

Configuring and Troubleshooting N-Port ID Virtualization

ESXi 5.0

N-Port ID Virtualization (NPIV) is an ANSI T11 standard that describes how a single Fibre Channel HBA port can register with the fabric using several worldwide port names (WWPNs). This allows a fabric-attached N-port to claim multiple fabric addresses. Each address appears as a unique entity on the Fibre Channel fabric.

The primary source of information on configuring NPIV in a vSphere environment is the *vSphere Storage* documentation, available on the VMware Web site.

This technical note provides additional details about certain specific NPIV configurations and information on diagnostic techniques that might be helpful as you configure NPIV. It also includes information to help you understand error messages that might appear as you are working with NPIV in a vSphere environment. This technical note includes the following topics:

- [“Enabling the NPIV Feature on a Fibre Channel Port on a Brocade Switch”](#) on page 1
- [“Identifying HBAs in a Host System”](#) on page 2
- [“Confirming That I/O Traffic Goes Through an NPIV HBA”](#) on page 2
- [“Confirming Creation of a VPORT”](#) on page 3
- [“Interpreting Error Messages”](#) on page 4

Enabling the NPIV Feature on a Fibre Channel Port on a Brocade Switch

If you are using a Brocade switch, be sure that the NPIV capability is enabled on the switch. You can determine the status with the following command:

```
admin> portcfgshow 0
```

If the NPIV capability is enabled, the results of the `portcfgshow 0` command include the following line:

```
NPIV capability          ON
```

If the NPIV capability is not enabled, enable it with the `portCfɡNPIVPort` command.

portCfɡNPIVPort Command Usage

```
portCfɡNPIVPort <PortNumber> <Mode>
```

Table 1. portCfɡNPIVPort Modes

Mode	Meaning
0	Disable the NPIV capability on the port
1	Enable the NPIV capability on the port

portCfgNPIVPort Command Example

```
portCfgNPIVPort 0 1
```

Identifying HBAs in a Host System

The sections that follow include sample commands with specific values identifying the HBAs in the example host systems. To determine what specific values to use in those commands, you can check the `/proc` file system. With ESXi, `/proc` nodes and commands are not available unless you have the ESXi Shell enabled.

To determine the types of HBAs in the system

Enter the following command:

```
# ls /proc/scsi
```

Depending on the version of your host, the two supported HBA brands, QLogic and Emulex, are identified differently.

Table 2. HBAs That Support NPIV

Supported HBAs	ESX/ESXi 3.5	ESX/ESXi 4.x and ESXi 5.0
Emulex	lpfc_740	lpfc820
QLogic	qla2300_707_vmw	qla2xxx

To determine instance numbers for the `cat` commands

Enter one of the following commands, depending on the type of HBA in your system.

The output of the `ls` command includes a number for each active HBA in the system. You use this number in the `cat` commands shown in the following sections.

Table 3. HBA `ls` Commands

Supported HBAs	Commands
QLogic	<code>ls /proc/scsi/qla2xxx</code>
Emulex	<code>ls /proc/scsi/lpfc820</code>

Confirming That I/O Traffic Goes Through an NPIV HBA

You can check to be sure I/O traffic is actually going through an NPIV HBA in various ways. For example, you can check the Fibre Channel switch traffic on the virtual port. Another approach is to use the ESXi Shell to check the `/proc` nodes of the HBA to get the `reqs` details.

For example, a QLogic HBA updates the `Total reqs` value when I/O goes through that HBA.

In the following command, replace the final number `6` with the number of the HBA you want to check. You can determine the number to use by following the instructions in [“Identifying HBAs in a Host System”](#) on page 2.

```
# cat /proc/scsi/qla2xxx/6
```

```
Virtual Port 1 SCSI LUN Information:
```

```
( 0:10): Total reqs 10, Pending reqs 0, flags 0x0, 2:0:1000,
```

Confirming Creation of a VPORT

You can use various methods to be sure a VPORT is created when an NPIV-enabled virtual machine is powered on.

To use the `esxcli` command

For ESXi 5.0 hosts, use the vSphere CLI to run the `esxcli storage core adapter list` command.

The following sample output indicates that VPORTs are present.

```
vmhba34 lpfc820 link-up fc.2024000c29889a28:2824000c29889a28 () virtual
vmhba35 lpfc820 link-up fc.2024000c29889a28:2824000c29889a29 () virtual
```

If you use ESXi Shell, you can run the `cat` command to check the `/proc` nodes of the physical HBA for the VPORT.

To check the `/proc` nodes of QLogic HBAs

For QLogic, run the following command.

Replace the final number with the number of the HBA you want to check. You can determine the number to use by following the instructions in [“Identifying HBAs in a Host System”](#) on page 2.

```
# cat /proc/scsi/qla2xxx/#
```

```
FC Port Information for Virtual Ports:
Virtual Port index = 1
Virtual Port 1:VP State = <ACTIVE>, Vp Flags = 0x0
scsi-qla2-port-3=500601609020fd54:500601601020fd54:a00000:1000: 1;
scsi-qla2-port-4=500601609020fd54:500601681020fd54:a10000:1000: 1;
```

```
Virtual Port 1 SCSI LUN Information:
( 0:10): Total reqs 10, Pending reqs 0, flags 0x0, 2:0:1000,
```

To check the `/proc` nodes of Emulex HBAs

For Emulex, run the following command.

Replace the final number with the number of the HBA you want to check. You can determine the number to use by following the instructions in [“Identifying HBAs in a Host System”](#) on page 2.

```
# cat /proc/scsi/lpfc820/3
```

```
SLI Rev: 3
  NPIV Supported: VPIs max 127  VPIs used 1
  RPIs max 512  RPIs used 13

Vports list on this physical port:
Vport DID 0x2f0901, vpi 1, state 0x20
  Portname: 48:19:00:0c:29:00:00:0d  Nodename: 48:19:00:0c:29:00:00:0b
```

You can also check the Fibre Channel switches for the WWPN of the VPORT.

To check the Brocade switch for the VPORT WWPN

1 Run the `switchshow` command:

```
admin> switchshow
...
17 17 id N2 Online F-Port 10:00:00:00:c9:36:5a:bb
18 18 id N4 Online F-Port 2 NPIV public
19 19 id N4 Online F-Port 2 NPIV public
20 20 id N2 Online F-Port 10:00:00:00:c9:4e:ab:da
```

- 2 Run the portloginshow command:

```
admin> portloginshow 18
Type  PID      World Wide Name      credit df_sz cos
=====
fe  4f1201  28:24:00:0c:29:88:9a:28  16  2048  c  scr=3
fe  4f1200  10:00:00:00:c9:50:ee:66  16  2048  c  scr=3
ff  4f1201  28:24:00:0c:29:88:9a:28  12  2048  c  d_id=FFFFFF
ff  4f1200  10:00:00:00:c9:50:ee:66  12  2048  c  d_id=FFFFFF
```

You can also use the nsshow command.

```
admin> nsshow

Type Pid    COS      PortName                      NodeName                      TTL(sec)
N    a00401;   3;28:d9:00:0c:29:00:02:97;28:d9:00:0c:29:00:01:97; na
Fabric Port Name: 20:04:00:05:1e:02:6e:0f
Permanent Port Name: 28:d9:00:0c:29:00:02:97 <== Vport WWPN
Port Index: 4
Share Area: No
Device Shared in Other AD: No
The Local Name Server has 4 entries }
```

If you are using QLogic switches, you can check the Fibre Channel switch for the WWPN of the VPORT.

To check the QLogic switch for the VPORT WWPN

Run the show ns command.

In the output of the show ns command, check the NodeWWN column. You can identify WWPNs of VMware VPORTs by the sequence 00:0c:29 in the third, fourth, and fifth fields.

```
#> show ns
```

Seq No	Domain ID	Port ID	Port Type	COS	PortWWN	NodeWWN
1	182 (0xb6)	b60200	N	3	21:00:00:e0:8b:88:e2:8b	20:00:00:e0:8b:88:e2:8b
2	182 (0xb6)	b60300	N	3	21:00:00:e0:8b:80:83:4b	20:00:00:e0:8b:80:83:4b
3	182 (0xb6)	b60700	N	3	21:00:00:1b:32:00:4b:34	20:00:00:1b:32:00:4b:34
4	182 (0xb6)	b60a00	N	3	50:06:01:60:41:e0:1d:98	50:06:01:60:c1:e0:1d:98
5	182 (0xb6)	b60b00	N	3	50:06:01:68:41:e0:1d:98	50:06:01:60:c1:e0:1d:98
6	182 (0xb6)	b60c00	N	3	21:01:00:1b:32:20:4b:34	20:01:00:1b:32:20:4b:34
7	182 (0xb6)	b60c01	N	3	28:33:00:0c:29:00:00:34	28:33:00:0c:29:00:00:33 <== VPort VWWN
8	182 (0xb6)	b60d00	N	3	21:01:00:1b:32:20:95:36	20:01:00:1b:32:20:95:36
9	182 (0xb6)	b60e00	N	3	21:00:00:e0:8b:92:0b:90	20:00:00:e0:8b:92:0b:90
10	182 (0xb6)	b60f00	N	3	21:01:00:e0:8b:a8:85:81	20:01:00:e0:8b:a8:85:81

Interpreting Error Messages

This section provides information to help you interpret the most common error messages seen in an NPIV environment.

Driver Load Time Errors

The VportGetInfo information in the server's log files provides information about the NPIV support at various levels in the system. On ESXi hosts, look for this information in `/var/log/messages`.

System with No Problems

When the system has no problems, the `VportGetInfo` messages are similar to the example below for each physical HBA in the system that has NPIV support.

Example Output

```
vmkernel: 0:00:00:57.292 cpu6:1040)SCSI: VportGetInfo:748: GetInfo for adapter vmhba2,
      [0x3f8ae380], max_vports=64, vports_inuse=0, linktype=0, state=1, failreason=0,
      rv=0, sts=0
```

HBA Without NPIV Support

When the HBA does not NPIV support, the `rv` value is less than 0. The typical value is -1. To resolve this problem, use an HBA with NPIV support.

Example Output

```
vmkernel: 0:00:00:41.167 cpu6:1039)SCSI: VportGetInfo:748: GetInfo for adapter vmhba0,
      [0x3f8a6980], max_vports=0, vports_inuse=0, linktype=0, state=0, failreason=0,
      rv=-1, sts=bad001f
```

Other Problems

If the value of `state` is 0 and the value of `failreason` is non-zero, see [Table 4](#) for details on the meaning of the message.

Example Output

```
vmkernel: 0:00:00:53.379 cpu6:1040)SCSI: VportGetInfo:748: GetInfo for adapter vmhba1,
      [0x3f8aca00], max_vports=64, vports_inuse=0, linktype=0, state=0, failreason=2,
      rv=0, sts=0
```

Table 4. Failure Codes

failreason Value	Reason for Failure	Remedy for the Failure
0	Unknown failures	Unknown reason, contact tech support
1	Fibre Channel link is down	Check the physical links from your machine to the Fibre Channel switch
2	Fabric does not support NPIV	Enable NPIV capability on the Fibre Channel switch port
3	Fabric does not have resources	Remove some VPORTs from the Fibre Channel fabric
4	Fabric LOGOUT initiated	
5	HBA does not have enough resources	Check and update the Fibre Channel HBA firmware

Virtual Machine Start Time Errors

In some cases, you might find an error recorded in the VMkernel log when an NPIV virtual machine is powered on.

Example Log Entry

```
vmkernel: 0:01:12:35.280 cpu3:1083)SCSI: VportDiscovery:489: NPIV vport rescan complete, [0:21]
(0x3f8e8268) [0x2bb66e8] status=0xbad0001
```

Generally, this means that the NPIV code in the VMkernel is not able to find any devices on the VPORT. There are a number of possible causes. To track down the cause, make the following checks:

- Check zoning in the configuration of the switch to be sure correct access is set for the NPIV WWN LUNs.
- Check the switch port to be sure it has NPIV capability enabled.
- Check the LUN's HostID to be sure it matches the physical HBA and virtual HBA in the storage array.

If you have comments about this documentation, submit your feedback to: docfeedback@vmware.com

VMware, Inc. 3401 Hillview Ave., Palo Alto, CA 94304 www.vmware.com

Copyright © 2007–2011 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Item: EN-000696-00
