

Wes Lee Roberts

Technical Account Manager - OpenVMS

Ambassador

HP Services - OpenVMS Platinum Support

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

Purpose



- Overview of setting up SANs in an OpenVMS environment
- Tips to aid in management and troubleshooting
- Common adventures when managing SANs with OpenVMS for the first time
- Not a course in SAN fundamentals or a tutorial
- Why? Fibre Channel is different than any other storage for OpenVMS

Overview




- Fibre Channel advantages
 - SANs – multi-vendor, multi-OS, multi-system detached storage – industry common components
 - High performance, cabling distance & redundancy
- Challenges
 - Learning a new technology
 - Can grow to be complex


SAN Configurations



- **Redundant** – more than one fabric
- **Multipath** – each disk unit available on multiple adapters and fabrics



Simple Configuration




More Complex Configuration

Setup – EVA/HSG cabling with 2 fabrics



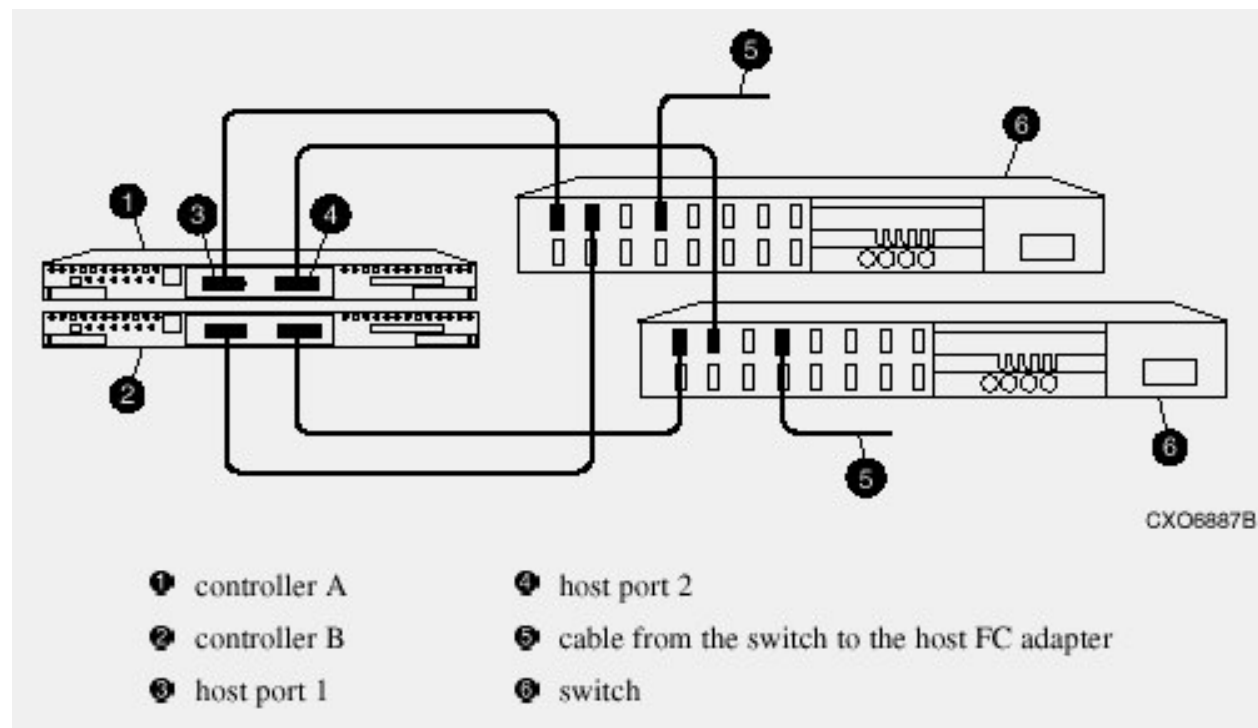
- Three possibilities – HSG Data Replication Mgr Style
 - Each HSG on a fabric – least redundant.
 - Not supported on EVAs, for CA or otherwise.
 - Use only for HSG DRM configurations.



Setup – EVA/HSG cabling with 2 fabrics



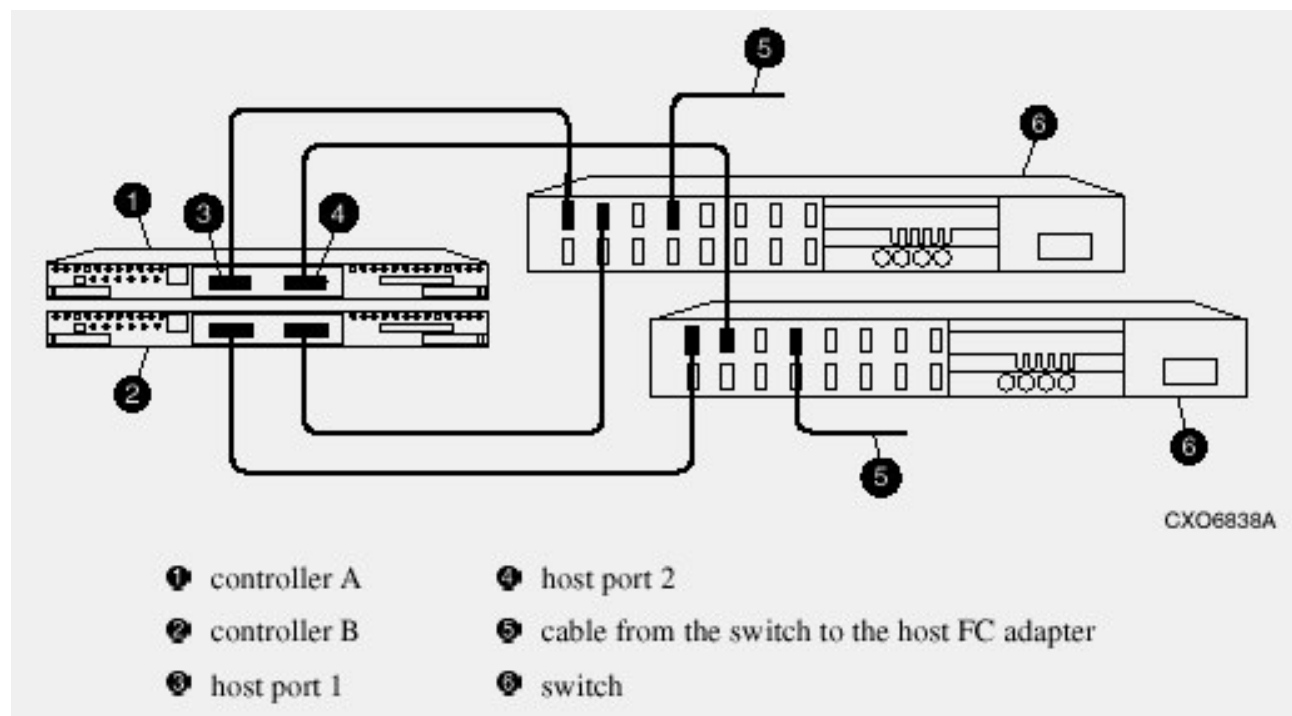
- Three possibilities – “1s & 2s” style
 - Each EVA/HSG on both fabrics – high redundancy
 - Use for non-DRM HSG and all EVA multipath and transparent failover configurations, but not CA



Setup – EVA/HSG cabling with 2 fabrics




- Three possibilities – “crossing” style
 - Each EVA/HSG on both fabrics – high redundancy
 - Use for non-DRM HSG and all EVA multipath and CA configurations



Setup – HSG80 cabling with 4 SANs



- 4 fabrics with an HSG80 port on each
 - highest redundancy/performance/capacity
- 4 fabrics are not currently supported with EVAs
 - A SAN Appliance must see all the ports on the EVA
 - Supported with added InterSwitch Links and zoning




Setup – MSA 1000 cabling with 2 fabrics - Using the fibre port

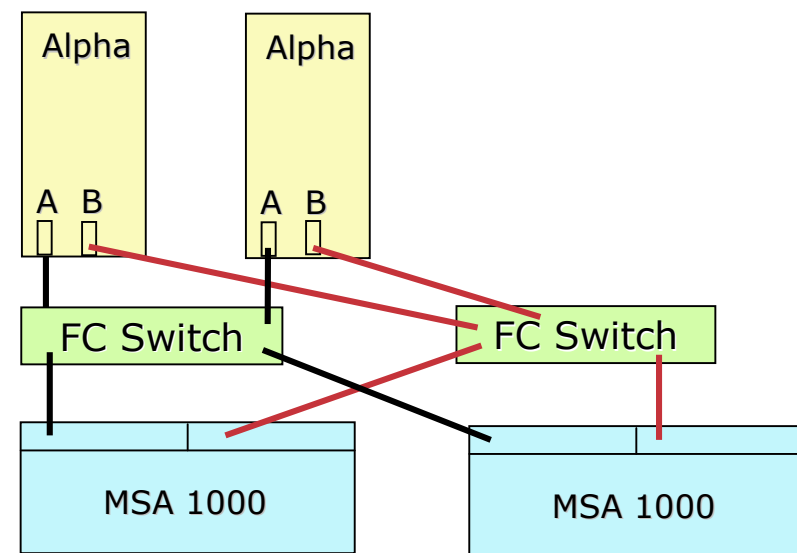


There are three FC options for the MSA1000.

- These figures demonstrate using the fibre port option to directly connect to another FC device.



Not a supported
OpenVMS configuration



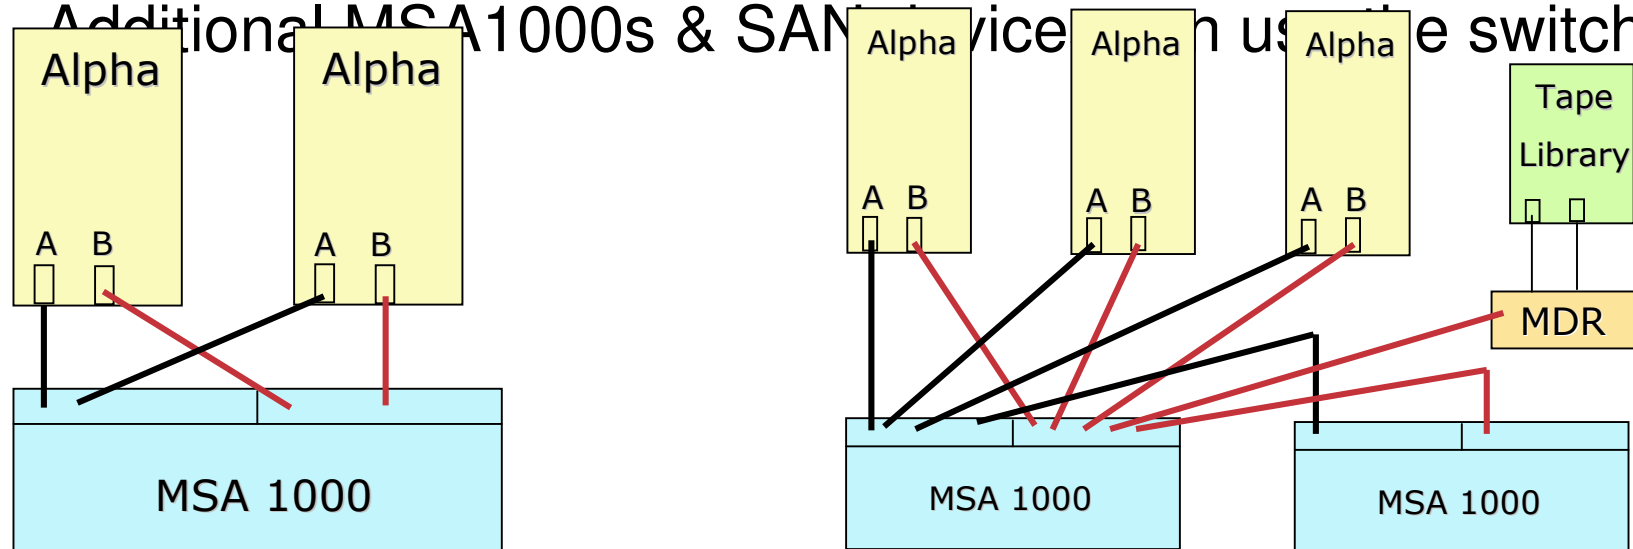
- MSA1000s shown here are dual redundant controllers. Single controller models are supported as well

Setup – MSA 1000 cabling with 2 fabrics - Using the embedded switch/hub



There are embedded 3 port hub & 8 port switch options

- In both cases one port is internal and the rest are available for external connections.
- Additional MSA1000s & SAN devices can use the switches.

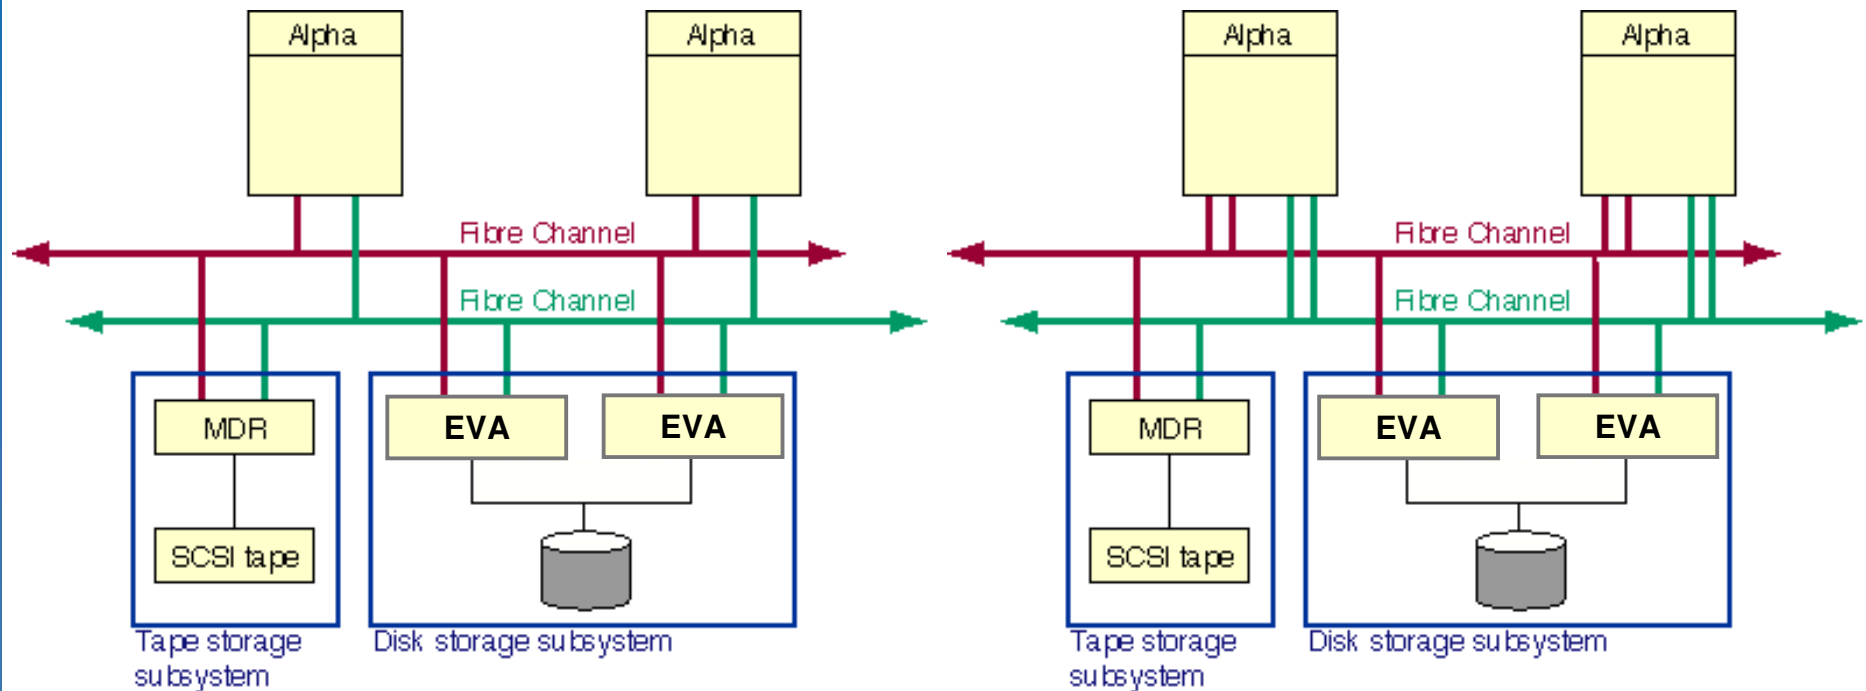


- The 3 port hub is **not** currently supported for OpenVMS, but support is planned.

Setup – Adapter cabling

Typically two fabrics with one or two adapters per fabric

- high **redundancy** and high **performance**



Setup – Cabling suggestions



- **Label** cables
- Be **consistent** – very important
 - whichever EVA/HSG cabling option you choose, use it for all EVA/HSGs (unless some HSG DRM & some not)
 - Cable the adapters the same from system to system
 - first adapters on the same fabric, etc.
- With GBIC style switches verify the GBIC type as they all look very similar
- **Document** everything and keep it current – use a common format to facilitate sharing via e-mail. This could significantly reduce time for future problem resolution.

Setup – Device naming



- An EVA/HSG/MSA presents two device types to OpenVMS
 - \$1\$GGAnnnnn: - console lun for the EVA/HSG pair
 - \$1\$DGAAnnnnn: - disk unit (HSG), virtual disk (HSV)
- There is no allocation class setting for an HSG/HSV/MSA
 - The 1 in \$1\$ is always 1 and cannot be changed
 - Recommend using the first digit or two of nnnnn to designate the EVA/HSG/MSA pair or the site – as in \$1\$DGAccnnn or \$1\$DGA1501:
- HSG80s must be set for Multibus failover and to SCSI-3 mode for OpenVMS

```
HSG80> SET MULTIBUS_FAILOVER
COPY={THIS|OTHER}
HSG80> SET THIS SCSI_VERSION=SCSI-3
```



Setup – Device naming



- The nnnnn in \$1\$xGAnnnnn: is set via an Identifier not by the controller name, or disk unit number dNNN (HSG) or LUN (EVA/MSA)
 - Identifiers have the range 0-32767
 - Note – MSCP only serves devices named 0-9999
- The console lun Identifier for a controller pair:
 - Is set via the SAN Appliance and Command View EVA for the EVA, on the HSG via the command SET THIS IDENTIFIER=nnnnn and on the MSA SET THIS_CONTROLLER_ID nnn
 - Is required for OpenVMS to see disks
 - Must be unique among controller pairs and disks visible to the OpenVMS system, and be non zero

Setup – Device naming



A disk unit Identifier (HSG), virtual disk OS Unit ID (EVA), and ID (MSA):

- Is set on the EVA via Command View EVA or SSSU
Is set on the HSG via the SET dNNN IDENTIFIER=nnnnn
Is set on the HSG via the SET UNIT_ID lun nnnnn
- Is required for each unit for OpenVMS to see the disk unit
- Must be unique SAN/Zone-wide. It is possible to give more than one unit the same identifier (doesn't work well, very bad).
- Recommend relating the unit, dNNN, or virtual disk, vdNNN, to the Identifier nnnn. Reduces confusion.
- Is viewed via the presentation tab in CommandView EVA, or SHOW dNNN or SHOW UNITS FULL on the HSG, or SHOW lun or SHOW UNITS on the MSA

EVA - Host setup notes



- Access is managed/configured per host using CommandView EVA or SSSU. Host port entries are automatically created when a system accesses the controller for the first time
 - “wwidmgr –show wwid”, booting OpenVMS, etc. - should create any entries as needed.
 - Host ports can be created manually – but that is not generally recommended
- A host must be defined before it can access any units. First the host is defined, which includes selecting one port (adapter). It is very important to set the operating system properly. Additional ports can then be added to a host.

EVA - Host setup notes




- VDisks must be presented to hosts to be accessible
 - When presenting a VDisk you may select a lun presentation number (1 – 255) or accept a default for each host.
 - This lun number is remembered by VMS to help identify the VDisk. So, if you want to recreate a VDisk and want to re-use the OS-ID (so the VMS device name and paths) it is important to use the same lun presentation number too or you will end up with additional paths.
- Recommend adopting a standard for VDisk names, lun presentation numbers and OS-ID settings.
- It is recommended to delete unneeded or inactive hosts.

WorldWide IDentifiers (WWIDs)

- Every adapter or controller connected to a SAN has a WWID that is unique – the form is nnnn-nnnn-nnnn-nnnn where the n's are hex digits.
- HS controller WWIDs (5000-nnnn-nnnn-nnnX) begin with 5000 and last digit denotes:

For HSGs

- 0 = the controller pair
- 1 = Port 1 bottom controller
- 2 = Port 2 bottom controller
- 3 = Port 1 top controller
- 4 = Port 2 top controller




- View HSG WWIDs with SHOW THIS/OTHER
 - Can also be viewed from OpenVMS (more later)

WorldWide IDentifiers (WWIDs)

- For EVAs the last digit (5000-nnnn-nnnn-nnnX) of a WWID for a controller pair denotes:

VCS V1.x V2.x+

- 0 0 = the controller pair
- D 9 = Port 1 of controller A
- C 8 = Port 2 of controller A
- 9 D = Port 1 of controller B
- 8 C = Port 2 of controller B



- Controller A/B designation is determined at system initialization time – Controller A is the one with the WWN prompt. Can use the locate command in CommandView EVA to determine the physical locations or follow cables.
- View HSV WWIDs from the SAN Appliance
 - Can also be viewed from OpenVMS (more later)

WorldWide IDentifiers (WWIDs)

- View KGPSA adapters & WWIDs at the AlphaServer SRM console with: “wwidmgr -show adapter”
 - KGPSA WWIDs typically begin with 1000 or 2000.

```
LP00>>>wwidmgr -show adapter
```


item	adapter	WWN	Cur. Topo	Next Topo
[0]	kgpsaa0.0.0.4.6	1000-0000-c920-05ab	FABRIC	FABRIC
[1]	kgpsab0.0.0.8.6	1000-0000-c921-0ce0	FABRIC	FABRIC
[9999]	All of the above			

- As of OpenVMS V7.2-2 the WWID of a KGPSA is viewable with the SHOW DEVICE/FULL command, ie:
\$ SHOW DEVICE/FULL FGA0:
 - Can be viewed from the EVA/HSV/MSAs (more later)
 - All WWIDs in the SAN can be viewed from the switches.
- Document all adapter and controller WWIDs.

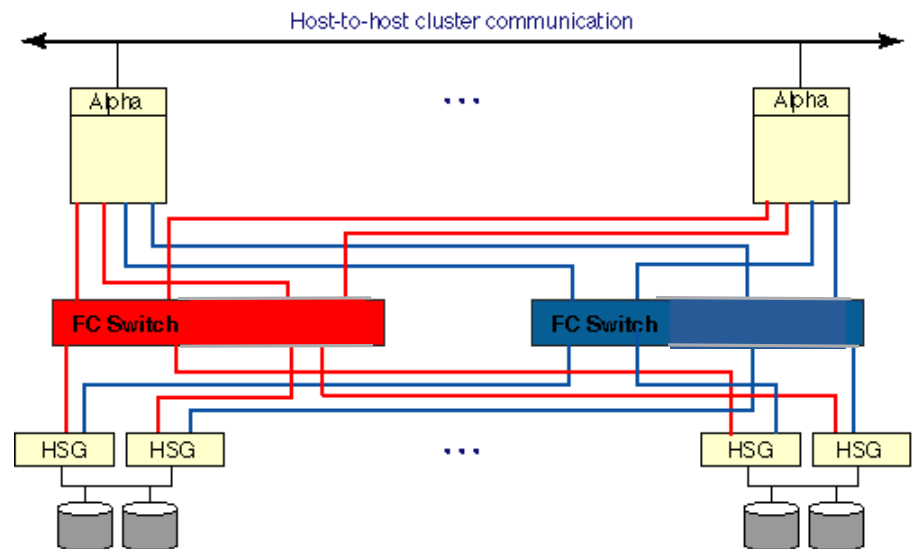
Paths / Connections



- A path or connection is defined by the two end points or WWIDs (the adapter & the EVA/HSG/MSA port)



4 paths per disk unit
on each system



VM0486A-AI

8 paths per disk unit on
each system

Path / Connection Status

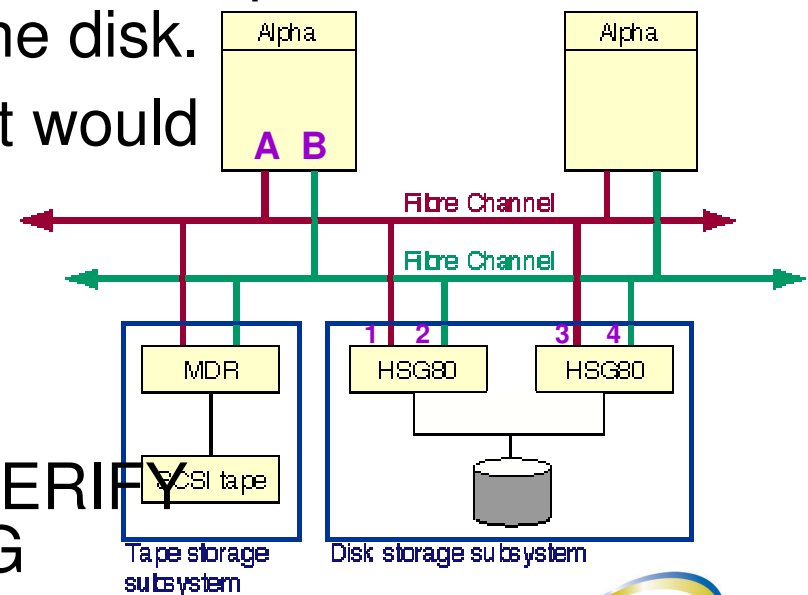


A path or connection can be considered

- Good or bad meaning it is functional or not
- Valid or invalid meaning it exists or not
- Connected (active) or not connected (inactive)
 - It is connected if the path's controller port is on the controller currently serving the disk.
 - It is not connected if disk unit would have to fail-over to the other EVA/HSG/MSA controller to be accessed on the path.

- To disable old paths use
`SYSMAN IO SCSI_PATH_VERIFY`

If the paths A-1 (adapter A to HSG port 1) & B-2 are connected then the paths on A-3 & B-4 are not connected, and visa-versa



AlphaServer Console Access to SAN Disks



- Due to the long WWIDs fibre channel devices cannot be incorporated into a console device name.
- Fibre Channel devices must have an association defined between the device name and WWID before they can be accessed from the console.
 - Only FC based system (with shadowset members) and dump disks require console access.
- Use the utility wwidmgr to manage the associations.
- 8 environment variables to store association info.
 - N1-N4 are used for the WWIDs to access
 - wwid0-wwid3 are used for device names

Using wwidmgr



- The users guide is on AlphaServer console CDs in the doc directory as wwidmgr.pdf & wwidmgr.ps
- Wwidmgr runs in console diag mode. Since the context is not the same as when the system boots an INIT may be necessary before entering diag mode.
 - On AlphaServer 8x00, 4x00, or 1200 systems, the console must be manually set in diagnostic mode.
 - P00>>> set mode diag
- After using wwidmgr the system must be re-initialized (with an INIT) before booting or to exit diag mode.
- Wwidmgr commands have the syntax:

Using wwidmgr



- Remember: EVA/HSG/MSA console lun & disk units must be assigned identifiers before OpenVMS can access them. These can be used with wwidmgr, too.
- Displaying FC devices
 - To poll each FC adapter and display devices found
 - wwidmgr –show wwid
 - wwidmgr –show wwid –full !displays all paths
 - (ev:none) means no environment variable is set for it
 - To display a specific device
 - wwidmgr –show wwid –full –udid *<identifier>*

Using wwidmgr

- For example:



```
P00>>>wwidmgr -show wwid
```

```
polling for units on kgpsa0, slot 4, bus 0, hose6...
```

```
kgpsaa0.0.0.4.6          PGA0          WWN 1000-0000-c920-05ab
```

```
polling for units on kgpsa1, slot 8, bus 0, hose6...
```

```
kgpsab0.0.0.8.6          PGB0          WWN 1000-0000-c920-5e75
```

```
[0] UDID:1 WWID:01000010:6000-1fe1-0000-04a0-ffff-fffe-0002-0000 (ev:none)
[1] UDID:2 WWID:01000010:6000-1fe1-0000-04a0-ffff-fffe-0003-0200 (ev:none)
[2] UDID:77 WWID:01000010:6000-1fe1-0000-04a0-ffff-fffe-0005-0300 (ev:wwid0)
[3] UDID:99 WWID:01000010:6000-1fe1-0000-04a0-ffff-fffe-0006-0100 (ev:none)
[4] UDID:40 WWID:01000010:6000-1fe1-0000-04a0-ffff-fffe-0005-0000 (ev:none)
```

```
P00>>>wwidmgr -show wwid -udid 99 -full
```

```
[3] UDID:99 WWID:01000010:6000-1fe1-0000-04a0-ffff-fffe-0006-0100 (ev:none)
```

```
- current_unit:2044 current_col: 0 default_unit:2044
```

	via adapter	via fc_nport	Con	DID	Lun
-	kgpsaa0.0.0.4.6	5000-1fe1-0000-04a3	No	210513	99
-	kgpsaa0.0.0.4.6	5000-1fe1-0000-04a1	Yes	210713	99
-	kgpsab0.0.0.8.6	5000-1fe1-0000-04a2	Yes	410513	99
-	kgpsab0.0.0.8.6	5000-1fe1-0000-04a4	No	410713	99



Using wwidmgr



- Define a device name and WWID association
 - `wwidmgr -quickset -udid <identifier>`
This sets the first available wwid# and typically all 4 of the N# environment variables.
 - Use also for additional devices on same controller pair
- Check out the association settings
 - `wwidmgr -show ev` !displays the wwid#s & N#s
 - `wwidmgr -show r[eachability]`
Shows path information & status for defined devices
- Resetting/clearing wwidmgr environment variables
 - `wwidmgr -clear all` !deletes all the wwid#s & N#s
 - `wwidmgr -clear wwid<0-3>` !delete a single wwid#
 - `wwidmgr -clear N<1-4>` !delete a single N#

Using wwidmgr - For example:



```
P00>>>wwidmgr -quickset -udid 77      !quickset also shows reachability info
P00>>>wwidmgr -quickset -udid 40      !which is not shown here to save space
P00>>>wwidmgr -show reachability
```

Disk assignment and reachability after next initialization:

6000-1fe1-0000-04a0-ffff-fffe-0005-0300

	via adapter:	via fc nport:	connected:
dga77.1001.0.4.6	kgpsaa0.0.0.4.6	5000-1fe1-0000-04a3	Yes
dga77.1002.0.4.6	kgpsaa0.0.0.4.6	5000-1fe1-0000-04a1	No
dgb77.1003.0.8.6	kgpsab0.0.0.8.6	5000-1fe1-0000-04a2	No
dgb77.1004.0.8.6	kgpsab0.0.0.8.6	5000-1fe1-0000-04a4	Yes

6000-1fe1-0000-04a0-ffff-fffe-0005-0000

	via adapter:	via fc nport:	connected:
dga40.1001.0.4.6	kgpsaa0.0.0.4.6	5000-1fe1-0000-04a3	No
dga40.1002.0.4.6	kgpsaa0.0.0.4.6	5000-1fe1-0000-04a1	Yes
dgb40.1003.0.8.6	kgpsab0.0.0.8.6	5000-1fe1-0000-04a2	Yes
dgb40.1004.0.8.6	kgpsab0.0.0.8.6	5000-1fe1-0000-04a4	No

Use the reachability information to assign the environment variables bootdef_dev & dump_dev all the devices/paths

```
>>> set bootdef_dev dga77.1001.0.4.6,dgb77.1003.0.8.6,dga77.1002.0.4.6,dgb77.1004.0.8.6
>>> set dump_dev dga40.1001.0.4.6,dgb40.1003.0.8.6,dga40.1002.0.4.6,dgb40.1004.0.8.6,d
ga77.1001.0.4.6,dgb77.1003.0.8.6,dga77.1002.0.4.6,dgb77.1004.0.8.6
```

Using wwidmgr



To now set up a device on a second EVA/HSG pair of controllers (for example shadowed system disks)

- Set a wwid# environment variable for the device without effecting the N# values.
 - wwidmgr -set wwid -udid <identifier>
- Since each N# is a WWID for a port of an EVA/HSG, 2 N#s can be used for each controller pair and retain redundancy for console access.
 - It is important to have the N#s set so the 2 ports to an EVA/HSG go to each controller over different SANs to maintain redundancy. You must pick which to keep.
 - If redundancy is desired for these console definitions then all disks to be defined must be restricted to 2 EVA/HSG controller pairs.
 - Not at issue for MSA as they only have two ports.

Using wwidmgr



The previous example results in environment variables

```
P00>>>wwidmgr -show ev
wwid0          77 1 WWID:01000010:6000-1fe1-0000-04a0-ffff-fffe-0005-0300
wwid1          40 1 WWID:01000010:6000-1fe1-0000-04a0-ffff-fffe-0005-0000
wwid2
wwid3
N1              50001fe1000004a3
N2              50001fe1000004a1
N3              50001fe1000004a2
N4              50001fe1000004a4
```

Note from the show reachability example, these HSG80s are in a 1s and 2s cabling configuration.

- That means N1 and N2 are HSG80 ports on the same SAN or adapter, and N3 and N4 are on the other.
- Keep either N1 & N3 or N2 & N4 so a path on each SAN and on each controller in a typical 2 SAN config.

Using wwidmgr



Set up a device on a second EVA/HSG pair of controllers - continued

- Find the paths for the EVA/HSG ports for the device.
 - `wwidmgr –show wwid –full –udid <identifier>`
Note the wwid of the EVA/HSG
 - `wwidmgr -show port`
List EVA/HSG port WWIDs in an item numbered list.
 - Note the desired EVA/HSG port WWIDs – be sure to identify one for each controller and fabric.
- Define the devices using the item # displayed in the show port output and the # from an N# to use
 - `wwidmgr -set port -item <i#> -node <# from N#>`

Using wwidmgr



- - For example:

```
P00>>>wwidmgr -set wwid -udid 399
P00>>>wwidmgr -show wwid -udid 399 -full
[9] UDID:399 WWID:01000010:6000-1fe1-0000-14d0-0009-8179-1057-002b (ev:wwid3)
- current_unit:3044 current_col: 0 default_unit:2044
      via adapter      via fc_nport      Con      DID      Lun
-      kgpsaa0.0.0.4.6  5000-1fe1-0000-14d1  No      210513    399
-      kgpsaa0.0.0.4.6  5000-1fe1-0000-14d3  Yes     210713    399
-      kgpsab0.0.0.8.6  5000-1fe1-0000-14d2  No      210513    399
-      kgpsab0.0.0.8.6  5000-1fe1-0000-14d4  Yes     210713    399
P00>>>wwidmgr -show port
[0] 1000-0000-c920-ca72
[1] 2006-0060-6900-0953
[2] 20fc-0060-6900-0953
[3] 5000-1fe1-0000-04a2 (N3)
[4] 5000-1fe1-0000-04a1 (N2)
[5] 5000-1fe1-0000-04a3 (N1)
[6] 5000-1fe1-0000-04a4 (N4)
[7] 5000-1fe1-0000-14d1
[8] 5000-1fe1-0000-14d3
[9] 5000-1fe1-0000-14d2
[10] 5000-1fe1-0000-14d4
P00>>>wwidmgr -set port -item 7 -node 2
P00>>>wwidmgr -set port -item 10 -node 4
```



OpenVMS & FibreChannel paths

- Remember – a path is defined by two endpoints on a SAN. OpenVMS uses the adapter name and the WWID of the EVA/HSG/MSA port. For example:

```
$ show device /fu $1$DGA46
```

```
Disk $1$DGA46:(LILDOG), device type DEC HSG80, is online, allocated, deallocate  
on dismount, mounted, file-oriented device, shareable, available to cluster,  
device has multiple I/O paths, error logging is enabled.
```

```
I/O paths to device          4  
Path PGA0.5000-1FE1-0005-35E3 (LILDOG), primary path.  
Error count          2    Operations completed          26  
Path PGA0.5000-1FE1-0005-35E1 (LILDOG).  
Error count          0    Operations completed          4  
Path PGB0.5000-1FE1-0005-35E2 (LILDOG), current path.  
Error count          2    Operations completed          8  
Path PGB0.5000-1FE1-0005-35E4 (LILDOG).  
Error count          0    Operations completed          4
```

- The **primary** path is the first one found by OpenVMS at I/O configuration.
- The **current** path is the one in use now.

OpenVMS & FibreChannel paths



- OpenVMS Multipath devices have an UCB for every path. These can be viewed with **ANALYZE/SYSTEM:**

SDA> SHOW DEVICE \$1\$DGA#####

- Changing the current path manually generates an IO if the device is mounted – and none if not mounted. When the I/O is attempted, if the current path is:
 - bad another will be chosen via a mount verification
 - not connected, then the controller will attempt to fail the disk over so this path will be connected – if the disk is mounted elsewhere in the cluster paths will be disrupted on each node when the next IO occurs there and new paths chosen via mount verification.

OpenVMS & Changing paths

- A path may be set manually via
 - \$ SET DEVICE/PATH=device-path/SWITCH
- Starting with OpenVMS V7.3-1, mounting a multipath disk can result in a path change.
- A path may change automatically by a mount verification
 - As of V7.3-1, the same algorithm is used as when mounting a disk.
- When a path is changed an OPCOM message is generated which includes whether the change was manual or automatic.
 - Note path changes may result in Mount Verifications, too.

OpenVMS & FibreChannel paths

The current path of FC devices can be viewed via
\$ SHOW DEVICE/MULTIPATH

```
$ show device/mult
```

Device Name	Device Status	Error Count	Paths Avl/Tot	Current path
\$1\$DGA1:	(LILDOG) Online	0	4/ 4	PGA0.5000-1FE1-0005-35E3
\$1\$DGA3:	(LILDOG) Mounted alloc	2	4/ 4	PGA0.5000-1FE1-0005-35E3
\$1\$DGA45:	(LILDOG) Online	2	4/ 4	PGA0.5000-1FE1-0005-35E3
\$1\$DGA46:	(LILDOG) Mounted	7	4/ 4	PGB0.5000-1FE1-0005-35E2

Device Name	Device Status	Error Count	Paths Avl/Tot	Current path
\$1\$GGA99:	Online	0	4/ 4	PGB0.5000-1FE1-0005-35E2

```
$
```

```
$ sho dev dg/mult
```

Device Name	Device Status	Error Count	Paths Avl/Tot	Current path
\$1\$DGA1:	(LILDOG) Online	0	3/ 4	PGA0.5000-1FE1-0005-35E3
\$1\$DGA3:	(LILDOG) Mounted alloc	2	4/ 4	PGA0.5000-1FE1-0005-35E3
\$1\$DGA45:	(LILDOG) Online	2	4/ 4	PGA0.5000-1FE1-0005-35E3
\$1\$DGA46:	(LILDOG) Mounted	8	4/ 4	PGB0.5000-1FE1-0005-35E2

OpenVMS & Path Balancing



- Path balancing is the selection of current paths so all adapters and controller ports available are utilized.
- Recall that a path is defined by two end points, adapter and controller port. Once set, the path used only changes if a mount verification occurs, or there is manual intervention.
- Beginning with OpenVMS V7.3-1 path balancing algorithms are part of device discovery, mounting and mount verification.
 - So, do you need your own script in addition?
 - OpenVMS does not take into account actual IO load when selecting paths.
 - Having a script makes recovery from a SAN fail-over event easier.

OpenVMS & Path Balancing



Prior to OpenVMS V7.3-1

The default path selection behavior was the current path would be the primary path.

Path balancing was a manual process and a system management responsibility.

- Path balancing could only be done from the OpenVMS system. Since OpenVMS did not distinguish whether or not a path was connected, HSG preferred path settings were meaningless.
- Question – If paths are manually balanced on all cluster nodes and one node was rebooted, by default, what happened to the path settings?

Answer – The primary path was chosen as the current path. When the disks were mounted some units failed-over on the HSG/HSV's, resulting in path changes on other nodes

OpenVMS & Path Balancing

Prior to OpenVMS V7.3-1



- Recommended procedure to balance paths:
- Create a command procedure for all SET DEVICE/PATH commands, so that all adapters and controller ports are used.
- Execute the procedure prior to mounting disks.
 - If the balance of paths get out of synch. due to some event, Execute the procedure again. Recall, doing this will result in mount verification messages.
 - If possible use the same procedure on all applicable cluster nodes. This reduces unit fail-over on controllers, but does requires consistent cabling from node to node.

OpenVMS & Path Balancing



As of OpenVMS V7.3-1

- When a new path to a multipath device is configured, the direct path with the fewest devices is chosen automatically as the current path.
- The MOUNT utility can, for a disk, trigger an automatic path selection search biased toward connected paths.
- Same algorithm used during Mount Verification if needed.
- Benefits include:
 - Minimizing path changes on other hosts in the cluster.
 - Tends to preserve any manual path setting done on other cluster nodes.
 - Enables the use of HSG/V capabilities for initial default distribution of devices between the two controllers.
 - Tends to balance use of available paths from a node to the disk units.
 - Prefers direct paths over MSCP served paths.



OpenVMS & Path Balancing

As of OpenVMS V7.3-1



Path selection during a multipath disk MOUNT:

1. If the current path is a direct path with a connected status it is used.
2. If the current path is an MSCP path selected manually by a path switch command it is used.
3. All direct paths are considered in order of increasing use as a current path for other devices. The first connected path found is selected as the current path. If none were connected, the first path considered is used which results in a controller fail-over attempt to make the path connected.
4. The MSCP served path is tried.
5. If the result was a current path change, an automatic path switch is performed and an OPCOM message is issued.

The MOUNT utility can repeat this algorithm a number of times until a working path is found. The number of retries depends on elapsed time & specified command qualifiers.



OpenVMS & Path Balancing



As of OpenVMS V7.3-1

Automatic path selection algorithms give meaning to preferring paths on the HSG/EVA controllers:

- Set the preferred HSG controller (top/bottom) to serve each unit. Note - only when both controllers are rebooted does this have an effect.

SET unit

PREFERRED_PATH={THIS|OTHER}_CONTROLLER

- EVAs have preferred path and preferred fail-back path selection via the SAN Appliance HSV Element Manager:
 - Setting preferred paths only on HSV has the same effect as on HSGs.
 - Setting preferred fail-back effects recovery when a member of an HSV pair fails and later recovers. Note – when the controller recovers and makes the fail-back settings, path changes (mount verifications) result



OpenVMS & Path Balancing

As of OpenVMS V7.3-1



Recommended procedure to balance paths if needed:

- Create a command procedure for all SET DEVICE/PATH commands, so that all adapters and controller ports are used. Match the settings to the controller preferred path settings.
- Execute the procedure after disks are mounted as needed. Recall, doing this will result in mount verification messages.
- If possible use the same procedure on all applicable cluster nodes. This reduces unit fail-over on controllers, but does requires consistent cabling from node to node.
- Setting the preferred HSG/HSV controller (top/bottom) on each unit can reduce mount verifications when a path balancing script is run. On an HSV, optionally set the fail-back paths.



OpenVMS & Adding FibreChannel Devices



FC utilizes SCSI protocols. So, devices appear as local devices to OpenVMS.

When a unit is added on a controller, and an (OS) identifier assigned, it does not automatically appear on OpenVMS.

- An **autoconfigure is required** for a new unit to appear on an OpenVMS node. Recommend **doing this cluster-wide**.

```
$ MCR SYSMAN
```

```
SYSMAN> SET ENVIRONMENT/CLUSTER
```

```
SYSMAN> IO AUTOCONFIGURE
```

If MSCP serving is disabled (or on a version prior to V7.3-1) it is very important to autoconfigure on all nodes before using a new device.

- Many sites disable MSCP serving for a variety of reasons, so use caution.



OpenVMS & FibreChannel paths

As of OpenVMS V7.3-1



Switching between MSCP served and direct paths for multipath disks. Great news for multi-site OpenVMS Clusters with extended SANs

- If a multipath disk is enabled for MSCP serving and its direct paths fail its current path will change to an MSCP path. If one or more direct paths are then restored it will fail back to a direct path.
- Use the MPDEV_REMOTE system parameter to enable this capability.
- To manually switch the current path to an MSCP served path, use
SET DEVICE device-
name/SWITCH/PATH-MSCP

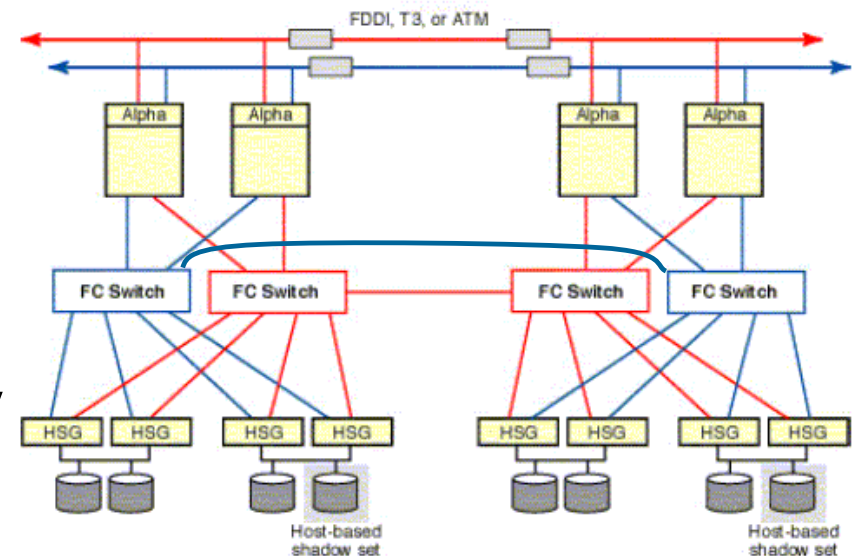


Two-site Host Based Volume Shadowing



OpenVMS supports HBVS over extended SANs

- Extended fibre channel distance limits are
 - 100 km for direct fibre
 - Limited only by latency/throughput requirements for FC over IP
- If using extended VMSClusters too:
 - use the V7.3-1 MSCP fail-over capability.
 - Distance may be limited by the cluster interconnect



Firmware and ECOs



- Fibre Channel is a continually evolving technology. So, staying **current on firmware** levels (for controllers, switches and adapters) and related OpenVMS ECOs is important. Check firmware versions on:
 - HSGs via SHOW THIS or SHOW OTHER, current is ACS V8.7-7 (V8.8 will be the final version). EVAs via CommandView EVA, current is VCS V3.014. MSAs via SHOW VERSION, current is 4.32
 - Brocade Switches with the telnet console “version” command.
 - Adapters via ANALYZE/SYSTEM - checking FG devices. Firmware is updated with the console. See the current Alpha console firmware CD, online at <ftp://ftp.digital.com/pub/DEC/Alpha/firmware/readme.html>
- On OpenVMS check for FIBRE-SCSI, SYS, SHADOWING and UPDATE ECOs with PRODUCT SHOW HISTORY. Current OpenVMS ECOs are available at <http://www.itrc.hp.com> or anonymous ftp at <ftp.itrc.hp.com>.

```
SDA> fc show device fga0
FGA0: operational firmware revision CS3.91A1
SDA> fc show device fgb0
FGB0: operational firmware revision CS3.91A1
```





New in OpenVMS V7.3-2

Dynamic Volume Expansion (DVE)

- Storage arrays today can dynamically expand units
 - Before V7.3-2 initialization was required to utilize the expanded capacity.
 - DVE allows allocation of extra bitmap space which enables later volume size expansion with the device mounted.
 - `$ INIT/LIMIT 1dga100: name ! Allocates 1TB bitmap`
 - `$ MOUNT/CLUSTER 1dga100: name ! Mount and use volume`
 - (expand the unit on the storage controller)
 - `$ SET VOLUME 1dga100:/SIZE=xxxxxx ! Set new size`
- /LIMIT sets the **Expansion Size Limit** (maximum growth limit)
- /SIZE sets the **Logical Volume Size** (current size of the volume)
- Current settings can be viewed with `SHOW DEVICE/FULL`



New in OpenVMS V7.3-2

Dynamic Volume Expansion

(DVE) For volumes initialized prior to OpenVMS V7.3-2

- `$ DISMOUNT/CLUSTER 1dga100: ! Dismount volume`
- `$ MOUNT 1dga100: name ! Privately mount volume`
- `$ SET VOLUME 1dga100: /LIMIT ! Allocate new bitmap`
- `$ MOUNT/CLUSTER 1dga100: name ! Cluster mount volume`
- (expand the unit on the storage controller)
- `$ SET VOLUME 1dga100: /SIZE=xxxxxx ! Set new size`
- DVE permits Dissimilar Device Shadowing (DDS)
 - Different sized volumes can now be shadowed
 - The **size of the smallest volume** determines the size of the shadow set
 - The shadow set can only be mounted on V7.3-2 nodes
 - The size of the shadow set can be grown using DVE

New in OpenVMS V7.3-2



Quieter Mount Verification

- In a SAN there are many reasons for “normal”, quick mount verifications, such as
 - Path switch by another cluster node
 - Rezoning a fabric
 - Transient command timeout to a busy storage device
- “Quieter” mount verification permits messages for infrequent, immediately recovered mount verification to be suppressed.
- Sysgen parameters (setting either of these to 0 disables this feature)
 - MVSUPMSG_INTVL (default is 1 hour)
 - MVSUPMSG_NUM (default 5)
- Mount verification messages are suppressed unless there are more than MVSUPMSG_NUM such mount verifications within MVSUPMSG_INTVL seconds to a single device

Some Challenges



- Mini-merge operations have always made use of a controller-assisted history log. This is not available with FC controllers.
 - A host-based Mini-Merge has been developed, scheduled for OpenVMS V8.2. A V7.3-2 release (via remedial kit) is planned for later this year.
- Sharing an EVA or MSA between systems with 1 Gb/s and 2 Gb/s adapters can result in poorer than expected performance.
 - Zoning or path balancing scripts can be used to work around this by isolating 1 Gb and 2Gb adapter access to different controller ports for an EVA.

Some Challenges



With 2 Gb adapters & the IO performance improvements in OpenVMS V7.3-1, systems & especially clusters can more easily produce more IOs than a single controller port can process.

- SAN controllers have command queues per port and use a **queue full** response as a flow control mechanism with a host.
 - EVA can queue up to 2,048 commands / port
 - MSA can queue up to 512 commands / port
 - HSG80 can queue up to 240 commands / port

HSG80s do not always handle this well.

- **Lower DIOLMs** reduce this, especially for VMS backup and database backup operations. Settings lower than 32-64 may hurt backup performance (100 is required for DECWindows).
- **Path balancing** also helps by spreading the IOs among the controller ports.
- Counter in QF Seen column in SDA> fc stdt /all

MSA 1000 & OpenVMS



- The MSA1000 requires at least OpenVMS V7.2-2 for fabric/switch topology.
 - In either case be sure to have the latest FIBRE_SCSI ECO installed for the version of OpenVMS.
 - Arbitrated loop (hub) technology is not supported on OpenVMS at this time. It will likely require at least OpenVMS V7.3-1 when added. If then using a hub, the KGPSAs must be set to LOOP topology in WWIDMGR.
`wwidmgr -show ada`
`wwidmgr -set adapter -item N -topo loop`
- The MSA1000 cannot be shared with other OSes, at least for now.
- Supported in a VMSCluster of up to 4 nodes.

MSA 1000 & OpenVMS

Redundant configurations

- In redundant configurations **only one** MSA1000 controller is **active** at a time. Since each controller has one FC port, only one port will be servicing all the disks at any one time.
- A disk fail-over is not as quick as with HS controllers since all the disks failover together.
 - Can only path balance between adapters (if you enough). If the current path for a disk is changed to a path on the other controller all the disks for that MSA1000 controller pair will fail over to the other controller. So set all paths to the same controller port.
 - Having only two ports per pair simplifies setting up a Shadowed System disk in WWIDMGR.

HSG & MSA Connections



- On an HSG/MSA a path is called a Connection. It is defined by the adapter WWID & controller port.
- Connections are automatically created when a system accesses the controller for the first time
 - “wwidmgr –show wwid”, booting OpenVMS, etc. - should create any connections as needed.
 - Connections can be created manually – but that is not generally recommended
- HSG80s have a limit of 96 defined connections. It is recommended to delete unneeded or inactive connections.

> DELETE connection-name

Note a connection cannot be deleted if it has access paths currently enabled

HSG/MSA Connections



```
HSG80TOP >show connection
```

Connection							Unit
Name	Operating system	Controller	Port	Address	Status		Offset
!NEWCON13	VMS	OTHER	1	011200	OL other		0
	HOST_ID=2000-0000-C922-4C03			ADAPTER_ID=1000-0000-C922-4C03			
!NEWCON14	VMS	OTHER	2	011600	OL other		0
	HOST_ID=2000-0000-C922-9C4C			ADAPTER_ID=1000-0000-C922-9C4C			
!NEWCON15	VMS	THIS	1	011200	OL this		0
	HOST_ID=2000-0000-C922-4C03			ADAPTER_ID=1000-0000-C922-4C03			
!NEWCON16	VMS	THIS	2	011600	OL this		0
	HOST_ID=2000-0000-C922-9C4C			ADAPTER_ID=1000-0000-C922-9C4C			

Very important the “operating system” for HSG or profile for MSA is set correctly – the default HSGs is WINNT or the setting of the last connection set, and for MSAs is Windows.

HSG80> SET connection-name OPERATING_SYSTEM=VMS

MSA> SET CONNECTION conn-name PROFILE=OpenVMS

Recommend re-naming connections (max. of 9 characters on HSG and 16 on MSA)

HSG80 or MSA> RENAME !NEWCON13 LildogA
or via ADD CONNECTION



Switch settings

Brocade switches - others have similar settings



Setting the switch domain_id can aid in troubleshooting

- To set the FC switch domain id at the telnet console, first disable the switch:

```
admin> switchDisable
```

```
admin> flashset !Must be unique in the fabric and  
domain_id [0]: 9 !if possible unique for the site.
```

```
admin> switchEnable
```

```
MAIN09:admin> switchShow
```

```
switchName:      MAIN09
```

```
switchType:      2.4
```

```
switchState:     Online
```

```
switchRole:      Principal
```

```
switchDomain:     9
```

```
switchId:        fffc01
```

```
switchWwn:        10:00:00:60:69:10:5c:de
```

```
port 0: sw Online F-Port 50:00:1f:e1:00:00:f7:03
```

```
port 1: sw Online F-Port 50:00:1f:e1:00:00:f7:01
```

```
port 2: sw No_Sync
```

```
port 3: sw Online F-Port 10:00:00:00:c9:24:36:2b
```

```
port 4: sw Online F-Port 50:00:1f:e1:00:05:63:d4
```

```
port 5: sw Online F-Port 50:00:1f:e1:00:05:d8:81
```

```
port 6: sw Online F-Port 10:00:00:00:c9:23:81:e1
```

```
port 7: id Online F-Port 50:00:1f:e1:00:05:d8:83
```

```
port 8: id Online F-Port 10:00:00:00:c9:22:47:5f
```

```
port 9: id Online F-Port 50:00:1f:e1:00:05:63:d1
```

```
port 10: id No Light
```

```
port 11: sw Online F-Port 10:00:00:00:c9:22:79:c3
```

```
port 12: sw Online F-Port 10:00:00:00:c9:23:01:4d
```

```
port 13: -- No_Module
```

```
port 14: -- No_Module
```

```
port 15: -- No_Module
```

HSG/MSA Console Access



- The only fully functional HSG/MSA management access is via the console port.
 - Scripting access for most HSG commands is available via the SANWorks Command Scripter product. It utilizes the SWCC agent to access the controller.
 - HSZTerm (Set Host/SCSI) is officially not supported. Using while a SWCC agent is running can crash the system.
 - A good console manager is the best solution, providing:
 - Console access (via the web and command line/terminal)
 - Scripting access
 - Console output logging. This is very important as this will capture all error output from the controller.
- HSGs do not have names - recommend setting the prompt



EVA Console Access



- Full management access for the EVA is provided via the SAN Management Appliance.
 - CommandView EVA provides a GUI interface with full management access.
 - The Enterprise Storage System Scripting Utility (SSSU) provides a limited but quite useful command line interface from the host system. It can be used to save and restore an EVA configuration.
 - The SMA can be set up with WEBES and/or SNMP traps to report errors and events.
- A second SAN Appliance can be configured for backup/fail-over.
- A separate management zone is recommended for the SAN Appliance(s) and the managed EVAs.

Additional Resources



- Section 6.6 & 6.7 and Chapter 7 of the Guidelines for OpenVMS Cluster Configurations manual contain excellent OpenVMS Fibre Channel configuration information.
 - The OpenVMS Documentation set is available on the web at http://h71000.www7.hp.com/doc/os732_index.html
- For the latest white papers, including “Fibre Channel in a Disaster-Tolerant OpenVMS Cluster System”, and information about fibre channel on OpenVMS visit the web page
<http://h71000.www7.hp.com/openvms/fibre/>
- For documentation on the EVA visit the web page
<http://h18000.www1.hp.com/products/storageworks/enterprise/index.html>
- For documentation on the MSA 1000 visit the web page
<http://h18000.www1.hp.com/products/storageworks/msa1000/index.html>

Summary



- **Cabling consistency** for systems & controllers is key
- Name HSG/MSA connections, set the O/S, & set switch host_ids
- Develop standard for EVA lun presentation numbers, OS_IDs, and VDisk names.
- Maintain SAN configuration documentation in common formats that can be shared via e-mail – Cable labels helps too
- Keep current on firmware (controller, adapter & switch) and ECOs. Do not mix versions of firmware.
- If path balancing, create a command procedure. As of OpenVMS V7.3-1 execute it after mounting disks during system startup.
- Path changes can cause mount verifications and mount verifications can cause path changes.
- **Remember to IO AUTOCONFIGURE** on all Cluster FC nodes when adding devices/units.



Questions?



Hear more about our service offerings by visiting us in the **Solutions Showcase!**



- HP Web Support Tools
- HP Active Savings Tool
- HP Education
- HP Business Continuity & Availability Solutions
- HP Adaptive Enterprise Agility Assessment
- HP Financial Services
- HP IT Consolidation Solutions
- HP Radio Frequency Identification (RFID)





Co-produced by:

