XgOS vNIC Switching Guide

Release 2.6.0 and Later

Xsigo Systems
70 West Plumeria Drive
San Jose, CA 95134
USA

http://www.xsigo.com
Tel: +1.408.329.5600
Part number: 650-20052-02 Rev A
Published: Oct 2012

EMI Statement, United States of America (Class A)

“NOTE: This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.”

EMI Statement, Canada (Class A)

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

EMI Statement, Europe and Australia (Class A)

“Warning - This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.”

EMI Statement, Japan (Class A)

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

“This is a Class A product based on the standard of the Voluntary Control Council For Interference by Information Technology Equipment (VCCI). If this equipment is used in a domestic environment, radio disturbance may arise. When such trouble occurs, the user may be required to take corrective actions.”

Lithium Battery - Replacement and Disposal

CAUTION!

Danger of explosion if the lithium battery is incorrectly replaced. Replace only with the same or equivalent type recommended by the manufacturer. Dispose of used batteries according to the manufacturer's instructions.

Laser Caution for I/O Cards (CDRH-US)

USE OF CONTROLS OR ADJUSTMENTS OR PERFORMANCE OF PROCEDURES OTHER THAN THOSE SPECIFIED HEREIN MAY RESULT IN HAZARDOUS RADIATION EXPOSURE.

Complies with 21 CFR Chapter 1, Subchapter J, Part 1040.10.

IEC 60825-1: 1993, A1: 1997, A2: 2001; IEC 60825-2: 2000

Replacement Laser Transceiver Modules

For continued compliance with the above laser safety Standards, only approved Class 1 modules from our approved vendors should be installed in the product. Contact Xsigo Customer Support (see Technical Support Contact Information) for approved-vendor contact information.

Power Cord Set Requirements – General

The requirements listed below are applicable to all countries:

The length of the power cord set must be at least 6.00 feet (1.8 m) and a maximum of 9.75 feet (3.0 m).

All power cord sets must be approved by an acceptable accredited agency responsible for evaluation in the country where the power cord set will be used.

The power cord set must have a minimum current capacity of 13A and a nominal voltage rating of 125 or 250 V ac~, as required by each country's power system.

The appliance coupler on the power cord must meet the mechanical configuration of an EN 60320 / IEC 60320 Standard Sheet C20 connector, which is the connector on the Fabric Director. The C20 connector supports a C19 plug as the mating part on the power cord that connects to the Fabric Director.

Power Cord Set Requirements – Specifics By Country

United States (UL), Canada (CSA)

The flexible power cord set must be UL Listed and CSA Certified, minimum Type SVT or equivalent, minimum No. 18 AWG, with 3-conductors that includes a ground conductor. The wall plug must be a three-pin grounding type, such as a NEMA Type 5-15P (rated 15A, 120V) or Type 6-15P (rated 15A, 250V).

Europe (Austria (OVE), Belgium (CEBEC), Denmark (DEMKO), Finland (SETI), France (UTE), Germany (VDE), Italy (IMQ), Netherlands (KEMA), Norway (NEMKO), Sweden (SEMKO), Switzerland (SEV), U.K. (BSI/ASTA)

The flexible power cord set must be <HAR> Type H03VV-F, 3-conductor, minimum 0.75mm^2 conductor size. Power cord set fittings, particularly the wall plug, must bear the certification mark of the agency responsible for evaluation in the country where it is being used, with examples listed above.

Australia (DFT/SAA)

Cord is as described under “Japan (PSE)” immediately below. Pins in the power plug must be with the sheathed, insulated type, in accordance with AS/NZS 3112:2000.

Japan (PSE)

The appliance coupler, flexible cord, and wall plug must bear a “PSE” Mark in accordance with the Japanese Denan Law. The flexible cord must be Type VCT or VCTF, 3-conductor, 0.75 mm^2 conductor size. The wall plug must be a grounding type with a Japanese Industrial Standard C8303 (15A, 125V) configuration.

Software Compliance – GPL (GPL v2) Licenses and Notices

Xsigo Systems, a wholly owned subsidiary of Oracle, uses certain elements of GNU Public License (GPLv2) code. Under the conditions of the GPL licensing agreement, you are entitled to request a copy of the open source/freeware code. For questions about Xsigo’s use of the GPL code, or to request a copy of the code, you can contact Xsigo by completing the web form at http://pages.xsigo.com/compliance.html. Afterward, Xsigo will contact you to assist you with your request.

Xsigo Systems, eine ganz besessene Tochtergesellschaft von Oracle, gewisse Elemente des GNU Public License (GPLv2) Code. Unter den Bedingungen von der GPL Lizenzvertrag werden Sie berechtigt, eine Kopie der offenen Quelle/ Freewarecodes zu erbitten. Für Fragen um den Gebrauch von Xsigo des GPL Codes oder eine Kopie des Codes zu erbitten, können Sie Xsigo durch Vollenden der Gewebesform an http://pages.xsigo.com/compliance.html kontaktieren. Nachher wird Xsigo Sie kontaktieren, Ihnen mit Ihrer Bitte zu helfen.

Xsigo Systens, une filiale entièrement possédée d'Oracle, utilisent de certains éléments de GNU Public License (GPLv2) le code. Sous les conditions du GPL autorise l'accord, vous êtes autorisé à demander une copie du code de code source libre/graticiel. Pour les questions de l'usage de Xsigo du code de GPL, ou demander une copie du code, vous pouvez contacter Xsigo en complétant la forme Web à http://pages.xsigo.com/compliance.html. Après, Xsigo vous contactera pour vous aider avec votre demande.

Documentation Purpose and Audience

This guide provides the basic information that you need to configure (vNIC Switching) also called vNIC-to-vNIC Switching on a single Oracle Fabric Director or multiple Oracle Fabric Directors. This guide also provides an overview of the vNIC-to-vNIC Switching feature as well as how it interacts with many of the Fabric Director's vNIC and networking-related features.

This guide is intended for data-center network administrators, and it assumes that its readers have knowledge and familiarity with common configuration and management tasks related to administering a data center. Although this guide does present some conceptual material about topics and technologies, it is not intended as a complete and exhaustive reference on those topics.

Document Overview

This guide contains the following sections:

- Overview — Documents an overview of what vNIC-to-vNIC Switching is and what benefits it provides.
- Single Chassis vNIC-to-vNIC Switching — Documents the configuration of vNIC-to-vNIC Switching on a single Fabric Director, including how to configure vNIC-to-vNIC Switching for use with many of the Fabric Director's other networking-related features. Also, additional information is provided for host-side considerations when configuring vNIC-to-vNIC Switching.
- Multi-Chassis vNIC-to-vNIC Switching — Documents the configuration of vNIC-to-vNIC Switching in a multiple Fabric Director environment, including how to configure vNIC-to-vNIC Switching for use with many of the Fabric Director's other networking-related features. Also, additional information is provided for host-side considerations when configuring vNIC-to-vNIC Switching.

Related Documentation

This document is part of a set of documentation for the Fabric Director. Table 1 on page i shows the other documents in the Oracle Fabric Director documentation set.

Table 1 Related Documentation for the Xsigo Systems Fabric Director

Document	Part Number	Revision Level and Date
Fabric Manager User Guide	650-30005-02	Rev A 10/2012
Fabric Director Quick Install Guide	650-20022-04	Rev A 10/2012
Fabric Director Hardware and Drivers Installation Guide	650-30008-03	Rev A 10/2012
Fabric Accelerator Quick Start Guide	650-20085-03	Rev A 10/2012
Fabric Performance Manager User Guide	650-20082-02	Rev A 10/2012
XgOS Software Upgrade Guide	650-20028-06	Rev A 10/2012
XgOS Remote Booting Guide	650-20029-08	Rev A 10/2012
XgOS Command-Line Interface User Guide	650-20052-02	Rev A 10/2012

Table 1 Related Documentation for the Xsigo Systems Fabric Director

Document	Part Number	Revision Level and Date
Installing Host Drivers on Windows Core Servers	650-20081-02	Rev A 10/2012
Hyper-V Setup Guide	650-20040-02	Rev A 10/2012
SAN Install for Windows 2008 Servers	650-20078-03	Rev A 10/2012

Release notes are also available with each major hardware or software release of Oracle's Fabric Director.

Revision Trail

Table 2 shows the revision history for this document.

Table 2 Revision History

Document Title	Document Number	Revision Level	Revision Date
XgOS vNIC Switching Configuration Guide	650-20052-02	A	10/2012
XgOS vNIC Switching Configuration Guide	650-20052-01	A	11/2009

Syntax Usage

Table 3 shows the typographical conventions used in this document.

Table 3 Syntax Usage

Syntax Marker	Means...	Example
bold text, courier font	a command	**`help`**
blank space	a delimiter for commands and arguments	**`system show version`**
- (dash)	you are specifying an argument	**`set ethernet-card 1 -type`**
= (equals sign)	you are specifying a parameter for an argument	**`set ethernet-card 1 -type=nwEthernet10Port1GbCard`**

Technical Support Contact Information

Xsigo Customer Support Services is willing to help solve any reported issues 24 hours a day, 7 days a week, 365 days a year. The Xsigo Technical Assistance Center (TAC) is open 9:00 a.m. to 6:00 p.m. PST Monday through Friday. If you need assistance, you can contact the Xsigo Technical Assistance Center (TAC) in any of the following ways:

- Email

  You can send an email to *support@xsigo.com* and we will respond within 24 hours (Monday through Friday).

- Web Access

  You can create a Service Request through the Support Web interface (http://www.xsigo.com/support) and we will respond within 24 hours (Monday through Friday). If you do not have a log-in, we will be more than happy to provide you with access to create, view, update and close Service Requests. You can also open RMA cases through the Web.

- Phone Contact

  In the event that you are in need of a faster response for any reason, Xsigo provides response to all phone calls in a maximum of 30 minutes (24 hours a day, 7 days a week, 365 days a year).

  — You can reach us through the Xsigo switchboard by dialing +1 408-329-5600 and selecting option “2”

  — You can reach us through a direct line, by dialing +1 408-736-3013 (24 hours a day, 7 days a week)

  — For our US customers, you can call us through our toll-free number by dialing 866-974-4647

Contents

Contents

Xsigo vNIC Switching

Oracle's XgOS operating system supports vNIC Switching (also called vNIC-to-vNIC Switching), which is a feature that enables two vNICs that are terminated on the same port on the same I/O module, to switch traffic between them without requiring an external Ethernet switch. The actual switching of traffic occurs just like in an external Ethernet switch, but instead, the switching occurs on the Ethernet module itself.

The vNIC-to-vNIC Switching is enabled by default, and provides the following benefits:

- network traffic can bypass the "physical network and just switch between the two vNICs. Traffic to and from the servers connected through these vNICs is seamlessly transmitted and received.
- cost, configuration, and management is reduced, because the vNICs are connected directly from the host server to the Oracle Fabric Director, thus bypassing the physical Ethernet switch.
- a high-bandwidth, low-latency private network is available on-demand for traffic supported by vNIC-to-vNIC Switching. One example of vNIC-to-vNIC Switching is for supporting VMware vMotion functions which have unique requirements for moving VMs.

In the simplest example, vNIC-to-vNIC Switching supports two host servers connected to the same Fabric Director through separate vNICs terminated on the same port on the same module. This configuration is shown in Figure 1.

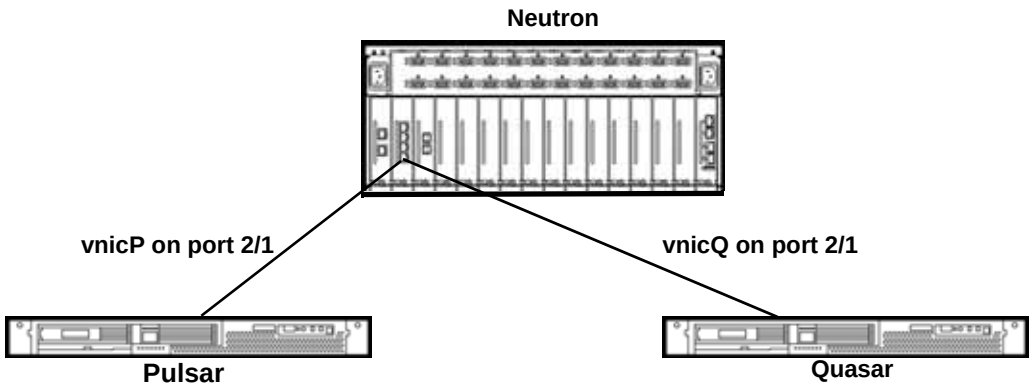

Figure 1 Single Chassis vNIC-to-vNIC Switching

In this example, traffic sent between Pulsar and Quasar can be seamlessly switched between them on the Ethernet module in slot 2 because both vNICs are terminated on port 1 in slot 2. In this example, the switching functions are handled by on-board logic embedded in the Ethernet module's processors. The need for an intermediary physical Ethernet switch is eliminated.

Additional considerations exist when considering different configurations and the various networking-related features available in the Fabric Director. These topics are discussed in the following chapters.

Hardware Support for vNIC-to-vNIC Switching

Xsigo's vNIC-to-vNIC Switching feature is supported on the following I/O modules:

- the 10 GE module
- the 10-Port GE module

On the server side, all supported Xsigo HCAs enable vNIC-to-vNIC switching.

Next Steps

Determine the environment in which you will be using vNIC-to-vNIC Switching (either on a single Oracle Fabric Director or multiple Oracle Fabric Directors), then go to the appropriate chapter:

- Chapter 2, "Single Chassis vNIC-to-vNIC Switching."
- Chapter 3, "Multi-Chassis vNIC-to-vNIC Switching."

Single Chassis vNIC-to-vNIC Switching 3

This chapter contains the following topics:

Standard vNIC Switching

With single-chassis vNIC-to-vNIC Switching, two vNICs for two different hosts are terminated on the same I/O port on the same I/O module, as shown in Figure 1.

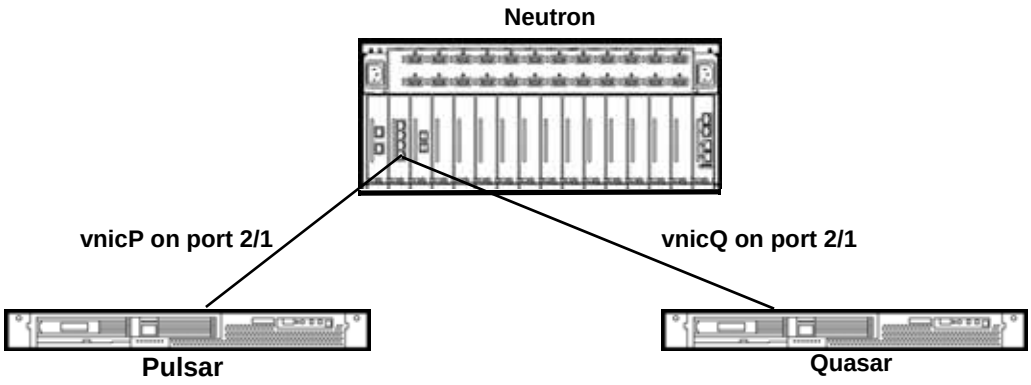

Figure 1 Standard vNIC-to-vNIC Switching (Single Chassis)

Configuration Details

For the example in Figure 1, the vNIC-to-vNIC Switching occurs for traffic sent between Quasar and Pulsar. The vNIC-to-vNIC Switching occurs automatically when the two vNICs are terminated on the same port. On the Oracle Fabric Director "Neutron," this configuration would be:

Step 1 Create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 2 Create a vnic in each server profile, and terminate it on the same I/O port and I/O module on the Fabric Director:

```
add vnic pulsar.vnicP 2/1
add vnic quasar.vnicQ 2/1
```

Step 3 Set the IP address allocation method for the vNICs. In this example, assume that addresses will be assigned through DHCP:

```
set vnic pulsar.vnicP 2/1 -addr-type=dhcp
set vnic quasar.vnicQ 2/1 -addr-type=dhcp
```

Step 4 Continue configuring the vNIC properties by issuing the **`set vnic`** command and setting the required parameters.

HA vNICs and vNIC Switching

For the example in Figure 2, the vNIC-to-vNIC Switching occurs for the traffic between Pulsar and Quasar. Any traffic between the two servers occurs on the Primary vNICs, which are the online and active vNICs. These vNICs must be terminated on the same I/O port. the Secondary vNICs are also configured on the same port, but cannot be configured on the same port as the primary vNICs or else you lose high availability. The secondary vNICs (shown as blue lines and named with “2”) are live standby vNICs are available to switch traffic between the two hosts if the primary vNICs go offline. With HA vNICs configured, you gain redundant data paths for your network traffic. HA vNICs do not load share.

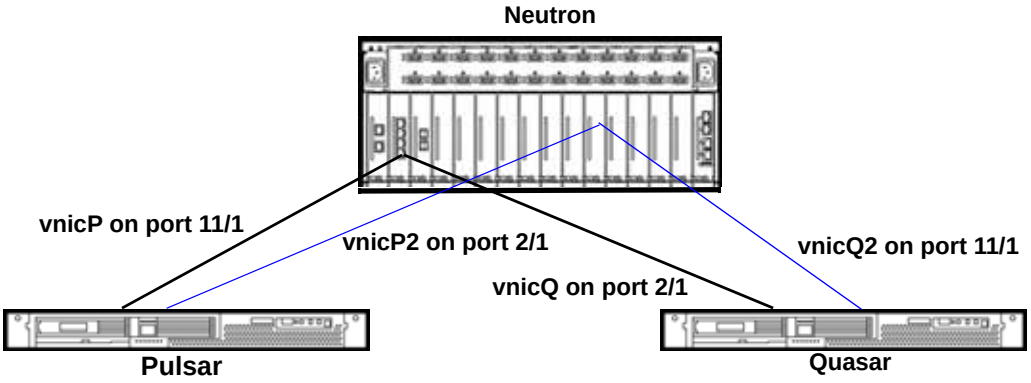

Figure 2 Example of HA vNIC

To configure HA vNICs to take advantage of vNIC-to-vNIC Switching as shown in this example, the configuration would be as follows:

Step 1 Create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 2 Create a primary vNIC in each server profile, and terminate it on the same I/O port and I/O module on the Fabric Director:

```
add vnic pulsar.vnicP 2/1 -primary
add vnic quasar.vnicQ 2/1 -primary
```

Step 3 Create the secondary vNICs. Make sure to:

- terminate them on a different port than the primary vNICs.
- terminate them on a different I/O module than the primary vNICs, as a best practice.
- specify the MAC address of the related primary vNIC.

For example:

```
add vnic pulsar.vnicP2 11/1 ha -secondary -mac=10:02:53:56:09
add vnic quasar.vnicQ2 11/1 ha -secondary -mac=06:13:53:12:04
```

Step 4 Set the IP address allocation method for the primary vNICs. In this example, DHCP will be used.

```
set vnic pulsar.vnicP 2/1 -addr-type=dhcp
set vnic quasar.vnicQ 2/1 -addr-type=dhcp
```

Step 5 Set the IP address allocation method for the secondary vNICs. The IP address allocation method should match what is configured for the primary vNICs.

```
set vnic pulsar.vnicP2 11/1 -addr-type=dhcp
set vnic quasar.vnicQ2 11/1 -addr-type=dhcp
```

Step 6 As an option, set the auto switchover parameters for the secondary vNICs. Auto switchover is true by default, but can be set to `false` if needed. The auto switchover setting (true or false) determines whether the secondary vNICs switchback from primary to secondary when the original primary vNICs come back online. In this example, auto switchover will be not enabled.

```
set vnic pulsar.vnicP2 11/1 -autoswitchover=false
set vnic quasar.vnicQ2 11/1 -autoswitchover=false
```


Note

In this example, module-level HA exists, but port-level HA can also be configured instead. You can configure port-level HA by terminating the secondary vNICs on the same module (but different port) as the primary vNIC. Even though this configuration is supported, Xsigo does not recommend it because both the primary and secondary vNICs are terminated on the same module, so in the unlikely event of a module failure, data paths between the servers are lost.

LAG and vNIC Switching

The Xsigo implementation of Link Aggregation Groups (LAG) treats all physical ports in the LAG as one logical port. Link aggregation groups (LAG) and vNIC Switching can work together to determine how vNIC Switching operates. When the vNICs are terminated within the same LAG, their traffic can be switched between across different physical ports on the same module, thereby relaxing the requirement that the vNICs are terminated on the same physical port. Be aware that the different physical ports must be in the same LAG to support this “cross module” method of vNIC Switching. Also, be aware this cross module LAG is supported on the 10-Port GE module only.

In Figure 3 on page 7, vnicP and vnicQ are terminated on port 2.1 and 2.10 respectively. The dot notation separating the slot and port identifiers indicates that the ports are members of a LAG. In this example, port 2.1 and 2.10 are configured in the same LAG (represented by the blue rectangle), so the on-board I/O processor can support vNIC Switching between the two ports but only if they are on the same I/O module and in the same LAG. If one of the ports is moved out of the LAG, vNIC Switching will no longer be supported.

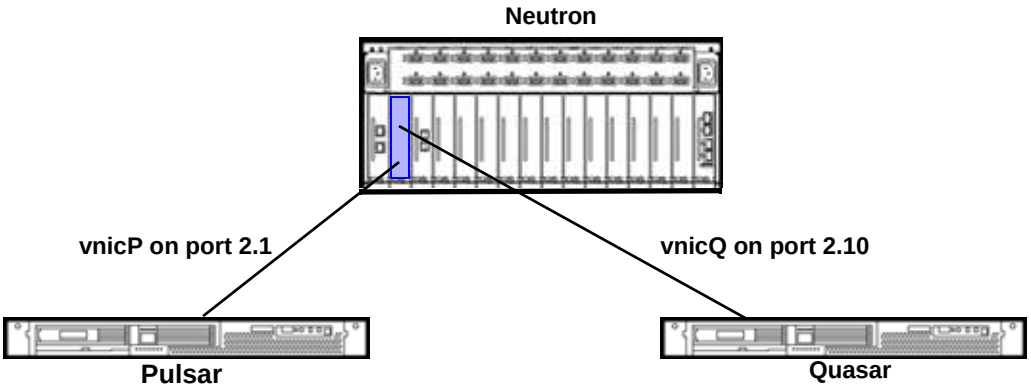

Figure 3 vNIC-to-vNIC Switching in Link Aggregation Group

To support vNIC-to-vNIC Switching between different physical ports in a LAG, you would configure the vNICs as follows:

Step 1 Create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 2 Create the LAG by specifying the LAG name and adding the first port:

```
add lag 2.1 port 2/1
```

Step 3 Add the second port to the LAG:

```
add lag 2.1 port 2/10
```

Step 4 Create a vNIC in each server profile, and terminate it on the LAG:

```
add vnic pulsar.vnicP 2.1
add vnic quasar.vnicQ 2.1
```

Step 5 Set the IP address allocation method for the vNICs. In this example, DHCP will be used.

```
set vnic pulsar.vnicP 2/1 -addr-type=dhcp
set vnic quasar.vnicQ 2/10 -addr-type=dhcp
```

Step 6 Complete the vNIC configuration for both vNICs as needed by using the **`set vnic`** command to assign the operational parameters required.

Note

This procedure shows how to configure cross-module vNIC Switching, in which two different ports are used, but vNIC Switching still occurs on the module. A LAG containing the two different ports is required to make this configuration work. Standard vNIC Switching, in which vNICs for two separate servers are terminated on the same physical port, is also supported. This configuration typically occurs outside of a LAG, but can also be configured in the same LAG.

Standard vNIC Switching and cross-module vNIC Switching are mutually exclusive. You cannot configure the same two vNICs for standard and cross-module Switching, nor can you configure two vNICs on the same physical port (standard vNICs Switching) and have that port tied to another port in the same LAG (cross-module vNIC Switching).

Network QoS and vNIC Switching

Network QoS allows for bandwidth usage parameters to be applied to vNICs to ensure that bandwidth is available to hosts. Network QoS is configurable on individual vNICs, and vNIC Switching can be affected by the Network QoS profiles that are configured on them.

In Figure 4, assume that Network QoS is configured on vnicP to guarantee a specific minimum bandwidth for the host Pulsar for vNIC Switching to be successful for these hosts. In this case, the same Network QoS profile must be configured on vnicQ. Without an exact duplicate Network QoS profile configured on each vNIC, when vNIC Switching occurs between the two hosts, the traffic can be altered by Network QoS Profiles in such a way that overruns or underruns can occur.

When configuring Network QoS in a vNIC Switching environment, always make sure to configure the same Network QoS Profile on both vNICs. Even a minor discrepancy between the two profiles can have an affect on how successfully vNIC-switched traffic is sent between two hosts.

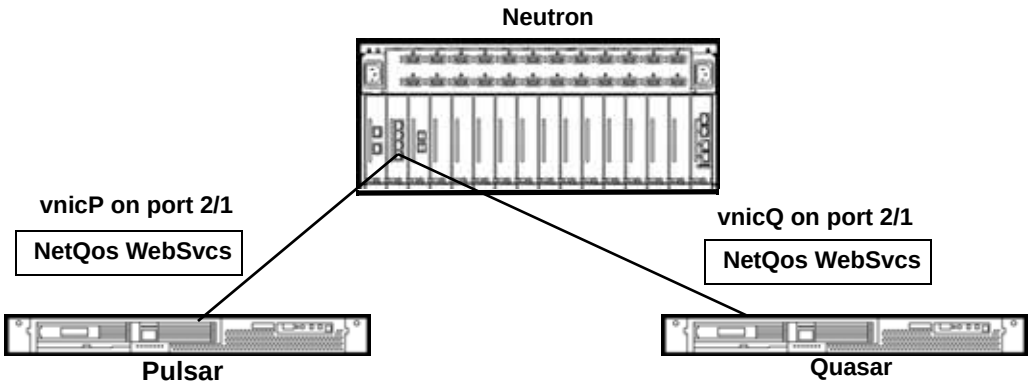

Figure 4 Network QoS and vNIC-to-vNIC Switching

To support Network QoS settings and vNIC-to-vNIC Switching, you would configure the vNICs as follows:

Step 1 Create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 2 Create a vNIC in each server profile:

```
add vnic vnicP.pulsar 2/1
add vnic vnicQ.quasar 2/1
```

Step 3 List all the default Network QoS Policer Profiles available:

```
show qos network policer
name                level  descr     cir          pir        cbs        pbs
-------------------------------------------------------------------------------
default/100m_1g     global 100m_1g   100Mbps      1Gbps      17.8814MB 35.7628MB
default/100m_250m   global 100m_250m 100Mbps      250Mbps    17.8814MB 35.7628MB
default/10g_10g     global 10g_10g   9.9297Gbps   9.9297Gbps 1.73395GB 1.73395GB
default/10m_100m    global 10m_100m  10Mbps       100Mbps    1.78814MB 3.57628MB
default/10m_1g      global 10m_1g    10Mbps       1Gbps      1.78814MB 3.57628MB
default/10m_50m     global 10m_50m   10Mbps       50Mbps     1.78814MB 3.57628MB
default/1g_10g      global 1g_10g    1Gbps        9.9297Gbps 178.814MB 357.628MB
default/1m_10m      global 1m_10m    1Mbps        10Mbps     182.617KB 365.234KB
default/250m_500m   global 250m_500m 250Mbps      500Mbps    44.7035MB 89.407MB
default/2g_10g      global 2g_10g    2Gbps        9.9297Gbps 357.628MB 715.256MB
default/3g_10g      global 3g_10g    3.00293Gbps  9.9297Gbps 536.965MB 1.04876GB
default/4g_10g      global 4g_10g    4Gbps        9.9297Gbps 715.256MB 1.39698GB
default/500m_750m   global 500m_750m 500Mbps      750Mbps    89.407MB  178.814MB
default/50m_100m    global 50m_100m  50Mbps       100Mbps    8.9407MB  17.8814MB
default/5g_10g      global 5g_10g    5.00122Gbps  9.9297Gbps 894.287MB 1.74665GB
default/64k_1m      global 64k_1m    64Kbps       1Mbps      11.7188KB 23.4375KB
default/6g_10g      global 6g_10g    6.00587Gbps  9.9297Gbps 1.04876GB 2.09752GB
default/750m_1g     global 750m_1g   750Mbps      1Gbps      134.11MB  268.221MB
default/7g_10g      global 7g_10g    7.00171Gbps  9.9297Gbps 1.22266GB 2.44532GB
default/8g_10g      global 8g_10g    8Gbps        9.9297Gbps 1.39698GB 2.79397GB
default/9g_10g      global 9g_10g    9.0022Gbps   9.9297Gbps 1.57199GB 3.14398GB
```

Step 4 Determine if you can use a default Network Policer or if you need to create a custom Policer:

- If you want to use a default Network QoS Policer, proceed to Step 5.
- If you want to create a custom QoS Policer, proceed to Step 8.

Step 5 To use a default QoS Policer, assign it to the vNIC for the first host. For example, assume the default Policer for 5 Gbps and 10 Gbps is on egress traffic only:

```
set vnicP.pulsar egress-qos -policer=default/5g_10g enable
```


Note

This command also has the **ingress-qos** qualifier to set the QoS variables for traffic leaving the I/O module and entering the network. Also, by using the **disable** qualifier in the command, you can selectively enable or disable the assigned Network QoS Policer without having to delete the Policer or unbind it from the vNIC.

Step 6 Assign a copy of the same QoS Policer to the vNIC for the second host. For example:

```
set vnic vnicQ.quasar egress-qos -policer=default/5g_10g enable
```

Note

The two QoS Policers must be identical for the vNICs attached to both hosts.

Step 7 When both default Policer is added, go to Step 11.

Step 8 To create a custom QoS Policer, first add the QoS Policer and specify the bandwidth variables. For example, create a custom Policer with 2 Mbps CIR, and optional CBS of 1Mbps and PBS of 2 Mbps:

```
add qos network policer WebSvcs/1 2m -cbs=1m -pbs=2m
```

Step 9 After the custom Policer is added, assign it to the vNIC for the first host. For example:

```
set vnic vnicP.pulsar egress-qos -policer=WebSvcs/1 enable
```

Note

This command also has the **`ingress-qos`** qualifier to set the QoS variables for traffic leaving the I/O module and entering the network. Also, by using the **`disable`** qualifier in the command, you can selectively enable or disable the assigned Network QoS Policer without having to delete the Policer or unbind it from the vNIC.

Step 10 Assign a copy of the same QoS Policer to the vNIC for the second host. For example:

```
set vnic vnicQ.quasar egress-qos -policer=WebSvcs/1 enable
```

Step 11 Set the IP address allocation method for the vNICs. In this example, DHCP will be used.

```
set vnic vnicP.pulsar 2/1 -addr-type=dhcp
set vnic vnicQ.quasar 2/1 -addr-type=dhcp
```

Step 12 Complete the vNIC configuration for both vNICs as needed by using the **`set vnic`** command to assign the operational parameters required.

vNIC Mirroring and vNIC Switching

Xsigo's vNIC Mirroring feature allows traffic received or transmitted on a vNIC to be copied on the I/O module that the vNIC is terminated on. The traffic is copied on an I/O module's network processing chip, and then can be redirected to an output, either a port or a vNIC. A configuration of vNIC Mirroring and vNIC Switching is supported on either a port or a vNIC. If you are sending the vNIC mirror data to an output device, you must specify the device's destination MAC address.

In Figure 5 on page 11, the example shows vNIC Mirroring to an external device a traffic analyzer (such as a sniffer). The mirrored vNIC traffic (shown as a blue line) is sent through an output port.

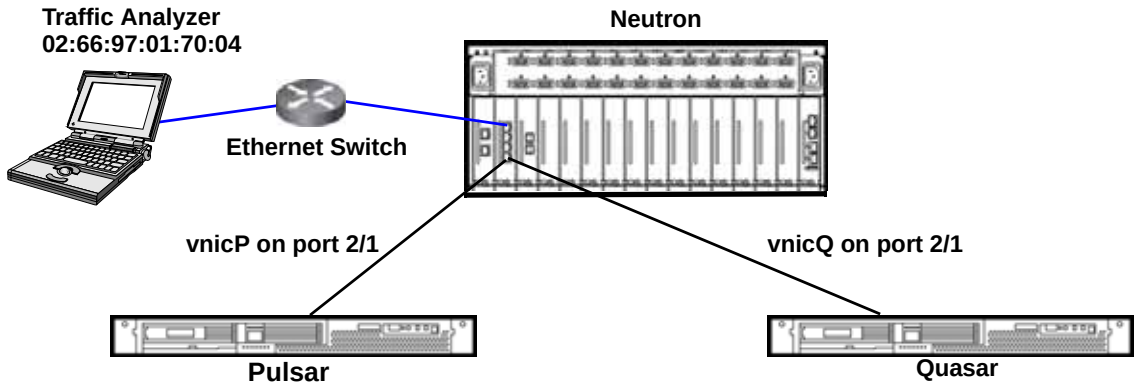

Figure 5 vNIC Mirroring to a Physical Port

To configure vNIC Mirroring for a vNIC-Switching port, you must know the MAC address of the device that will be receiving mirror data. To have bidirectional traffic between Pulsar and Quasar vNIC switched, and also mirror bidirectional traffic to the traffic analyzer, the configuration would be:

Note

The vNIC Mirroring feature is typically used in troubleshooting situations. In multichassis configuration, it is unlikely that both chassis will experience the same failure, so vNIC Mirroring is typically not required on both Fabric Directors. Instead, vNIC Mirroring is typically only used on one chassis, which is terminating the vNIC that contains the error. However, vNIC Mirroring can be configured on both Oracle Fabric Directors if needed.

Step 1 Create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 2 Create a vNIC in each server profile, and terminate it on the same physical port:

```
add vnic pulsar.vnicP 2/1
add vnic quasar.vnicQ 2/1
```

Step 3 Create the vNIC Mirror and assign it to an I/O port that is not supporting the vNIC switched hosts, and assign destination MAC address for the mirrored traffic:

```
set vnic vn1.pubstest1 -mirror=2/5 -mirror-mac=02:66:97:01:70:04
```

Note

The mirror must be configured on the same I/O module as the source traffic.

Step 4 Set the direction of traffic that will be mirrored on the vNIC:

```
set vnic vn1.pubstest1 -mirror-direction=both
```

Note

Traffic can be mirrored unidirectionally for transmit or receive (`txOnly` or `rxOnly` respectively) by setting the **`-mirror-direction=`** qualifier appropriately. Also, you can leave the mirror configured but disable it by specifying none for the direction. (None is the default).

Step 5 Set the IP address allocation method for the vNICs. In this example, DHCP will be used.

```
set vnic vnicP.pulsar 2/1 -addr-type=dhcp
set vnic vnicQ.quasar 2/1 -addr-type=dhcp
```

Step 6 Complete the vNIC configuration for both vNICs as needed by using the **`set vnic`** command to assign the operational parameters required.

VLAN Tagging and vNIC Switching

VLAN tagging supports isolation of different networks and hosts through the application of a VLAN tag, which can occur at the port level or the vNIC level. If the VLAN tagging occurs on the I/O port (as opposed to the network), then the way that VLAN tagging occurs can affect vNIC Switching.

In Figure 6, the vNICs must remain in the same VLAN in order to allow for successful vNIC Switching of traffic between the two hosts. As a result, to support vNIC Switching in a VLAN environment, you can should set the VLAN tagging mode on the I/O module to trunk. In trunk mode, the VLAN tags are not stripped at either the host or the I/O module, which allows traffic to be successfully vNIC switched between Pulsar and Quasar provided that both hosts are part of the same VLAN.

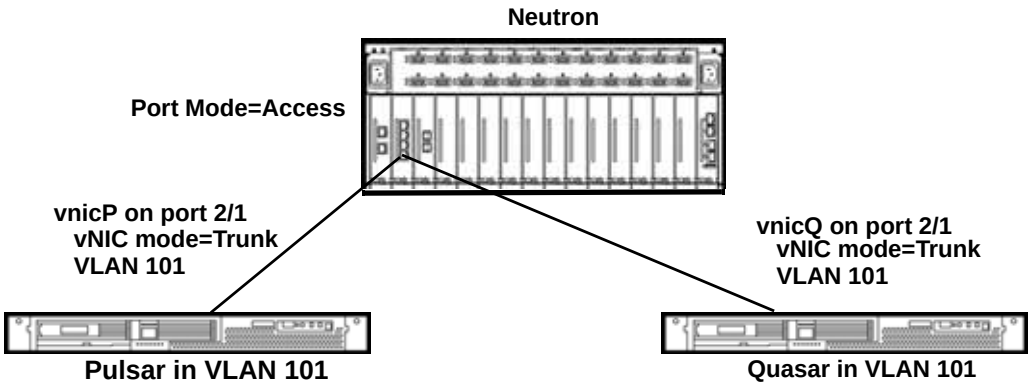

Figure 6 VLAN Tagging with vNIC Switching

In Figure 6, the hosts are VLAN aware and will use only VLAN 101 for communication. The port is set to trunk mode, so the VLAN tags are preserved for all vNICs terminated on that port, even VLAN tags not associated with the hosts. Because VLAN tags are preserved at the port level, traffic isolation is enforced for all VLANs. On the vNIC level, the tag is also preserved so that traffic is delivered to the correct host. In this example, the configuration should be as follows:

Step 1 Configure the Ethernet port mode:

```
set ethernet-port 2/1 -mode=trunk
```

Step 2 Add vNICs to the hosts:

```
add vnic vnicP.pulsar 2/1 -mode=trunk -access-vlan=101
add vnic vnicQ.quasar 2/1 -mode=trunk -access-vlan=101
```

Step 3 On the hosts, configure VLAN 101, assign an IP address to the VLAN, and configure additional VLAN properties (as needed) for your network.

Host Considerations

Some host considerations exist for vNIC Switching. Be aware of the following:

- For ESX 3.5, local ID numbers are required for vNICs. The vNICs that will be vNIC Switching cannot have the same local ID number.
- On ESX hosts, do not use vNIC Switching. Instead, use ESX's native NIC Teaming feature.
- On Solaris hosts, for HA and vNIC Switching, use Solaris's native IPMP for HA.

Multi-Chassis vNIC-to-vNIC Switching 15

This chapter contains the following topics:

Standard vNIC Switching

With multi-chassis vNIC-to-vNIC Switching, two vNICs for two different hosts are terminated on the same I/O port on the same I/O module. This configuration will exist on both Oracle Fabric Directors.

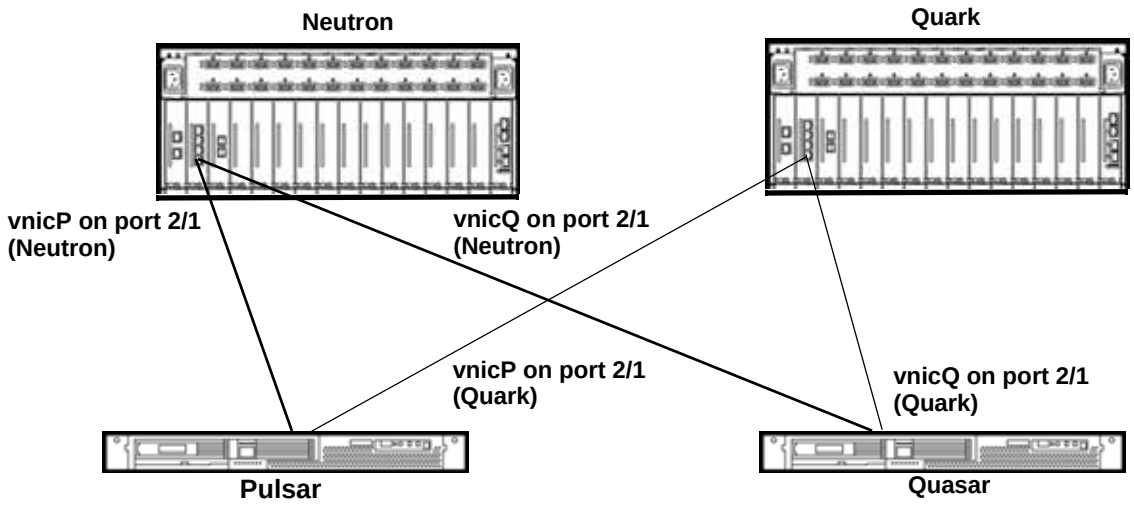

Figure 1 Standard vNIC-to-vNIC Switching

Configuration Details

For the example in Figure 1, the vNIC-to-vNIC Switching occurs for traffic sent between Pulsar and Quasar. The vNIC-to-vNIC Switching occurs automatically when the two vNICs are terminated on the same port. This configuration would be:

Step 1 On Neutron, create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 2 On Quark, create server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 3 On Neutron, create a vnic in each server profile, and terminate it on the same I/O port and I/O module on the Fabric Director:

```
add vnic pulsar.vnicP 2/1
add vnic quasar.vnicQ 2/1
```

Step 4 On Quark, create a vNIC in each server profile, and terminate it on the same I/O port and I/O module on the Fabric Director:

```
add vnic pulsar.vnicP 2/1
add vnic quasar.vnicQ 2/1
```

Step 5 Set the IP address allocation method for the vNICs. In this example, assume that addresses will be assigned through DHCP:

```
set vnic pulsar.vnicP 2/1 -addr-type=dhcp
set vnic quasar.vnicQ 2/1 -addr-type=dhcp
```

Step 6 Continue configuring the vNIC properties by issuing the **`set vnic`** command and setting the required parameters.

HA vNICs and vNIC Switching

In Figure 2, the vNIC-to-vNIC Switching occurs for traffic between Pulsar and Quasar. Any traffic between the two servers occurs on the Primary vNICs (indicated by (P)), which are the online and active vNICs. These vNICs must be terminated on the same I/O port on each Fabric Director. the Secondary vNICs are also configured on the same port, but cannot be configured on the same port as the primary vNICs or else you lose high availability. The secondary vNICs (shown as blue lines and indicated by (S)) are live standby vNICs and they are available to switch traffic between the two hosts if the primary vNICs go offline. With HA vNICs configured, you gain redundant data paths for your network traffic. HA vNICs do not load share.

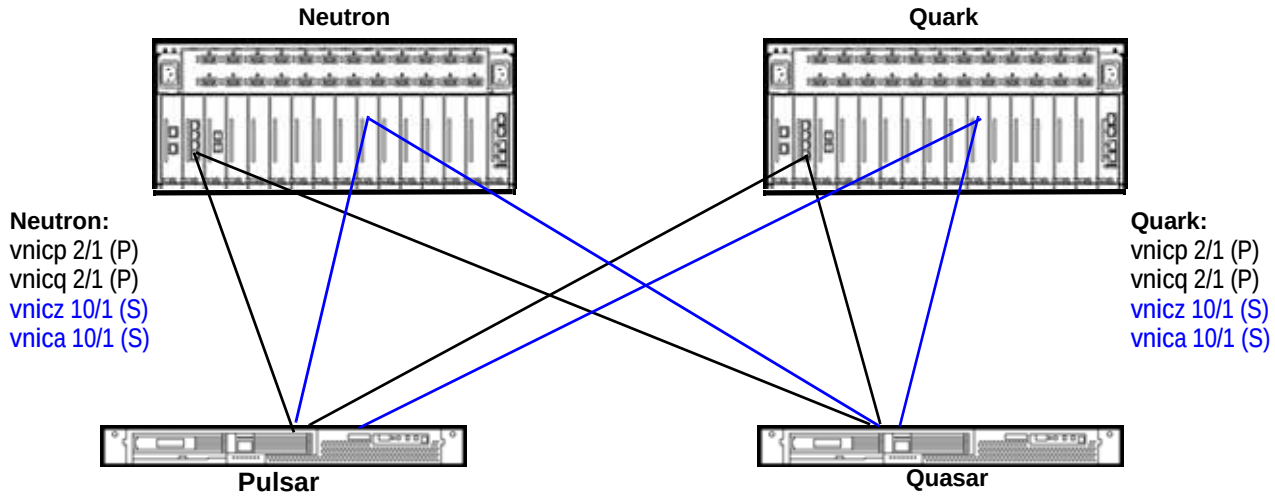

Figure 2 Example of HA vNIC

To configure HA vNICs to take advantage of vNIC-to-vNIC Switching as shown in this example, the configuration would be as follows:

Step 1 On Neutron, create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 2 On Quark, create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 3 On Neutron, create a primary vNIC in each server profile, and terminate it on the same I/O port and I/O module:

```
add vnic pulsar.vnicP 2/1 -primary
add vnic quasar.vnicQ 2/1 -primary
```

Step 4 On Quark, create a primary vNIC in each server profile, and terminate it on the same I/O port and I/O module:

```
add vnic pulsar.vnicP 2/1 -primary
add vnic quasar.vnicQ 2/1 -primary
```

Step 5 On Neutron, create the secondary vNICs. Make sure to:

- terminate them on a different port than the primary vNICs.
- terminate them on a different I/O module than the primary vNICs, as a best practice.


Note

You can terminate the primary and secondary vNICs on the same I/O module if you choose, but be aware that this configuration protects against port-level failures only.

Xsigo recommends configuring your secondary vNICs with the MAC address of a primary vNIC that is on a different I/O module whenever possible, to provide protection in the unlikely case of a module failure.

For example:

```
add vnic pulsar.vnicZ 10/1 ha -secondary -mac=10:02:53:56:09
add vnic quasar.vnicA 10/1 ha -secondary -mac=06:13:53:12:04
```

Step 6 On Quark, create the secondary vNICs. Follow the guidelines documented in Step 5:

```
add vnic pulsar.vnicZ 10/1 -secondary -mac=10:02:53:56:09
add vnic quasar.vnicA 10/1 -secondary -mac=06:13:53:12:04
```

Step 7 On Neutron, set the IP address allocation method for the primary vNICs. In this example, DHCP will be used.

```
set vnic pulsar.vnicP 2/1 -addr-type=dhcp
set vnic quasar.vnicQ 2/1 -addr-type=dhcp
```

Step 8 On Quark, set the IP address allocation method for the primary vNICs. In this example, DHCP will be used.

```
set vnic pulsar.vnicP 2/1 -addr-type=dhcp
set vnic quasar.vnicQ 2/1 -addr-type=dhcp
```

Step 9 On Neutron, set the IP address allocation method for the secondary vNICs. The IP address allocation method should match what is configured on the primary vNICs.

```
set vnic pulsar.vnicZ 10/1 -addr-type=dhcp
set vnic quasar.vnicA 10/1 -addr-type=dhcp
```

Step 10 Set the IP address allocation method for the secondary vNICs. The IP address allocation method should match what is configured for the primary vNICs.

```
set vnic pulsar.vnicZ 10/1 -addr-type=dhcp
set vnic quasar.vnicA 10/1 -addr-type=dhcp
```

Step 11 As an option, set the auto switchover parameters for the secondary vNICs on Neutron. Auto switchover is true by default, but can be set to false if needed. Auto switchover true or false determines whether the secondary vNICs switchback from primary to secondary when the original primary vNICs come back online. In this example, auto switchover will not be enabled.

```
set vnic pulsar.vnicZ 10/1 -autoswitchover=false
set vnic quasar.vnicA 10/1 -autoswitchover=false
```

Step 12 As an option, set the auto switchover parameters for the secondary vNICs on Quark. If you choose to set the auto switchover parameters, the parameters must be set to the same value on both Fabric Directors.

```
set vnic pulsar.vnicZ 10/1 -autoswitchover=false
set vnic quasar.vnicA 10/1 -autoswitchover=false
```


Note

In this example, module level HA exists, but port level HA can also be configured instead. You can configure port-level HA by terminating the secondary vNICs on the same module (but different port) as the primary vNIC. Even though this configuration is supported, Xsigo does not recommend it because both the primary and secondary vNICs are terminated on the same module, so in the unlikely event of a module failure, data paths between the servers are lost.

LAG and vNIC Switching

The Xsigo implementation of Link Aggregation Groups (LAG) treats all physical ports in the LAG as one logical port. Link Aggregation Groups (LAG) and vNIC Switching can work together to determine how vNIC Switching operates. When the vNICs are terminated within the same LAG, their traffic can be switched between different physical ports on the same module (called “cross module” vNIC Switching), thereby relaxing the requirement that the vNICs are terminated on the same physical port. Be aware that the different physical ports must be in the same LAG to support this method of vNIC Switching. Also, cross module vNIC Switching is supported on the 10-Port GE module only.

In Figure 3 on page 20, vnicP and vnicQ are terminated on port 2.1 and 2.10 respectively. The dot notation separating the slot and port identifiers indicates that the ports are members of a LAG. In this example, port 2.1 and 2.10 are configured in the same LAG (represented by the white rectangular background), so the on-board I/O processor can support cross module vNIC Switching between the two ports (shown by the red line) as long as:

- the ports are on the same I/O module
- in the same LAG

If one of the ports is moved out of the LAG, vNIC Switching will no longer be supported.

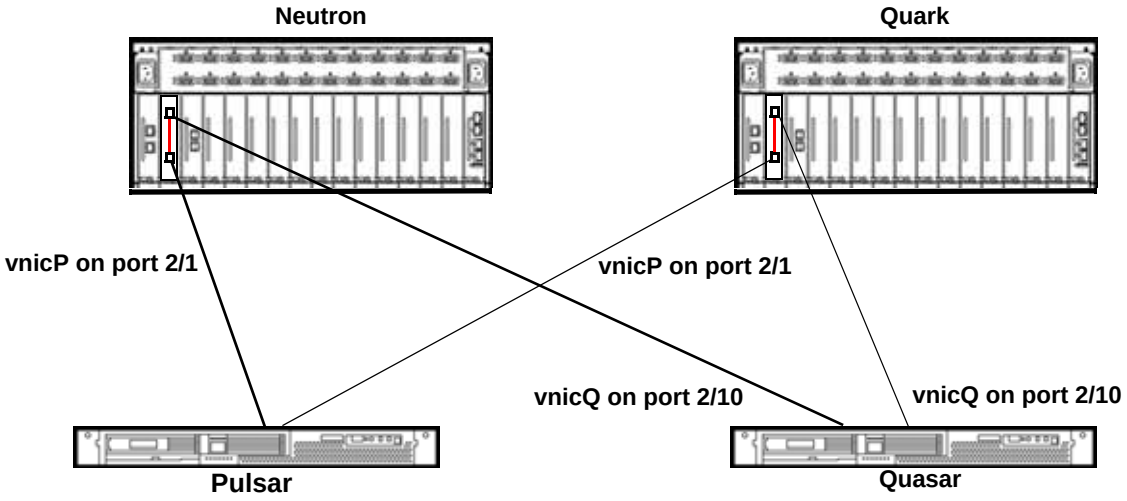

Figure 3 vNIC-to-vNIC Switching in Link Aggregation Group

To support vNIC-to-vNIC Switching between ports in a LAG, you would configure the vnICs as follows:

Step 1 On Neutron, create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 1 On Quark, create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 2 On Neutron, create the LAG by specifying the LAG name and adding the first port:

```
add lag 2.1 port 2/1
```

Step 3 On Neutron, add the second port to the LAG:

```
add lag 2.1 port 2/10
```

Step 4 On Quark, create the LAG by specifying the LAG name and adding the first port:

```
add lag 2.1 port 2/1
```

Step 5 On Quark, add the second port to the LAG:

```
add lag 2.1 port 2/10
```

Step 6 On Neutron, create a vNIC in each server profile, and terminate it on a port in the LAG:

```
add vnic pulsar.vnicP 2.1
add vnic quasar.vnicQ 2.1
```

Step 7 On Quark, create a vNIC in each server profile, and terminate it on a a port in the LAG:

```
add vnic pulsar.vnicP 2.1
add vnic quasar.vnicQ 2.1
```

Step 8 On Neutron, set the IP address allocation method for the vNICs. In this example, DHCP will be used.

```
set vnic pulsar.vnicP 2.1 -addr-type=dhcp
set vnic quasar.vnicQ 2.1 -addr-type=dhcp
```

Step 9 On Quark, set the IP address allocation method for the vNICs. In this example, DHCP will be used.

```
set vnic pulsar.vnicP 2.10 -addr-type=dhcp
set vnic quasar.vnicQ 2.10 -addr-type=dhcp
```

Step 10 On both Fabric Directors, complete the vNIC configuration for the vNICs as needed by using the **`set vnic`** command to assign the operational parameters required.


Note

This procedure shows how to configure cross-module vNIC Switching, in which two different ports are used, but vNIC Switching still occurs on the module. A LAG for the two different ports is required to make this configuration work. Standard vNIC Switching, in which vNICs for two separate servers are terminated on the same physical port, is also supported. This configuration typically occurs outside of a LAG, but can also be configured in the same LAG.

Standard vNIC Switching and cross-module vNIC Switching are mutually exclusive. You cannot configure the same two vNICs for standard and cross-module Switching, nor can you configure two vNICs on the same physical port (standard vNICs Switching) and have that port tied to another port in the same LAG (cross-module vNIC Switching).

Network QoS and vNIC Switching

Network QoS allows for bandwidth usage parameters to be applied to vNICs to ensure that bandwidth is used or available to hosts based on the needs of your network. Network QoS is configurable on individual vNICs, and vNIC Switching can be affected by the Network QoS profiles that are configured on them.

In Figure 4 on page 22, assume that Network QoS is configured on vnicP to guarantee a specific minimum bandwidth for the host. In this case, the same Network QoS profile must be configured on vnicQ. Without an exact duplicate Network QoS profile configured on each vNIC, when vNIC Switching occurs between the two hosts, the traffic can be altered by Network QoS Profiles in such a way that overruns or underruns can occur.

When configuring Network QoS in a vNIC Switching environment, always make sure to configure the same Network QoS Profile on both vNICs. Even a minor discrepancy between the two profiles can have an affect on how successfully vNIC-switched traffic is sent between two hosts.

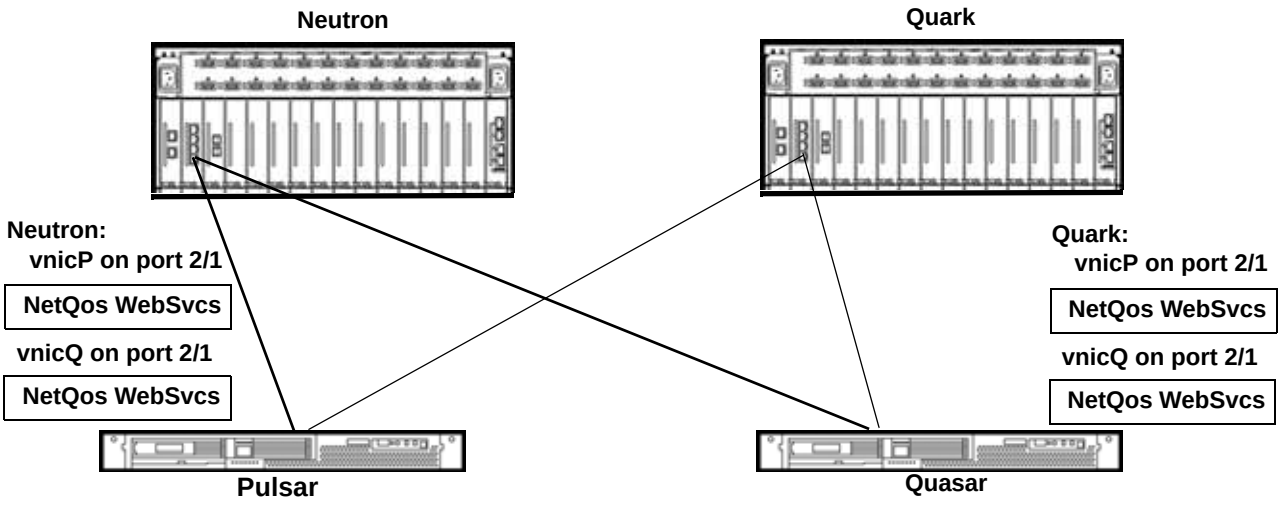

Figure 4 Network QoS and vNIC-to-vNIC Switching

To support Network QoS settings and vNIC-to-vNIC Switching, you would configure the vNICs as follows:

Step 1 On Neutron, create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 1 On Quark, create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 2 On Neutron, create a vNIC in each server profile and terminate it on a port:

```
add vnic vnicP.pulsar 2/1
add vnic vnicQ.quasar 2/1
```

Step 3 On Quark, create a vNIC in each server profile, and terminate it on a port:

```
add vnic vnicP.pulsar 2/1
add vnic vnicQ.quasar 2/1
```

Step 4 On both Fabric Directors, list all the default Network QoS Policer Profiles available:

```
show qos network policer
name                level  descr     cir         pir       cbs       pbs
-------------------------------------------------------------------------------
default/100m_1g     global 100m_1g   100Mbps     1Gbps     17.8814MB 35.7628MB
default/100m_250m   global 100m_250m 100Mbps     250Mbps   17.8814MB 35.7628MB
default/10g_10g     global 10g_10g   9.9297Gbps  9.9297Gbps 1.73395GB 1.73395GB
default/10m_100m    global 10m_100m  10Mbps      100Mbps   1.78814MB 3.57628MB
default/10m_1g      global 10m_1g    10Mbps      1Gbps     1.78814MB 3.57628MB
default/10m_50m     global 10m_50m   10Mbps      50Mbps    1.78814MB 3.57628MB
default/1g_10g      global 1g_10g    1Gbps       9.9297Gbps 178.814MB 357.628MB
default/1m_10m      global 1m_10m    1Mbps       10Mbps    182.617KB 365.234KB
```

```
default/250m_500m   global 250m_500m 250Mbps     500Mbps    44.7035MB 89.407MB
default/2g_10g      global 2g_10g    2Gbps       9.9297Gbps 357.628MB 715.256MB
default/3g_10g      global 3g_10g    3.00293Gbps 9.9297Gbps 536.965MB 1.04876GB
default/4g_10g      global 4g_10g    4Gbps       9.9297Gbps 715.256MB 1.39698GB
default/500m_750m   global 500m_750m 500Mbps     750Mbps    89.407MB  178.814MB
default/50m_100m    global 50m_100m  50Mbps      100Mbps    8.9407MB  17.8814MB
default/5g_10g      global 5g_10g    5.00122Gbps 9.9297Gbps 894.287MB 1.74665GB
default/64k_1m      global 64k_1m    64Kbps      1Mbps      11.7188KB 23.4375KB
default/6g_10g      global 6g_10g    6.00587Gbps 9.9297Gbps 1.04876GB 2.09752GB
default/750m_1g     global 750m_1g   750Mbps     1Gbps      134.11MB  268.221MB
default/7g_10g      global 7g_10g    7.00171Gbps 9.9297Gbps 1.22266GB 2.44532GB
default/8g_10g      global 8g_10g    8Gbps       9.9297Gbps 1.39698GB 2.79397GB
default/9g_10g      global 9g_10g    9.0022Gbps  9.9297Gbps 1.57199GB 3.14398GB
```

Step 5 Determine if you want to use a default Network Policer or create a custom Policer:

- If you want to use a default Network QoS Policer, proceed to Step 6.
- If you want to create a custom QoS Policer, proceed to Step 11.

Step 6 To use a default QoS Policer, assign it to the vNIC for the first host on Neutron. For example, assume the default Policer for 5 Gbps and 10 Gbps is on egress traffic only:

```
set vnic vnicP.pulsar egress-qos -policer=default/5g_10g enable
```


Note

This command also has the **ingress-qos** qualifier to set the QoS variables for traffic leaving the I/O module and entering the network. Also, by using the **disable** qualifier in the command, you can selectively enable or disable the assigned Network QoS Policer without having to delete the Policer or unbind it from the vNIC.

Step 7 On Neutron, assign a copy of the same QoS Policer to the vNIC for the second host. For example:

```
set vnic vnicQ.quasar egress-qos -policer=default/5g_10g enable
```


Note

The two QoS Policers must be identical for the vNICs attached to both hosts.

Step 8 To use a default QoS Policer, assign it to the vNIC for the first host on Quark. For example, assume the default Policer for 5 Gbps and 10 Gbps is on egress traffic only:

```
set vnic vnicP.pulsar egress-qos -policer=default/5g_10g enable
```

Step 9 On Quark, assign a copy of the same QoS Policer to the vNIC for the second host. For example:

```
set vnic vnicQ.quasar egress-qos -policer=default/5g_10g enable
```

Step 10 When default Policers for both hosts are added to both Fabric Directors, go to Step 17.

Step 11 To create a custom QoS Policer on Neutron, first add the QoS Policer and specify the bandwidth variables. For example, create a custom Policer with 2 Mbps CIR, and optional CBS of 1Mbps and PBS of 2 Mbps:

```
add qos network policer WebSvcs/1 2m -cbs=1m -pbs=2m
```

Step 12 On Neutron, after the custom Policer is added, assign it to the vNIC for the first host. For example:

```
set vnic vnicP.pulsar egress-qos -policer=WebSvcs/1 enable
```


Note

This command also has the **ingress-qos** qualifier to set the QoS variables for traffic leaving the I/O module and entering the network. Also, by using the **disable** qualifier in the command, you can selectively enable or disable the assigned Network QoS Policer without having to delete the Policer or unbind it from the vNIC.

Step 13 On Neutron, assign a copy of the same QoS Policer to the vNIC for the second host. For example:

```
set vnic vnicQ.quasar egress-qos -policer=WebSvcs/1 enable
```

Step 14 To create a custom QoS Policer on Quark, first add the QoS Policer and specify the bandwidth variables. For example, create a custom Policer with 2 Mbps CIR, and optional CBS of 1Mbps and PBS of 2 Mbps:

```
add qos network policer WebSvcs/1 2m -cbs=1m -pbs=2m
```

Step 15 On Quark, after the custom Policer is added, assign it to the vNIC for the first host. For example:

```
set vnic vnicP.pulsar egress-qos -policer=WebSvcs/1 enable
```

Step 16 On Quark, assign a copy of the same QoS Policer to the vNIC for the second host. For example:

```
set vnic vnicQ.quasar egress-qos -policer=WebSvcs/1 enable
```

Step 17 On Neutron, set the IP address allocation method for the vNICs. In this example, DHCP will be used.

```
set vnic vnicP.pulsar 2/1 -addr-type=dhcp
set vnic vnicQ.quasar 2/1 -addr-type=dhcp
```

Step 18 On Quark, set the IP address allocation method for the vNICs. In this example, DHCP will be used.

```
set vnic vnicP.pulsar 2/1 -addr-type=dhcp
set vnic vnicQ.quasar 2/1 -addr-type=dhcp
```

Step 19 Complete the vNIC configuration for both vNICs as needed by using the **set vnic** command to assign the operational parameters required.

vNIC Mirroring and vNIC Switching

Xsigo's vNIC Mirroring feature allows traffic received or transmitted on a vNIC to be copied on the I/O module that the vNIC is terminated on. The traffic is copied at the network processing chip, and then can be redirected to an output, either a port or a vNIC. To support vNIC mirroring and vNIC Switching, the mirrored data can be sent to an output port, vNIC on the same I/O port, or a destination MAC address.

In Figure 5 on page 25, the example shows vNIC Mirroring to an external device a traffic analyzer (such as a sniffer). The mirrored vNIC traffic (shown as a blue line) is sent through an output port.

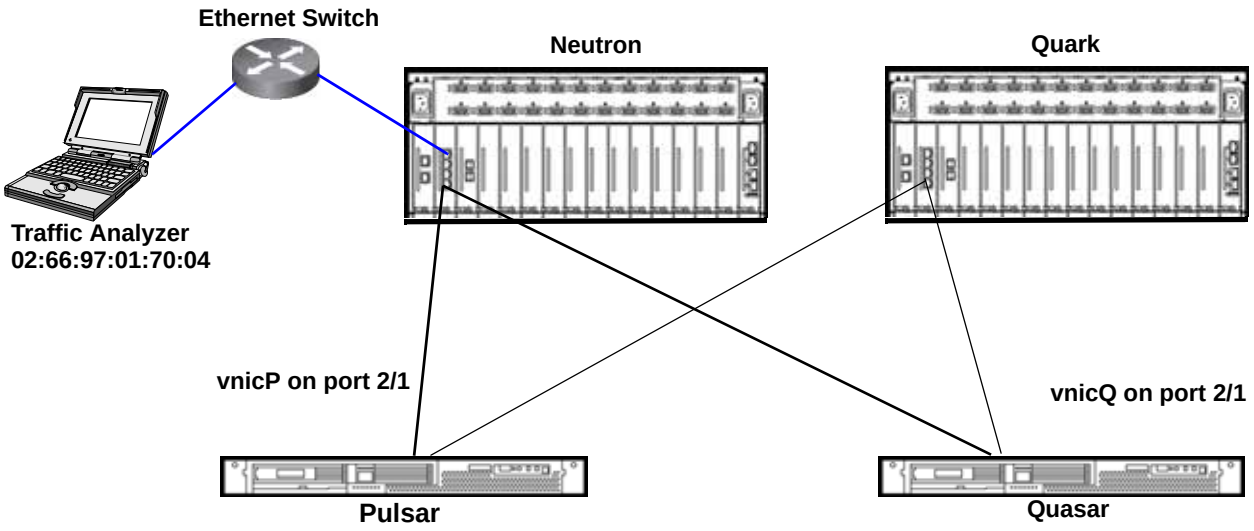

Figure 5 vNIC Mirroring to a Physical Port

To configure vNIC Mirroring for a vNIC-Switching port, you must know the MAC address of the device that you will be attaching. To have bidirectional traffic between Pulsar and Quasar vNIC switched, and also mirror bidirectional traffic to the traffic analyzer, the configuration would be:

Note

The vNIC Mirroring feature is typically used in troubleshooting situations. In multichassis configuration, it is unlikely that both chassis will experience the same failure, so vNIC Mirroring is typically not required on both Fabric Directors. Instead, vNIC Mirroring is typically only used on one chassis, which is terminating the vNIC that contains the error. However, vNIC Mirroring can be configured on both Fabric Directors if needed.

Step 1 On Neutron, create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 2 On Quark, create a server profile on each server:

```
add server-profile pulsar
add server-profile quasar
```

Step 3 On Neutron, create a vNIC in each server profile, and terminate it on the same physical port:

```
add vnic pulsar.vnicP 2/1
add vnic quasar.vnicQ 2/1
```

Step 4 On Quark, create a vNIC in each server profile, and terminate it on the same physical port:

```
add vnic pulsar.vnicP 2/1
add vnic quasar.vnicQ 2/1
```

Step 5 On the chassis that needs the troubleshooting session, create the vNIC Mirror and assign it to a MAC address (an external device) on a specific port:

```
set vnic vn1.pubstest1 -mirror=2/5 -mirror-mac-address=02:66:97:01:70:04
```

Step 6 Set the direction of traffic that will be mirrored on the vNIC:

```
set vnic vn1.pubstest1 -mirror-direction=both
```


Note

Traffic can be mirrored unidirectionally for transmit or receive (`txOnly` or `rxOnly` respectively) by setting the -mirror-direction qualifier appropriately. Also, you can leave the mirror configured but disable it by specifying none for the direction. (None is the default).

Step 7 If needed, you can configure vNIC Mirroring on the second chassis to copy traffic from a vNIC terminated on that Oracle Fabric Director. The traffic can be copied to the same destination MAC address, or a different one.

VLAN Tagging and vNIC Switching

VLAN tagging supports isolation of different networks and hosts through the application of a VLAN tag, which can occur at the port level or the vNIC level. If the VLAN tagging occurs on the I/O port, then the way that VLAN tagging occurs can affect vNIC Switching.

In Figure 6 on page 27, the vNICs must remain in the same VLAN in order to allow for successful vNIC Switching of traffic between the two hosts. As a result, to support vNIC Switching in a VLAN environment, you can should set the VLAN tagging mode on the I/O module to trunk. In trunk mode, the VLAN tags are not stripped at either the host or the I/O module, which allows traffic to be successfully vNIC switched between Pulsar and Quasar provided that both hosts are part of the same VLAN.

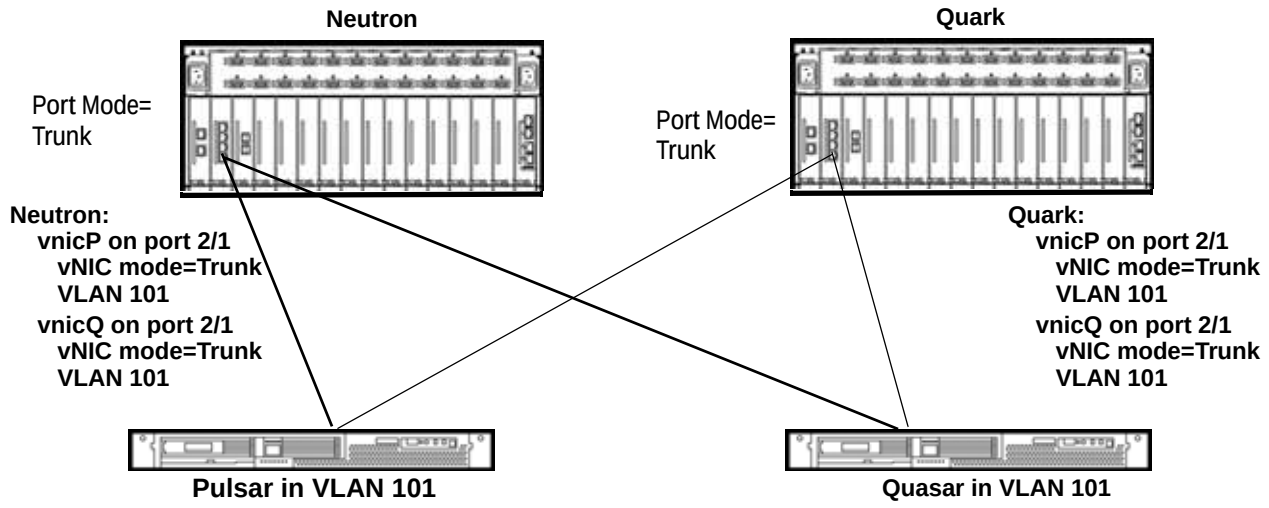

Figure 6 VLAN Tagging with vNIC Switching

To support vNIC Switching in the VLAN tagged environment, you will need to make sure that packets between the hosts are not stripped of their VLAN tags, or retagged with a different VLAN ID. Use trunk mode on the vNIC and on the port. In Figure 6, the hosts are VLAN aware and will use VLAN 101 only for communication. The port is set to trunk mode, so the VLAN tags are preserved for all VLANs, even VLANs not associated with the hosts. Because VLAN tags are preserved at the port level, traffic isolation is enforced for all VLANs. In this example, the configuration should be as follows:

Step 1 On Neutron, configure the Ethernet port mode:

```
set ethernet-port 2/1 -mode=trunk
```

Step 2 On Quark, configure the same Ethernet port mode:

```
set ethernet-port 2/1 -mode=trunk
```

Step 3 On Neutron, add vNICs to the hosts:

```
add vnic vnicP.pulsar 2/1 -mode=trunk -access-vlan=101
add vnic vnicQ.quasar 2/1 -mode=trunk -access-vlan=101
```

Step 4 On Quark, add vNICs to the hosts:

```
add vnic vnicP.pulsar 2/1 -mode=trunk -access-vlan=101
add vnic vnicQ.quasar 2/1 -mode=trunk -access-vlan=101
```

Step 5 On the hosts, configure VLAN 101, assign an IP address to the VLAN, and configure additional VLAN properties (as needed) for your network.

Host Considerations

Some host considerations exist for vNIC Switching. Be aware of the following:

- For ESX 3.5, local ID numbers are required for vNICs. The vNICs that will be vNIC Switching cannot have the same local ID number.
- On ESX hosts, do not use vNIC Switching. Instead, use ESX's native NIC Teaming feature.
- On Solaris hosts, for HA and vNIC Switching, use Solaris's native IPMP for HA.

C

E

H

L

N

P

S

V