



Power Systems SAN Multipath Configuration Using NPIV

V1.2

Bejoy C Alias

IBM India Software Lab

Revision History

Date of this revision: 27-Jan-2011	Date of next revision : TBD
---	------------------------------------

Revision Number	Revision Date	Summary of Changes	Changes marked
V1.0	23-Sep-2010	1. Initial creation	N
V1.1	27-Oct-2010	1. Included bookmarks for figures 2. Modified section 2.4 to include details about Live Partition Mobility (Thanks to Bob Foster for pointing out this missing part)	N
V1.2	27-Jan-2011	1. Updated section 2.3.2 with review comments from Dan Braden (IBM US)	N

Preface

This article explains the steps for Power Systems SAN multipath configuration using the SAN port virtualization feature NPIV. Using this virtualization feature, the SAN HBA port utilization can be increased and based on the bandwidth requirements it can reduce the overall HBAs required for SAN access. This document shows the configuration done using IBM AIX Operating System and shows the NPIV configuration part only. It does not explain the SAN switch configuration in detail or the details like VIOS installation, Virtual SCSI mapping, etc.

Table of Contents

1. Introduction	4
2. SAN Multipath Configuration using NPIV.....	4
2.1. Pre-requisites.....	4
2.2. VIOS Configuration	5
2.3. AIX LPAR Configuration.....	7
2.4. Configuring the SAN Storage.....	8
3. References	9

1. Introduction

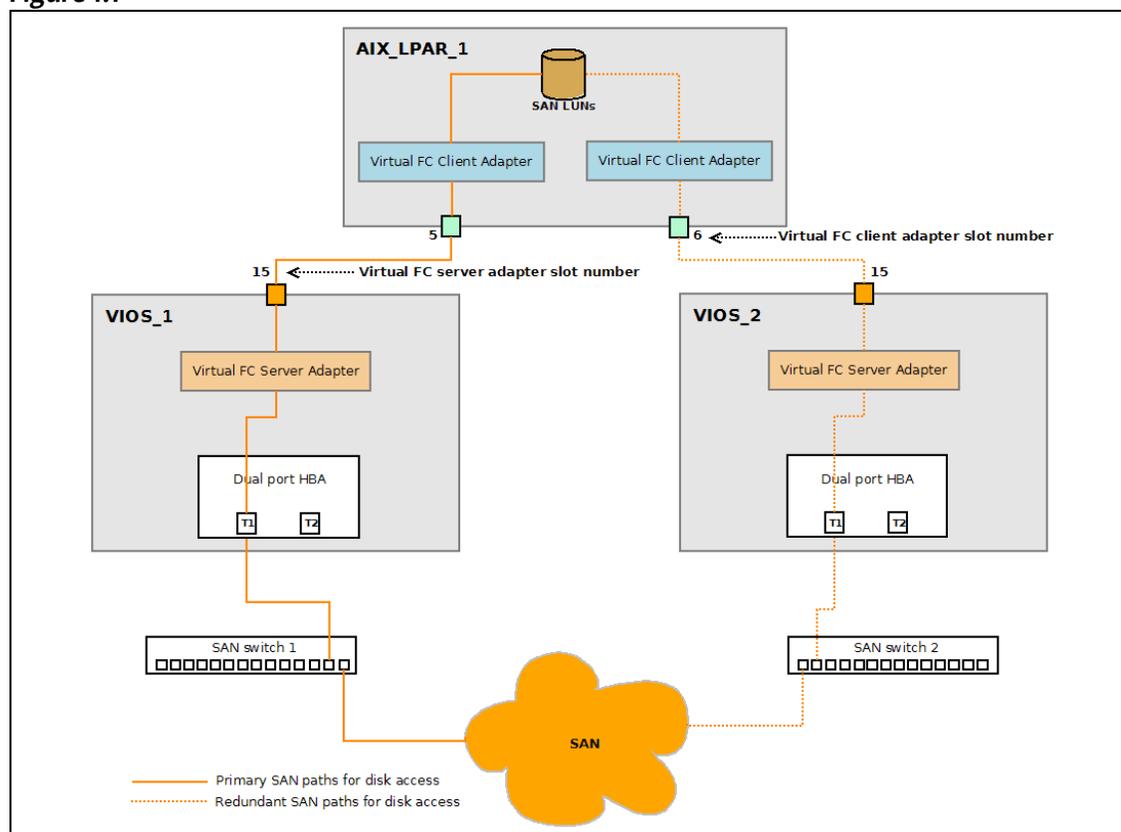
N_Port ID Virtualization or NPIV is a Fibre Channel facility allowing multiple N_Port IDs to share a single physical N_Port. This allows multiple Fibre Channel initiators to occupy a single physical port, easing hardware requirements in Storage Area Network design. An NPIV-capable fibre channel HBA can have multiple N_Port IDs, each with a unique identity and world wide port name.

On System p the Virtual I/O Server (VIOS) allows sharing of physical resources between LPARs including virtual SCSI, virtual Fibre Channel (NPIV) and virtual networking. NPIV allows System p logical partitions (LPARs) to have dedicated N_Port IDs, giving the OS a unique identity to the SAN, just as if it had a dedicated physical HBA(s).

For availing maximum redundancy for the SAN paths of the LPARs, the paths has to go through two different HBAs and those HBAs should be assigned to two different VIO Servers.

The below given diagram ([Figure 1.1](#)) shows the logical representation of the configuration explained in this document.

Figure 1.1



2. SAN Multipath Configuration using NPIV

This section explains the pre-requisites and the steps involved in configuring the multipath SAN connectivity for the Power Systems using NPIV.

2.1. Pre-requisites

This lists only those pre-requisites for the AIX Operating System which are required for doing the Multipath NPIV configuration.

- IBM POWER6 or above processor based hardware

- 8Gb PCIe Dual-port FC Adapter (Feature Code #5735)
- SAN switch supporting the NPIV feature
- VIOS version 2.1 or above
- AIX 5.3 Technology Level 09 or above
- AIX 6.1 Technology Level 02 or above

2.2. VIOS Configuration

You have to create dual VIOS instances on the System p server to avail maximum redundancy for the paths. Each VIOS should own at least 1 number of 8Gb FC Adapter (FC#5735) and one Virtual FC Server Adapter to map it to the client LPAR.

NOTE: Before doing any NPIV related configuration, the VIOS and AIX LPAR profile configuration has to be completed with all the other details required for running the VIOS and the AIX LPAR, like proper Virtual SCSI adapter and/or Virtual Ethernet Adapter mapping, physical SCSI/RAID adapter mapping to the VIOS and/or AIX LPAR etc. This is because, for doing a Virtual FC Server Adapter to Client Adapter mapping both the VIOS and AIX LPAR should exist in the system, NOT like Virtual SCSI adapter mapping where you have to provide only the slot numbers and it does not check for the existence of the client.

2.2.1. Creating Virtual FC Server Adapters

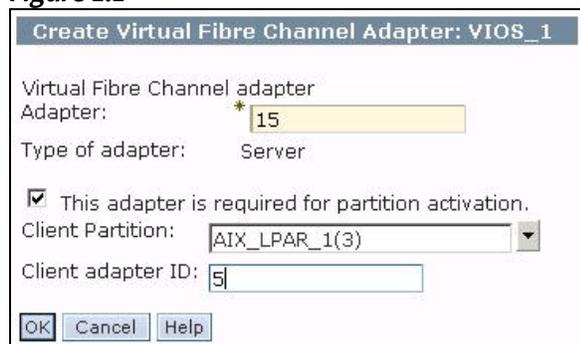
Open the first VIOS's profile for editing and navigate to the Virtual Adapters tab, click on Actions > Create > Fibre Channel Adapter as shown in the [figure 2.1](#)

Figure 2.1



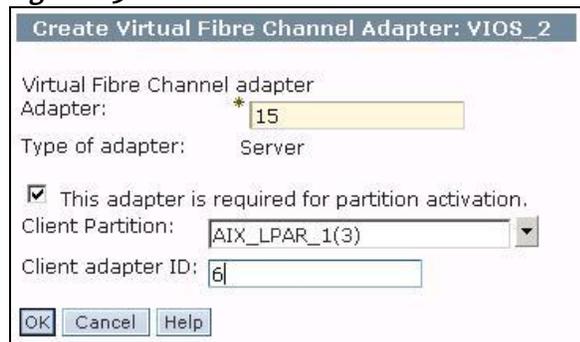
Input the Virtual FC Server Adapter slot number, select the client LPAR from the dropdown which will be connected to this server adapter and input its planned Virtual FC Client Adapter slot number as shown in [figure 2.2](#)

Figure 2.2



Same way, edit the profile of the second VIOS and add the Virtual FC Server Adapter details as shown in [figure 2.3](#)

Figure 2.3



If you have any more LPARs which you want to configure for NPIV, repeat the above mentioned steps with those LPAR details.

2.2.2. Mapping Virtual FC Server Adapters to the physical HBA ports

Once the Virtual FC Server Adapters are created, it has to be mapped to the physical HBA ports for connecting it to the SAN fabric. This will connect the Virtual FC Client Adapters to the SAN fabric through Virtual FC Server Adapter and physical HBA port.

For mapping the Virtual FC Server Adapters to the HBA port, the HBA port should be already connected to the SAN switch enabled with NPIV support. You can run the command `\lsnports` in VIOS `\$` prompt to know whether the connected SAN switch is enabled for NPIV. If the value for the `\fabric` parameter shows as `\1` that means that HBA port is connected to a SAN switch supporting the NPIV feature. If the SAN switch does not support NPIV, this value will be `\0` and if there is no SAN connectivity there won't be any output for the `\lsnports` command.

```
$ \lsnports
name          physloc          fabric tports aports swwpns awwpns
fcs0          U5802.001.00B0557-P1-C7-T1      1    64    64    2048    2048
fcs1          U5802.001.00B0557-P1-C7-T2      1    64    64    2048    2048
```

If the value for `\fabric` is `\1`, we can go ahead mapping the Virtual FC adapter to those HBA ports. In the VIOS, the Virtual FC adapters will be shown as `\vfchost` devices. We can run the `\lsdev -vpd | grep vfchost` command to know which device represents the Virtual FC adapter on any specific slot.

```
$ \lsdev -vpd |grep vfchost
vfchost0      U9117.MMA.6583005-V1-C15          Virtual FC Server Adapter
```

In our setup, it shows that the `\vfchost0` is the server adapter device which is on slot 15. If there is more number of Virtual FC adapters, we can easily identify the device name if you know the server adapter slot number.

Now we need to map the device `\vfchost0` to the physical HBA port `\fcs0` using the `\vfcmmap -vadapter vfchost0 -fcp fcs0` command. Once it is mapped, check the status of the mapping using the command `\lsmmap -vadapter vfchost0 -npiv`. Please note that the status of the port is showing as `\NOT_LOGGED_IN`, this is because the client configuration is not yet complete and hence it cannot login to the fabric.

```

$ vfcmap -vadapter vfchost0 -fcp fcs0
$ lsmap -vadapter vfchost0 -npiv
Name          Physloc          ClntID ClntName      ClntOS
-----
vfchost0      U9117.MMA.6583005-V1-C15          3

Status:NOT_LOGGED_IN
FC name:fcs0          FC loc code:U5802.001.00B0557-P1-C7-T1
Ports logged in:0
Flags:4<NOT_LOGGED>
VFC client name:          VFC client DRC:

```

Repeat the above mentioned steps in the second VIOS LPAR also. If you have more client LPARs, repeat the steps for all those LPARs in both the VIOS LPARs.

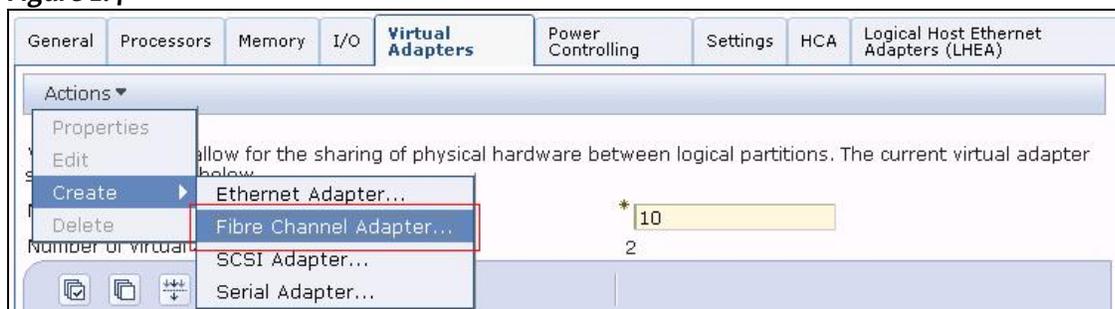
2.3. AIX LPAR Configuration

Now we need to create the Virtual FC Client Adapters in the AIX LPAR and configure necessary drivers in the AIX OS to use multipathing.

2.3.1. Creating Virtual FC Client Adapters

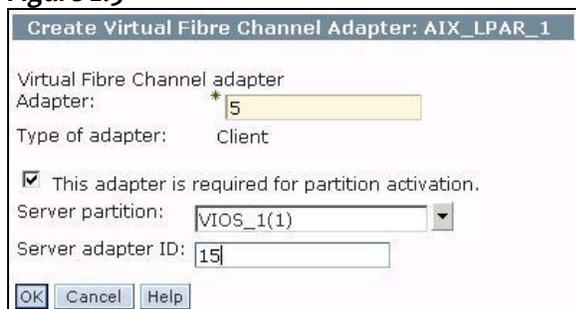
Open the LPAR profile for the AIX LPAR for editing and add the Virtual FC Client Adapters by navigating to the Virtual Adapters tab and selecting Actions > Create > Fibre Channel Adapter as shown in the [figure 2.4](#)

Figure 2.4

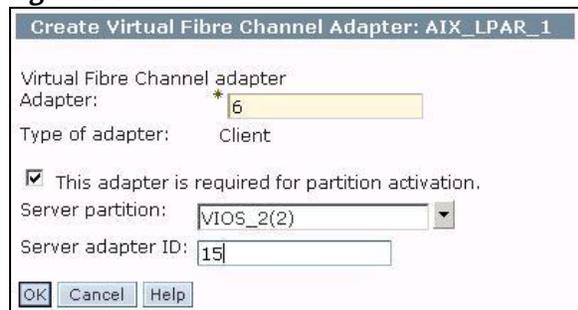


Input the first Virtual FC Client Adapter slot number details as shown in the [figure 2.5](#). You need to make sure that the slot numbers entered is exactly matching with the slot numbers entered while creating the Virtual FC Server Adapter in the first VIOS LPAR. This adapter will be the first path for the SAN access in the AIX LPAR.

Figure 2.5



Create the second Virtual FC Client Adapter with the slot number details as shown the [figure 2.6](#). Make sure the slot numbers match with the slot numbers we have entered in the second VIOS LPAR while creating the Virtual FC Server Adapter.

Figure 2.6


2.3.2. Configure AIX Operating System to use multipath

Now activate the AIX LPAR and install AIX in it, note that the minimum version required to support the NPIV feature is AIX 5.3 TL9 or AIX 6.1 TL2. Once the AIX installation is complete, depending on the SAN Storage box, you need to install the necessary subsystem driver and configure it.

To attach the storage, the storage vendor will specify what multi-path code must be used with their storage, and there may be options. We generally recommend multi-path code compliant with the MPIO architecture provided the storage supports it. For ESS, SVC, Storwize, DS6000, DS8000, DS3950, DS4000 (most models) and DS5000, you can use SDDPCM which is compliant with the MPIO architecture. This requires a "host attachment" fileset which populates the ODM to recognize and configure the storage correctly, and the SDDPCM code itself. The URL where you can download the SDDPCM is given below in References section.

If you are using IBM's XIV storage, native XIV support is included in many of the latest AIX releases and separate driver installation may not be necessary, separate XIV Host Attachment Kit is required for older version of AIX. Consult your storage administrator if you are in any doubt of the driver requirements in AIX.

You can now check the status of the Virtual FC Server Adapter ports in both the VIOS to check whether the ports are successfully logged in to the SAN fabric.

```
$ lsmap -vadapter vfchost0 -npiv
```

Name	Physloc	ClntID	ClntName	ClntOS
vfchost0	U9117.MMA.6583005-V1-C15	3	AIX_LPAR_1	AIX

```

Status:LOGGED_IN
FC name:fcs0          FC loc code:U5802.001.00B0557-P1-C7-T1
Ports logged in: 1
Flags:a<LOGGED_IN,STRIP_MERGE>
VFC client name:fcs0  VFC client DRC:U9117.MMA.6583005-V3-C5-T1
    
```

You can use the `'lsmap -vadapter vfchost0 -npiv'` command. You can see that the port is logged into the SAN fabric and it shows the client LPAR details too.

2.4. Configuring the SAN Storage

Once the AIX LPAR is activated and the Virtual FC ports are showing as logged in, you can configure your SAN storage box to do the proper Zoning and Access Control by providing the WWNs shown in the AIX LPAR. Note that the Virtual FC Client Adapter holds the WWNs for the clients and Virtual FC Server Adapters are just a connecting adapter used for passing through the SAN fabric data to the physical HBA port.

You can get the WWNs of the Virtual FC Client Adapter ports either from the AIX operating system or from the Hardware Management Console. From AIX, find out the Virtual FC Adapter which is mapped to the first and second VIOS by seeing the slot numbers from the command output `'lscfg -vp1 fcs*'`. Once you have found the device name for the specific adapter, get the WWN for that adapter by using the command `'lscfg -vp1 fcs0 |grep Address'` as shown below.

```
# lscfg -vpl fcs0 |grep Address
Network Address.....C0507602505C003A
```

For seeing the WWNs of the Virtual FC Client Adapter from the HMC, open the properties for the LPAR and navigate to the Virtual Adapters tab, then select the Virtual FC Client Adapter for which you want to see the WWN and click on properties.

NOTE: When viewing the properties of the Virtual FC Client Adapter from the HMC, it will show two WWNs for each Virtual FC Client Adapter as shown in [figure 2.7](#). The second WWN shown here is not used until there is a live migration activated on this LPAR through Live Partition Mobility.

Figure 2.7



Similar way, find out the WWN for the second adapter which is mapped through the second VIOS. Create the SAN LUNs, do proper Zoning and map the LUNs to these two adapters.

If you are planning to use Live Partition Mobility on this LPAR, the Zoning for SAN access has to be done using both the WWNs shown in the [figure 2.7](#). When a live migration happens for this LPAR, the new migrated hardware will be accessing SAN storage using the second WWN of the Virtual FC Client Adapter, so you have to make sure the second WWN is also configured in Zoning and Access Control.

Once it's done run the '`cfgmgr`' command in AIX LPAR and check whether the disks are getting detected. Once it is detected you can check the status of multipathing.

If the storage is IBM DS series, use the SDD commands '`datapath query adapter`', '`datapath query device`', '`lsvpcfg`' etc.

If the storage is XIV, you can check the status of multipathing using the '`lspath`' command. It will show the output as shown below. You can see that there are two separate paths for the disk '`hdisk2`' which is through two separate virtual FC adapters.

```
# lspath
Enabled hdisk0 vscsi0
Enabled hdisk1 vscsi1
Enabled hdisk2 fscsi0
Enabled hdisk2 fscsi1
```

3. References

S/N	Item	Link
1	IBM TechDocs	http://w3-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/FQ128819
2	SDDPCM for IBM Storage	http://www-01.ibm.com/support/docview.wss?rs=540&context=ST52G7&uid=ssg1S4000201&loc=en_US&cs=utf-8&lang=en