



Tips for setting up SANs for hp OpenVMS

Manfred Kaser
HP Services

Vortrag 3D04

Overview

Fibre channel advantages

- SANs – multi-vendor, multi-OS, multi-system detached storage – industry common components
- High performance, cabling distance & redundancy

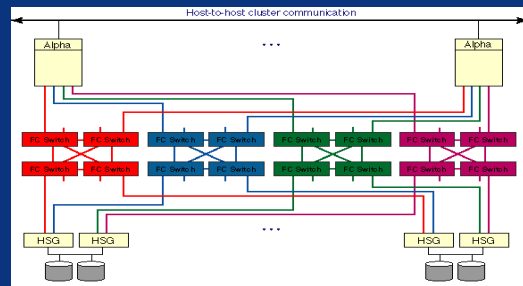
Challenges

- ~~Mini-merge with ACS 8.7 (HSG80 firmware), OpenVMS V7.3-1 plus appropriate ECOs !~~
- Learning a new technology
- Can grow to be complex

Setup – HSG80 cabling with 4 SANs

One option -- each HSG80 port on a SAN with 4 SANs

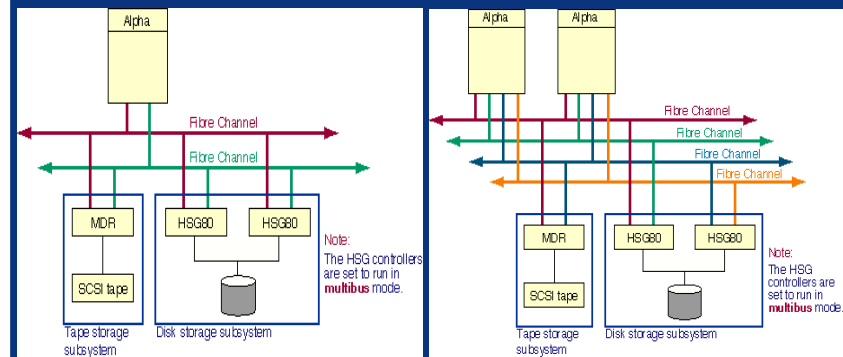
- Highest redundancy/performance/capacity
- Not currently supported with HSV110s - the SAN Appliance must see all the ports on the HSV



Setup – Adapter cabling

Typically one adapter per SAN

- high redundancy and high performance



Setup – Cabling suggestions

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

Setup – Device naming

- An HSG/HSV presents two device types to OpenVMS
 - \$1\$GGAnnnnn: - console lun for the HSG/HSV pair
 - \$1\$DGAnnnnn: - disk unit (HSG), virtual disk (HSV)
- There is NO allocation class setting for an HSG/HSV
 - the 1 in \$1\$ is always 1 and cannot be changed
 - recommend using the first digit or two of nnnnn to designate the HSG/HSV 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\$xDGAnnnnn: is set via an Identifier, not by the controller name, or disk unit number dNNN (HSG) or virtual disk name vdNNN (HSV)
 - identifiers have the range 0-32767
 - note – MSCP only serves devices named 0-9999
- The console lun Identifier for an HSG/HSV pair:
 - is set via the SAN Appliance and the HSV Element Manager for the HSV and via the command SET THIS IDENTIFIER=nnnnn on the HSG
 - is required for OpenVMS to see HSG/HSV disks
 - must be unique among HSG/HSV pairs and disks, and be non zero

Setup – Device naming

A disk unit Identifier (HSG), virtual disk OS Unit ID (HSV):


- is set on the HSV via the SAN Appliance and the HSV element manager
- is set on the HSG via the SET dNNN IDENTIFIER=nnnnn
- is required for each unit for OpenVMS to see the disk or virtual 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 SAN Appliance on the HSV or via SHOW dNNN or SHOW UNITS FULL on the HSG

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

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

- HSG HSV
- 0 0 = the controller pair
- 1 9 = Port 1 bottom controller
- 2 8 = Port 2 bottom controller
- 3 D = Port 1 top controller
- 4 C = Port 2 top controller



View HSV WWIDs from the SAN Appliance, and HSG WWIDs with SHOW THIS/OTHER

- Can also be viewed from OpenVMS (more later)

WorldWide IDentifiers (WWIDs)

View KGPSA adapters and WWIDs at the AlphaServer SRM console with: "set wwidmgr -show adapter"

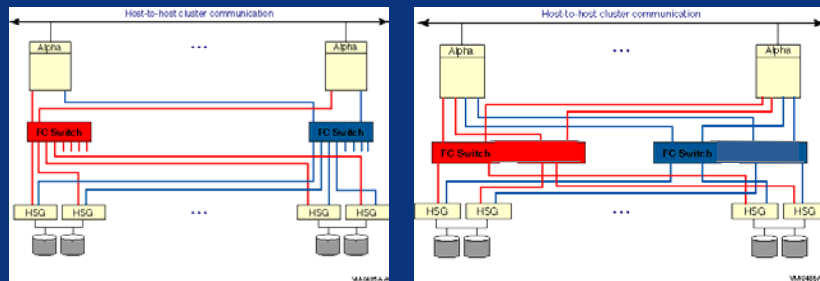
- KGPSA WWIDs typically begin with 1000 or 2000

```
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 on the adapter device, ie: \$ SHOW DEVICE/FULL FGA0:
- Can also be viewed from the HSG/HSVs (more later)
- All WWIDs in the SAN can be viewed from the switches
- Document all adapter and controller WWIDs as they will be needed for troubleshooting

Paths / Connections

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



4 paths per disk unit
on each system

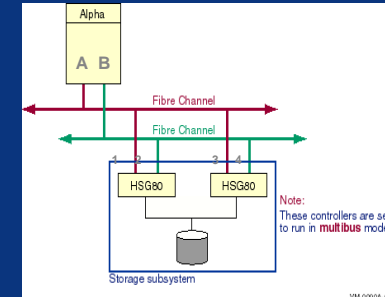
8 paths per disk unit
on each system
*HSGs may only be connected to 2
SANs

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
 - it is connected if the path's HSG/HSV port is on the controller currently serving the disk
 - it is not connected if disk unit would have to fail-over to the other HSG/HSV to be accessed on the path

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



Note:
These controllers are set
to run in **multibus** mode.

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
 - only 4 devices and 4 paths can be assigned

Using `wwidmgr`

- The `wwidmgr` users guide is on the AlphaServer console CD in the doc directory as `wwidmgr.pdf` and `wwidmgr.ps`
- On AlphaServer 8x00, 4x00, or 1200 systems, the console must be in diagnostic mode before using `wwidmgr`.
 - P00>>> set mode diag
- Since the context is not the same as when the system boots an INIT may be necessary before entering diag mode.
- After using `wwidmgr` the system must be re-initialized before booting or to exit `wwidmgr` mode.
 - P00>>> init
- `Wwidmgr` commands have the syntax:
 - P00>>> `wwidmgr command [-option [parameter]]`

Using wwidmgr

Remember: HSG/HSV 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>

OpenVMS & FibreChannel paths

Remember – a path is defined by two endpoints, such as two WWIDs. OpenVMS uses the adapter name and the WWID of the HSG/HSV port. For example:

```
$ show device /fu $1$DGA706
Disk $1$DGA706: (MUN01), device type HSG80, is online, file-oriented device,shareable,
device has multiple I/O paths, served to cluster via MSCP Server, error logging is enabled.
. . . I/O paths to device                    5
Path MSCP (MUNAXP), primary path.
Error count                                0 Operations completed                0
Path PGA0.5000-1FE1-0004-EDB1 (MUN01).
Error count                                0 Operations completed                0
Path PGA0.5000-1FE1-0004-EDB3 (MUN01), current path.
Error count                                0 Operations completed                52304
Path PGA0.5000-1FE1-0004-EDB4 (MUN01).
Error count                                0 Operations completed                0
Path PGA0.5000-1FE1-0004-EDB2 (MUN01).
Error count                                0 Operations completed                0
```

- 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#####
```

The current path (or UCB) is the one to be used for the next I/O for the device.

- Changing the current path manually has no effect on the HSG/HSV controllers or the rest of the OpenVMS Cluster until the next I/O is issued.
- When an I/O is attempted, if the current path is:
 - bad another will be chosen via a mount verification
 - not connected, then the HSG/HSVs 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 and new paths chosen via mount verification

OpenVMS & Changing paths

A path can be set manually via

```
$ SET DEVICE/PATH=device-path/SWITCH
```

A path can be changed automatically by a mount verification.

- Starting with OpenVMS V7.2-2 (or earlier with VMS721_SYS-V0100 or VMS721H1_SYS-V0600) mount verification attempts the current path first.
- Before OpenVMS 7.2-2 mount verification attempts the primary path first

Starting with OpenVMS V7.3-1, mounting a multipath disk can result in a path change.

When a path is changed an OPCOM message is generated which includes whether the change was manual or automatic. Note that 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      Device      Error Paths Current
Name        Status      Count Avl/Tot 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      Device      Error Paths Current
Name        Status      Count Avl/Tot path
$1$GGA99:   Online     0   4/ 4  PGB0.5000-1FE1-0005-35E2
$
$ sho dev dg/mult
Device      Device      Error Paths Current
Name        Status      Count Avl/Tot 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.

Observation is most sites choose to do path balancing. But, is it always needed?

- With fibre channel's high performance characteristics, bottlenecks are less common. So, path balancing may not improve performance.
- However, it does enhance failure detection.

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

New with OpenVMS V7.3-1

Automatic path balancing for multipath disk and tape devices occurs during certain events.

- 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 the connected path presently used the least as a current path.
- Benefits include:
 - Minimizing the disruption on other hosts in the cluster.
 - Tends to preserve any manual load balancing 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 the use of available paths from this host to the disk devices.
 - Prefers direct paths over MSCP served paths.

© 2003 hp

DECUS Bonn 2003

page 23

OpenVMS & Path Balancing

New with OpenVMS V7.3-1

Path selection during the MOUNT of a multipath disk, is as follows:

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 checked. They 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, and an attempt is made to fail over the device on the HSG/HSV pair 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 trigger this path selection algorithm a number of times until a working path is found. The number of retries depends on the elapsed time and the qualifiers specified with the command.

© 2003 hp

DECUS Bonn 2003

page 24

OpenVMS & Path Balancing

New with OpenVMS V7.3-1

Automatic path selection algorithms give meaning to preferring paths on the HSG/HSV 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

- HSVs have preferred path and preferred fail-back path selection via the SAN Appliance and the 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 on the OpenVMS systems.

OpenVMS & Path Balancing

New with OpenVMS V7.3-1

Recommended procedure to balance paths:

- Set the preferred HSG/HSV controller (top/bottom) to serve each unit. On an HSV, optionally set the fail-back paths.
- 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 require consistent cabling from node to node.

OpenVMS & Adding FibreChannel Devices

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

When a unit is added on an HSG/HSV, 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. Recommended to do this cluster-wide.

```
$ MCR SYSMAN
SYSMAN> SET ENVIRONMENT/CLUSTER
SYSMAN> IO AUTOCONFIGURE
```

Prior to OpenVMS V7.3-1 (or in V7.3-1 with MSCP serving disabled), in an OpenVMS Cluster, it was very important to autoconfigure on all nodes before using a new device.

- Disabling MSCP serving (if not needed) of SAN devices was also recommended prior to OpenVMS V7.3-1

OpenVMS & FibreChannel paths

New with 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
```

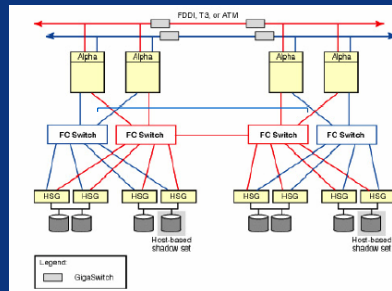
Two-site Host Based Volume Shadowing

With OpenVMS V7.2-2, HBVS over extended SANs is supported. It is also available for OpenVMS V7.2-1 and V7.2-1H1 via a Shadowing ECO, but the "Enhanced Volume Shadowing functionality" must be chosen during installation.

- Adds new member timeout qualifier to force manual intervention in the event the SAN inter-site link is lost, but not the SCSI link. This avoids the possibility of split shadow sets.

Not necessary with OpenVMS V7.3-1 if disks are MSCP served.

- New /SITE and /READ_COST qualifiers were added since all the disks appear local.
- See the ECO release notes and/or the current Volume Shadowing for OpenVMS manual, section 4.7 - available on the web at <http://www.openvms.compaq.com:8000/index.html>



Firmware and ECOs

Fibre Channel is a relatively new and 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 V8.7-1, or ACS V8.6x-12. HSVs via the SAN Appliance, current is VCS V2.0.
- 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.

```
SDA> fc show device fga0
FGA0: operational firmware revision DS3.03X2
SDA> fc show device fgb0
FGB0: operational firmware revision DS3.03X2
```

- On OpenVMS check for FIBRE-SCSI, SYS, SHADOWING and UPDATE ECOs with PRODUCT SHOW HISTORY. Current OpenVMS ECOs are available at <http://ftp1.support.compaq.com/public/vms/>

HSG80 Connections

On an HSG a path is called a Connection. It is defined by the adapter WWID and the HSG 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

HSG80 Connections

```
HSG80TOP_>show connection
```

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

Very important the “operating system” is set correctly – the default is WINNT or the setting of the last connection set.

```
HSG80> SET connection-name OPERATING_SYSTEM=VMS
```

Recommend re-naming connections (max. of 9 characters)

```
HSG80> RENAME !NEWCON13 LildogA
```

HSG80s do not have names - recommend setting the prompt

```
HSG80> SET THIS PROMPT="HSG80_99Top> "
```


HSV Connections

On an HSV a path is called a Connection. It is defined by the Host adapter WWID and the HSV port.

Host adapter access is managed/configured using the HSV Element Manager from the SAN Appliance. 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 connections as needed.
- Hosts and host ports can be created manually – but that is not generally recommended

A host must be add/configured on the HSV before it can access any units. First you add a host, which includes one port (adapter). It is very important to set the operating system properly in the pull-down menu. Additional ports can then be added to a host.

It is recommended to delete unneeded or inactive hosts.

> DELETE connection-name

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

Switch settings

Setting the switch domain_id can aid in troubleshooting.

To set the FC switch host_id at the telnet console, the switch has to be disabled first:

```
admin> switchDisable
```

```
admin> flashset
```

```
domain_id [0]: 9 !Must be unique in the fabric and  
: !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/HSV Console Access

The only fully functional HSG management access is via the console port.

- Scripting access for most commands is available via the SANWorks Command Scripter product. It utilizes the SWCC agent to access the controller.
- HSZTerm (Set Host/SCSI) works via the GGA device. It is officially not supported. Using while a SWCC agent is running can crash the system. Provides access to all commands and utilities except VTDPY.

A good console manager is the best solution for an HSG 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.

Full management access for the HSV is provided via the SAN Appliance. Both GUI (SANWorks Element Manager) and character-cell (Enterprise Storage System Scripting Utility (SSSU)) interfaces are available from the Appliance.

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://www.openvms.compaq.com:8000/index.html>

For details on the fibre channel related enhancements in OpenVMS V7.3-1, see sections 4.9 and 4.16 of the OpenVMS Alpha Version 7.3--1 New Features and Documentation Overview

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://www.openvms.compaq.com/openvms/fibre/>

For documentation on the StorageWorks enterprise virtual array (HSV) visit the web page

<http://www.compaq.com/products/storageworks/enterprise/documentation.html>

Summary

Cabling consistency for systems & controllers is key

Name HSG80 connections, set the O/S, & set switch host_ids

Maintain SAN configuration documentation in common formats that can be shared via e-mail

Keep current on firmware (controller, adapter & switch) and ECOs.
Do not mix versions of firmware.

If path balancing, create a command procedure. Prior to OpenVMS V7.3-1 execute it before 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.

Console management is recommended for HSG80s.

Questions?



i n v e n t