

Configuring and Troubleshooting N-Port ID Virtualization

ESX Server 3.5, ESX Server 3i version 3.5

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 VMware Infrastructure 3 environment is the *Fibre Channel SAN Configuration Guide*, available on the VMware Web site.

This technical note provides additional details about certain specific NPIV configurations and information on diagnostic techniques that may be helpful as you configure NPIV. It also includes information to help you understand error messages that may appear as you are working with NPIV in a VMware Infrastructure environment. It covers 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 is Going 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 `portCfGNPIVPort` command.

Usage

```
portCfgNPIVPort <PortNumber> <Mode>
```

Table 1. portCfgNPIVPort modes

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

Example

```
portCfgNPIVPort 0 1
```

Identifying HBAs in a Host System

The sections that follow include examples 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 using the service console on a traditional ESX Server host.

To determine the types of HBAs in the system, enter the following command:

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

QLogic HBAs are listed as `qla2300`. Emulex HBAs are listed as `lpfc`.

To determine the instance numbers you should substitute in the `cat` commands shown in the following sections, enter one of the following example commands, depending in the type of HBA in your system. The output of the `ls` command includes a number for each active HBA in the system.

QLogic Example

```
# ls /proc/scsi/qla2300
```

Emulex Example

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

Confirming That I/O Traffic is Going 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, using the service console on traditional ESX Server, is checking 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/qla2300/6
```

```
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. This section describes methods appropriate for QLogic, Emulex, and Brocade switches.

If you are using QLogic or Emulex switches and are using traditional ESX Server, you can use the **cat** command in the service console to check the `/proc` nodes of the physical HBA for the VPORT.

QLogic Example

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/qla2300/6
```

```
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,
```

Emulex Example

In the following command, replace the final number **3** 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/lpfc/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
```

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

Brocade Example

```
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. 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.

QLogic Example

```
#> 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 traditional ESX Server, look for this information in the `/var/log/vmkernel*` files. On ESX Server 3i, 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 2](#) 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 2. 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 may 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

© 2007 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,880,022, 6,944,699, 6,961,806, 6,961,941, 7,069,413, 7,082,598, 7,089,377, 7,111,086, 7,111,145, 7,117,481, 7,149,843, 7,155,558, 7,222,221, 7,260,815, 7,260,820, and 7,269,683; patents pending. VMware, the VMware "boxes" logo and design, Virtual SMP and VMotion are registered trademarks or trademarks 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.

Revision 20071017