Storage Device	Supported Connections	Supported OS	Supported Servers	Non-HA Support	HA Software Support		
					MD/UX	MC/SG	SG-OPS
SureStore E Disk Array XP512 <sup>1</sup>	Direct Connect	MPEix 6.0, 6.5	E(9x8), K(9x9KS), T(9x9)	Yes	No	No	No
		HP-UX 10.20	A, D, K, R	Yes	No	No	No
		HP-UX 11.0	A, D, K, R, V, N, L	Yes	No	No	No
		W2K	HP and Non HP W2K servers <sup>2</sup>	Yes	N/A		
		NT EE4.0, SP6+	HP and Non HP NT servers <sup>2</sup>	Yes	N/A		

**Notes**  
1. The XP256 is not supported with the SHAC product and there are no plans to support this configuration in the future. The XP256 has SCSI CHIP pairs available for SCSI host connectivity.  
2. XP512 will support all server configurations which meet the following conditions:  
a. The server is listed on the Microsoft HCL under the system/server field; and  
b. The OS/HBA/Driver combination specified is supported by the server or the HBA vendor.



N S S O

Stress-Free Storage...Guaranteed



### SCSI to FC Bridge Support (SHAC)

The following table details the supported configurations for SCSI hosts connected to the SCSI to FC Bridge product (project name: SHAC). The SCSI to FC Bridge product number is the A5814A. Direct Connect configurations refer to direct connections to the Bridge. No SAN components are supported between the bridge and the XP512. The host mode type should be set to 08.

## Properties

- SCSI-FC Router appears to the host as a SCSI device and attaches directly to the SCSI host adapter
- Target device on the Fibre Channel loop appears to the host as a SCSI device
- SHAC cannot run in the XP512 high speed mode, where there are 2 loop ids on a port. The SCSI-FC Router requires only one loop id per port available in the standard mode.



N S S O  
Stress-Free Storage...Guaranteed



## Review of Key Points

- What is SHAC?
- Offerings?
- SHAC Benefits?
- Connectivity Options
- Documentation



N S S O  
Stress-Free Storage...Guaranteed



## Module Wrap



N S S O  
Stress-Free Storage...Guaranteed



# Module 18

## Host Storage Map & XPPF



**ESBU**  
Stress-Free Storage...Guaranteed



## Module Agenda



➤ HP Storage Map

➤ XPPF



➤ Wrap Up





**ESBU**  
Stress-Free Storage...Guaranteed




# Host Storage Map

- Host Storage Map allows you view and manage host system files and their locations on devices associated with disk arrays
- Runs on a host system independent of the Command View web-based application.



Stress-Free Storage...Guaranteed



## Student Notes

The HP Surestore Host Storage Map XP application allows you view and manage host system files and their locations on devices associated with disk arrays. This application runs on a host system independent of the Command View web-based application.

Host Storage Map provides you with a graphical map to the device where data is stored on the disk array. This host-mapping application is linked to Command View, where you can also access Device Management for the disk array.

Invoke the Host Storage Map application whenever you want to view how the logical files are mapped to the physical devices, and to see how these files are distributed over the array. Viewing file distribution is particularly useful to ensure that any host that has dual paths is mapped through different hardware. You invoke the Host Storage Map when you want to study hardware performance when files are distributed over various arrays.

### Installing Host Storage Map

Host Storage Map is an optional utility in Command View. This application can run on host systems with various operating system environments, such as HP-UX, Windows NT, AIX, and Solaris.

To activate Host Storage Map, you must first download and install this feature from the Command View Management Server, using the Support page.

## HP Host Storage Map

- Ships with Command View XP
- Definition: Application that allows you to view and manage host system files & their locations on devices associated with the disk arrays
  - Runs **in-band** on connected host on HP-UX, Solaris, NT, AIX boxes (Command View independent)
  - Data collection/display is based on XPINFO.EXE
  - Provides a graphical map to the device where data is stored on the disk array
  - Requires Java Runtime Environment (JRE)



ESBU  
Stress-Free Storage...Guaranteed





NOTE: Although the user's guide goes through the entire installation process, the guide does not mention that JRE is a requirement for installation. Using a browser, please reference the SUPPORT tab from within Command View XP server to obtain the JRE required for the Host Storage Map.




# Starting Host Storage Map

- Execute the HostStorageMap command from the host system command line interface



Stress-Free Storage...Guaranteed



## Student Notes

To start Host Storage Map, you execute the HostStorageMap command from the host system command line interface.

The HostStorageMap command also provides optional arguments, which enable you to view additional information from the command line interface.

Host Storage Map does not require Command View to be active. The Host Storage Map application running on the host system formulates the required URL for each array that is connected to the host and managed by Command View.

To start Host Storage Map:

1. Use the command line or command dialog box of your operating system environment.
2. Enter the HostStorageMap command:  
HostStorageMap

The Host Storage Map application starts and the Startup window appears. By default, this command displays the Startup window. From the Startup window, you can enter configuration information to view and manage host storage resources. Commands selected from the Startup window display additional information in web browser pages.

## Syntax

The syntax of the command is

HostStorageMap [ optional arguments ]

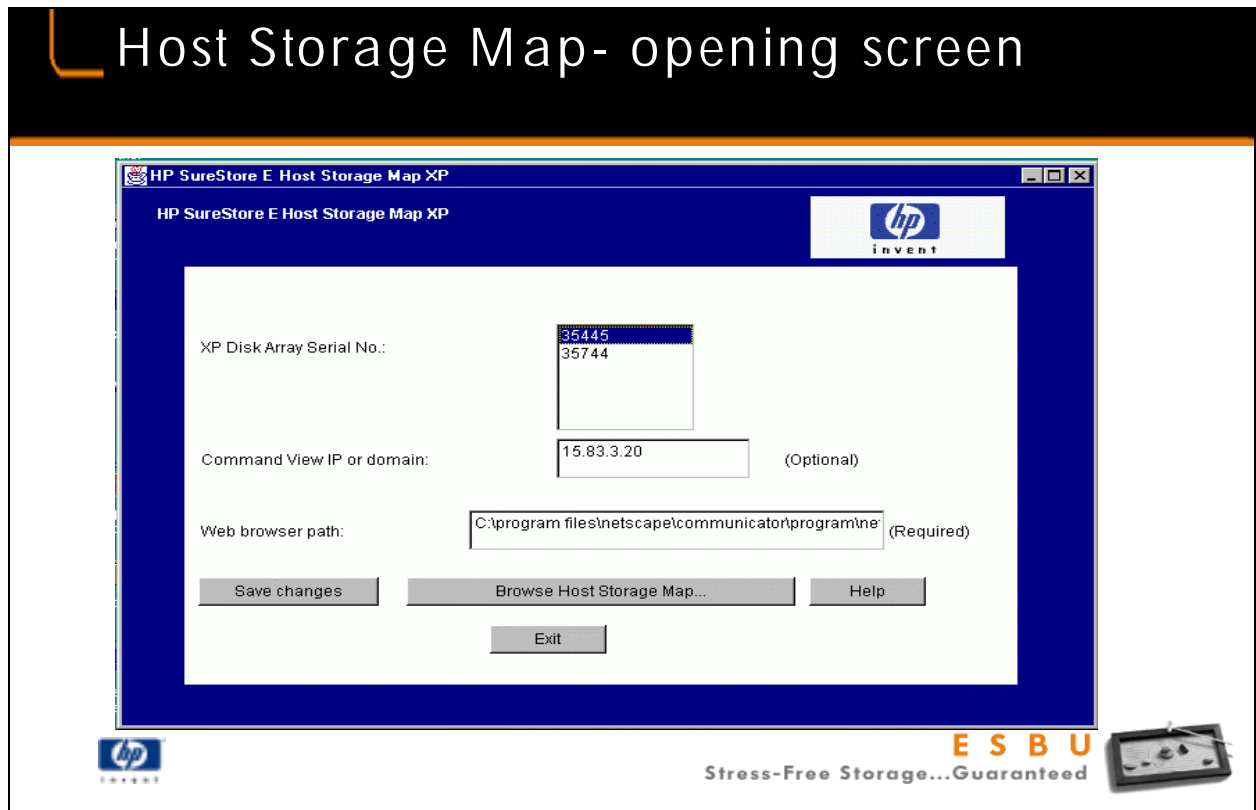
## Optional Arguments

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Host Based Software For the XP Array

- a Attempts to report information even for devices that are not OPEN type devices. For these devices, it reports only the target ID, LUN ID, product ID (type), device size, and firmware revision.
- b This command displays array information from the latest xpinform files that reside in the same directory with this command. The information displayed is not always current. Therefore, use this command option when you are confident that there has been no change since the last HostStorageMap operation. You can use the Browse Host Storage Map button to view more detailed configurations of disk arrays.
- c Displays information about Continuous Access and Business Copy volume designation. Data fields include:
  - Subsystem number
  - CT group number
  - CA volume designation
  - BC volume designation for each of the three MU Numbers
- d[char] Display a single line of data per logical device. Each data field is delimited by a comma or by a single char if specified. The data fields appear in the following order:
  - device file
  - target ID
  - LUN ID
  - port ID
  - CU:Ldev
  - type
  - device size
  - serial#
  - code revision
  - subsystem
  - CT group
  - CA volume
  - BC0 volume
  - BC1 volume
  - BC2 volume
  - ACP pair
  - RAID level
  - RAID group
  - disk1
  - disk2
  - disk3
  - disk4
- e Provide a name cross-reference of all LUNs connected to the array.
- f raw\_device Report information only for the specified logical device.
- h Display the Help message.
- i Display the identification information. The data fields include:

SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Host Based Software For the XP Array

- LUN type (product ID)
  - device size
  - serial number
  - firmware revision
  - l Sort output by LDEV number.
  - p Sort output by Port number.
  - r Display RAID configuration information. Data fields include:
    - ACP pair number
    - RAID level
    - RAID group
    - disk mechanisms
  - u Create report using updated ioscan list. By default, the report is created using the kernel image of the last ioscan.
  - v Display the xpinfo version. The -v option has precedence over all other options. If you specify the -v option along with other output options, only the xpinfo revision displays.
- Verbose Format
- If you do not specify options –i, –c, –r, or –d, the results are displayed in verbose format.



## Student Notes

### Fields

**XP Disk Array Serial No.** Displays serial numbers that correspond to the disk arrays.

**Command View IP or domain** Displays the network address or name of the system where Command View resides. This value is required to hyperlink devices to the Command View management station. If the IP address of Command View is provided but the device is not added into Command View through Device Administration, the launch of Command View will fail. If a value is not provided for the Command View IP, no hyperlink is created.

**Web browser path** - Displays the location of the web browser. You must provide a path to one of the browsers in this field.

### Buttons

**Save changes** - Saves all information you enter into the fields.

**Browse Host Storage Map** - Opens the web browser and displays a list of device files and related information.

**Help** - Provides access to the online Help information about how to use this application.

**Exit** - Closes the Host Storage Map page. If the web browser is still open, you must close it before choosing the Exit button.

The host storage map requires the Java Runtime Environment (JRE) to be installed. Upon launching HostStorageMap.Bat from the

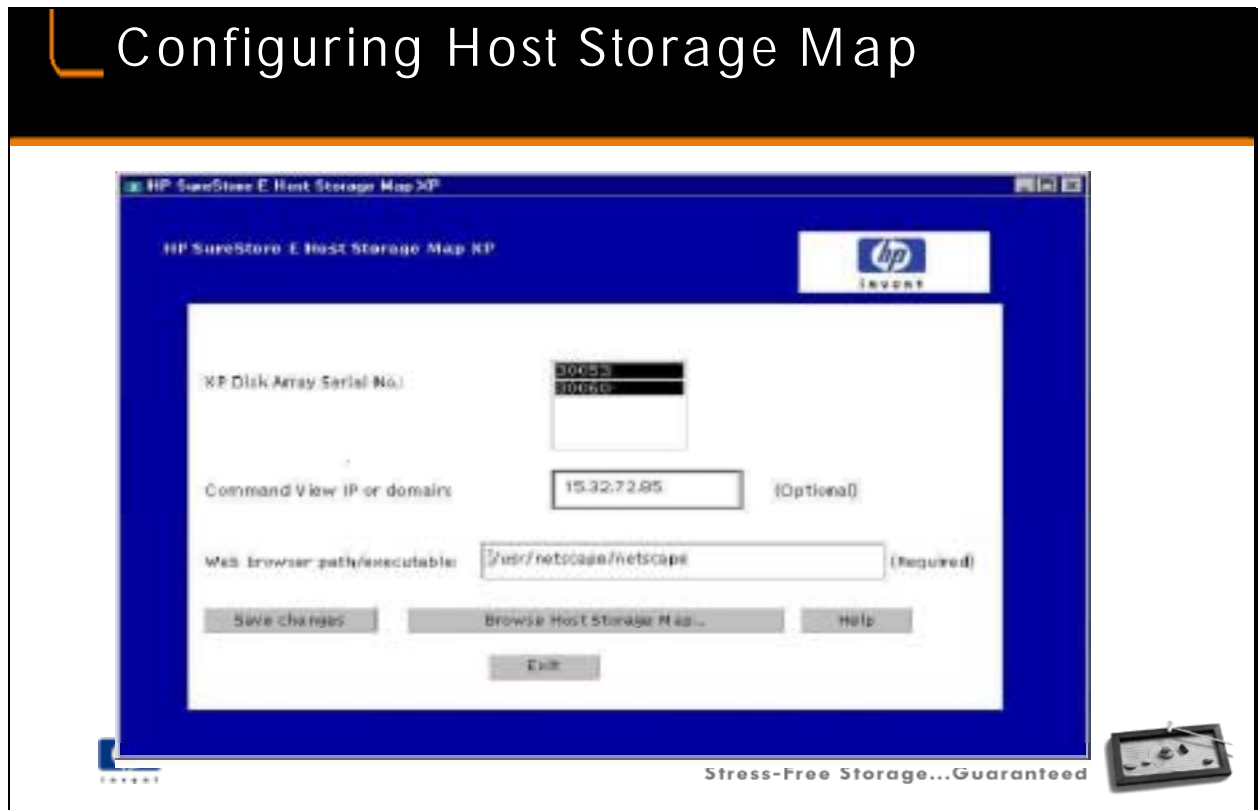
SR26013 HP SureStore XP Family Technical Pre-sales HP Channel Partner Training  
Host Based Software For the XP Array

\HPSS\HostStorageMap directory, the program runs XPINFO.EXE in the background and displays the screen shown above.

Field Definitions:

XP Array Disk Serial No: The array serial numbers of all arrays found during XPINFO.EXE's scan of the in-band connected arrays (see next slide for information on XPINFO).

Command View IP or Domain: The IP address of the machine that is running Command View XP server. If filled in, the program should allow you to click on the array's serial number (shown in following slide) and launch into Command View XP for management of that array.



### Student Notes

Configure the Host Storage Map application by entering values in each of the fields on the Startup window. The configuration values you enter and save remain intact the next time you open Host Storage Map.

Whenever you want to change the Command View IP address for any particular array, you must first close the web browser and then return to the Startup window.

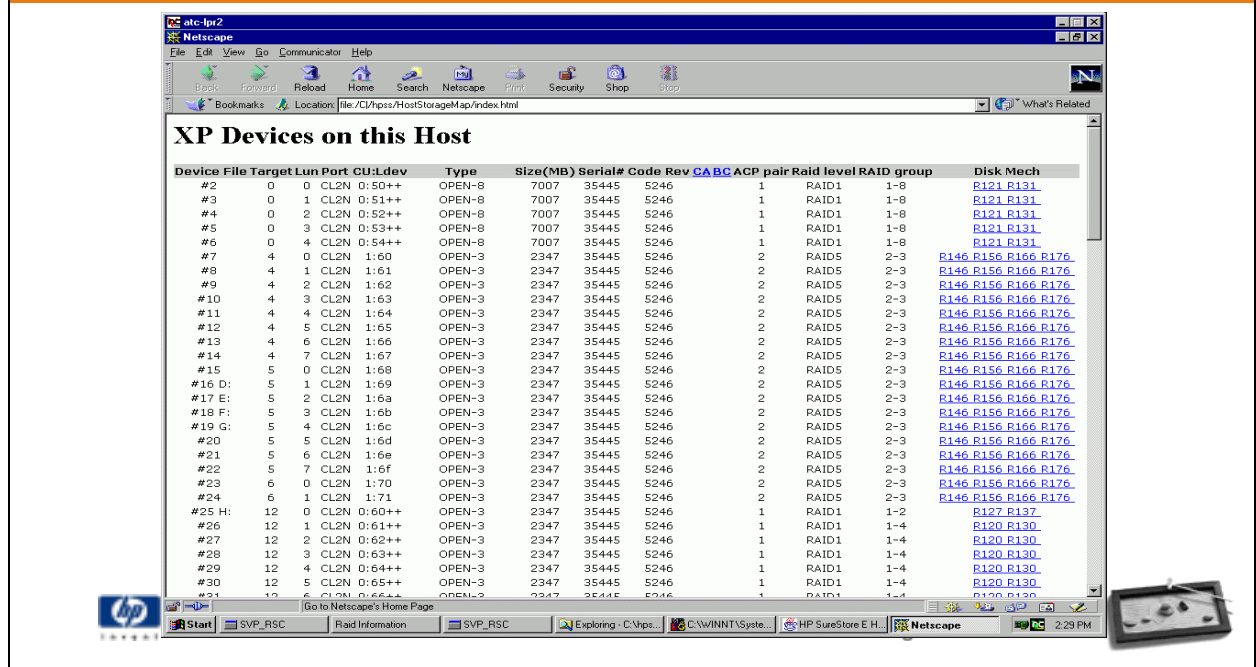
When the Host Storage Map application initially starts, it searches for connections to disk arrays. Host Storage Map then assembles a list of the system files and their locations that are stored on all devices that are associated with the disk arrays.

#### Adding an IP address

The Command View IP address does not appear for all of the disk arrays in the list the first time. You can add an IP address for the disk array by following these steps.

1. Select one or more serial numbers from the XP Disk Array Serial No. list box.
2. Type the IP address or the host name of the Command View that manages the selected array in the Command View IP or domain field.
3. Click on the Save Changes button. The system saves the values you enter for configuration.

## Browse Host Storage Map UI



### Student Notes

When you click on the Browse Host Storage Map button from the Startup window, the web browser is launched and displays the XP Devices on this Host page.

In the CU:Ldev column, ++ indicates that there are multiple paths to the device. From this page, you can view any arrays and related devices for which you have assigned an IP address.

### Links

To view additional information, select any of the underlined links. The information displayed is determined by the item you selected:

Serial # Displays the Command View login page for the selected device.

CA and BC Displays a summary of details for all of the host device files.

Y Displays details for a single device file.

Disk Mech Displays an illustration of the disk array configuration.

Displayed above is the first screen of the host storage map.

# BC/CA Details

The top screenshot displays a 'BC/CA Details' window with a table listing various volumes and their properties. The table has columns for 'Device', 'LUN', 'Volume', 'LUN Name', 'LUN ID', 'LUN Size', 'LUN Type', 'LUN Format', and 'LUN Status'. The bottom screenshot shows a similar window but with a different set of data, likely for a single device file.

HP  
Stress-Free Storage...Guaranteed

## Student Notes

Example This first screen shows a list of all CA and BC volumes.

Example This second screen shows BC/CA details for a single device file.

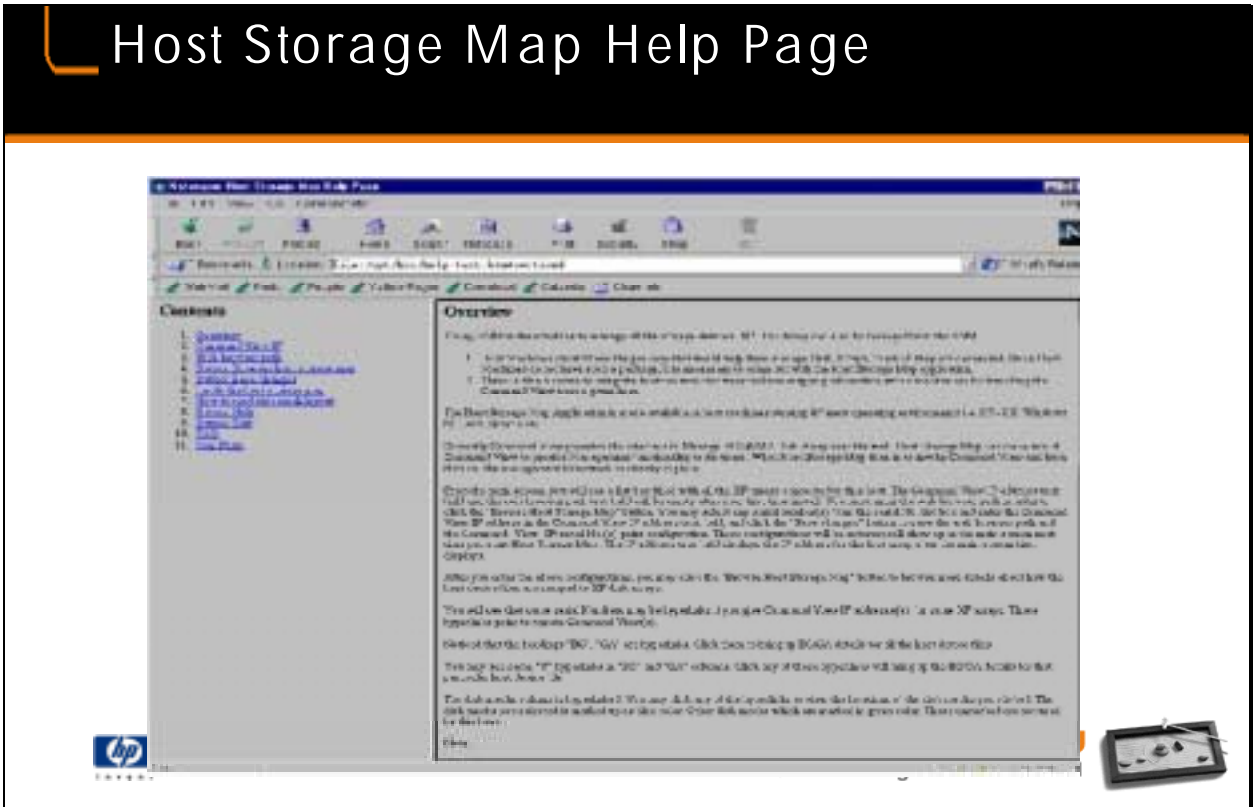


## Disk Array Layout Map




By clicking on a disk mechanism, the storage map displays the disk canister number that was selected. In this example, R146 was selected.  
Student Notes

You can see how files are distributed among the disks in the arrays you monitor. From the devices list, select a Disk Mech item to view the Disk Array Layout Map. An image illustrates how all relevant files are distributed.







## Student Notes

You can select Help from the Startup window to get information about how best to use the Host Storage Map application.

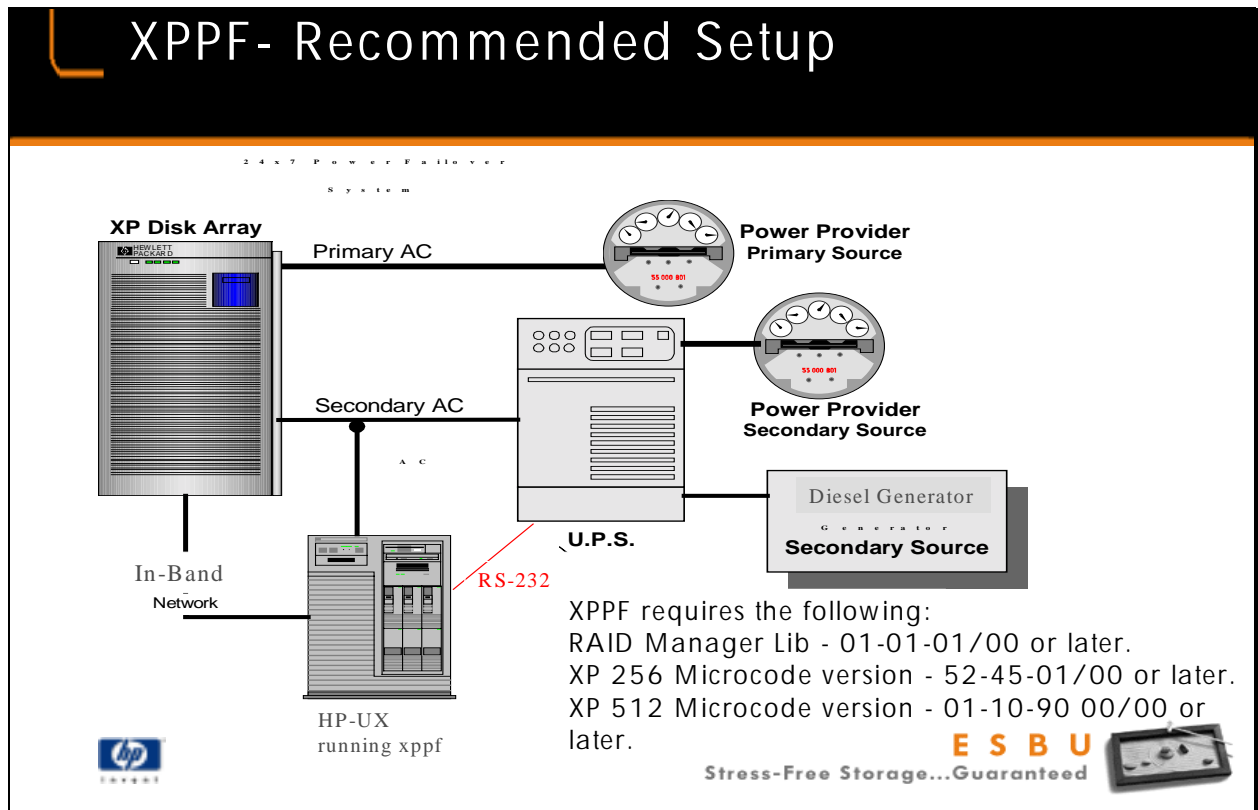
XPPF

**Power Fail Cache-** destage of cache to disk in event of immanent power failure

- XPPF Program File (free!)  
[http://essd.boi.hp.com/essdatc/hw\\_prods/xp256/xppf/xppf.zip](http://essd.boi.hp.com/essdatc/hw_prods/xp256/xppf/xppf.zip)
- XPPF White Paper  
[http://essd.boi.hp.com/essdatc/hw\\_prods/xp256/ups\\_pf/ups\\_pf.pdf](http://essd.boi.hp.com/essdatc/hw_prods/xp256/ups_pf/ups_pf.pdf)



The internal architecture of the XP256 and XP512 has no single point of failure, but the XP unit(s) is reliant upon AC power. In the event of total power loss to the array, the XP will maintain data in cache via internal battery units for a total of 48 hours maintain shared memory for 96 hours. If power is not restored data held in cache and system configuration information will be lost. System configuration information can only be restored by reloading from floppy (assumes disk #31 has been updated with each change to the array) or from flash memory when the system is restarted. To extend this interval before data loss, an Uninterruptable Power Supply (UPS) is required. However, in the event of a total power failure, the array is now at mercy of the UPS thus compromising cache data. XPPF is an HP program provided free of charge to customers that will help eliminate the loss of cache data in the event of eminent power outage.



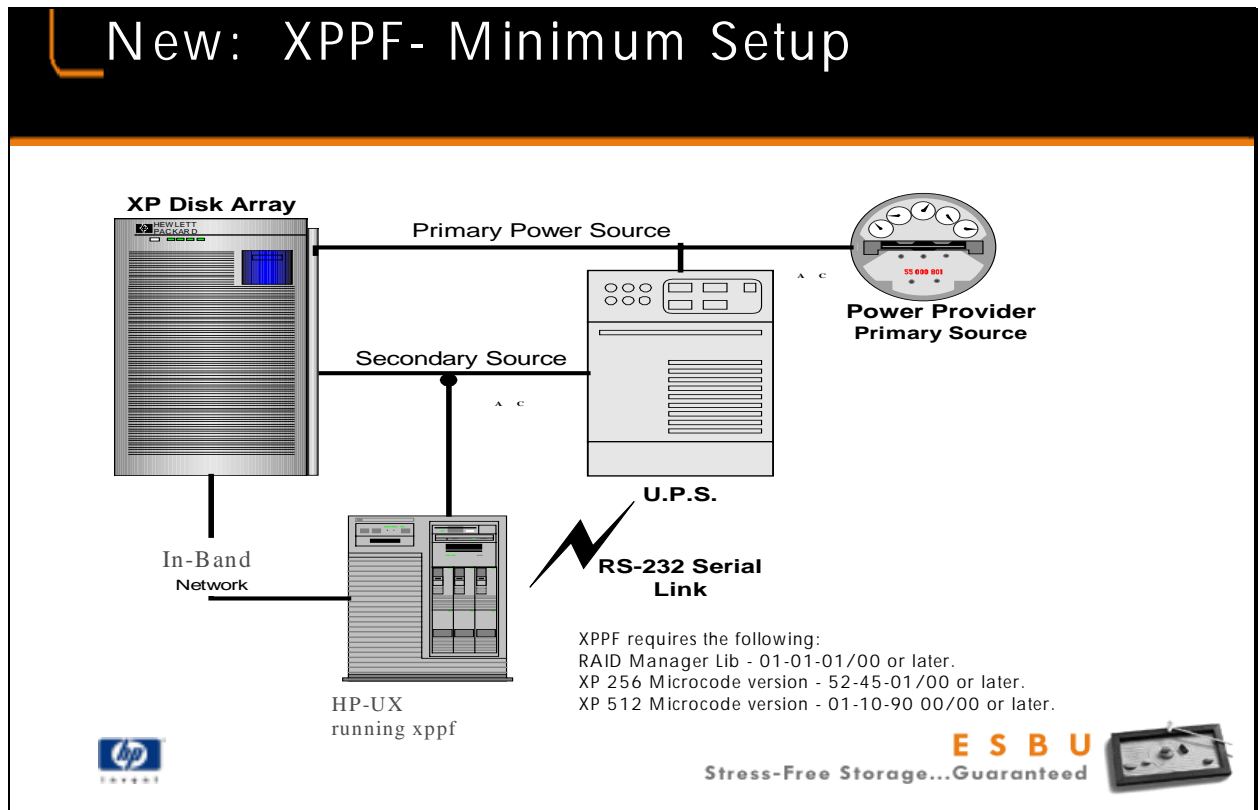
To ensure no loss of data, HP developed Power Fail Cache. This implementation requires the following:

#### UPS

HP-UX in-band connected host connect to both the UPS and the array. This host will run a free program called "xppf". XPPF received notification from the UPS of eminent power outage and initiates the following from the host to the array:  
destaging operation of all writes held in cache  
force all successive writes to bypass cache and write straight to disk. The UPS should allow up to 30 minutes uptime for the array to complete all destaging. If power is restored during this destaging operations, xppf will set the array back to normal, non-powerfail operation.

RS232 connection from the host to the UPS. This serial connection allows the UPS to send notice to the host of an eminent power failure is occurring.

Any UPS that has software that allows command file execution for events will be able to use the xppf program (see American Power Conversion Corp <http://www.apcc.com>).



Illustrated above is the minimum configuration setup possible for use with the xppf program. Although this is not the full power redundant configuration that HP would recommend, xppf would be a smart choice in this implementation.

## Module Wrap-up



**ESBU**  
Stress-Free Storage...Guaranteed

