

Release Notes for Oracle Fabric Interconnect and XgOS 3.9.2

These release notes document information about the Oracle Fabric Interconnect F1-15 and F1-4 and version 3.9.2 of Oracle's XgOS command-line interface. Additional release notes exist for:

- Oracle Virtual Networking host drivers for Windows hosts
- Oracle Virtual Networking host drivers for Linux hosts
- Oracle Virtual Networking host drivers for ESX Server Classic 4.1 and ESXi Server 4.1 hosts
- Oracle Virtual Networking host drivers for ESXi Server 5.0 host drivers and ESXi 5.1 hosts
- Oracle Virtual Networking host drivers for Solaris hosts
- Oracle Fabric Manager



Note

If you will be upgrading the Fabric Interconnect to the latest version of XgOS, as part of the upgrade:

- all I/O modules will reboot
- the I/O service running on the Fabric Interconnect will restart, which will cause a service interruption
- the SCP might reboot based on conditions related to the XgOS kernel or drivers

Be aware that I/O will be interrupted temporarily, but will resume. Depending on changes in the new version of XgOS, the Fabric Interconnect might be rebooted as part of the upgrade.

These notes contain the following sections:

- [What's New in this Release](#) on page 2
- [System Requirements for Fabric Interconnects](#) on page 2
- [System Limitations and Restrictions](#) on page 3
- [User Guides](#) on page 8
- [Documentation Additions](#) on page 9
- [Downloading XgOS](#) on page 15
- [Known Problems](#) on page 17
- [Fixed Problems](#) on page 20
- [Accessing Oracle Support](#) on page 21

Overview

The Oracle Fabric Interconnect is a service-oriented platform that interconnects data-center resources based on application demands.

Customers and partners are requested to send comments and report bugs to Oracle by filing a customer case through the My Oracle Support web portal.

The Fabric Interconnect can be managed using any of the following interfaces:

- the XgOS Command Line Interface (CLI)
- Fabric Manager, which is a GUI for managing multiple Fabric Interconnects, or Fabric Manager with the VMware Extension, which is an extension that registers Fabric Manager for interoperability with VMware VirtualCenter or vSphere.

What's New in this Release

This release contains fixes. For more information, see [Fixed Problems](#) on page 20.

System Requirements for Fabric Interconnects

This section documents system requirements for the Fabric Interconnect, such as host operating systems, servers, Ethernet network switches and routers, SAN switches, and storage that the Fabric Interconnect supports.

Operating Systems

The following platforms are supported by the Fabric Interconnect:

- Oracle Solaris 10 Update 11 and 11 Update 1
- Oracle Linux 5 Update 1 and later, 6 Update 2 (32-bit and 64-bit), and 6 Update 3 (32-bit- and 64-bit)
- Red Hat Enterprise Linux 4 (RHEL) Update 8
- RHEL 5 GA (Update 0) and Updates 1 through 7
- RHEL 6 Update 1 and Update 2
- CentOS 5 GA (Update 0) and Updates 1 through 6
- Citrix XenServer 5.6 FP1 (fully supported) and 5.6 SP2 (experimental support)
- Microsoft Windows Server 2003 R2 with SP2
- Microsoft Windows Server 2008 SP2
- Microsoft Windows Server 2008 R2
- Oracle Virtual Machine (OVM) 3.1.1
- VMware ESX Server Classic 4.1 GA and Update 1, and ESXi Server 4.1 GA and Update 1
- VMware ESXi Server 5 GA (Update 0) and Update 1
- VMware ESXi Server 5.1 GA (Update 0) and Update 1

Separate release notes exist for the various host drivers for each platform. For information about host drivers, see the release notes for the relevant platform.

Servers and HCAs

The Oracle Fabric Interconnect can operate with any industry-standard server that supports x86- and x86_64-based CPUs from Intel or AMD. The Fabric Interconnect supports PCI Express servers only.

Version 2.7.0 and later requires that each server be connected to the Oracle Fabric Interconnect using a host channel adapter (HCA) installed in a PCI Express slot. Server HCAs must be mem-free. The Fabric Interconnect does not support memory-based HCAs.

Supported Firmware Version for ConnectX Family of HCAs

XgOS 3.9.2 supports Oracle HCAs as well as third-party manufacturers' HCAs with the requirement that the HCAs use the required version of firmware. Oracle HCAs include:

- Sun InfiniBand Dual Port 4x QDR PCIe Low Profile Host Channel Adapter M2. For information, including the required firmware version, see <http://docs.oracle.com/cd/E19241-01/index.html>
- Sun InfiniBand Dual Port 4x QDR PCIe ExpressModule Host Channel Adapter M2. For information, including the required firmware version, see <http://docs.oracle.com/cd/E19157-01/index.html>
- Oracle Dual Port QDR InfiniBand Adapter M3. For information, including the required firmware version, see http://docs.oracle.com/cd/E40985_01/index.html

For third-party HCAs, consult the manufacturer's documentation for the required firmware version.

Ethernet Network Switches and Routers

The Fabric Interconnect I/O ports operate with any 1 Gbps or 10 Gbps network switch that supports established IEEE standards for Ethernet.

Storage and SAN Switches

The Oracle Fabric Interconnect can operate with industry-standard Fibre Channel SAN storage devices, and any SAN switches that support NPIV login through the Fibre Channel SAN. For Oracle vHBAs to successfully connect to a Fibre Channel switch, NPIV must be enabled on the switch.

System Limitations and Restrictions

This section documents system limitations and restrictions for InfiniBand Fabric Interconnects and XgOS. Be aware of the following system limitations and restrictions for the Fabric Interconnect and XgOS 3.9.2.

Before Burning Firmware on an HCA, You Must Manually Ensure It Is Correct

When using the in-band upgrade feature to upgrade an HCA's firmware (`set physical-server <name> upgrade-hca <lid> firmware <filename>`), the command does not perform any compatibility check to ensure that the correct version or type of firmware is about to be burned onto the HCA. Therefore, you must manually ensure that the firmware is correct for an HCA before burning that firmware onto the HCA. This requirement is especially important for most blade servers, which use system specific HCAs known as mezzanine cards that are often flashed with a custom firmware image created by the blade server's manufacturer.



The in-band upgrade feature is not currently supported for ConnectX-3-based HCAs.

Burning firmware on a server HCA or mezzanine card has the following main phases:

1. Identifying the HCA to determine which firmware is required.
2. Finding the firmware on the firmware vendor's site. Be aware that firewalls or other security restrictions might prevent you from connecting to the vendor's website.
3. Downloading the correct firmware to the Fabric Interconnect so that the firmware can be "pushed" from the Fabric Interconnect to the HCA(s) or mezzanine cards.
4. Upgrading the firmware on the HCA or mezzanine card.

Follow this procedure, which assumes you are burning Mellanox firmware. Although this procedure uses Mellanox as an example, the procedure should be applicable to other vendors firmware.

Step 1 Issue the `show physical-server` command for each server and find the board ID. In the following example, the `board-id` field is highlighted.

```
show physical-server dellblade13.lab.xsigo.com hca
lid          node-guid          board-id          device-id
firmware-version  hardware-version  option-version
-----
6           0024e89097ffac03      DEL08F0120009    26428           2.7.710
unknown
69          f04da290977932e4      DEL08F0120009    26428           2.7.710
unknown
2 records displayed
```

Step 2 Make a note of the board ID listed in the command. You will need it to search for the compatible firmware.



Mellanox HCA board IDs start with `mt_` but other manufacturers may start with different prefixes (for example, `HP` for Hewlett Packard, or `DEL` for Dell).

Step 3 Go to <http://www.mellanox.com>. (If you are using a different vendor's firmware, log on to that vendor's website.)

Step 4 At the Mellanox (or other vendor's) website, enter the board ID into the *Search* field on the home page and press **Enter** (or click the magnifying glass icon) to start the search.

Typically the query returns only 1 link and in that link you will be able to find the board-id and FW image.



Due to firewall rules and security restrictions, you might not be able to connect to this site. If you can connect, you can copy this URL for use in the next step.

If you cannot connect to this site, you will need to find an alternate method to download the firmware to a “safe” location in your network. For example, download the firmware to an SCP server then install the firmware on the Fabric Interconnect from that SCP server. (In addition to HTTP shown in this step, you can also use SCP or FTP.)

- Step 5** When you have located the correct firmware, install it on the Fabric Interconnect by issuing the **system install hca-firmware** <full-path-to-file-name> command:. For example:

```
system install hca-firmware http://www.mellanox.com/downloads/firmware/file.bin
```

The <full-path-to-file-name> argument requires the full file path to the firmware. The firmware file name is a .bin file, and FTP, SCP, or HTTP are supported file transfer protocols.

When the command completes, the firmware file is downloaded to the Fabric Interconnect at which point it can be pushed from the Fabric Interconnect to the server HCAs or mezzanine cards.

- Step 6** When the download is complete, verify that the correct firmware is present on the Fabric Interconnect by issuing the **show system hca firmware** command and finding the firmware file you just downloaded. You will use the filename in the next step.
- Step 7** When you are certain that each firmware image is correct for the HCA(s), issue the **set physical-server** <name> **upgrade-hca** <lid> **firmware** command to burn the correct firmware onto the HCA(s).



If you are burning multiple firmware versions on multiple HCAs or mezzanine cards, make sure not to mix up the firmware versions! You must get the correct firmware version onto the correct HCAs or cards.

- Step 8** Repeat this procedure as needed for all HCAs on which you will be burning firmware.

To Upgrade Fabric Interconnects Connected to ESXi 5.0 and 5.1 Hosts with PVI vNICs, Set Server Profiles “Down”

If your Fabric Interconnect(s) are connected to ESXi 5.0 hosts which have one or more PVI vNICs, use this procedure to upgrade the Fabric Interconnect. This upgrade procedure is required only for Fabric Interconnects that are connected to ESXi 5.0 hosts that have PVI vNICs.

For each server profile connected to an ESXi 5.0 host:

- Step 1** Before upgrading chassis, set the server profile(s) connected to ESXi 5.0 hosts to “down” state.
- ```
set server-profile <profile-name> down
```
- Step 2** Perform the Fabric Interconnect reboot or XgOS upgrade.
- Step 3** After the reboot (or upgrade) is complete, set the server profile(s) connected to ESXi 5.0 hosts to “up”:
- ```
set server-profile <profile-name> up
```

For Dynamic LAGs, Set LACP on the Fabric Interconnect I/O Port and Its Peer

With the current implementation of LACP, if passive-mode LACP will be used on a LAG, both the Fabric Interconnect's Gigabit Ethernet I/O ports and the peer device at the end of a port in the LAG must be configured with LACP. If LACP is not enabled at both ends of the links in the LAG, a loop can occur.

For VLAN Tagging, Make Sure Port and vNIC Modes Are Congruent

The Fabric Interconnect does not currently check the validity of VLAN tagging modes. For example, the Fabric Interconnect allows the following configuration:

- an Ethernet port or Link Aggregation Group (LAG) set to `trunk` mode and a vNIC terminating on that Ethernet port or LAG set to `access` or `trunk` mode. This configuration is valid and the Fabric Interconnect allows it.
- an Ethernet port or LAG set to `access` mode and a vNIC terminating on that Ethernet port or LAG set to `trunk` mode. This configuration is *invalid*, but the Fabric Interconnect still allows it.

The listed settings are a misconfiguration that can negatively affect the traffic on the VLAN, and can cause hosts to not receive traffic. Be aware that this misconfiguration is not prevented on the Fabric Interconnect, and no error message is displayed if this configuration is specified.

Consideration for Enabling and Disabling IGMP

Be aware of the following consideration regarding enabling and disabling IGMP:

The ability to enable or disable IGMP is supported at the I/O module-level only. You cannot control the state of IGMP traffic on a per-port basis.

Subnet for a Server Profile's Default Gateway Must Match at Least One vNIC in the Same Server Profile

Due to the current implementation of subnet masking on the Fabric Interconnect, the default gateway's IP subnet must match the subnet of at least one vNIC configured in the same Server Profile. Be aware of this limitation when configuring the default gateway for a Server Profile.

Use `-allow-lun-masking` at vHBA Creation to Dynamically Add or Delete LUN Masks

LUN Masks cannot be configured after the vHBA is created unless the vHBA was initially created with a LUN Mask or the `-allow-lun-masking` option is set when the vHBA is initially configured. If your vHBAs need a LUN Mask, or might need them in the future, use the `-allow-lun-masking` option the first time you create each vHBA. For example:

```
add vhma vh2.profile1 3/2 -allow-lun-masking
```

With the `-allow-lun-masking` option set when the vHBA is initially created, LUN Masks can be attached to or unattached from the vHBA either now or at any time in the future.

Merging InfiniBand Fabrics Causes a Service Interruption

The Fabric Interconnect's InfiniBand fabric supports both of the following operational modes:

- standalone in which a single Fabric Interconnect has a self-contained fabric
- shared, in which multiple Fabric Interconnects share the same IB fabric

If you have multiple standalone Fabric Interconnects, and will merge them together into a shared fabric, be aware that merging the IB Fabrics is not a seamless process. The multiple separate IB fabrics will successfully merge, but this process does cause service interruptions. Because merging IB fabrics is disruptive, Oracle recommends that you perform this procedure during a maintenance window, or a similar period of time when service interruptions will be kept to a minimum. A procedure for merging individual fabrics can be found in the *Fabric Accelerator Quick Start Guide*.

The system downgrade Command Does Not Revert the Existing Configuration

The **system downgrade** command is implemented to allow you to back out of an upgrade to a configuration that is not functioning correctly. As a result, the **system downgrade** command does not save the “new” (and possibly non-functional) config and convert it after the downgrade completes. This design is intentional and prevents the possibility of bringing non-functional configs into other software versions after the downgrade.

Be aware that the **system downgrade** command does not convert the latest configuration by default. However, you can explicitly bring the latest configuration along with a downgrade by exporting the latest configuration, performing the downgrade, then re-importing the latest config.

Cannot Specify a Netmask for an Inbound ACL Set Up for IP Address Range

For ACLs, using the ‘in’ condition for IP address ranges does not work predictably because this direction automatically attaches a netmask to the IP address of each IP address range. You cannot manually add a netmask to the IP address range; the ACL assumes the network class on its own. This issue will be addressed in a future release of software.

Deleting In-Use vHBAs Can Cause Server Instability

Attempting to delete an in-use vHBA can cause serious server instability. To avoid such instability, follow the instructions for vHBA removal in the *XgOS Command-Line Interface User Guide*.

Virtual NICs Might Not Be Displayed in Ethereal Interfaces

A vNIC might not be displayed in Ethereal interfaces. If this situation occurs, restart the Ethereal network packet-filtering driver by issuing the **net stop npf** command followed by the **net start npf** command.

Limitations on the Length of Virtual Resource Names

Be aware of the following limitations for virtual resource names:

- With Linux operating systems, the names of virtual resources are restricted to the following lengths:
 - vNICs: 10 characters
 - vHBAs: 15 characters
 - Server Profiles: 31 characters

When creating these entities on the Fabric Interconnect, comply with the name string limits. If you violate the length restrictions, VLANs will not work.

- For Windows HA vNICs, the vNIC name, which includes `vnic_name` plus `server_profile_name`, or `<vnic_name>.<server_profile_name>`, must be 31 characters or less. This limit occurs because the host allows for only 16 characters in the name string, and counts the dot (.) that separates `<vnic_name>` from `<server_profile_name>` as one of the characters. As a result:
 - The maximum character string length for `<vnic_name>` is 14 characters
 - The maximum supported character string length for `<server_profile_name>` is 17 characters

When creating HA vNICs on the Fabric Interconnect for Windows hosts, comply with the name string limits.

- When creating vNICs on the Fabric Interconnect, the XgOS CLI restricts HA vNICs to a maximum of 7 characters, but Fabric Manager allows a maximum of 9 characters. Be aware of this discrepancy, and when creating HA vNICs through Fabric Manager, make a best effort to name them with a maximum of 7 characters to have predictable HA vNIC names between the CLI and Fabric Manager.
- For ESX Classic ESXi servers, vNIC names must be limited to a maximum of 9 characters.

Limitation on Restoring Configuration

Anytime you import a backed-up Fabric Interconnect configuration, if that configuration is not in sync with the current hardware configuration on the Fabric Interconnect, the import can fail. You might experience this failure when importing a configuration that was exported prior to some hardware or virtual I/O resource changes, or when a software feature like vNIC Mirroring was configured in the exported config, but is not configured on the chassis where you re-import.

User Guides

User guides for the Fabric Interconnect are available on CD for shipments to new customers, and can be downloaded from the My Oracle Support web portal.

Oracle provides the following Fabric Interconnect product documentation in PDF format:

- *Oracle Fabric Interconnect Hardware and Host Drivers Installation Guide*
- *XgOS Software Upgrade Guide*
- *XgOS Command-Line User Guide*
- *Remote Boot Guide*
- *Oracle Fabric Manager User Guide*
- *Oracle Fabric Manager VMware Integrator User Guide*
- *Oracle Software Defined Network Controller User Guide*
- *Fabric Performance Monitor User Guide*
- *XgOS vNIC Switching Configuration Guide*

You can access these manuals from the Oracle Technical Network (OTN) at http://docs.oracle.com/cd/E38500_01/index.html. On that web page, scroll down to find the document that you need, then click the **Download** button to pull a copy of the PDF to your web browser.

Documentation Additions

The following sections contain additional text that does not appear in the *Fabric Interconnect Hardware and Host Drivers Guide, 3.8.0*.

Upgrading to the FP-2

XgOS 3.9.0 and later supports a new Front Panel, called the FP-2. This hardware component contains the SCP, and the main difference between the FP-2 and its predecessor is that the SCP is embedded on the FP-2 and is no longer a field-replaceable part. Existing Front Panels are still supported in the field, but will no longer be an orderable item.



Note

It is not a requirement that you upgrade to the FP-2. The previous version of Front Panel is supported in XgOS 3.9.0 and later. These instructions are provided in cases where you want or need to upgrade the Front Panel Assembly to an FP-2.

Consideration for Upgrading to the FP-2

The minimum version that supports the FP-2 is XgOS 3.9.0. As a result, you cannot downgrade a Fabric Interconnect that has an FP-2 to XgOS version 3.8.x or earlier and have a functional FP-2.

For existing customers with the previous Front Panel (not new customers who received the FP-2 with their shipment from the factory), Oracle recommends keeping the previous Front Panel to allow flexibility in downgrading to XgOS 3.8.x if needed.

If you are upgrading a Fabric Interconnect that has ESX hosts connected, make sure to set the Server Profiles for the ESX hosts to “down” before attempting this upgrade.

Upgrade Procedure

The FP-2 is a field-upgradeable part (the SCP on the FP-2 is not). For the upgrade procedure you will need:

- #2 Phillips (cross-head) screwdriver
- an ESD wrist strap
- an ESD-safe work surface
- an ESD bag for the old Front Panel Assembly



Note

This upgrade procedure requires you to power down the Fabric Interconnect, so a service interruption will occur. Oracle recommends that you perform the upgrade during a scheduled service window.

The upgrade procedure has the following main phases:

- [Backing Up the Configuration](#) (optional, but recommended if you want to preserve your existing config after the hardware upgrade. If you do not save your existing config, you will need to recreate your entire config after the upgrade.)
- [Removing the Front Panel Assembly](#)

- [Installing the FP-2](#)
- [Restoring the Configuration](#) (This is conditional depending on whether the config was saved or not.)

Backing Up the Configuration

If you wish to retain your current setup after the replacement is finished you will need to follow these steps to save the configuration to a storage device other than the chassis. For illustrative purposes, this procedure will assume you are using a USB token to store the config and user files while the Fabric Interconnect is powered down.

- Step 1** Export your config to an external storage device. For example, insert the USB token into the USB port on the Management Module.
- Step 2** Issue the **system export <file-name>** command to save the config and all user files that you want to keep.
- The **<file-name>** argument is alphanumeric character string that names the current config. You can assign any name to the configuration you are saving. You do not need to specify the file extension or type as part of the file-name string.
 - The system config is saved to XML format by default.
 - The user files are typically located in the user home directory in `/var/fsroot/` on the Fabric Interconnect. So, for example, for the default “admin” user, the user files would be located in `/var/fsroot/admin`

For example, to save the current running configuration to a file `xgos-old-config`, you would issue the command as follows:

```
system export /usb/xgos-old-config.xml
```



You can also save the config in two steps, by:

- 1.) Saving the config to the `/home` directory of the account you logged in with—for example
system export xgos-old-config.xml
 - 2.) Copying the config to the USB token by using either the **file copy** or **cp** commands, for example:
file copy xgos-old-config.xml /usb/xgos-old-config.xml
or **cp xgos-old-config.xml /usb/old-config.xml**
-

- Step 3** When files are on the USB token, remove it.
- Step 4** Proceed to [Removing the Front Panel Assembly](#).

Removing the Front Panel Assembly

- Step 1** Power down the Fabric Interconnect.
- Step 2** Using the screwdriver, release the two spring-loaded screws (1) securing the release levers to the chassis. See [Figure 1](#).




Figure 1 Loosening the Fabric Interconnect F1-15 Front Panel Assembly Screws

- Step 3** Simultaneously grasp both of the release levers and slowly raise them until they are horizontal (2), then use the levers to carefully slide the assembly out of the chassis (3). See [Figure 2](#).




Figure 2 Removing the Fabric Interconnect F1-15 Front Panel Assembly

Installing the FP-2

To install the FP-2:

- Step 1 Carefully align the FP-2 with the Front Panel bay in the chassis. Make sure that the FP-2 is facing up, and the handles are pointing straight towards you as shown in (3) in [Figure 2](#).
- Step 2 Keeping the FP-2 level and holding it by the edges, gently slide the FP-2 into the chassis until you feel some resistance when the FP-2 is almost all the way inserted.



Note

This resistance is normal and occurs when the FP-2 plug contacts the power socket inside the chassis. When you feel this resistance, stop pushing on the FP-2.

- Step 3 When the FP-2 plug has contacted the socket inside the chassis, simultaneously grasp both of the release levers and slowly lower them until they are vertical and flush with the chassis. This step is the reverse of (2) in [Figure 2](#).
- Step 4 Using the screwdriver, tighten the two spring-loaded screws securing the release levers to the chassis. This step is the reverse of [Figure 1](#). *Do not overtighten the screws!*
- Step 5 Power on the system.
- Step 6 As an option, you can attach to serial port 1 (SER-1) at 115200 baud and leave the session open to watch the progress of the installation and start up while the Fabric Interconnect is brought back online.
- Step 7 When the chassis is up, log on as `admin/admin`.



Note

If it does not immediately allow you to log in, wait until it does. During startup, the Fabric Interconnect waits for a specific Ethernet interface called `ethMgmt` to become ready before allowing admin logon the first time.

Also, if you are not using the serial port, you will not be able to ssh to the chassis until the network is ready. In this case you can log on as `admin` as soon as it will accept your connection.

Also, be aware that the management Ethernet interface address is assigned by static addressing by default. The default address for the management interface is 192.168.1.1. This address is the default management interface address for the FP2 also, so when you install the FP2, it will use 192.168.1.1. Be aware that if you have manually set the management interface address to something other than 192.168.1.1, that non-default address will not be used after the FP2 is installed. You will not be able to connect to the non-default address, and you will need to use 192.168.1.1 instead.

After the FP-2 is correctly installed and the system powers up, the Fabric Interconnect recognizes itself as a “new” system so the First Boot Wizard will run to assist you in configuring basic functionality that gets the Fabric Interconnect back online.

- Step 8 Complete the FirstBoot wizard by entering the system name, subnet manager settings, phone home settings (if desired) and so on.
- Step 9 Proceed to [Restoring the Configuration](#).

Restoring the Configuration

If you have saved a configuration and want to reload the configuration, you can import the saved configuration. The configuration file was saved to an off-SCP device (such as a USB token), so when you import the configuration, make sure you specify the path and file name when you import the saved configuration.



Caution

When a saved configuration is being imported, the current I/O module configuration on the Fabric Interconnect must match the I/O Module configuration in the saved config. Do not add or remove I/O cards during the FP-2 hardware replacement procedure. If you change the I/O module configuration during the replacement procedure, the I/O module configuration will not match the saved configuration, and importing the config will abort.

To restore the configuration from a backup:

- Step 1 Insert the USB token into the USB port on the Management Module.
- Step 2 Copy all user files off of the USB token into the correct directory.
- Step 3 Import the configuration from the USB token by issuing the **system import <file-name>** command.

For example, to load the saved configuration called `xgos-old-config.xml` from a USB device, you would issue:

```
system import /usb/xgos-old-config.xml
```



Note

You can also import the config in two steps, by:

1.) Copying the config from the USB token to the /home directory of the current account, by using either the **file copy** or **cp** commands, for example at the XgOS prompt:

```
file copy /usb/xgos-old-config.xml xgos-old-config.xml  
or cp /usb/xgos-old-config.xml xgos-old-config.xml
```

2.) Importing the config from the /home directory (the default location)—for example, **system import xgos-old-config.xml**

When the saved config is successfully re-imported, the Fabric Interconnect will prompt you for a reboot.

- Step 4 When prompted for the reboot, enter **Y** to accept the reboot. You must enter **Y** to re-import the saved config. Do not press N or the saved config will not be imported and the Fabric Interconnect will be put into an indeterminate state.
- Step 5 Allow the reboot to complete, then remove the USB token from the USB port.



Note

Remember that the FP2 is supported in XgOS 3.9.0 and later. If the config you re-imported (`xgos-old-config.xml`, in this example) is less than 3.9.0, you will need to upgrade to XgOS 3.9.0 or later after the configuration is re-imported.

You can download XgOS 3.9.0 by following the instructions in [Downloading XgOS](#) on page 15 and running the **system upgrade** command to upgrade to XgOS 3.9.0 or later.

Consideration for the watch Command

Be aware that the **watch fc-card <slot>** command always shows utilization at zero instead of the actual utilization.

Some MIB Objects Are Not Mappable from Enterprise MIBs

The following list shows the MIB objects that remain in the Fabric Interconnect's Enterprise MIB, and are not mappable into the ENTITY-MIB, ENTITY-STATE-MIB, ENTITY-SENSOR-MIB, and IF-MIB.

- Chassis MIB Objects
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoChassisMfgDate.0 = STRING: 2007-4-6,0:0:0.0
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoChassisBaseMacAddr.0 = STRING: 00:13:97:05:70:00
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoChassisLastUpTime.0 = STRING: 2012.01.12.18.46.13.0
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoChassisCpuUsage.0 = STRING: 0
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoChassisMemoryUsage.0 = STRING: 36.3608
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoChassisWWN.0 = STRING: 50:01:39:70:00:04:70:00
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoChassisCpuUsageInt.0 = INTEGER: 0
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoChassisMemoryUsageInt.0 = INTEGER: 36
- Card MIB Objects
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoCardMfgDate.10 = STRING: 2006-2-24,0:0:0.0
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoCardAdminState.10 = INTEGER: up(1)
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoCardOperState.10 = INTEGER: up(2)
- Fan MIB Object
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoFanExpSpeed.1 = INTEGER: 4200
- Temperature MIB Object
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoTemperatureProbeMaxValue.1 = STRING: 55
- Power Supply MIB Object
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsignoPowerSupplyMfgDate.1 = STRING: 2007-3-9,0:0:0.0

- Port MIB Objects
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoPortType.1 = INTEGER: sanFcPort(121)
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoPortAdminSpeed.1 = STRING: AutoNegotiate
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoPortWWNN.1 = STRING: 50:01:39:71:00:04:70:5B
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoPortWWPN.1 = STRING: 50:01:39:70:00:04:70:5B
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoPortFullDuplex.1 = INTEGER: true(1)
- vNIC MIB Objects
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVnicLocation.1 = STRING: 5/7
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVnicEgressQoS.1 = STRING:
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVnicEgressQoS.1 = STRING:
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVnicIngressQoS.1 = STRING:
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVnicIngressQoS.1 = STRING:
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVnicServerProfileIndex.1 = INTEGER: 2
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVnicHAPreferences.1 = STRING:
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVnicTrunkMode.1 = INTEGER: noTrunkMode(2)
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVnicAccessVLANId.1 = STRING: 1
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVnicVLANIds.1 = STRING:
- vHBA MIB Objects
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVhbaLocation.1 = STRING: 10/1
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVhbaQoS.1 = STRING:
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVhbaQoS.1 = STRING:
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVhbaQoS.1 = STRING:
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVhbaQoS.1 = STRING:
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVhbaQoS.1 = STRING:
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVhbaWWNN.1 = STRING: 50:01:39:71:00:04:72:22
 - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVhbaWWPN.1 = STRING: 50:01:39:70:00:04:72:22

Downloading XgOS

This section documents where to get software and documentation for this version of XgOS.



Note

If your Fabric Interconnect contains an FP-2, be aware that the Fabric Interconnect must be running XgOS 3.9.0 or later to support the hardware.

Downloading Software

You can get software for this version of XgOS from My Oracle Support:

- Step 1 Point your browsers to support.oracle.com and log in to display the MOS Dashboard.
- Step 2 On the Dashboard, click the Patches and Updates tab.
- Step 3 In the Patch Search tab, click Search->Product or Family to display the search dialog. See [Figure 3](#).

Figure 3 My Oracle Support — Searching to Find Product

- Step 4 In the Product is field, enter “Oracle Fabric Interconnect.” This field will auto-fill when enough characters are entered to make the string unique.
- Step 5 In the Release is field, enter “3.9.2”. This field will auto-fill when enough characters are entered to make the string unique.
- Step 6 When these fields contain the product and release, click Search.
- Step 7 On the resulting table, click the row containing the 3.9.2 software and select **Download**.
- Step 8 When the software is downloaded, get it onto a network node the Fabric Interconnect can access.
- Step 9 From the Fabric Interconnect, run the system upgrade command to upgrade the XgOS on the Fabric Interconnect to version 3.9.2.

Getting Documentation

Oracle Fabric Interconnect documentation is available as PDF files on the Oracle Technical Network (OTN) at http://docs.oracle.com/cd/E38500_01/index.html.

On that web site, scroll down until you find the document that you need, then click the **Download** button to pull the document to your web browser.

Known Problems

The following tables list known problems with the Fabric Interconnect and XgOS:

- [Table 1 Known Problems in Version 3.9.2 Specific to the Oracle I/O Modules](#)
- [Table 2 Known Problems in Version 3.9.2 Specific to the Fabric Interconnect](#)

Table 1 Known Problems in Version 3.9.2 Specific to the Oracle I/O Modules

Number	Description
16330195	<p>When using link aggregation groups (LAGs) on the 10-Port GE I/O module, attempting to remove multiple LAG ports with a single command fails.</p> <p>To work around this issue, remove one port at a time by using the remove lag <LAG-number> port <port-to-be-removed> command.</p>
16335073	MAC-based QoS is not yet supported on the 4-Port 10 GE Module.
16335216	<p>When performing changes to the Fibre Channel storage network, a <code>RD_REG_DWORD_FAIL</code> error can cause one or more paths to storage to become dead or not seen by one or more of the Fabric Interconnects. If this problem occurs, FC I/O cards can reset (even FC I/O cards that would seem to be unrelated to the SFC topology change) which can cause a service interruption.</p> <p>If traffic stops and FC I/O cards do not reset, you can work around this problem by resetting the FC card. For example, assume the errored FC I/O card is in slot 11.</p> <p>Step 1 Issue the set system fc-card 11 down command</p> <p>Step 2 Wait for the card to be brought down and traffic to halt</p> <p>Step 3 Issue the set system fc-card 11 up command</p>
16336829	<p>When a large number of vHBAs are configured on a Fiber Channel module, a problem can cause the connected storage target to be lost if the fibre channel cable is repeatedly removed and reinserted. The problem occurs due to a timing issue in which “dead” messages are processed after RSCNs. This problem is seen only on ESX and Windows Server 2008 hosts. This problem also has been observed only if the Fibre Channel module is supporting a large number of vHBAs—for example, 36 or more.</p> <p>If you are using a Line Rate Fibre Channel module, you might encounter this issue if your I/O module has more than 36 vHBAs.</p>

Table 2 Known Problems in Version 3.9.2 Specific to the Fabric Interconnect

Number	Description
16330154	Some IBM X-series servers do not support HCAs with the Oracle Option ROM installed. Such servers are not able to perform PXE boot or iSCSI boot over a vNIC, nor are they able to SAN Boot over a vHBA. Replacement HCAs without the incompatible Option ROM are available from Oracle.

Table 2 (continued) Known Problems in Version 3.9.2 Specific to the Fabric Interconnect

Number	Description
16330295	The display of ACL Rule Sets and rules through the CLI can display the address mask portion of the IP address condition unpredictably. In cases where the ACL Rule is configured for a range (with the <> option), ACL addresses can sometimes receive a /24 address mask, and sometimes not.
16330478	<p>The Oracle implementation of LACP for Link Aggregation Groups does not contain a comprehensive set of LACP properties and statistics. As a result, you can display only some of the pertinent LACP information through the XgOS CLI.</p> <p>You can work around this issue by logging in to the peer device (Ethernet switch), and displaying the available LACP information on the switch. You should be able to combine the information available on the Fabric Interconnect and the peer device to obtain comprehensive information about LACP properties.</p>
16330523	When the same user is configured as a local user (user configured on the Fabric Interconnect) and a RADIUS user, a conflict occurs when the role is assigned during log in. For example, assume <i>paulw</i> is created as a local user with the operator role, and a <i>paulw</i> is also created as a RADIUS user with the network role. When the IMS search order is set to use internal authentication, and <i>paulw</i> logs in with the user and password for the network role, the operator role is enforced instead.
16330603	In the Oracle implementation of SNMP, an SNMP Agent problem prevents the vNIC and vHBA location information from being saved in the <code>xsigovnicremove</code> and <code>xsigovhbaremoved</code> trap. The problem occurs because no previous value for the location is saved. As a result, if vNICs or vHBAs are removed, the object is deleted at that time and the location cannot be retrieved for the trap.
16331758	If duplicate MAC addresses occur on two separate vNICs terminated on the same port, a constant series of RnR NACKs is generated, which can lead to heavy network flooding or broadcasts. The heavy network flooding or broadcasts, in turn, can cause a reboot of the Gigabit Ethernet I/O Module where the vNICs are configured.
16333475	<p>When the MTU set on an HA vNIC is different than the MTU set on a Windows host, the Windows host rejects traffic on the HA vNIC. In addition, the Fabric Interconnect does not recognize the mismatch when the Windows host has rejected the MTU mismatch. The Fabric Interconnect shows the vNIC state as <code>up/Failed</code> instead of <code>up/MTUMismatch</code>. For additional details about the state of the HA vNIC, see the Windows system event log.</p> <p>If you suspect this problem is occurring, make sure to check the HA vNIC MTU against the Windows Host MTU. If they are different, make the HA vNIC MTU match the Windows host's MTU.</p>

Table 2 (continued) Known Problems in Version 3.9.2 Specific to the Fabric Interconnect

Number	Description
16333708	<p>Sometimes a problem with LACP LAGs can cause a lockup after an SCP reboot. When the problem occurs, error messages similar to the following are displayed:</p> <pre>Apr 21 15:40:19 iop-11 vn2_agent_cp[608]: [ERR] VN2_CP Failed to send sync message: timeout: expected 30000, real: 29980 Apr 21 15:40:19 iop-11 vn2_agent_cp[608]: [ERR] l2m::L2M eth10g-11 HW Error: ERROR Cookie 5491 end (Command timeout)</pre> <p>When the error messages are displayed, commands no longer complete successfully when issued for the LACP LAG and the I/O Module where the LACP LAG is configured.</p> <p>For static LAGs, the error message are displayed if this problem occurs, but the lockup does not occur and commands for static LAGs and the I/O Module still complete successfully.</p> <p>To work around this problem, reboot the I/O module with the affected LACP LAG. For example, if the affected LACP LAG is on the I/O module in slot 6, you would issue set iocard 6 reset</p>
16333713	<p>The allowed VLAN range feature is supported on a trunk-mode vNIC or I/O port that is to specify the VLANs that are valid on that vNIC or port. This feature is not applicable to access-mode VLANs.</p> <p>However, a software problem causes the allowed VLAN range feature to be configurable on access-mode VLANs. The XgOS does not prevent this misconfiguration. When configuring the allowed VLAN range feature, make sure to do so only on trunk-mode vNICs or ports.</p>
16334326	<p>Using the in-band HCA upgrade feature to upgrade the firmware (set physical-server <name> upgrade-hca <lid> firmware <filename>) also causes the Option ROM to be removed from select HCAs. This problem occurs due to the type of firmware shipped with certain HCAs. The following HCAs are affected:</p> <ul style="list-style-type: none"> • HP_08B0420001 • HP_08B0000001 • HP_0AD0000003 • HP_0160000009 • DEL08F0120009 • DEL08F0120002 <p>If your server use any of these HCAs, and you will be using in-band update to flash new firmware on the HCA, be aware that you will need to re-flash the Option ROM after flashing the HCA firmware.</p>
16336217	<p>When using the in-band upgrade feature to upgrade an HCA's firmware (set physical-server <name> upgrade-hca <lid> firmware <filename>), the command does not perform any compatibility check to ensure that the correct version or type of firmware is about to be burned onto the HCA. As a result, it is possible to burn incompatible firmware onto an HCA, which, in the worst case, can cause the HCA to become unusable.</p> <p>You can avoid this problem by ensuring that you are burning the correct firmware for the HCAs in your servers. For information, see Before Burning Firmware on an HCA, You Must Manually Ensure It Is Correct on page 3.</p>

Table 2 (continued) Known Problems in Version 3.9.2 Specific to the Fabric Interconnect

Number	Description
16336437	<p>When hard zoning is used, if you add a second array and rezone to include that array, a problem can prevent the vHBAS from showing all targets in the new zone. When this problem occurs, the Fabric Interconnect shows an inconsistent view of the targets available through the vHBAs, and the target shows most vHBAs, but often shows no LUNs on the available targets. This problem has been observed in the following conditions:</p> <ul style="list-style-type: none"> • only when access to all LUNs is removed, which is an atypical procedure • only with HP targets
16336454	<p>Performing a chassis reboot or upgrading the XgOS sometimes causes ESXi 5.0 hosts to crash to purple screen (PSOD) when one or more Fabric Accelerator PVI s are configured. Be aware that this condition can occur.</p> <p>You can work around this problem by following the procedure documented in To Upgrade Fabric Interconnects Connected to ESXi 5.0 and 5.1 Hosts with PVI vNICs, Set Server Profiles “Down” on page 5.</p>

Fixed Problems

Table 3 shows the fixes in this release. This release might contain fixes from previous releases. If so, such information is indicated in text.

Table 3 Fixed Problems in XgOS 3.9.2

Number	Description
17362408	<p>A problem caused vHBA chip to experience data corruption issues.</p> <ul style="list-style-type: none"> • In one instance, an internal disk-write data RAM error caused one byte in 16 to have a bit flip condition. • In one instance, an software error caused a spontaneous link list not initialized condition. <p>These problems are fixed in Oracle’s XgOS 3.9.2.</p>
16745501	<p>A problem in chip-level microcode sent a message that the Fabric Interconnect’s 1 port, 10 Gigabit Ethernet modules sometimes treated as a fatal error and reset itself to recover. This problem is fixed in XgOS 3.9.1 by changing the message handling.</p>
16335175	<p>When a large number of vHBAs are configured on a host, the vHBA code sometimes missed or improperly handled RSCNs. As a result, one vHBA or target might not have been recognized. This is problem is fixed in XgOS 3.9.1.</p>

Accessing Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/us/corporate/acquisitions/xsigo/support-1849142.html> or visit <http://www.oracle.com/us/corporate/accessibility/support/index.html> if you are hearing impaired.

