Release Notes for Fabric Director and XgOS 3.9.0

These release notes document information about the Fabric Director (VP780 and VP560) and version 3.9.0 of Oracle's XgOS command-line interface. Additional release notes exist for:

- Oracle's Xsigo Windows host drivers
- Oracle's Xsigo Linux host drivers
- Oracle's Xsigo ESX Server Classic 4.1 and ESXi Server 4.1 host drivers
- Oracle's Xsigo ESXi Server 5.0 host drivers and ESXi 5.1 host drivers
- Oracle's Xsigo Fabric Manager


Note

If you will be upgrading the Fabric Director 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 Director 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 Director might be rebooted as part of the upgrade.

These notes contain the following sections:

Xsigo Systems
70 West Plumeria Drive
San Jose, CA 95134
USA
http://www.xsigo.com
Tel: +1.408.329.5600

2

Overview

Oracle's Xsigo Fabric Director 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 Xsigo by filing a customer case through the Xsigo Technical Support web portal (http://support.xsigo.com). Xsigo is fully committed to responding to all feedback regarding our product and greatly appreciates customer involvement. If you need to contact Xsigo Customer Support, you can facilitate your interaction with Customer Support by first gathering some troubleshooting information. See page 24.

The Fabric Director 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 Directors, or Fabric Manager with the VMware Extension, which is an extension that registers Fabric Manager for interoperation with VMware VirtualCenter or vSphere.

What's New in this Release

This release contains the following new features and functionality:

- In this release, the maximum number of simultaneous CLI sessions that can be open is 10. Exceeding this limit can cause the Fabric Director to behave unpredictably. The 10-session limit is not currently enforced through software, so you must make sure to limit the number of simultaneous sessions yourself.
- A new Front Panel assembly, called the FP-2, is available. This new part offers a new processor and a streamlined design which includes an embedded SCP which is no longer a field-replaceable part. In addition, the entire FP-2 is RoHS compliant. The FP-2 is installed in all new shipments from the factory.


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 also.

  The FP-2 is supported only with XgOS 3.9.0 and later. Earlier versions of XgOS do not support the FP2.

- Fixes have been added. For more information, see Fixed Problems on page 23.

System Requirements for Fabric Directors

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

Operating Systems

The following platforms are supported by the VP780 and VP560:

- Oracle Linux 5 Update 1, 6 Update 2, and 6 Update 3
- 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)
- Debian Ubuntu 8.10 Server (initial-release 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)

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

Servers and HCAs

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

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

Supported Firmware Version for ConnectX Family of HCAs

XgOS 3.9.0 supports ConnectX and ConnectX-2 HCAs with the requirement that the HCA firmware version is 2.7.x and later.

ConnectX-3 HCAs are supported when connected to FDR fabrics or IB switches, if the HCAs are running firmware version 2.10.700.

Supported Firmware Version for InfiniHost HCAs

XgOS 3.9.0 supports InfiniHost HCAs with requirement that the HCA firmware version is 5.3.0.

Ethernet Network Switches and Routers

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

4

System Limitations and Restrictions

This section documents system limitations and restrictions for InfiniBand Fabric Directors and XgOS. Be aware of the following system limitations and restrictions for the Fabric Director and XgOS 3.9.0.

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.


Note

The in-band upgrade feature is not currently supported for ConnectX-3 HCAs. For more information, see 23143 in the Known Problems section.

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 Director so that the firmware can be "pushed" from the Fabric Director 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                   unknown

69      f04da290977932e4      DEL08F0120009      26428        2.7.710
unknown                   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.

Note

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.

Note

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 Director 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 Director 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 Director at which point it can be pushed from the Fabric Director to the server HCAs or mezzanine cards.

Step 6 When the download is complete, verify that the correct firmware is present on the Fabric Director 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).

Caution

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 Directors Connected to ESXi 5.0 and 5.1 Hosts with PVI vNICs, Set Server Profiles “Down”

If your Fabric Director(s) are connected to ESXi 5.0 hosts which have one or more PVI vNICs, use this procedure to upgrade the Fabric Director. This upgrade procedure is required only for Fabric Directors 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 Director 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 Both the Fabric Director 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 Director’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 Director does not currently check the validity of VLAN tagging modes. For example, the Fabric Director 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 Director 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 Director 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 Director, 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 Director, 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 vhba 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 Director's InfiniBand fabric supports both of the following operational modes:

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

If you have multiple standalone Fabric Directors, 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, Xsigo 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 Director, 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 Director for Windows hosts, comply with the name string limits.
- When creating vNICs on the Fabric Director, 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 Director configuration, if that configuration is not in sync with the current hardware configuration on the Fabric Director, 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 Director are available on CD for shipments to new customers, and can be downloaded from the Xsigo Technical Support web portal.

Xsigo Systems provides the following Fabric Director product documentation in PDF format:

- *Fabric Director Hardware and Host Drivers Installation Guide*
- *XgOS Software Upgrade Guide*
- *XgOS Command-Line User Guide*
- *XgOS Remote Boot Guide*
- *Fabric Manager User Guide*
- *Fabric Accelerator Quick Start Guide*
- *Fabric Performance Monitor User Guide*
- *XgOS vNIC Switching Configuration Guide*

You can download these manuals by logging in to the Xsigo Support page (http://support.xsigo.com) and clicking the "Documentation" tab on toolbar at the top of the page. You will need a login and password to access the Xsigo Support page. See page 24.

Documentation Additions

The following sections contain additional text that does not appear in the Xsigo Technical Documentation.

Upgrading to the FP-2

The release 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. 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 Director 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), Xsigo recommends keeping the previous Front Panel to allow flexibility in downgrading to XgOS 3.8.x if needed.

If you are upgrading a Fabric Director 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 Director, so a service interruption will occur. Xsigo 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 assumes you are using a USB token to store the config and user files while the Fabric Director 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 Director. 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
```

Note

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 Director.

Step 2 Using the screwdriver, release the two spring-loaded screws (1) securing the release levers to the chassis. See Figure 1.

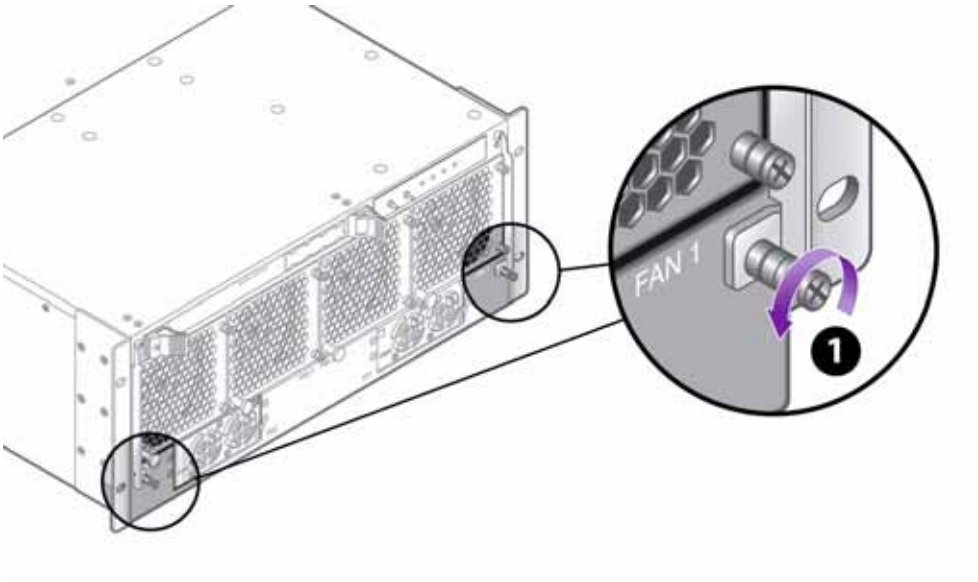

Figure 1 Loosening the VP780 Front Panel Assembly Screws

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.

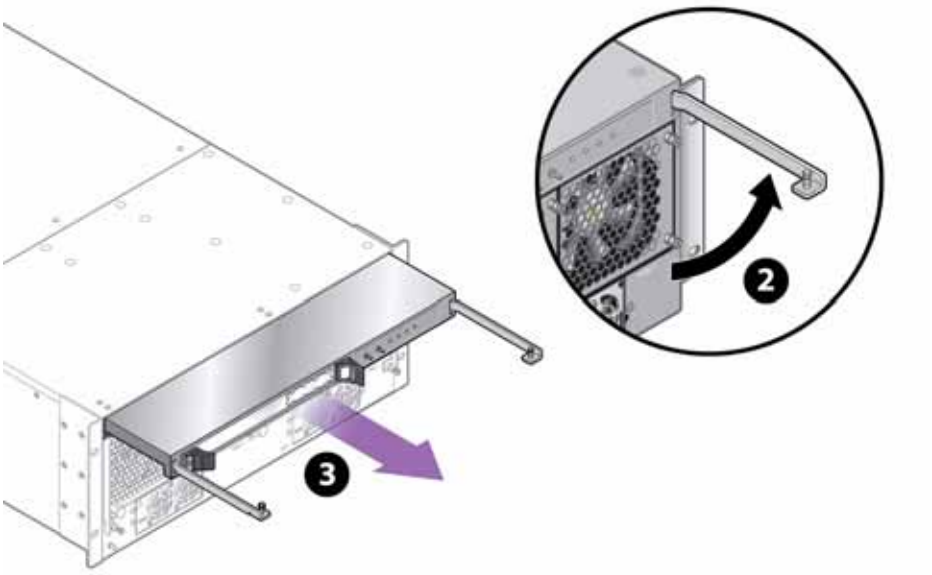

Figure 2 Removing the VP780 Front Panel Assembly

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 insides 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 Director 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 Director 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 Director recognizes itself as a "new" system so the First Boot Wizard will run to assist you in configuring basic functionality that gets the Fabric Director 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 a 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 Director 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 Director 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 Director 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 after the configuration is re-imported.

You can download XgOS 3.9.0 by following the instructions in Downloading XgOS on page 16 and running the **`system upgrade`** command to upgrade to XgOS 3.9.0.

Considerations for the watch Command

Be aware of the following considerations:

- The **`watch fc-card`** `<slot>` command always shows utilization at zero instead of the actual utilization.
- On Ethernet Fabric Directors, if you set multiple vNICS to **`-if=none`** then back to the original interface (or a new one) the **`watch vnic`** command shows zero utilization for those vNICs even when they are actively supporting traffic.

Some MIB Objects Are Not Mappable from Xsigo Enterprise MIBs

The following list shows the MIB objects that remain in the Xsigo 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::xsigoChassisMfgDate.0 = STRING: 2007-4-6,0:0:0.0
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoChassisBaseMacAddr.0 = STRING: 00:13:97:05:70:00
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoChassisLastUpTime.0 = STRING: 2012.01.12.18.46.13.0
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoChassisCpuUsage.0 = STRING: 0
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoChassisMemoryUsage.0 = STRING: 36.3608

  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoChassisWWN.0 = STRING: 50:01:39:70:00:04:70:00
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoChassisCpuUsageInt.0 = INTEGER: 0
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoChassisMemoryUsageInt.0 = INTEGER: 36
- Card MIB Objects
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoCardMfgDate.10 = STRING: 2006-2-24,0:0:0.0
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoCardAdminState.10 = INTEGER: up(1)
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoCardOperState.10 = INTEGER: up(2)
- Fan MIB Object
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoFanExpSpeed.1 = INTEGER: 4200
- Temperature MIB Object
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoTemperatureProbeMaxValue.1 = STRING: 55
- Power Supply MIB Object
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoPowerSupplyMfgDate.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::xsigoVnicEgressQoSCir.1 = STRING:
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVnicEgressQoSPir.1 = STRING:
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVnicIngressQoSCir.1 = STRING:
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVnicIngressQoSPir.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::xsigoVhbaQoSCir.1 = STRING:
  - XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVhbaQoSPir.1 = STRING:

— XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVhbaQoSCbs.1 = STRING:

— XSIGO-IODIRECTOR-ENTITY-MIB::xsigoVhbaQoSPbs.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 Director contains an FP-2, be aware that XgOS 3.9.0 is the minimum version that supports that hardware.

Downloading Software

You can get software for this version of XgOS from either of the following locations:

- Oracle's Technical Network (OTN)
- Oracle's Software Delivery Cloud

Oracle's Technical Network (OTN)

You can download this version of XgOS through Oracle's Technical Network (OTN), which is available without a user account or password. Software is available through this method, but not documentation. For release notes for this release, see Getting Documentation.

To get the software:

Step 1 Point your browser to http://www.oracle.com/technetwork/indexes/downloads/index.html

Step 2 Scroll down to the *Drivers* section.

Step 3 Click the Xsigo Drivers link.

Step 4 In the *Fabric Interconnect* section, find this version of XgOS, and click the link to download the software to a network-accessible node in your network.

Step 5 Using file copy, SCP, or another file transfer protocol, copy the XgOS software from the network-accessible node to the Fabric Director.

Step 6 When the new XgOS is on the Fabric Director, run the **`system upgrade`** command on the Fabric Director to upgrade it to the appropriate version.

Oracle's Software Delivery Cloud

You can download this version of XgOS through Oracle's Software Delivery Cloud (also called "edelivery").

Note

To access the Oracle Software Delivery Cloud, you will first need to get a user account and password. To get a user account and password, please follow the instructions on https://edelivery.oracle.com which is the Oracle Software Delivery Cloud website.

Software is available through this method, but not documentation. For release notes for this release, see Getting Documentation.

To get the software:

Step 1 Point your browser to https://edelivery.oracle.com.

Step 2 Log in with your account name and password.

Step 3 Accept the terms and conditions of usage (if prompted) and click ***Continue*** to display the Media Pack Search page as shown in Figure 3.

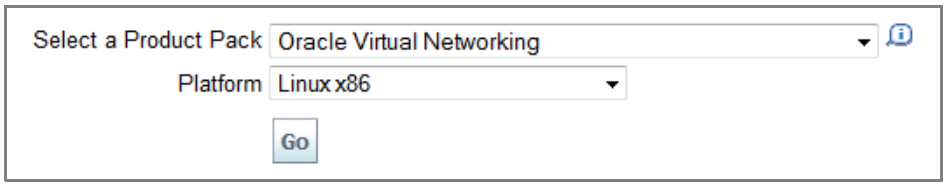

Figure 3 Searching for Oracle Virtual Networking Software

Step 4 From the *Select a Product Pack* dropdown menu, select Oracle Virtual Network.

Step 5 From the *Platform* dropdown menu, select any value. For XgOS software the platform does not matter. However, you can use the *Platform* dropdown to search for host drivers by OS or hypervisor type.

Step 6 Click ***Go*** to display the result of your search.

Step 7 On the resulting page, locate `Oracle Fabric Interconnect Firmware - XgOS 3.9.0` and click the ***Download*** button to get the XgOS on a network-accessible node in your network.

Step 8 Using file copy, SCP, or another file transfer protocol, copy the XgOS software from the network-accessible node to the Fabric Director.

Step 9 When the new XgOS is on the Fabric Director, run the **`system upgrade`** command on the Fabric Director to upgrade it to the appropriate version.

Getting Documentation

Xsigo documentation for this release is on the Xsigo ServiceWorks website. If you are looking for software, you will need to download it by using the procedure documented in either of the previous sections.

To download documentation:

Step 1 Point your browser to http://www.xsigo.com/supportdlc/

Step 2 On the left navigation panel, click the *Oracle – Xsigo Product Documentation* link as shown in figure Figure 4.

Figure 4 Getting Documentation for Software

Step 3 On the resulting page, scroll down to the *Release Notes* section.

Step 4 Find the documentation for this version of software, and click the link to display the PDF.

Known Problems

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

Table 1 Known Problems in Version 3.9.0 Specific to the Xsigo I/O Modules

Number	Description
9918	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. To work around this issue, remove one port at a time.
12498	On the 10 GE module, the `io link down alarm` is inconsistently reported when the InfiniBand cable is removed: • When the InfiniBand cable is unplugged, the `io link down alarm` is sent for some (but not all) the configured vNICs. • When the InfiniBand cable is plugged back in, the `io link down alarm` does not clear. Instead, it clears only when traffic occurs on the path affected by the cable pull. To workaround this problem, switch the traffic back to the affected IB link.
19942	MAC-based QoS is not yet supported on the 4-Port 10 GE Module.
21698	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. 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.

Table 2 Known Problems in Version 3.9.0 Specific to the Fabric Director

Number	Description
7989	Some IBM X-series servers do not support HCAs with the Xsigo 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 Xsigo Systems.
9505	With XgBoot installed, IBM X-series of servers will fail to continue past Xgboot in the boot order if vHBA or vNIC booting fails. This is due to a limitation in the X-server BIOS. When prompted, press ***q*** to quit XgBoot and the system will boot normally.

Table 2 (continued) Known Problems in Version 3.9.0 Specific to the Fabric Director

Number	Description
10317	If you have autocommit switched off and have many pending configuration changes when you begin a system upgrade or downgrade, you will be asked if you want to commit your changes. Infrequently, this can elicit an error message about database corruption. Xsigo strongly recommends leaving the autocommit option enabled. During an upgrade, this is not a problem because a new database is created as part of the upgrade. However, if you receive this error message during a downgrade, it is possible that the database is in an error state and the Fabric Director will not come up after the downgrade. You can avoid this problem by not disabling the autocommit option. To work around this issue, always back up your configuration, following the instructions in the *Xsigo Software Upgrade Guide*. If you have your configuration saved to an XML file, you can restore it if necessary after the downgrade.
11768	There is a known limitation in the IB components that limits the IB counters to a maximum value of 4294967295. This value is not necessarily the actual counter value, but is displayed whenever counters meet or exceed the maximum value. The actual value can be higher than this counter indicates. If you need to clear the IB counters you can: Step 1 Issue **`show diagnostics ib-topo`** Step 2 Get the port name from the `portname` column in the resulting output (for example, `taffy:serverPort9`) Step 3 Issue the **`set diagnostics ib-clear-counters name`** `<portname>` command where `<portname>` is the string `<chassisname>`**`:ServerPort`**`<portnum>` (for example, `taffy:serverPort9`)
11809	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.
12065	When creating individual ACL rules with the same condition, the CLI processes those rules as one rule. By default, if no action is specified for an ACL, the CLI applies “any”. This can lead to situations where you might think you’re creating two separate rules, but the CLI processes the rules as one. For example, assume you create ACL rule 1 with source IP address set to ANY, and ACL Rule 2 with destination IP address set to ANY. These two rules are actually processed as one rule for ANY source and destination address. Also, by default, if you do not set a specific value for options in the ACL rules, the CLI defaults to “ANY”. This can add more confusion to the ACL rules and what end-users expect. Continuing with the example, with rule 1 set to a source IP address of “ANY” and no destination IP address specified, the CLI assumes “ANY” for the destination rule. In rule 2, with the destination address set to “ANY” and no source IP address set, the CLI defaults to ANY source IP address. The net result is that these two rules end up duplicating each other because they are both set to ANY source IP address and ANY destination address. Then, because these are two of the same rules, the CLI processes them as one rule. This consideration occurs only when two rules are the same, and use non-specific values (either ANY, or no value, which defaults to ANY). To workaround this problem, instead of leaving an option in the ACL rule blank or unspecified, enter a value. For example, instead of rule 1 being ANY source IP address and nothing specified for destination address, enter an IP address (for example, 0.0.0.0 to apply to all destination addresses) or an unused IP address to set the address to none in the network.

Table 2 (continued) Known Problems in Version 3.9.0 Specific to the Fabric Director

Number	Description
13662	The Xsigo implementation of SNMP does not have complete MIB entries for System Info and 10 GE interfaces.
13918	The Xsigo 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. 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 Director and the peer device to obtain comprehensive information about LACP properties.
14634	When the same user is configured as a local user (user configured on the Fabric Director) and a RADIUS user, a conflict occurs when the role is assigned during log in. For example, assume *paulw* is created as a local user with the operator role, and a *paulw* is also created as a RADIUS user with the network role. When the IMS search order is set to use internal authentication, and *paulw* logs in with the user and password for the network role, the operator role is enforced instead.
15372	In the Xsigo implementation of SNMP, an SNMP Agent problem prevents the vNIC and vHBA location information from being saved in the `xsigovnicremove` and `xsigovhbaremoved` 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.
16627	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.
18344	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 Director does not recognize the mismatch when the Windows host has rejected the MTU mismatch. The Fabric Director shows the vNIC state as `up/Failed` instead of `up/MTUMismatch`. For additional details about the state of the HA vNIC, see the Windows system event log. 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.

18577	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:
	
	`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)`
	
	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.
	
	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.
	
	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`**

Table 2 (continued) Known Problems in Version 3.9.0 Specific to the Fabric Director

Number	Description
18582	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. 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.
19195	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: • HP_08B0420001 • HP_08B0000001 • HP_0AD0000003 • HP_0160000009 • DEL08F0120009 • DEL08F0120002 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.
20044	When a large number of vHBAs are configured on a host, the vHBA code can sometimes miss or improperly handle RSCNs. As a result, one vHBA or target might not be recognized. For example, if 63 or more vHBAs are configured on a Fibre Channel module, sometimes an RSCN can be missed when a target comes back online. If you notice that a vHBA or target is not seen by the Fabric Director when it should be, it is likely you've hit this condition. You can rescan the Fibre Channel topology from the Fabric Director to discover the missing vHBA or target: **`set vhba <vhba-name> rescan`**
21086	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. 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 4.
21306	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 Director 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: • only when access to all LUNs is removed, which is an atypical procedure • only with HP targets

Table 2 (continued) Known Problems in Version 3.9.0 Specific to the Fabric Director

Number	Description
21323	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 PVIs are configured. Be aware that this condition can occur. You can work around this problem by following the procedure documented in To Upgrade Fabric Directors Connected to ESXi 5.0 and 5.1 Hosts with PVI vNICs, Set Server Profiles "Down" on page 6.
23143	Due to a problem in the `mst_flint` tool used in XgOS, the firmware and Option ROM on ConnectX-3 HCAs currently cannot be upgraded using the in-band upgrade feature. In this release, if you issue either of the following commands on a ConnectX-3 HCA, they do not complete successfully: • **`set physical-server`** `<name>` **`upgrade-hca`** `<LID>` **`optionrom`** • **`set physical-server`** `<name>` **`upgrade-hca`** `<LID>` **`firmware`** When these commands are issued, they fail and display the following error message: `Upgrade of HCA for (lid 7) failed: Failed to execute command (exit status 2) [xg_flint]` If you want to upgrade ConnectX-3 HCA firmware or Option ROM, you can download updated Mellanox Firmware Tools (MFT) for ConnectX-3 cards from the Mellanox website and upgrade the HCA locally. Additional assistance with locating and using MFT for ConnectX-3 can be found through Mellanox support.

Fixed Problems

Table 3 shows the fixes in this release. This release also contains fixes from previous releases. Such fixes are indicated in text.

Table 3 Fixed Problems in XgOS 3.9.0

Number	Description
21732	Due the way that Compellent array moved the targets from one Port ID to another during failover, numerous aborts occurred for the old Port ID, which caused the vHBA drivers to delay the failover for several minutes. This problem is fixed in XgOS 3.9.0.
22434	In a previous version of XgOS, tagged IGMP queries were not successfully passed to trunk vNICs. This problem is fixed in XgOS 3.8.3.
22363	During an upgrade, a difference between the MTU setting on the port and the vNIC consistently prevented the config from being imported. When the config wasn't imported, the upgrade could not successfully complete. This problem is fixed in XgOS 3.8.3.
22326	If you disabled an LACP LAG, the operational state of the LAG was listed as `down` even though one of the ports of the LAG were still `up/up`. In this case, the vNIC on the port was `up/up`, but the interface state was `down/down` and traffic was still passed on the vNIC even though you set the LAG down. This problem is fixed in XgOS 3.8.3.

Table 3 (continued) Fixed Problems in XgOS 3.9.0

Number	Description
22310	When one or more vHBAs bound to the EMC RecoverPoint application, a problem occurred during the procedure of logging the target out of the Fabric Director, and logging back in. While the target was logged out, the `lost-to-dead` timer was started, and soon after the target logged back in, the `lost-to-dead` timer expired. As a result of the timer expiring, the target was counted as "dead," and all associated paths were lost. This problem is fixed in XgOS 3.8.2.
21843	A problem interrupted the traffic on vNICs or vHBAs on Windows Server 2008 hosts when out of order sequence errors occurred. When this problem occurred, traffic interruptions of approximately 50 to 60 seconds occurred. This problem is fixed in XgOS 3.8.2.
21753	A problem in XgOS 3.7.0 and later prevented a hybrid LAG from being set when you used the -tag-native= false option with the **`set ethernet-port`** `<port num>` **`-mode=trunk -tag-native=false`** command. When you issued this command, the tagging conditions for the native VLAN on the LAG port were not enforced. This problem is fixed in XgOS 3.8.2.
21666	If a flush bit error was encountered, a problem caused improper handling of a Fibre Channel card reset. This problem is fixed in XgOS 3.8.2. In XgOS 3.8.2, and the Fibre Channel card does not reset.
21629	During CT pass through handling, a SIGSEGV error occurred that eventually caused a buffer overflow. As a result, the vHBA software crashed, and the all vHBAs on the affected module were brought down. This problem is fixed in XgOS 3.8.2.
19433	A problem sometimes prevented the Xsigo Fibre Channel modules from properly handling a `flush_bit` error. When the error occurred, all vHBAs on the I/O module changed to `up/down`, and hosts lost connectivity to targets in the SAN. Also, an error-level message similar to the following was posted to the Fabric Director's Syslog utility: `Flush bit not getting reset` This problem is fixed in XgOS 3.8.2.

Technical Support Contact Information

Xsigo Systems is owned by Oracle. Xsigo customers may contact support via the Xsigo website, telephone or e-mail. In order to expedite troubleshooting, all new support requests must be submitted via the Xsigo self-service portal at: http://support.xsigo.com. In addition to opening cases, the Xsigo Support Portal will allow you to update your support cases, download software, search for and view knowledge-base articles, and access technical documentation.

In order to access the customer support portal, you will need to have a Xsigo Support Portal login. Your account team will provide you with the necessary login information to access the support portal. If you need additional logins for your staff, please contact your account team for assistance.

For all Critical (P1) cases, please call the Xsigo support center at either of the following phone numbers:

- **1 866-974-4647** (toll free)
- **1 408-736-3013** (international).

Alternatively, you can email `supportP1@xsigo.com` and you will be responded to within 30 minutes.

Gathering Information for Xsigo Technical Support

If the Xsigo Fabric Director encounters problems, please gather the following information from the Fabric Director before contacting Xsigo Technical Support or filing a case through the support website:

- Type and number of servers connected (brand, model, number of CPUs, size and type of memory)
- Output from the **`get-log-files -all`** command (for Oracle's Xsigo Fabric Director experiencing problems), which will gather the **`show tech-support`** information plus all log files, and place this information into `xsigo-logs.<director-name>.tar.gz`
- Optionally, you can get just the output from the xsigo-support by issuing the **`show tech-suppor`**t command, then copy the output to a file.