

A Guide to Success with the A5814A-003 SCSI-FC Router

(or)

A bunch of stupid things I've done and seen that you should avoid

BY: Jim Hawkins -- MPE/iX Lab
April 21st, 2003

The A5814A-003 SCSI-FC Router (a.k.a. "Router") is a unique piece of hardware. In our collective experience in the MPE/iX Lab we have found it to be a very reliable, high performing, and trustworthy part of the data center environment. That being said a A5814A-003 can cause you a lot of pain if you don't carefully plan your approach. 99.999% of the time "a problem with the Router" is caused by other components in the FC-SAN or, most often, a lack of understanding of the capabilities of the Router and of underlying FC-SAN principles.

This document is broken into sections which, at a high-level, provide all the information you need. They are:

- I) Rule #0: Routers don't handle change well.
- II) Really Surprising Fact #1: Routers "talk" to other Routers, Avoid this!
- III) Tip #1: FC Switch Zoning, your best friend.
- IV) Tip #2: LUN Security, tricky and best avoided.
- V) Really Surprising Fact #2: Windows systems can now STOMP on MPE/iX disks.
- VI) Rule #1: Plan, Plan, Plan, then implement carefully.
- VII) Examples, good and bad.

Exercising Good Judgment requires Experience.
Experience comes from exercising Bad Judgment.

Rule #0: Routers don't handle change well.

Any change to the FC environment visible to the Router will likely result in a Router marking its device map invalid. This will prevent you from accessing the device from MPE/iX which may result in System Hangs or Aborts. If a System Hang or Abort isn't scary enough for you the ONLY way to reset the Router Map is via a Double Power Cycle of the device while flipping switches too small for adult human fingers (and hard to see to boot)!

Activities which can cause a Router to Invalidate its Map include but are not limited to:

- A) Connection of an FC Cable in the FC-SAN.
- B) Adding/Removing another router (or powering up another router).
- C) Adding/Removing a LUN to device PORT.
- D) FC-Switch Zone changes.
- E) Change in LUN Security settings.

Due to this sensitivity the WORST possible strategy is "plug some Routers and see what happens" -- this is the path to Chaos or at least a headache. Close behind in the "creating a big mess sweepstakes" is the "I'll just add a few more LUNs to this Device Port or reZone this switch to add another Router" strategy. Likely this will cause existing Routers connected to the Device Port to stop working. Why is this?

Really Surprising Fact #1:

Routers "talk" to other Routers.

Routers are cautious and like to form a consensus of how to "map" all the devices visible in the FC-SAN. If two Routers can communicate they WILL exchange maps. If the maps do not agree then what the Routers think is the "older" map will be invalidated. You will then be forced to MANUALLY reset the Router with the "invalid" map. If a Router can "see" a device port through any set of FC connections than it can and will communicate with any other Router that can "see" the port. This is true even if there are no LUNs configured on that port! Therefore your best path to success is to limit the number of Routers connected to any device port and the best mechanism for this is FC Switch "Zoning".

Tip #1:

FC Switch Zoning, your best friend.

Any FC Switch worth anything allows you to "Zone" the ports on the switch. That is logically group a sub-set of ports such that devices connect to ports outside the group/Zone cannot communicate with the devices inside this Zone and vice versa. (A "Switch" without Zoning is really a "Hub". Hubs, though they work, should be avoided as they leave you open to too many problems.) Your absolute safest and best Zone configuration is one with two Ports, one to connect to a Router and one to connect to a Device Port. This will prevent this Router from talking with any other Router or Device Luns on any other device port.

"Ah," you say, "but, I can have multiple routers in my zones because I'll use LUN Security to virtually link the Router to my Device LUN". WRONG!

Tip #2:

LUN Security: tricky and best avoided.

LUN Security is a feature of most Storage Devices that allows one to limit the visibility and availability of a particular Device LUN to a single accessing device. This is generally accomplished by entering an FC-SAN hardware unique address (a.k.a. WWN) into a Device dependent configuration tool. Routers do have WWN values which can be used. Sounds good, but don't forget 'Really Surprising Fact #1: Routers "talk" to other Routers.'

Consider this scenario:

You create six LUNs on a Device Port.

You connect two Routers via a Switch to the Port.

You configure three LUNs to Router One WWN and three to Router Two WWN.

If you turn on Router One it will see its three LUNs. Turn it off.

If you turn on Router Two it will see its three LUNs. Leave it on.

Turn on Router One. It will see its three LUNs AND talk with Router Two.

Router Two will notice it's map is out of date BECAUSE it CANNOT "see" the LUNs that Router One can see.

Router Two map is marked invalid.

If you "Manually Reset" Router Two it sees its three LUNs and Router One map will be marked invalid.

Repeat ad nauseum.

Simply put: You CANNOT use LUN security to limit the view of LUNs to Routers that connect to the same device port. The ONLY time LUN Security can be used is if there is ONE and only ONE Router connected to a Device Port (either directly or via Zoning).

Really Surprising Fact #2:

Windows systems can now STOMP on MPE/iX disks.

Various versions of "Windows" type systems have a "Disk Administrator" tool (windisk.exe, diskmgmt, etc.) During boot if "new" disks are discovered it will give you the option to "claim" the disk. NEVER do this unless you are willing to risk losing data! If given the option disable this "feature."

The BEST option is to use Device Port settings, FC Switch Zoning and even LUN Security to prevent a "Windows" system from seeing your MPE/iX Disk Volumes. Any O.S. can cause this to happen but "Windows" systems tend to be the most aggressive. Your best protections are:

- A) Limit each Device Port to only one Type of O.S.
- B) Use FC Switch Zoning to limit the number of ports that are Zoned together.
- C) Where possible, LUN Security can help prevent problems.

"Multiple OS Access" can cause the following problems:

- o ":DSTAT ALL" suddenly shows device as UNKNOWN.
- o Session or System Hang (due to inability to complete I/O).
- o System Abort (due to Disk I/O failures in critical places like XM).
- o "Random" data corruption, especially in Volume Label areas.

Rule #1: Plan, Plan, Plan.

As with ANY data center deployment the best path to success is a careful design your FC-SAN in advance. Planning really should be "Rule #0" but most people already know that so their eyes glaze over and they stop paying attention. Hopefully, if your eyes haven't glazed over I have fully impressed upon you the simple fact that this ain't so simple. Good planning will really truly save you major headaches.

Following Rule #0 your goal is to create a STABLE "view" for a Router before you start trying to configure MPE/iX disks volumes. If you are FORCED to expand an existing configuration it is much easier and safer to expand into virgin territory of previously unused Device and FC Switch Ports. It can be much harder to accommodate "just one more LUN on a port" than you might think.

Make sure you understand the number of Systems, SCSI HBAs, Routers, Disk LUNs, Device Ports, Storage Devices and the resulting number FC Port and FC Switches to tie it all together (limit of no more than two levels of switches between Router and Device Port). Ideally you should also allow additional Device and Switch Ports to handle expansion or to correct a mistake in your original design. (A Spreadsheet is a really great tool to accomplish this)

- A) Create a list of all Systems and HBA/Routers that will be connected.
- B) Create a list of the desired number/size of LUNs per HBA/Router.
(recommend Device LUNs <8GB in size no more than 15 per HBA/Router)
- C) Create a list of Device Ports -- minimize the number of Routers per Port!
- D) Create list to allocate each Disk LUN to a device PORT
- E) "Merge" A-D into one list.
(Where you have more than on HBA/Router per Device Port you must then have an FC Switch to "join" these connections).
- F) Merge into "E" the list of all FC-Switch Ports.
- G) Note the Zoning required to minimize the number of Routers on each connection.
- H) Determine the number of FC Switches base upon number of Ports required.
- I) From G-H Document the Zoning for each switch.

Your result should be a document that can tell you which MPE/iX LDEV corresponds to which Storage Device LUN AND what path you follow to get there AND who else shares this same path.

Once the planning activities are complete. Implementation can follow:

- J) Make sure you have all Device, Switch, Router Documentation.
(NSS SPOCK)
(New router document)
- K) BEFORE YOU MAKE A SINGLE CONNECTION Make sure ALL devices have up-to-date firmware. Check NSS Streams device for current recommendations for Router,

Switch and Storage firmware levels.

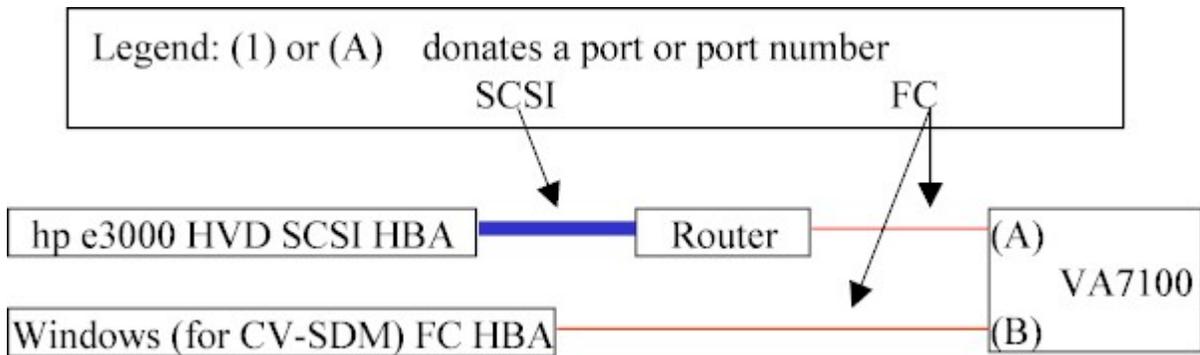
- L) Create all Storage Device LUNs and assign to Ports according to Plan.
- M) Create all FC Switch Zone according to Plan.
- N) Connect FC Cables from Switches to Devices according to Plan.
Confirm good connections.

You now have a "stable" environment that Routers are likely to be happy in.

- O) Connect Routers to System HBAs (Power OFF) according to Plan.
- P) Individually Turn Router Power On and Manually Rest Map.
- Q) One at a time connect the Router FC cables according to Plan.
- R) Verify System Connections via ODE Mapper/Mapper2.
Check that ALL expected LUNs are preset according to Plan.

Good Example #1:

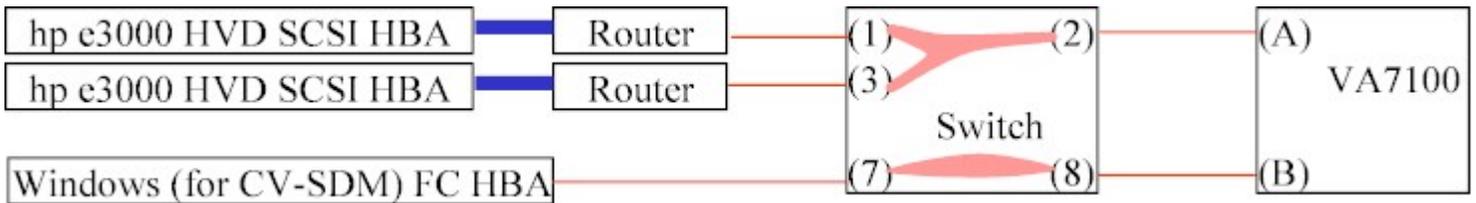
Minimum configuration, ONE HP e3000 HBA hooked directly to VA7100:



Notes: Recommend that HP e3000 be connected to the "performance port" of the VA7100, generally this is the port on Controller 1 of 2.

Good Example #2.1(OK):

TWO HP e3000 HBA hooked to a VA7100 (HBAs could be on different systems).

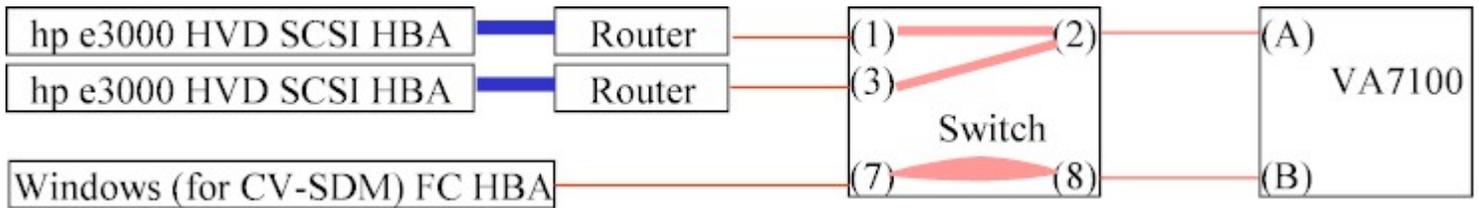


Notes:

- a) As you have two Routers and only one port you must use a FC Switch to multiplex some connections.
- b) To avoid "surprising fact #2" you MUST use "Zoning" on the Switch to isolate the HP e3000 Routers from the Windows System. Zone "HP" would contain Ports 1, 2, 3. Zone "Windows" Ports 7 and 8. .
- c) Do NOT attempt to use LUN Security on LUNs on Device Port A!

Good Example #2.2(Recommended):

TWO HP e3000 HBA hooked to a VA7100 (HBAs could be on different systems).

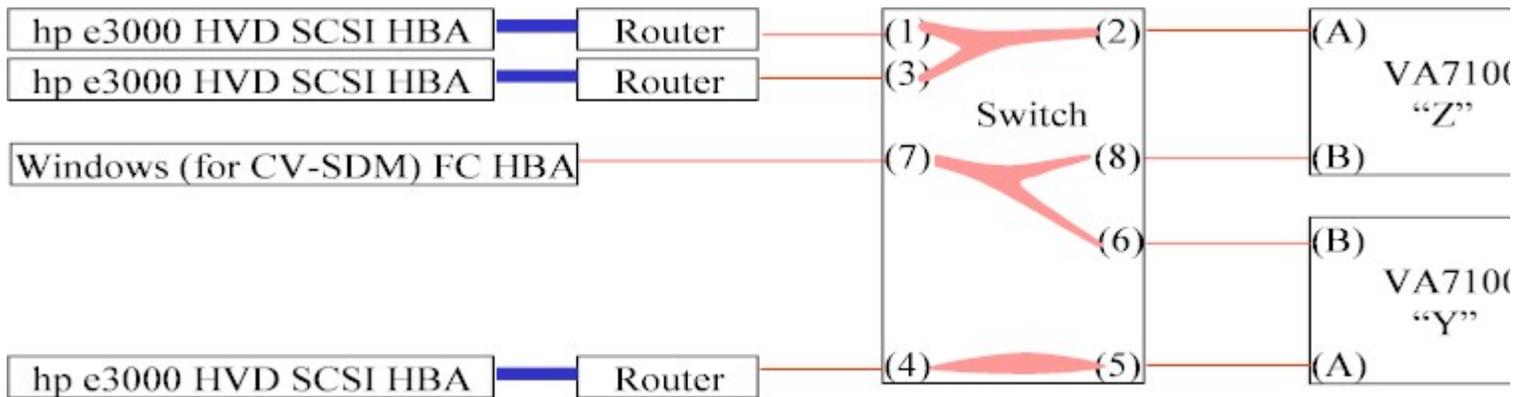


Notes:

- As you have two Routers and only one port you must use an FC Switch to multiplex these connections.
- To avoid "surprising fact #2" you MUST use "Zoning" on the Switch to isolate the HP e3000 Routers from the Windows System.
- Use "**Single Initiator Zoning**"; Zones: "Router1" (1,2), "Router2" (3,2), "Windows" Ports 7 and 8. .
- With Current Brocade products you should be able to use LUN Security!

Good Example #3

Three HP e3000 HBA/Routers hooked to two VA7100s which are managed by a single CV-SDM workstation.

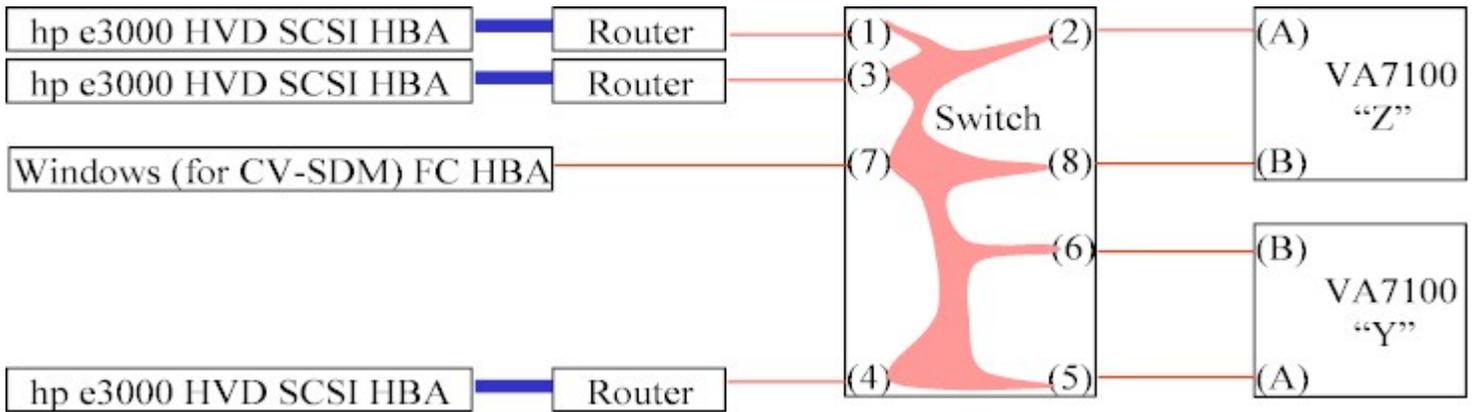


Notes:

- Zones are (1,2,3) (7,8,6) and (4,5)
- LUN Security could be used on VA7100 "Y" Port A,B and "Z" Port B.
- LUN Security NOT allows on VA7100 "Z" port A.

Bad Example #3a

Three HP e3000 HBA/Routers hooked to two VA7100 which are managed by a single CV-SDM workstation.

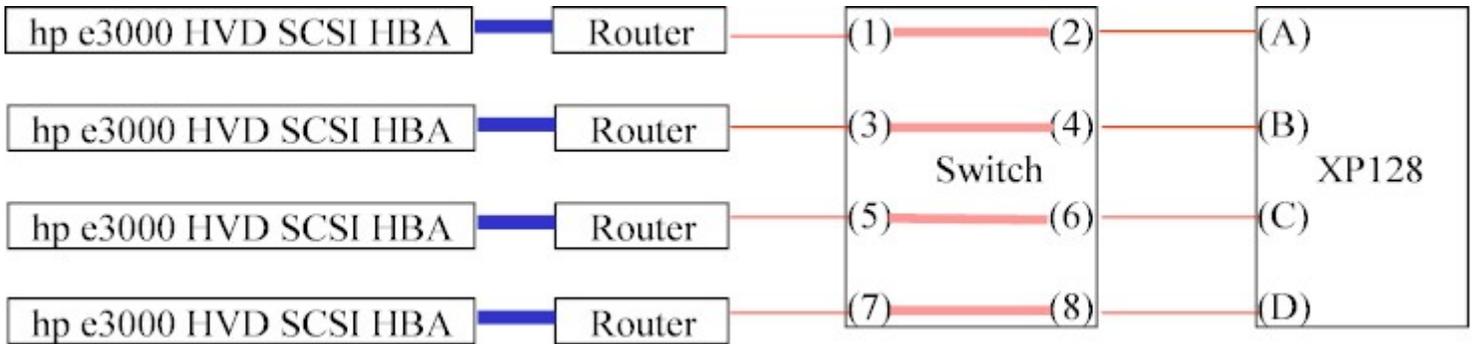


Notes:

- a) Zoning NOT use so (1,2,3,4,5,6,7,8) are all together.
- b) LUN Security cannot be used on ANY Device Port.
- c) Windows can easily write on all MPE LUNs (and vice versa)
- d) this is BAD don't do it!!!

Good Example #4

Four HBA/Routers connected to XP48/512 or XP128/1024 with lots of ports:

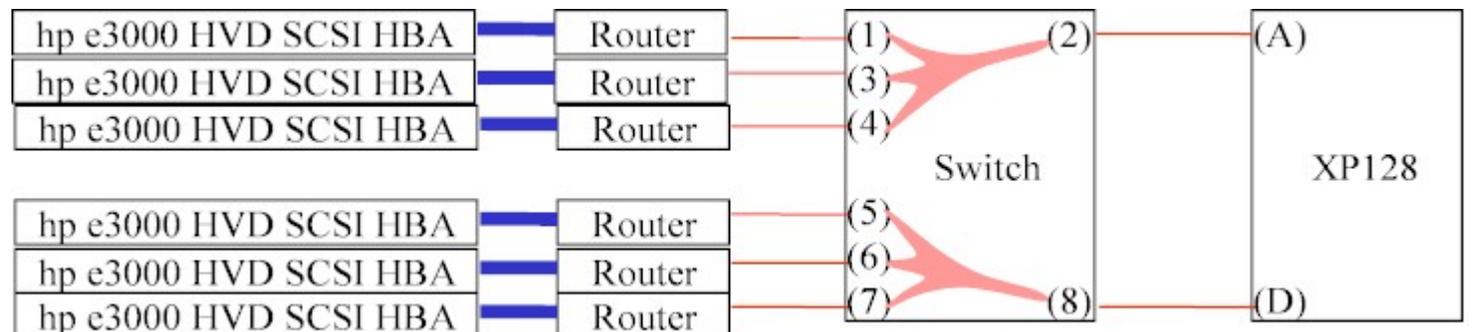


Notes:

- a) Zoning gives each Router exclusive access to one XP Port.
- b) LUN Security could be used
- c) A one Router, one Port configuration doesn't actually require a switch. . .
- d) (This could be expensive). . .

OK Example #5

Six HP e3000 HBA Routers connected to XP48/512 or XP128/1024 with few ports.



Notes:

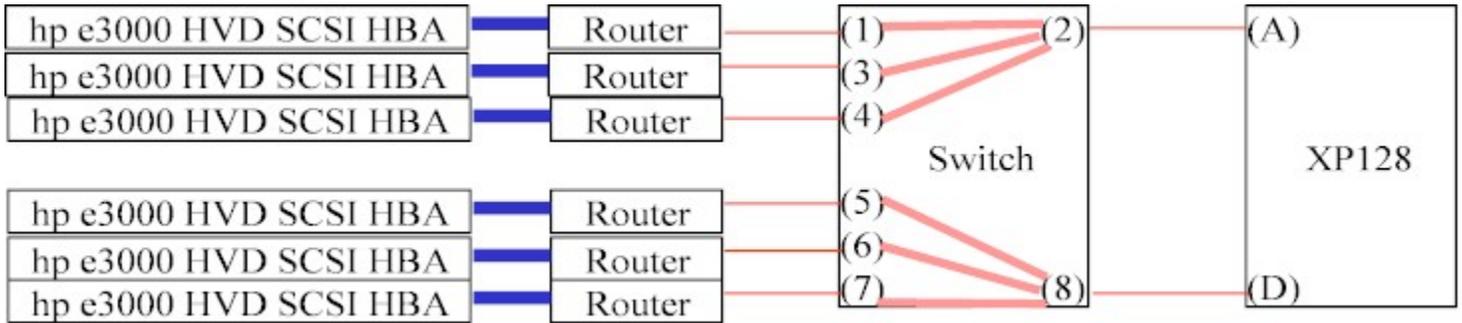
- a) Zoning (1,2,3,4) (5,6,7,8)
- b) LUN Security can NOT be used as Routers (1,3,4) share maps (as do 5,6,7)
- c) Lots of Routers grouped together.
- d) Switch is "single point of failure"

Great Example #6

Six HP e3000 HBA Routers connected to the XP48/512 or XP128/XP1024 with few ports (port consolidations)

This technique uses a feature in the switch called "Single Initiator Zoning". The switch is configured so that each zone contains only the router's port and the storage port. The first zone would be port (1,2) the second zone would be port (3,2) and the third zone is port (4,2).

This allows each router to "talk" to a single storage array port (for port consolidation) but doesn't allow the routers to talk to each other. With this configuration it is possible to implement Lun Security.



Notes:

- a) Zoning (1,2) (3,2) (4,2) (5,8) (6,8) and (7,8)
- b) LUN Security can be used because the Routers can't ("see" each other) and share maps
- c) Lots of Routers grouped together.
- d) Switch is "single point of failure"