



Session 3968

Troubleshooting continuous access EVA solutions



Don Fraser
Disaster Tolerant Solutions & Test Architect
NSS CA EVA Systems Integration & Test
Hewlett-Packard

© 2004 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice





Agenda

- Part one – Is it supported?
 - Definitions & overview
 - Supported configurations

- Part two – Troubleshooting
 - Understand what tools are available
 - Understand how to use those tools



A quick review of terms and concepts

→ replication direction




virtual disk



copy set member



log disk



disk group




EVA controller



data replication groups:

- source DR group



- destination DR group




managed set

EVA without replication

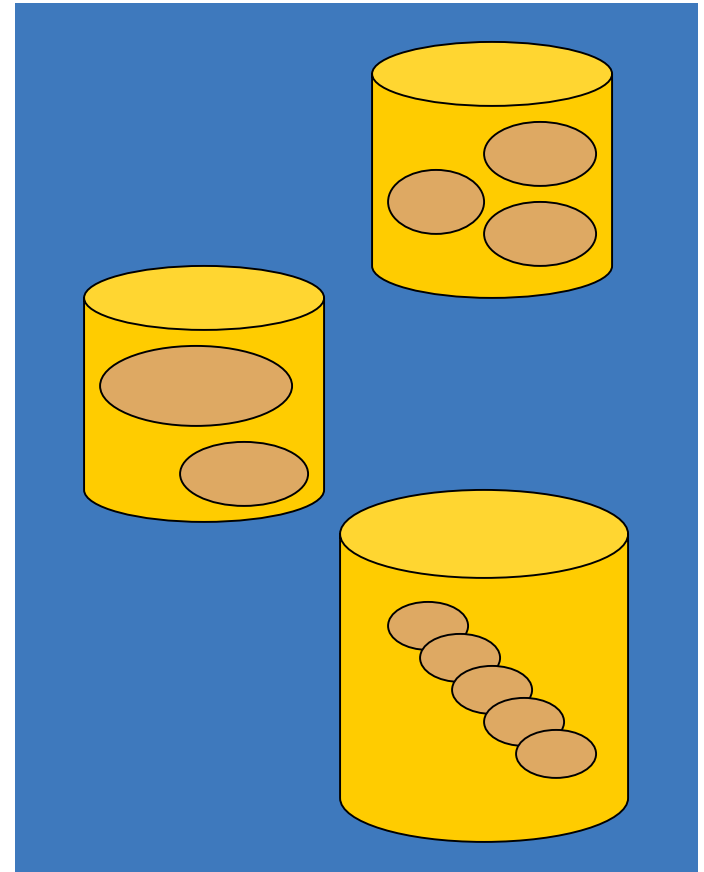
1 EVA



3 disk groups



2-5 virtual disks within
disk groups



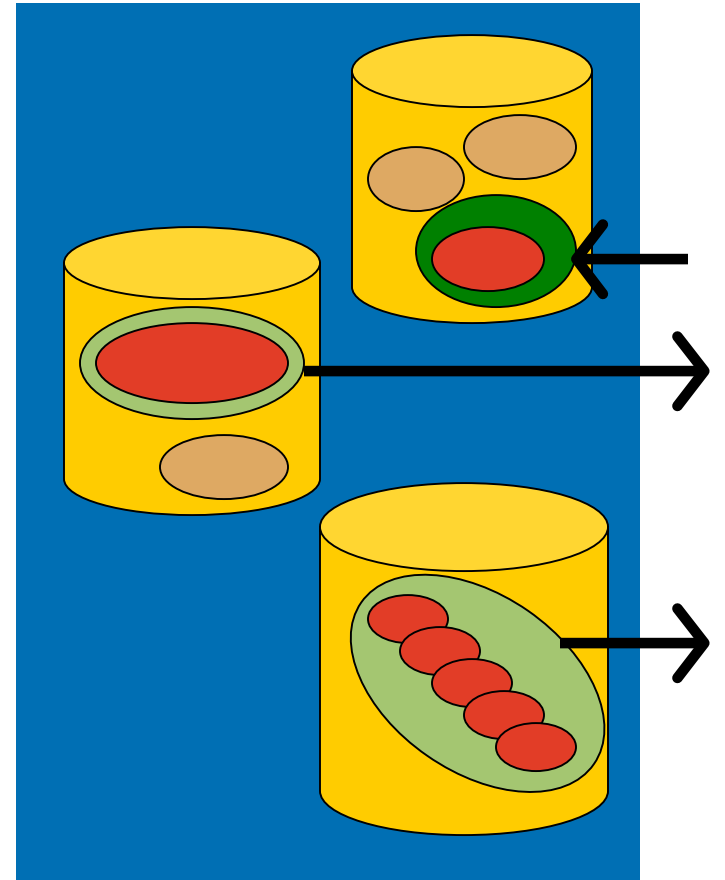
EVA with replication

3 DR groups

• 2 sources 

• 1 destination 

1-5 copy sets 

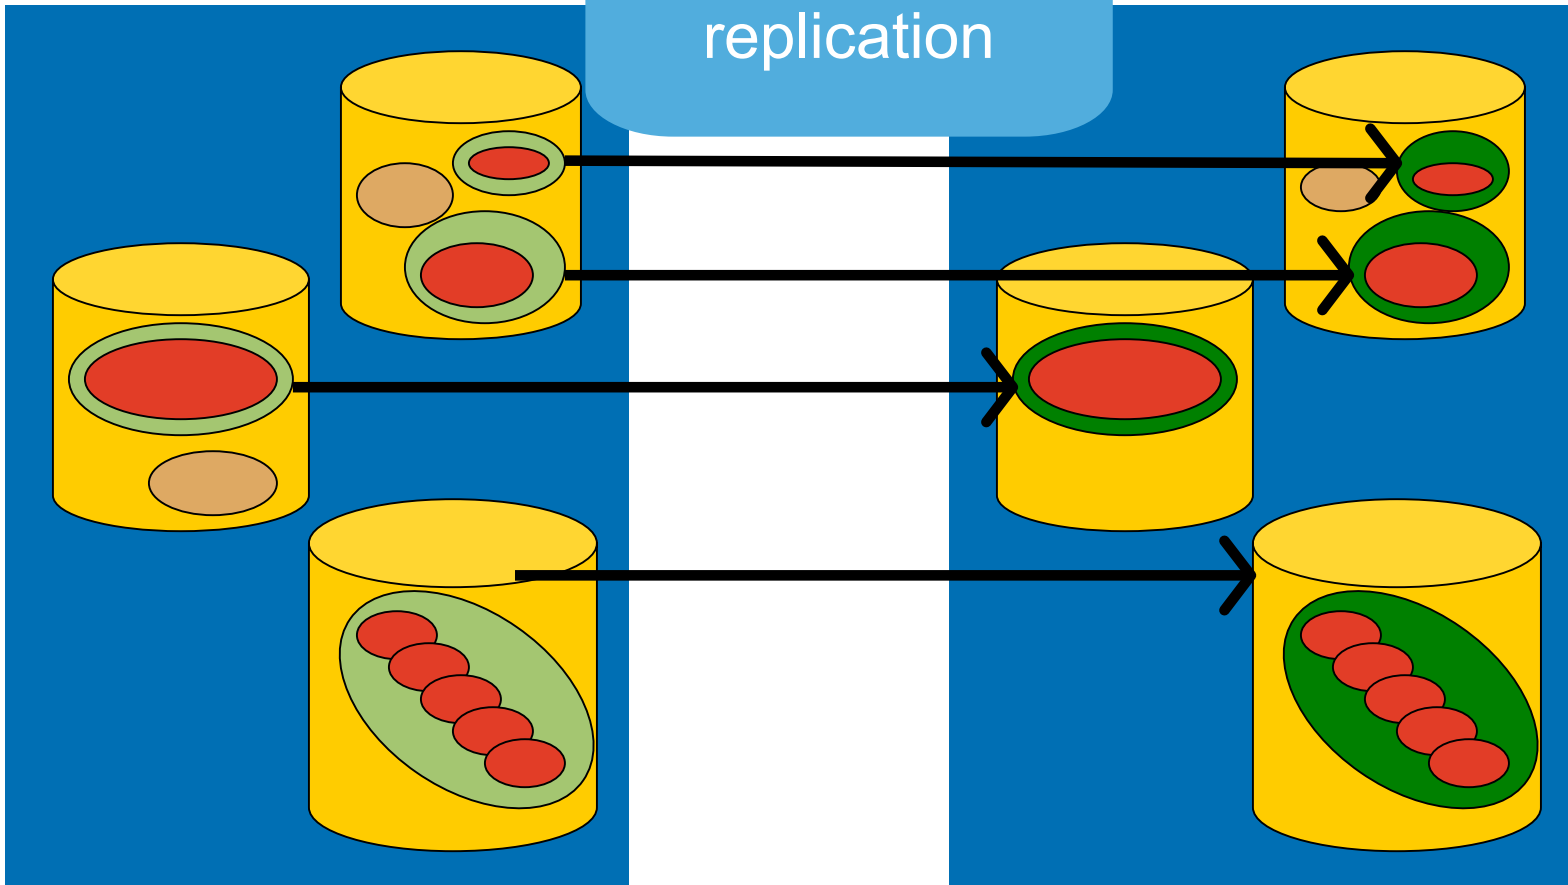


source

destination



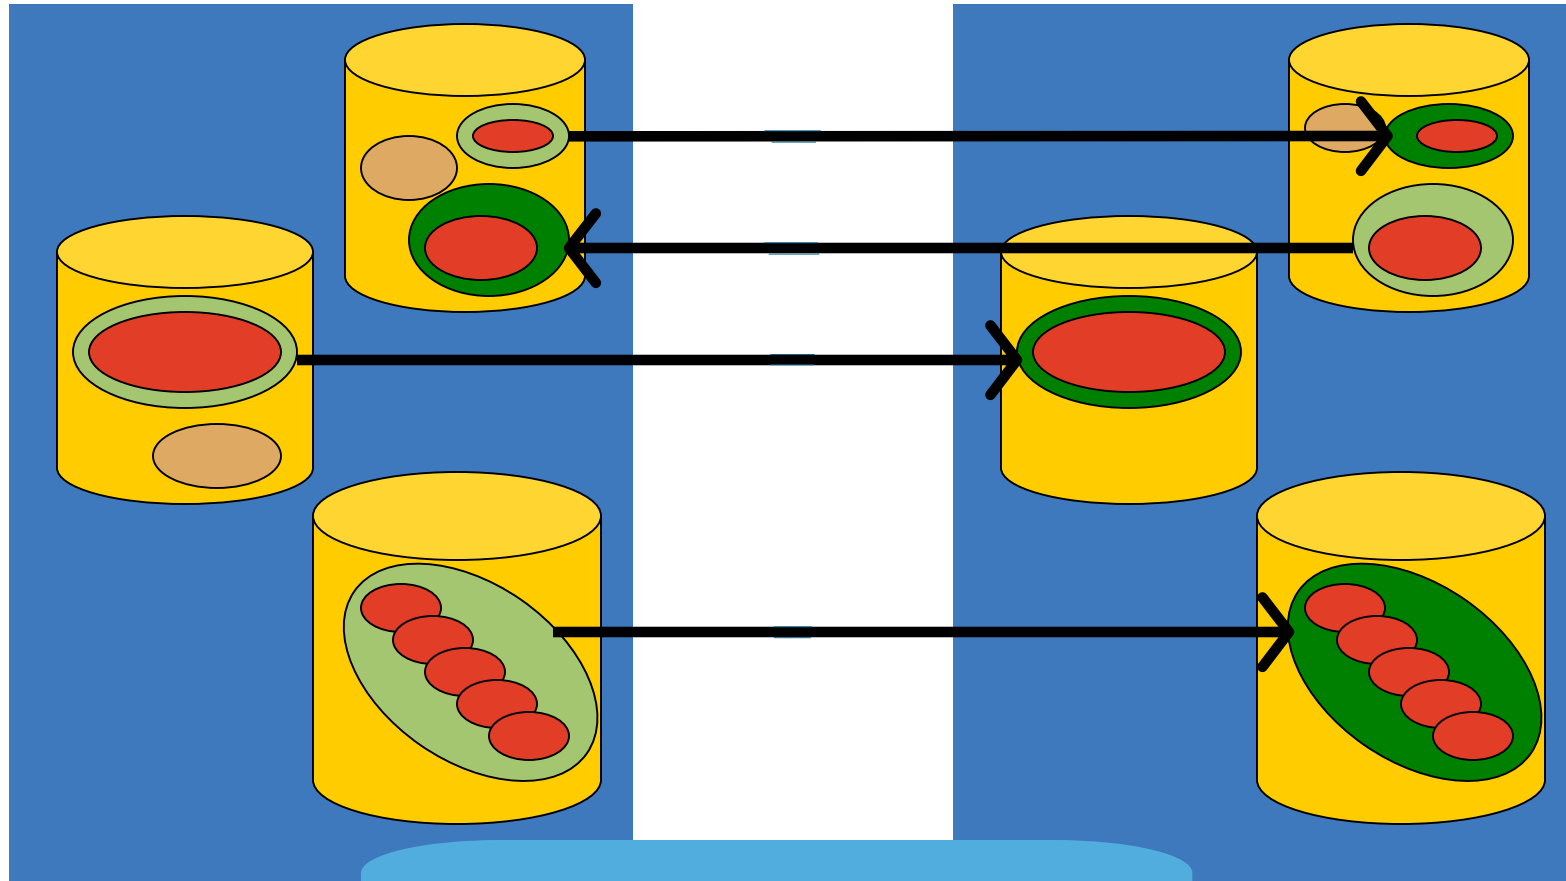
unidirectional replication



2004


Solutions and Technology Commitment & Experts

Both are source & destination



bidirectional replication

Managed set



- Collection of DR groups
- For purpose of managing together

Example: managed set of all DR groups sourced on this EVA



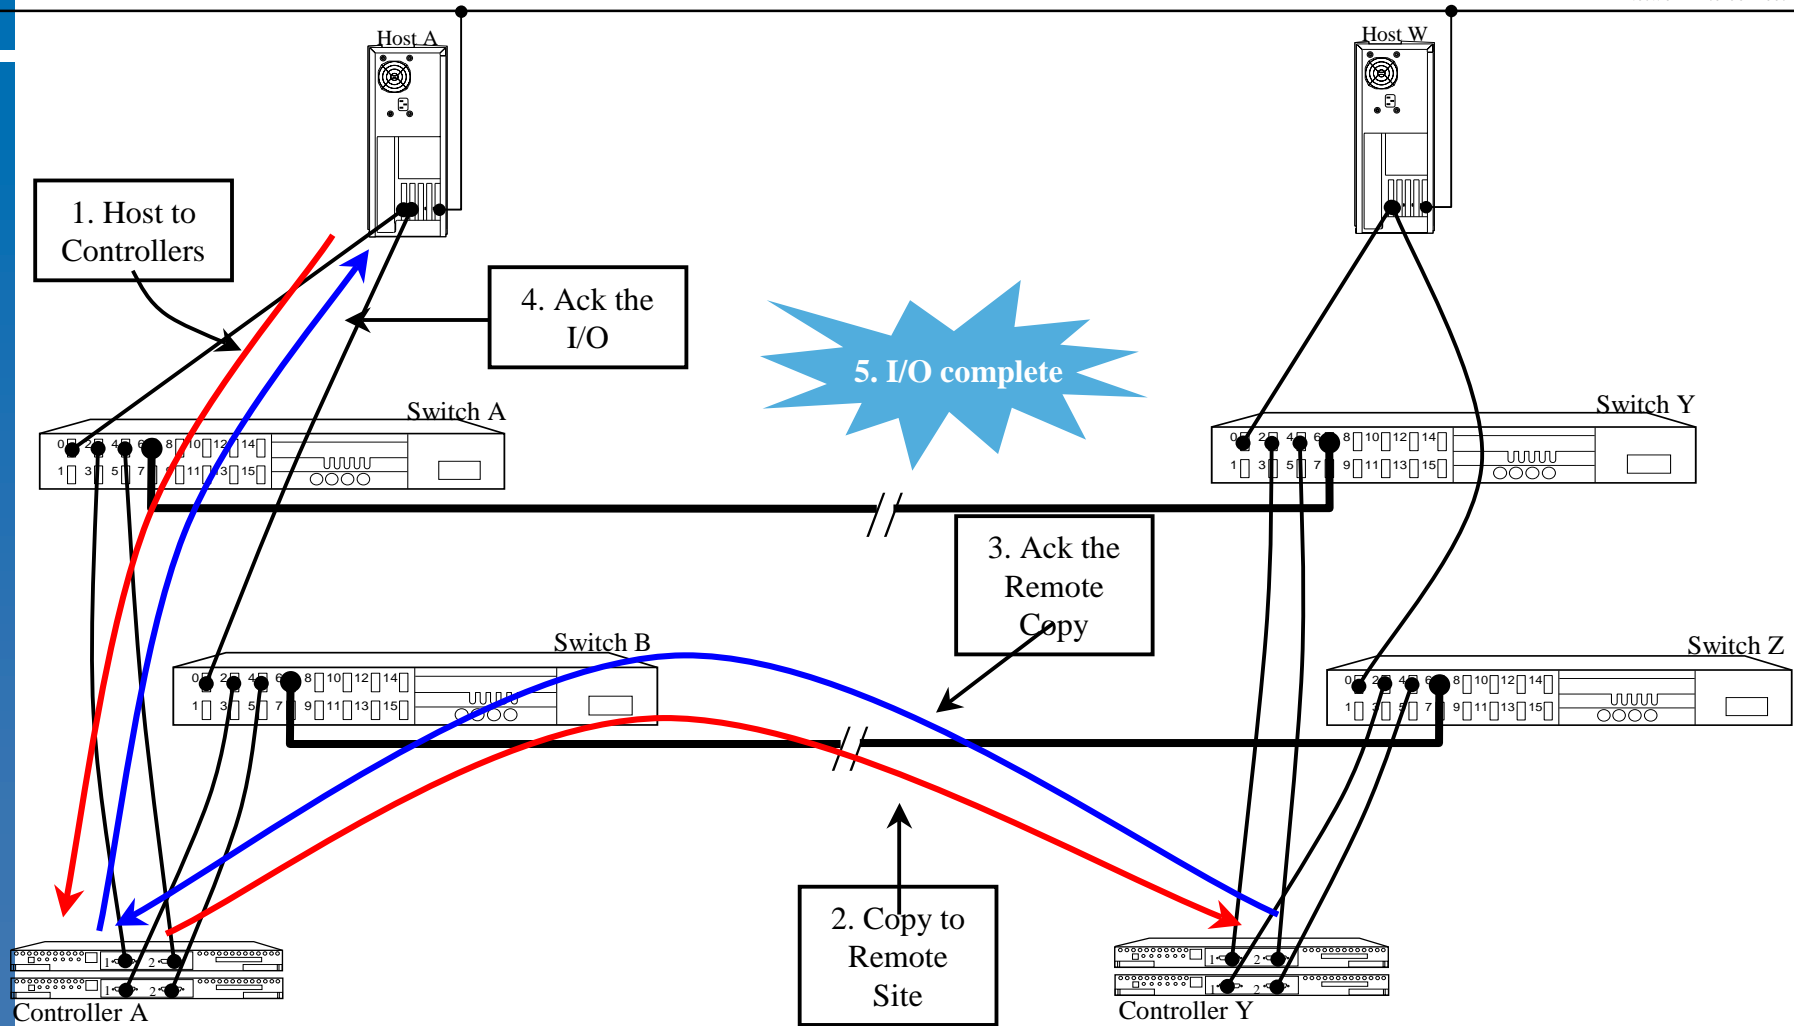
Synchronous replication

- Application can use either synchronous or asynchronous I/O to array and is independent of replication mode
- Source controller waits for acknowledgment that replication write has reached cache of destination controller, before sending acknowledgment to host
- Safest method of replication
 - Write order preserved
 - No completed write is at risk

Synchronous replication



invent
Network Interconnect



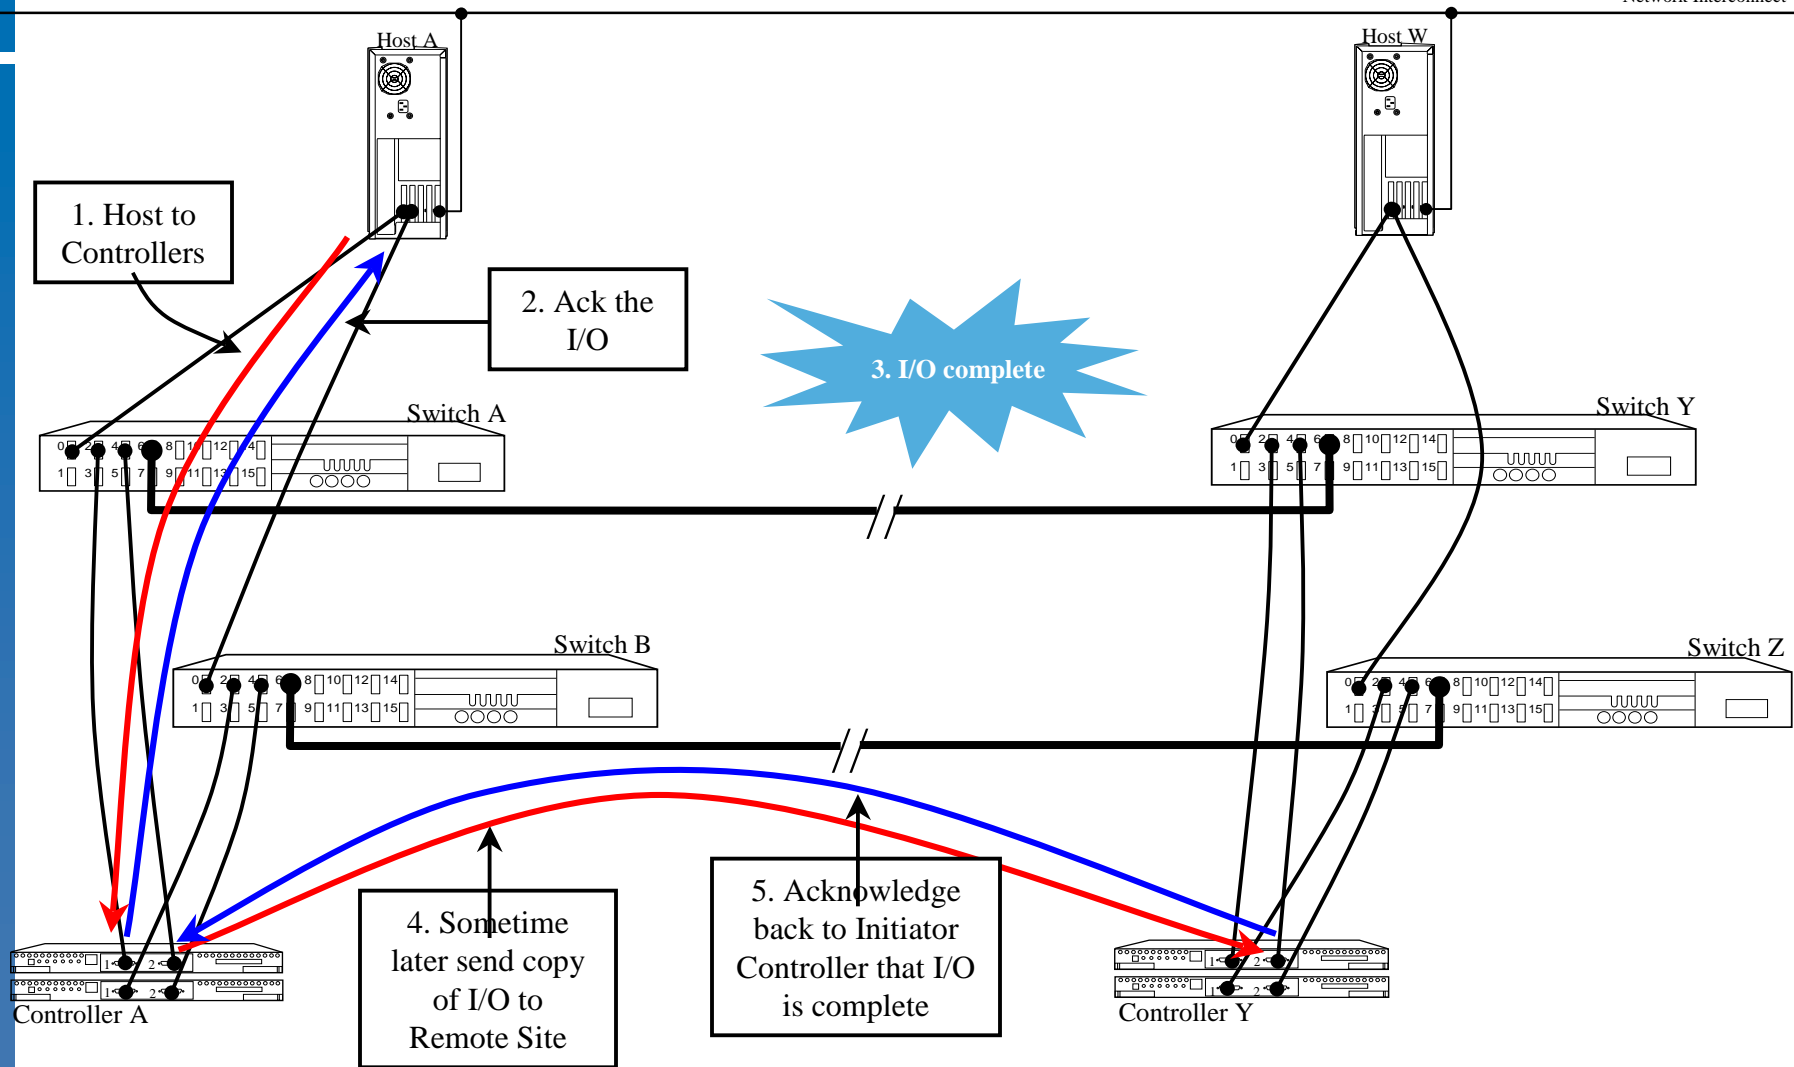
Asynchronous replication

- Application can use either synchronous or asynchronous I/O to array and is independent of replication mode
- Controller acknowledges the host write as complete, as soon as the write is in local cache
- Not the safest method of replication
 - Must be willing to lose data because I/O in flight may be lost if link broken and not restored

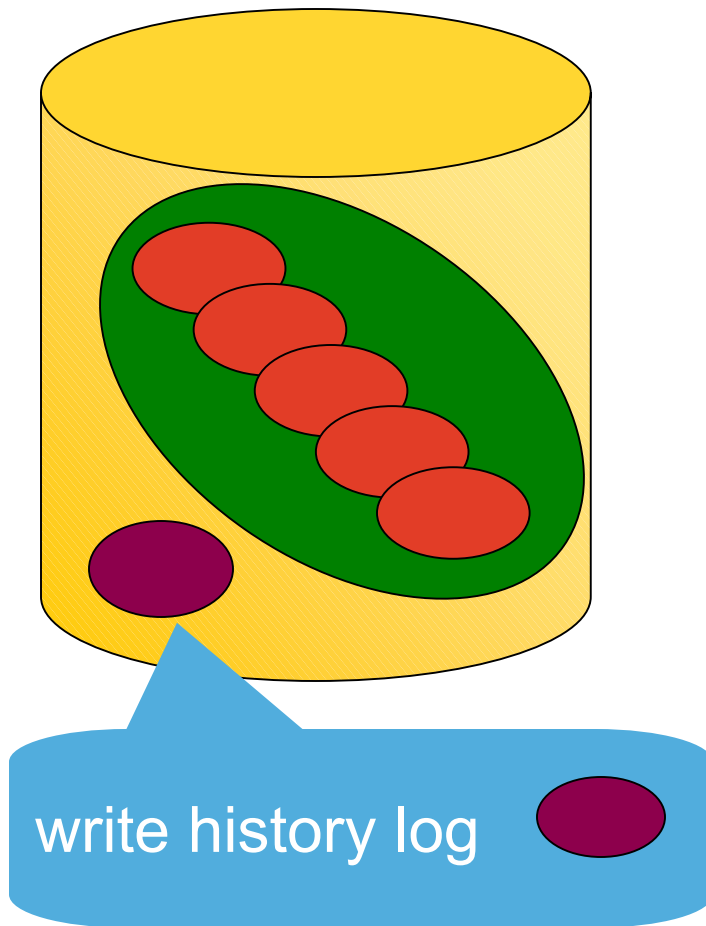
Asynchronous replication



invent
Network Interconnect




DR group



- Group of replicated copy sets
 - DR groups on EVAs are equivalent to association sets on HSGs
 - Share write history log
 - In-order application of writes is guaranteed
- When failsafe locked: if unable to write to one member, will not write to any member
- Best practice: use one copy set per DR group unless required by application

Write history log



- Used when all links are down
- Used when DR group is suspended
- Used when destination EVA pair is down
- Best practice: place write history log in dedicated disk group

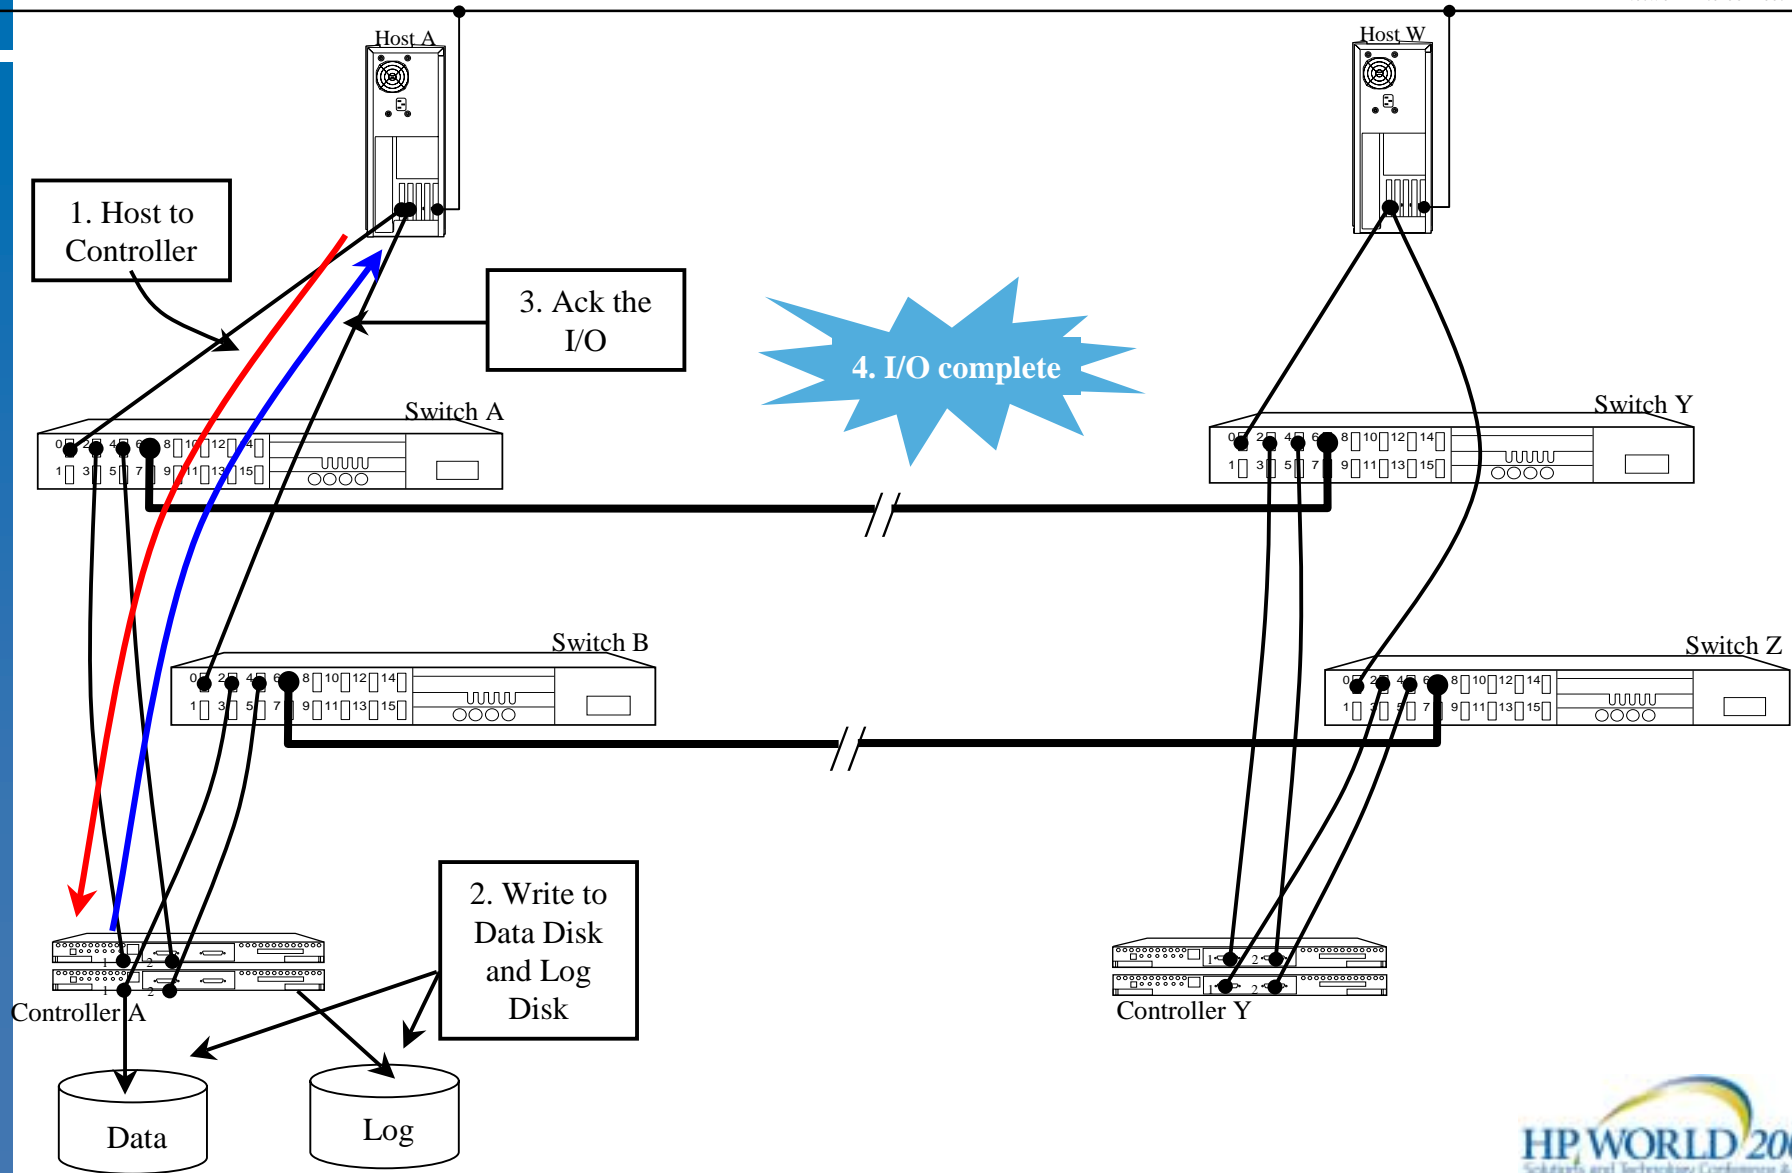
Write history log

- Automatically created at creation of DR group, log is placed into Disk Group with most free space
 - Starts with 136 MBytes of Vraid 1 space
- Expands as needed, to lesser of
 - 2 times capacity of DR group
 - No free space in disk group
 - 2 TB
- Maintains write order across all members of DR group
- If log becomes full, members of DR group are marked for full copy

Write history logging



invent
Network Interconnect




Write history log merge

- Replays contents of write history log in write order to remote site
- Its write for write, not changed blocks, or tracks
- Note: this is not the same write history log used in OpenVMS Host Based Volume Shadowing

Write history log merge



invent
Network Interconnect





Agenda – part one

- A quick review
- Definitions & overview
- Review of supported configurations



Supported solutions

Today, each EVA 3000/5000 array pair is limited to:


- Maximum of 128 copy sets per array pair
- Maximum of 128 DR group per array
- Maximum of 8 copy sets per DR group
- Minimum of 1 disk group per storage system
 - maximum of 16 or
 - number of drives divided by 8, which ever is lower
- Maximum of 512 Vdisks per array
 - 1-GB to 2-TB copy set size
- 16 storage systems per pair of SAN Mgt Appliances
- Maximum one way link latency of 100 mSec
 - Or cable length of 20,000 km or 12,500 miles

Supported solutions (cont)

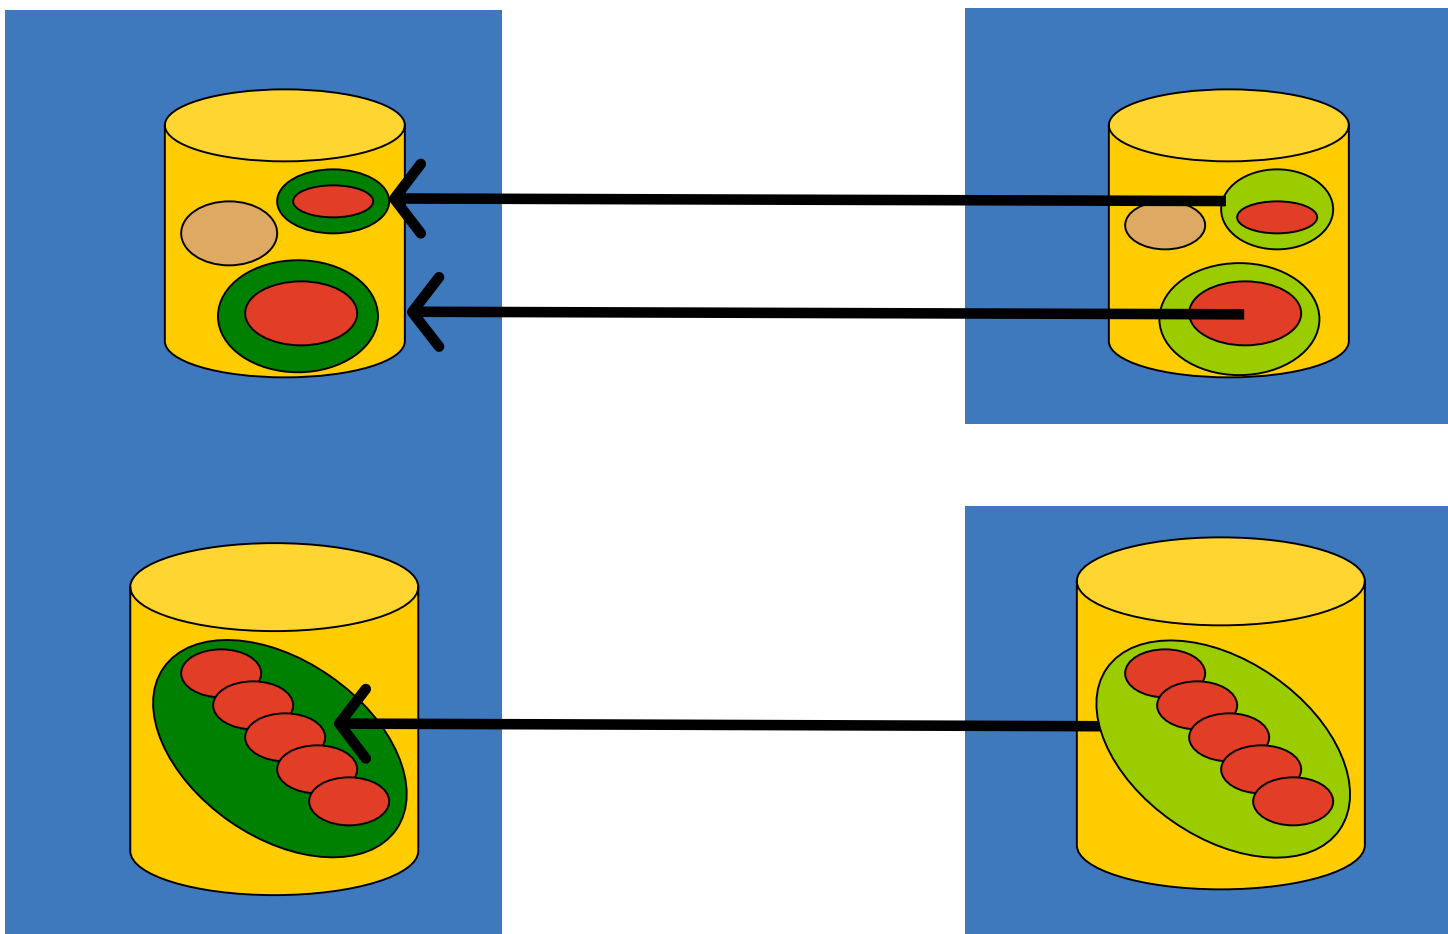
Today, each array pair is limited to (continued):

- Up to two replication relationships per array pair
 - Array to array fan-out or fan-in
 - Different DR groups, LUNs
- Members of a DR group can be from multiple disk groups
- Source and destination disk groups need not have the same geometry (numbers of disks) and should be the same size
- Maximum of 8 snapshots or snapclones per DR group at the source or destination site and up to 7 per source vdisk

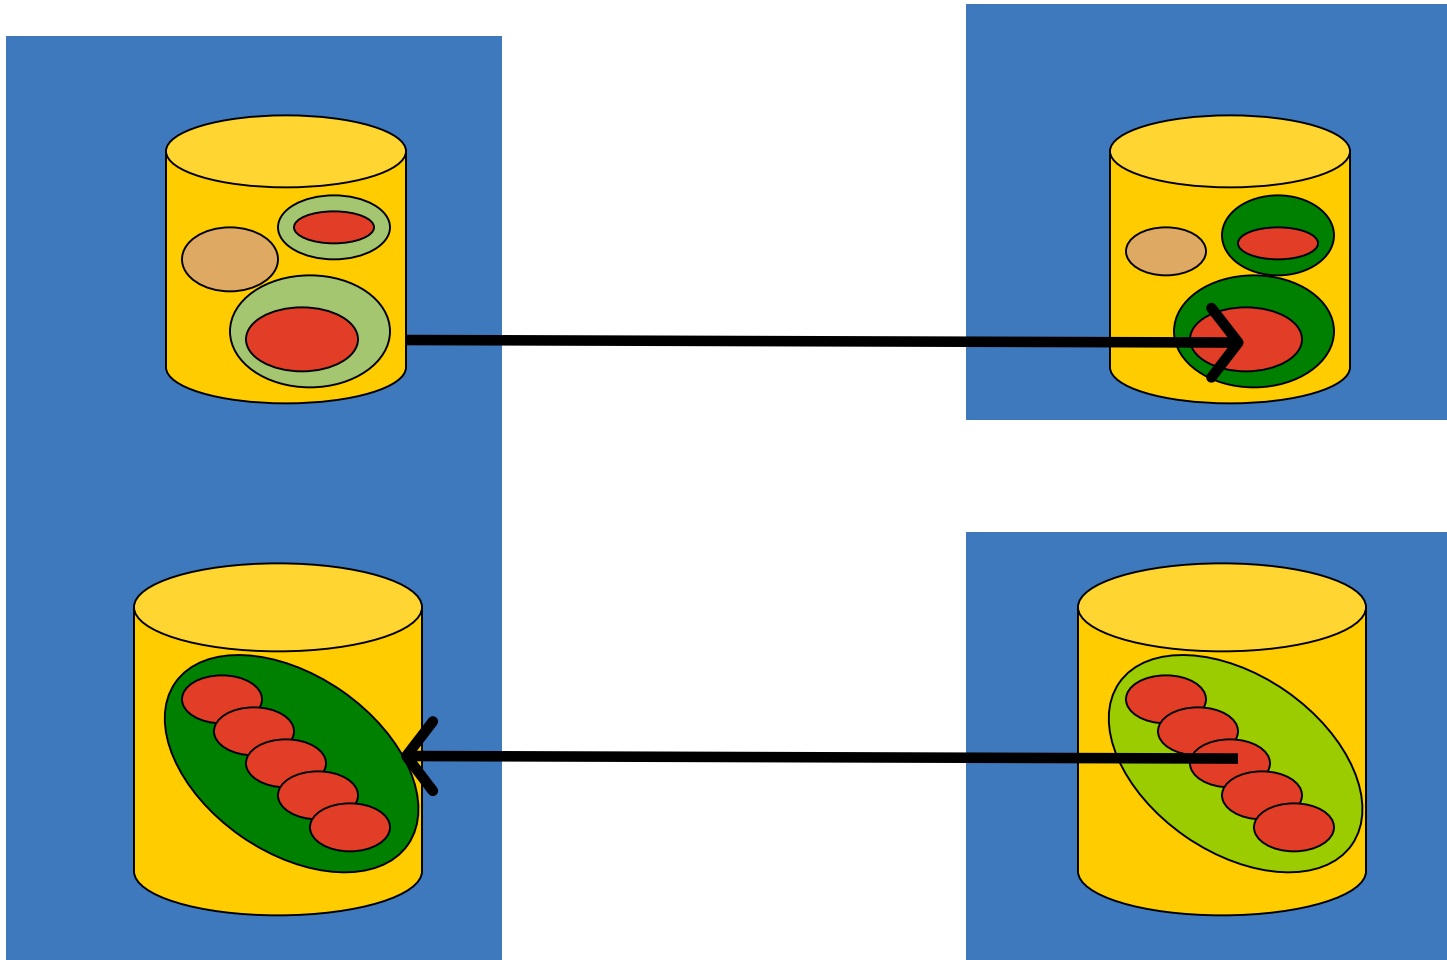
Multiple relationship fan-out



Multiple relationship fan-in



Multiple relationship cascade



Supported solutions (cont)

Today's Supported Operating Systems:

- HP HP-UX 11.0, 11i, 11i v2
- HP OpenVMS 7.2-2, 7.3-1, 7.3-2
- HP Tru64 5.1a, 5.1b
- IBM AIX 4.3.3, 5.1, 5.2
- Microsoft Windows: NT; 2000; 2003 (32 & 64 bit)
- Novell Netware 5.1, 6.0, 6.5
- Red Hat AS 2.1 (32 & 64 bit)
- SUN Solaris 2.6, 7, 8, 9
- SuSE SLES 7, 8 (32 bit & 64 bit)
- United Linux V1.0 (32 bit)

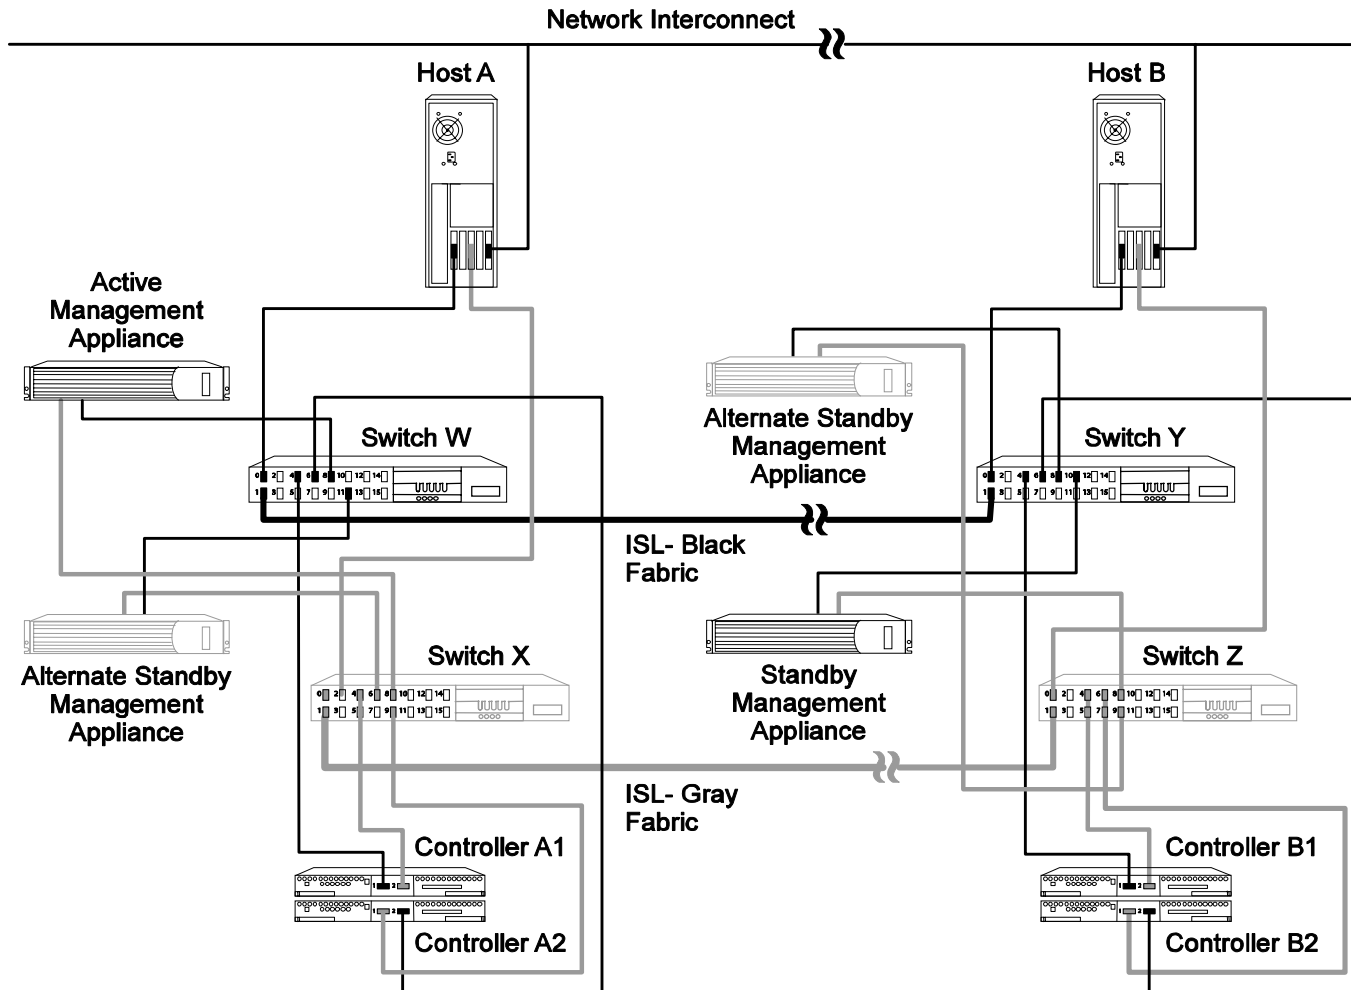


Supported solutions descriptions

- Basic solution
 - Direct fiber
 - Using WDM
 - Using FCIP
 - Stretched cluster
- Optional solutions
 - Single HBA
 - Single switch
 - Single fabric



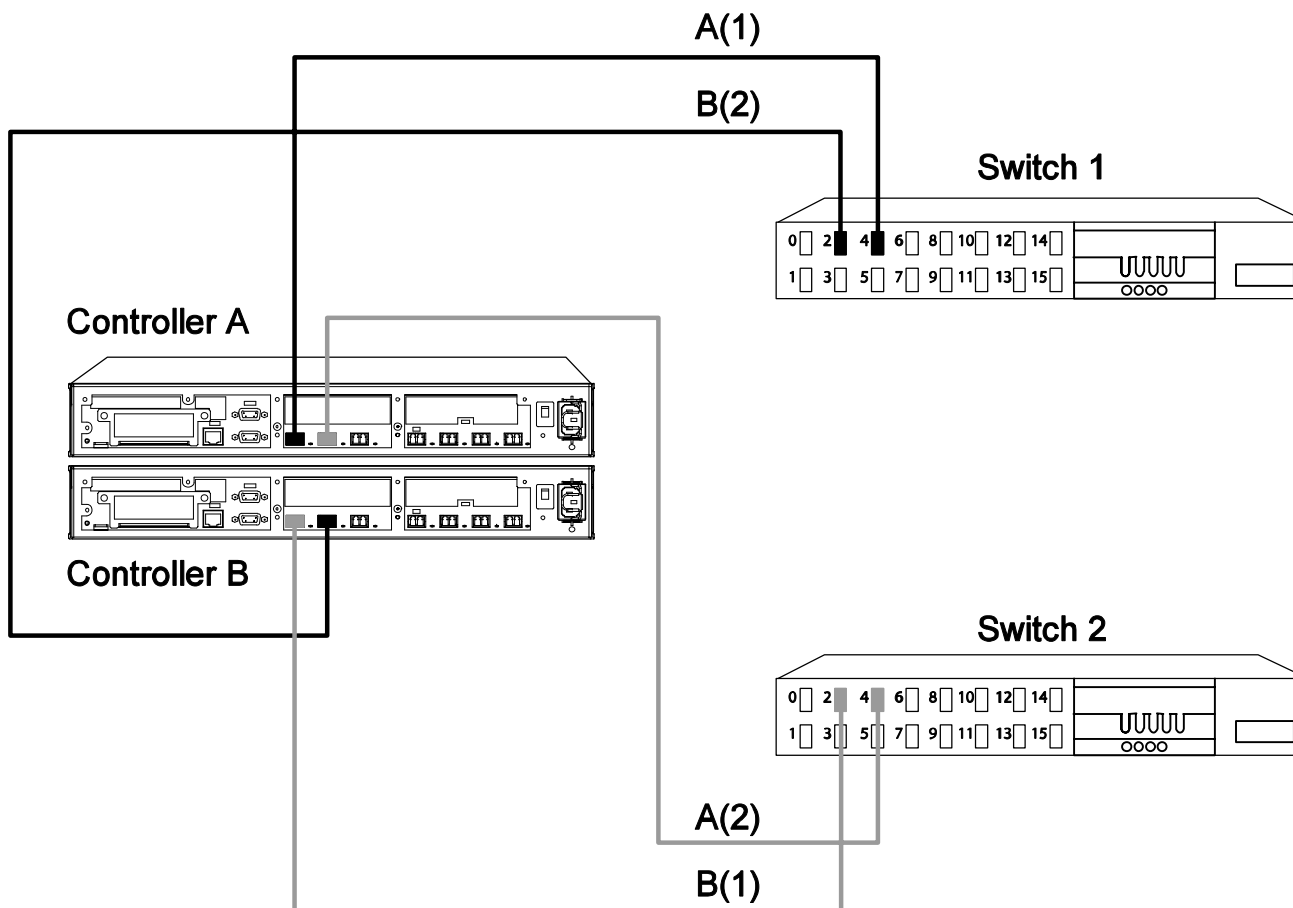
Basic Continuous Access Over Fiber configuration



CXO8009A



Supported cabling




CXO8092B

Example 1 of cabling that is not supported



THIS CABLING NOT SUPPORTED




CXO8094A



Example 2 of cabling that is not supported



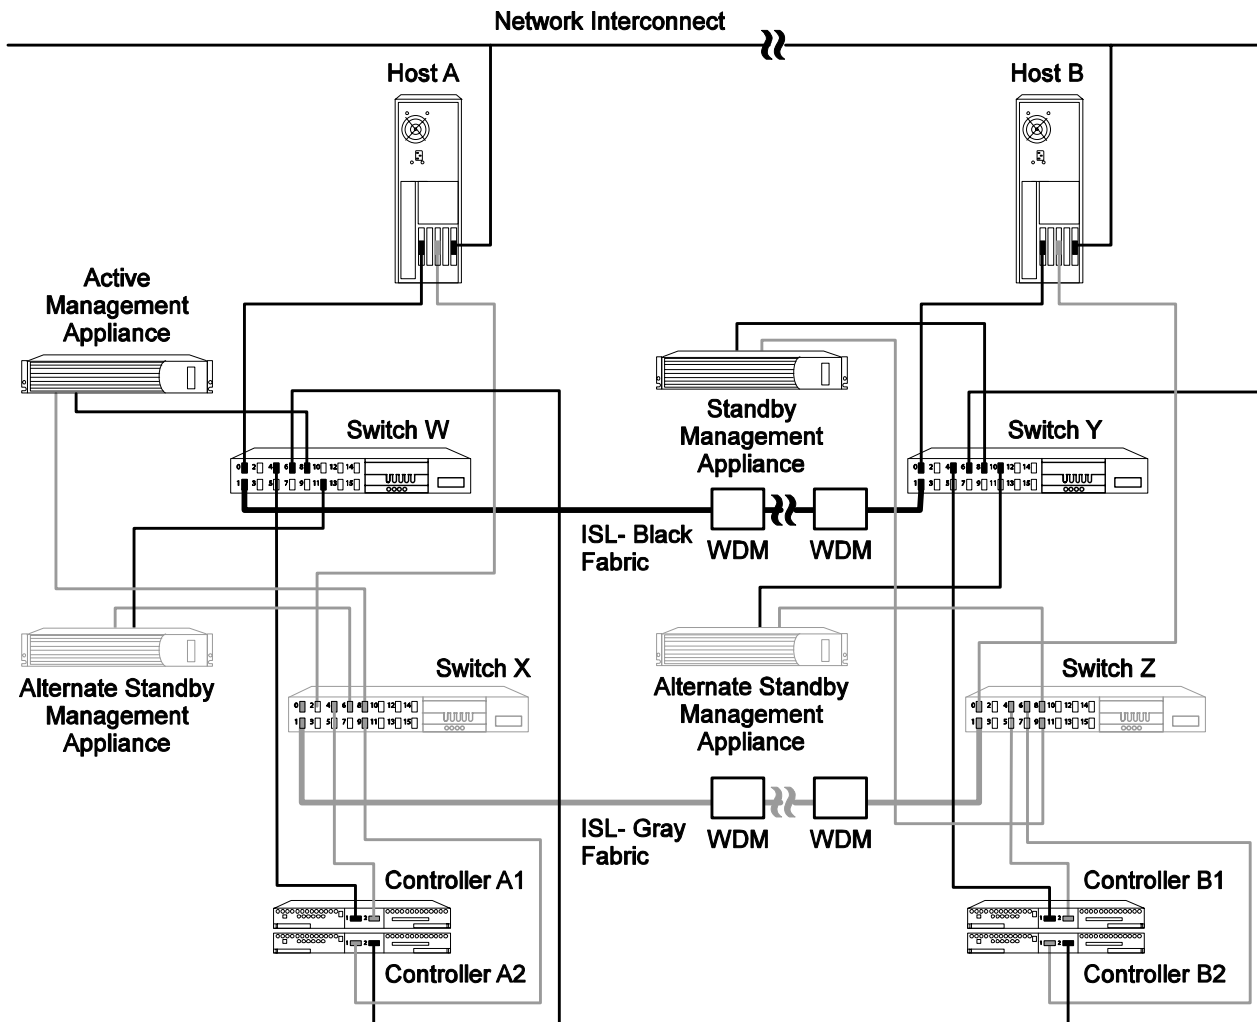
THIS CABLING NOT SUPPORTED



CXO8093A



Continuous Access EVA over WDM configuration



CXO8010A

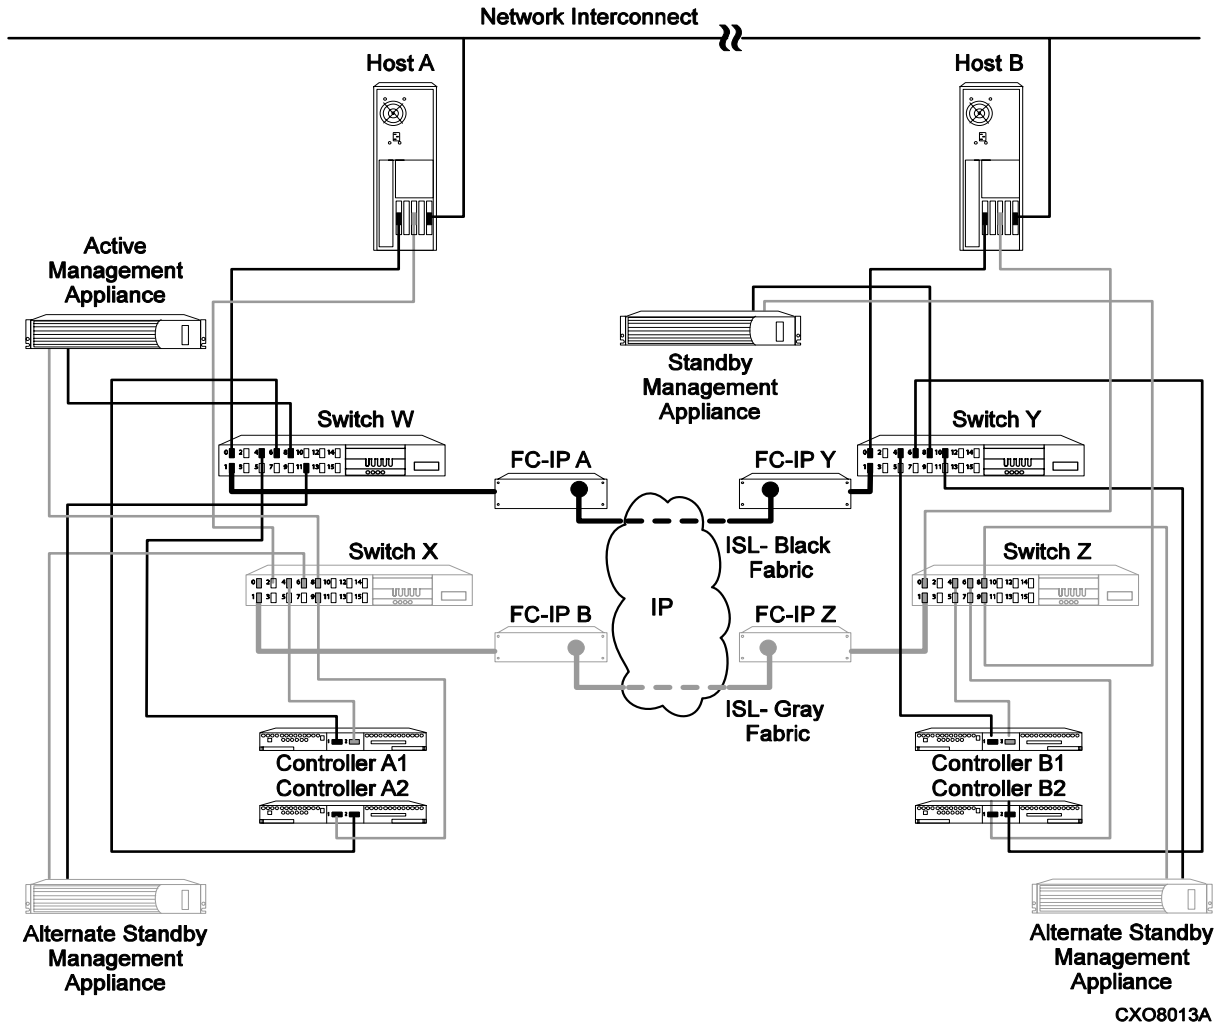


Continuous Access EVA over WDM configuration (cont)



- All WDM extensions regardless of vendor, or type of technology are supported.
- Quick check after installing any WDM is to insure that switches build a fabric

Continuous Access EVA Over IP or SONET configuration





Continuous Access EVA Over IP or SONET configuration (cont)

- Supported Switch and Gateways:
 - HP
 - B-Series & SR2122-2 FC to 1 GbE
 - Cisco
 - C-series & PA-FC-1G 1 GbE Fibre Channel Port Adapter for IP routers
 - C-series & MDS IP8 1 GbE IP adapter for FC MDS switch
 - CNT
 - B-Series or M-series & 1000, 1001, 1101 FC to GbE IP
 - SAN Valley
 - B-series and SL 700 single channel FC to GbE IP
 - B-series and SL1000 dual channel FC to GbE IP
 - B-series and SV2000 dual channel FC to GbE IP
 - Akara
 - B-series & Optical Utility Services Platform (OUSP) 2000 family
 - FC to SONET (OC3 to OC48)

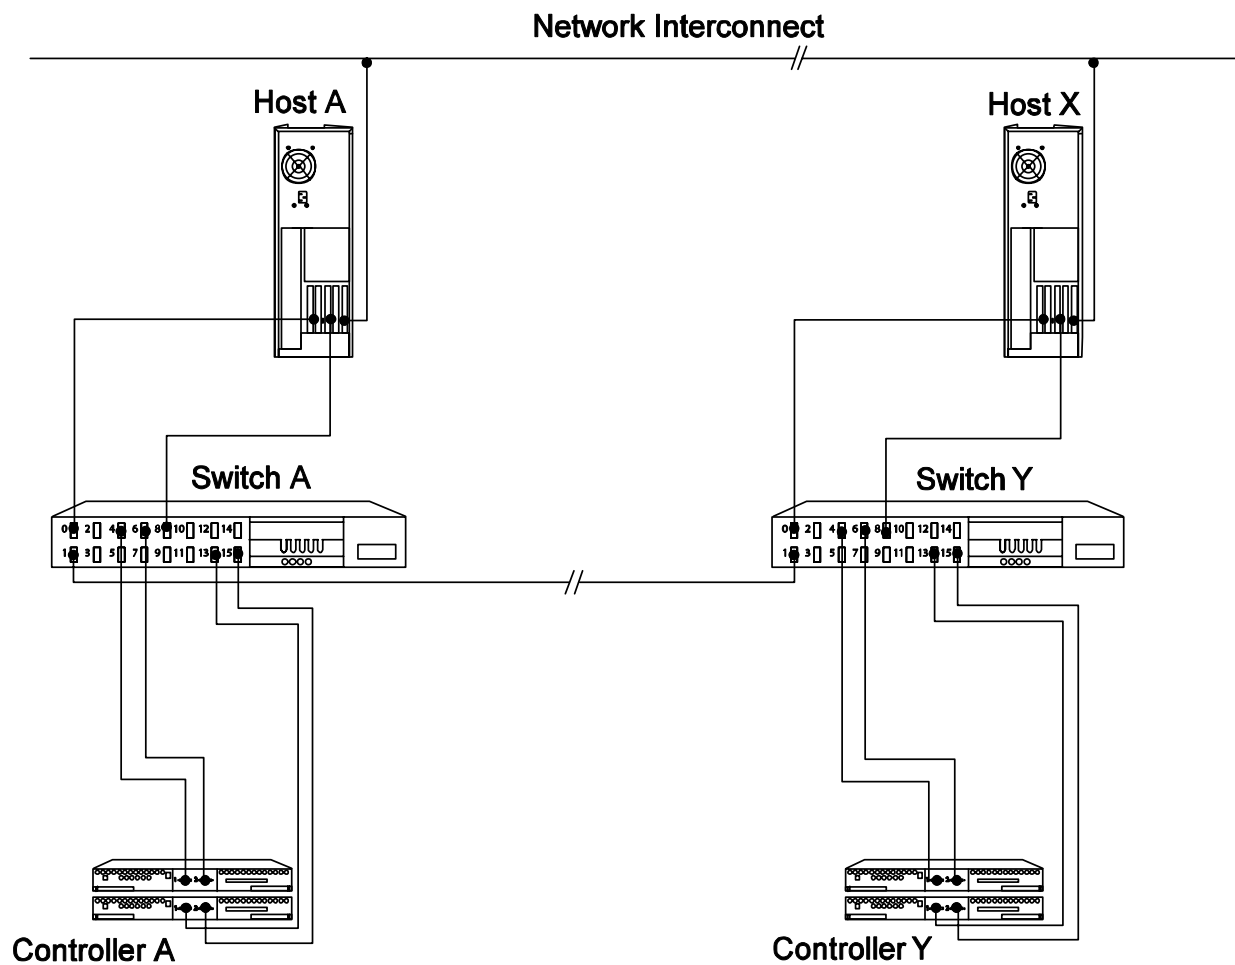




Optional configurations

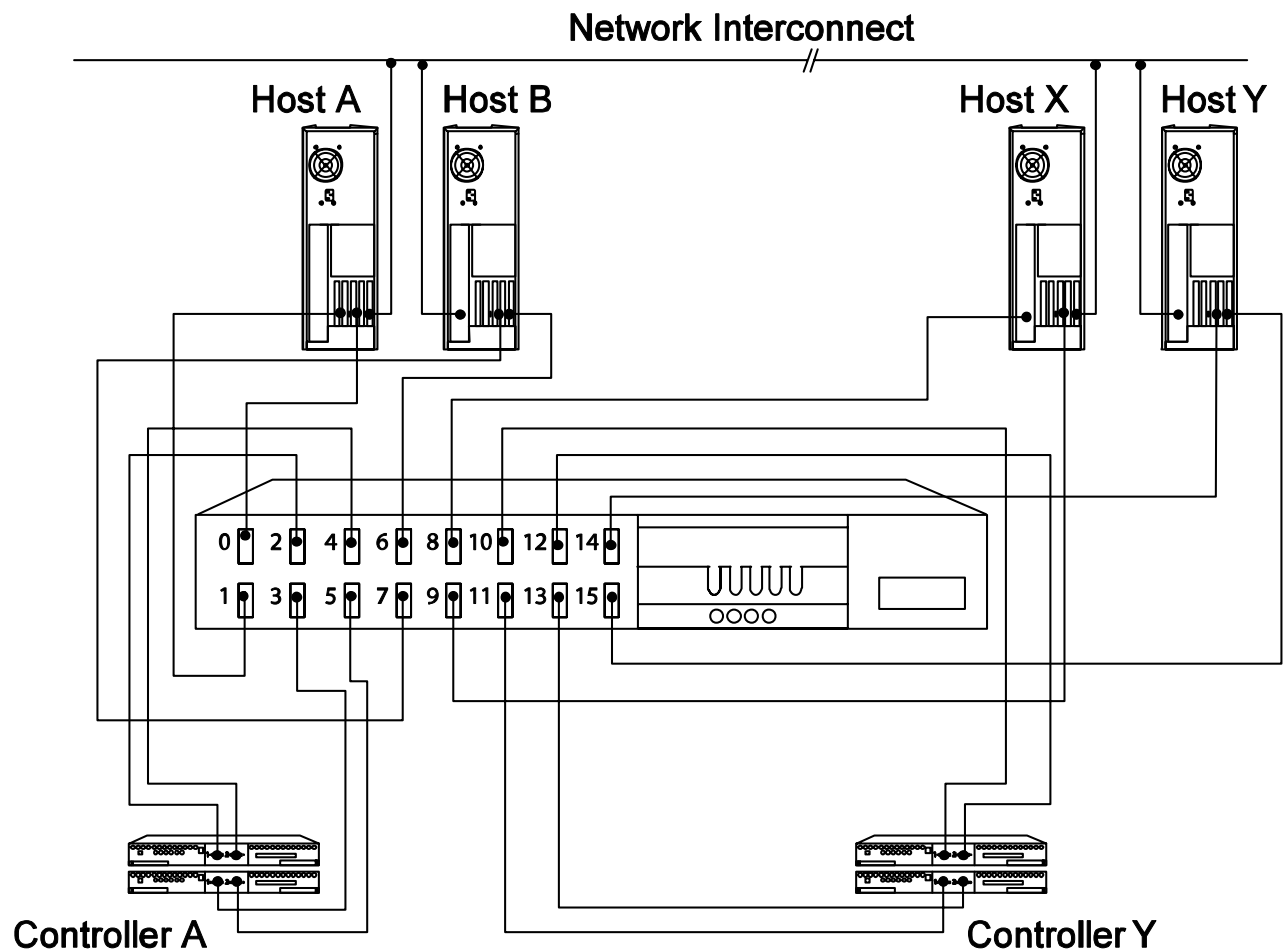
- Single HBA
 - Because there are two paths from any HBA to the EVA, Secure Path multi-path driver required (not required on OpenVMS and Tru64)
- Single fabric
- Single switch

Single-fabric configuration



CXO7866A

Single-switch configuration



CX08217A



But my solution is not supported

Either:

- To quote Picard: “Make it so”
- Or contact local account team for waiver





Agenda – part one

- A quick review
- Definitions & overview
- Review of supported configurations
- Questions???



Agenda

- Part one – Is it supported?
 - Definitions & overview
 - Supported configurations

- Part two – Troubleshooting
 - What changed? Did it work and now its not?
 - Understand what tools are available
 - Understand how to use those tools

What changed?

- Is this a new installation
 - If so, there are additional issues to consider such as was it set up correctly
- No, it has been working and now its not
 - Was there a change in the intersite network
 - Performance
 - Bandwidth
 - Was there a change in the application
 - More users
 - New version

What tools are available?

- Command View EVA
- Continuous Access User Interface
- Switch
- Fibre Channel Analyzers
- If using FCIP, then Network Analyzers

- Others?



Command View EVA or CA UI

- Is the appliance able to see both arrays?
 - Yes, then press on
- No; Then need to understand why not, as it could be something other than the array.

Command View EVA or CA UI

- Look first at source array
 - Is it working?
 - All ports active (host and backend)?
 - Are DR groups replicating?
 - Presentations correct?
- Then look at destination array
 - Is it working?
 - All ports active (host and backend)?
 - Are destination copies of Vdisks and DR Groups available?



Switch

- Use either CLI or GUI interface
- For each switch in fabric perform a “Show Fabric”, and confirm that all host ports are on correct fabrics
 - All ports 8 & D on one fabric
 - All ports 9 & C on other fabric
- In other words confirm correct wiring per CA EVA Design Guide

Switch (cont)

- Now look at performance by switch port
 - Identify host to array data paths
 - Identify array to array replication data paths
 - Look for choke points and over subscription of network
- Other
 - Any missing interswitch links?
 - Confirm switch code version is current

IP Network Analyzers

- Look at performance across IP network
 - Is it steady or bursty?
 - How does cross network performance compare to host to array performance?
 - Should be similar profiles
- Also look at end to end latency
 - Should be steady
 - If not, ask network vendor why not
 - Could it be dropped packets leading to re-transmissions?
 - How does it match expected latency versus distance?
 - Rough guess is $5 \text{ microsecond /km} * \text{distance (km)} * 2$

Potential network issues

- Maximum latency – 100 mSec one way
- Maximum latency jitter – 10 mSec
 - Difference between minimum and maximum
- Packet loss ratio
 - Less than 0.2% over 24 hours
 - Not to exceed 0.5% for more than 5 minutes in a two hour window
- Use PING to discover
 - > ping -n 3600 -l 2048
 - - n is number
 - - l (lower case L) is size



It helps if

- Use the tools during setup to get comfortable with them, and to know what the displays should look like when everything is working correctly
 - During peaks in workload
 - During valleys in workload
- Repeat as needed to look for trends
 - Avoid running out of disk space
 - Avoid over subscription of links
- Specifics in CA EVA Operations Guide



Where to find documentation and software



- The external StorageWorks Continuous Access (CA EVA) web page
- <http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html>
 - Click on "technical documentation" to find Design Guide and Operations Guides
 - Click on "software downloads" to find Continuous Access User Interface kit
 - Click on "related info" to find the Performance Estimator spreadsheet





HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:



RECOMMENDED TRAINING VENUE FOR THE
HP Certified Professional

