



Upgrading the Graphics Card in an SGI® Graphics Cluster™ System

Contributors

Written by Tom Dye

Edited by Allison Gosbin

Illustrated by Dan Young

Production by Rhonda Kunsman

Engineering contributions by Eric Wang, Armando Serrato, Robert Sanders, Ken Jones, Mark Schwenden, Carolyn Curtis, and others.

COPYRIGHT

This document contains private information of Silicon Graphics, Inc. The contents of this document may not be disclosed to third parties, copied, or duplicated in any form, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

LIMITED AND RESTRICTED RIGHTS LEGEND

The electronic (software) version of this document was developed at private expense; if acquired under an agreement with the USA government or any contractor thereto, it is acquired as "commercial computer software" subject to the provisions of its applicable license agreement, as specified in (a) 48 CFR 12.212 of the FAR; or, if acquired for Department of Defense units, (b) 48 CFR 227-7202 of the DoD FAR Supplement; or sections succeeding thereto. Contractor/manufacturer is Silicon Graphics, Inc., 1600 Amphitheatre Pkwy 2E, Mountain View, CA 94043-1351.

TRADEMARKS

OpenGL Performer, SGI, SGI DataSync, SGI Graphics Cluster, and SGI ImageSync are trademarks of Silicon Graphics, Inc.

Linux is a registered trademark of Linus Torvalds.

NVIDIA and NVIDIA QUADRO are trademarks of NVidia Corporation.

Red Hat is a trademark of Red Hat, Inc.

Windows NT is a trademark of Microsoft Corporation.

X Windows System is a trademark of The Open Group.

Contents

Upgrade Summary	1
Checking Contents	2
Following Electrostatic Discharge Precautions.....	2
Preparing a Node for the Upgrade	2
Sliding a Node Out of the Rack	2
Removing a Node From a Rack	8
Removing a Node Cover.....	9
Replacing a Graphics Card	10
Inserting a Node in a Rack.....	11
Preparing the System for Operation.....	11
Installing the Graphics Drivers	13
Windows NT Systems	13
Linux SGI Graphics Cluster Systems	14
Sources of Graphics Driver Installation Failures.....	15
Recompiling the NVIDIA Kernel on Linux Systems	16
Installing the Driver in Command-line Mode -- Linux Systems	17
Returning NV16 Graphics Cards	17

Upgrading the Graphics Card in an SGI Graphics Cluster System

This document contains the following sections:

- “Upgrade Summary”
- “Checking Contents”
- “Following Electrostatic Discharge Precautions”
- “Preparing a Node for the Upgrade”
- “Replacing a Graphics Card”
- “Inserting a Node in a Rack”
- “Preparing the System for Operation”
- “Installing the Graphics Drivers”
- “Returning NV16 Graphics Cards”

Upgrade Summary

This upgrade replaces the graphics card and the associated driver. Specifically, the SGI VR9 (NV20) graphics card replaces the SGI VR7 (NV16) graphics card; and NVIDIA graphics driver 1.0-1551 replaces 1.0-1420.

For systems running the Linux operating system, you must install the new graphics drivers **before** you install the new graphics cards.

For systems running the Windows NT operating system, install the new graphics cards before you install the new graphics drivers.

Pack the NV16 boards in the NV20 shipping materials. Return the boards to SGI.

Checking Contents

Ensure that none of the upgrade materials are damaged or missing. The upgrade includes the following materials:

- SGI VR9 graphics cards
- Software bundle
 - Graphics driver CD
 - Return address label, P/N 024-1707-001

Note: Do not damage or discard the shipping containers, packing materials, or antistatic bags. You will use these items to return the old VR7 graphics cards to SGI.

Following Electrostatic Discharge Precautions

Electrostatic discharge (ESD) can irreparably damage electronic equipment. Always follow preventative measures when handling system components:

- Remove a component from its antistatic bag only when you are ready to install that component.
- If you must handle a component before installation, do not place it on a surface that produces ESD (carpeting, for example) or near devices that create static electricity.
- Wear a grounding wrist strap and connect it to a grounded connection on your system when you install or remove a component.

Preparing a Node for the Upgrade

This section explains:

- “Sliding a Node Out of the Rack”
- “Removing a Node From a Rack”
- “Removing a Node Cover”

Note: For Linux systems, you must install the new NVIDIA graphics driver (1.0 1551) for every channel node, including the master node, before you replace the old graphics cards with the new cards.

Sliding a Node Out of the Rack

To replace the graphics card, it is usually sufficient to slide the node out from the rack on its extensible slide rails, where it locks into place. To do so, follow these steps:



Warning: When adding or servicing equipment in the rack, do not pull out more than two nodes at a time.

1. Make sure all users are off the affected system(s) and that data is backed up, if necessary. Shut down the system with the appropriate administrative tools.
2. If desired, power off the rack PDU with its circuit breaker.
3. Unlock the node bezel, pull the bezel out and away from the node, and swing the bezel down, as shown in Figure 1.

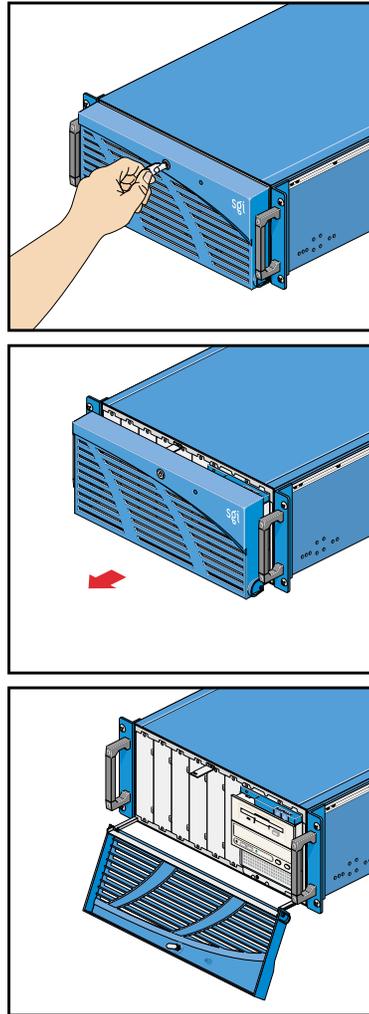


Figure 1 Opening a Node Bezel

4. Power off the node; refer to Figure 2 for the button location.

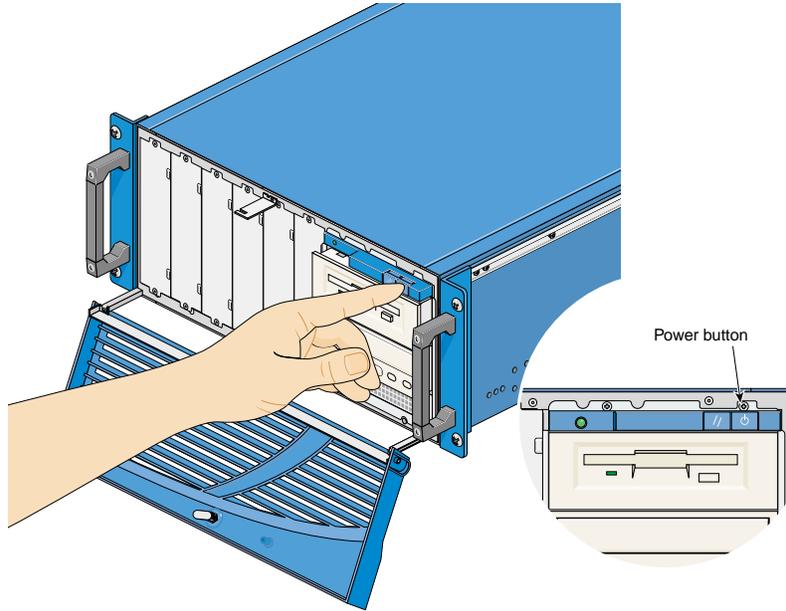


Figure 2 Powering Off the Node

5. At the back of the rack, detach the power cable from the node. Refer to Figure 3.

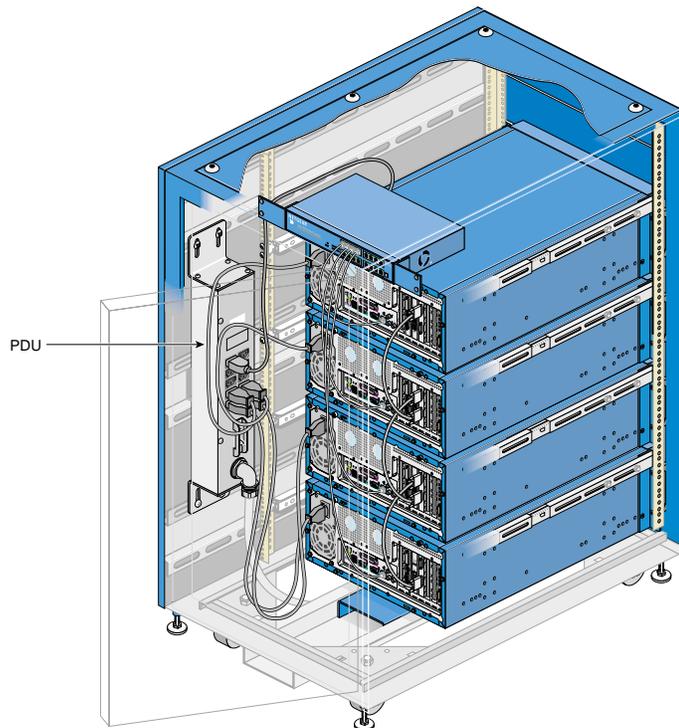


Figure 3 Node Power Cables

- For Series 12 systems, disconnect the daisy-chain cable from the node's external SGI ImageSync connector; refer to Figure 4. A Series 12 master node has one SGI ImageSync cable; channel nodes are daisy-chained to the nodes above and below.

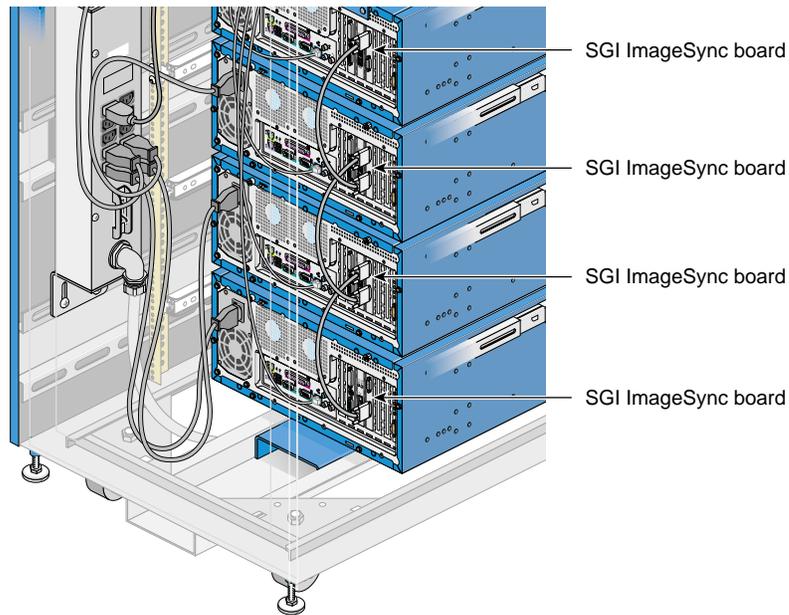


Figure 4 Disconnecting the External SGI ImageSync Cable, SGI Graphics Cluster Series 12 System

- Disconnect other cables, such as to peripherals, Ethernet, and external options.
- Pull out the rack's anti-tip shelf; refer to Figure 5.

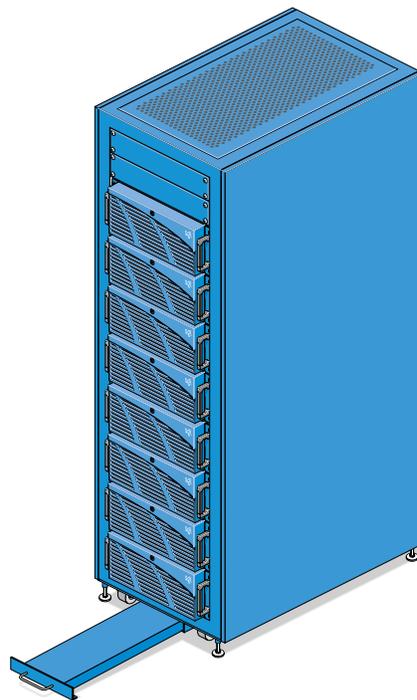


Figure 5 Pulling Out the Anti-Tip Shelf



Warning: Be sure to pull out the anti-tip shelf before you pull out a node. Failure to do so can result in the rack tipping over. Never pull out more than two nodes at a time.

9. Remove the four screws that secure the node to the rack chassis. These screws are above and below each handle; refer to Figure 6. Retain the screws.

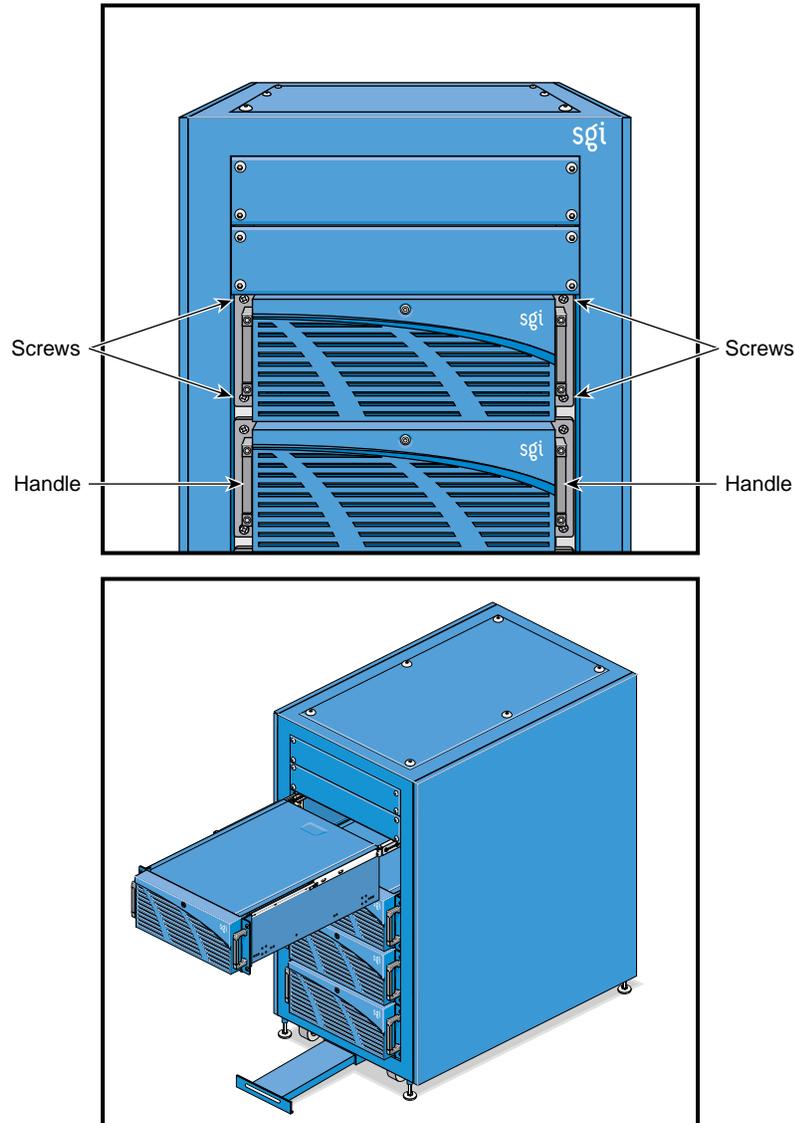


Figure 6 Freeing a Node from a Rack

- Using the node's handles, slide the node out of the rack; refer to Figure 7.

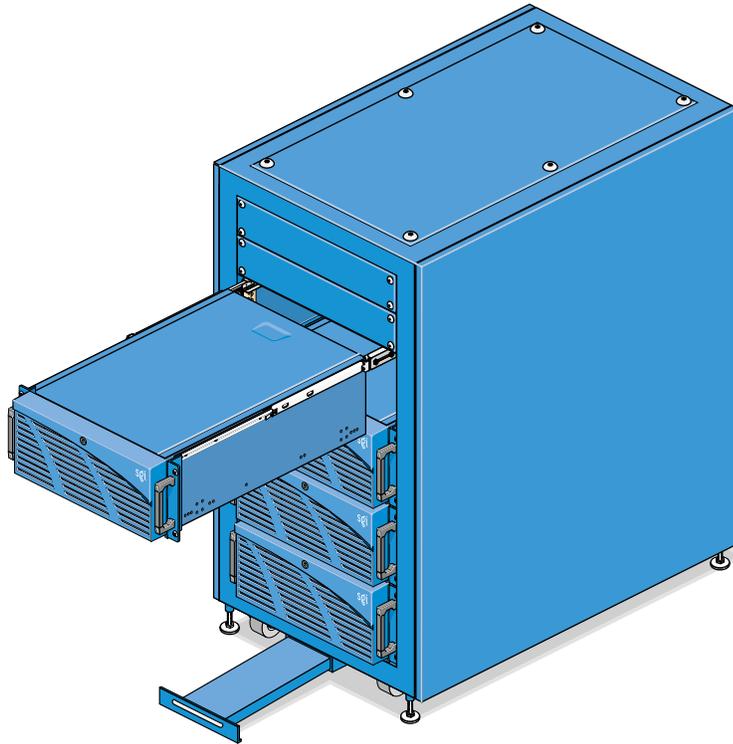


Figure 7 Sliding a Node Out of a Rack

When the node's slide rails are fully extended, they lock in place.

Note: To push the node back into the system, push in the lever on each slide rail; refer to Figure 8 on page 8.

- If you need to completely remove a node from a rack to gain access to the internal components, complete the following procedure, "Removing a Node From a Rack".

Removing a Node From a Rack

To completely remove a node from a rack and its slide rails, follow these steps:

Caution: Have a second person assist you when you remove and add nodes. Most nodes weigh between 32 and 34 lbs (14.4 and 15.5. kg).

1. Prepare a work surface to accommodate the node.
2. Pull the node out of the rack following all the steps in “Removing a Node From a Rack”.
3. Using your thumbs, push in the locking lever on each rail and slide the node forward—off the rails. Refer to Figure 8. Place the node on the work surface.

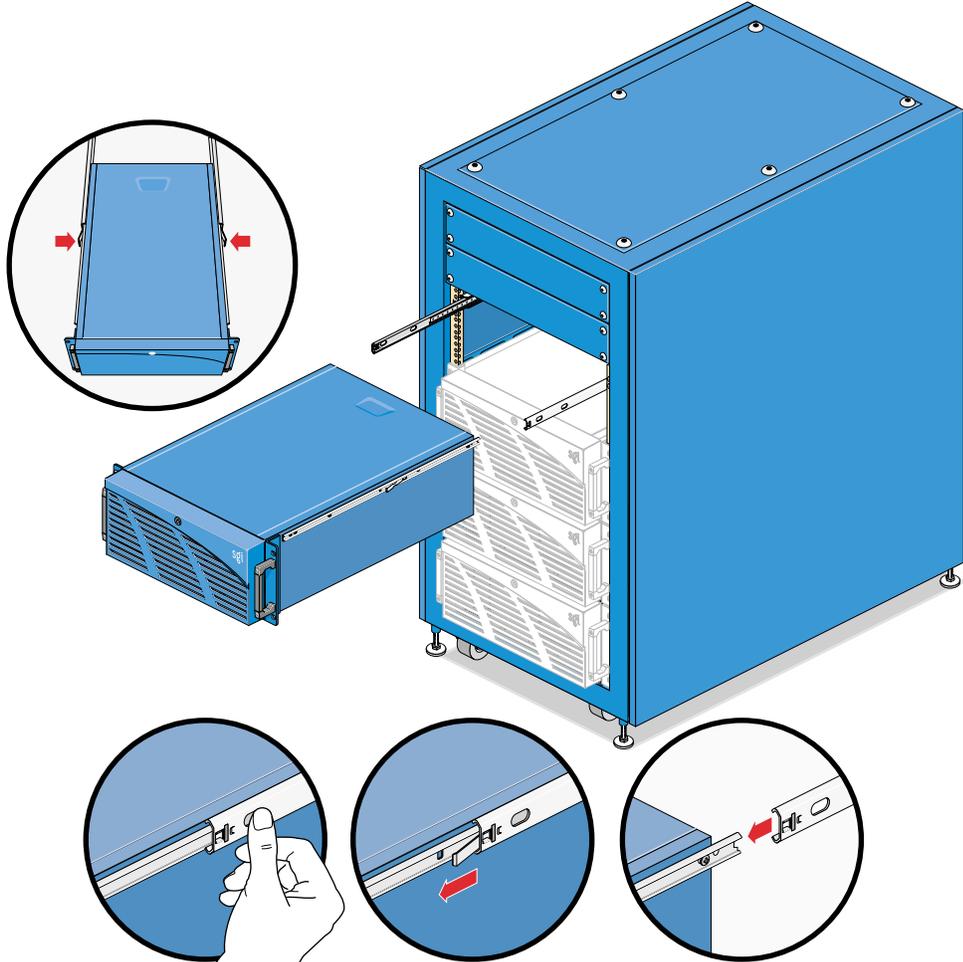


Figure 8 Slide-rail Locking Levers

Removing a Node Cover

Follow these steps to remove the node's cover:

1. Loosen the two captive screws at the back of the node; refer to Figure 9.

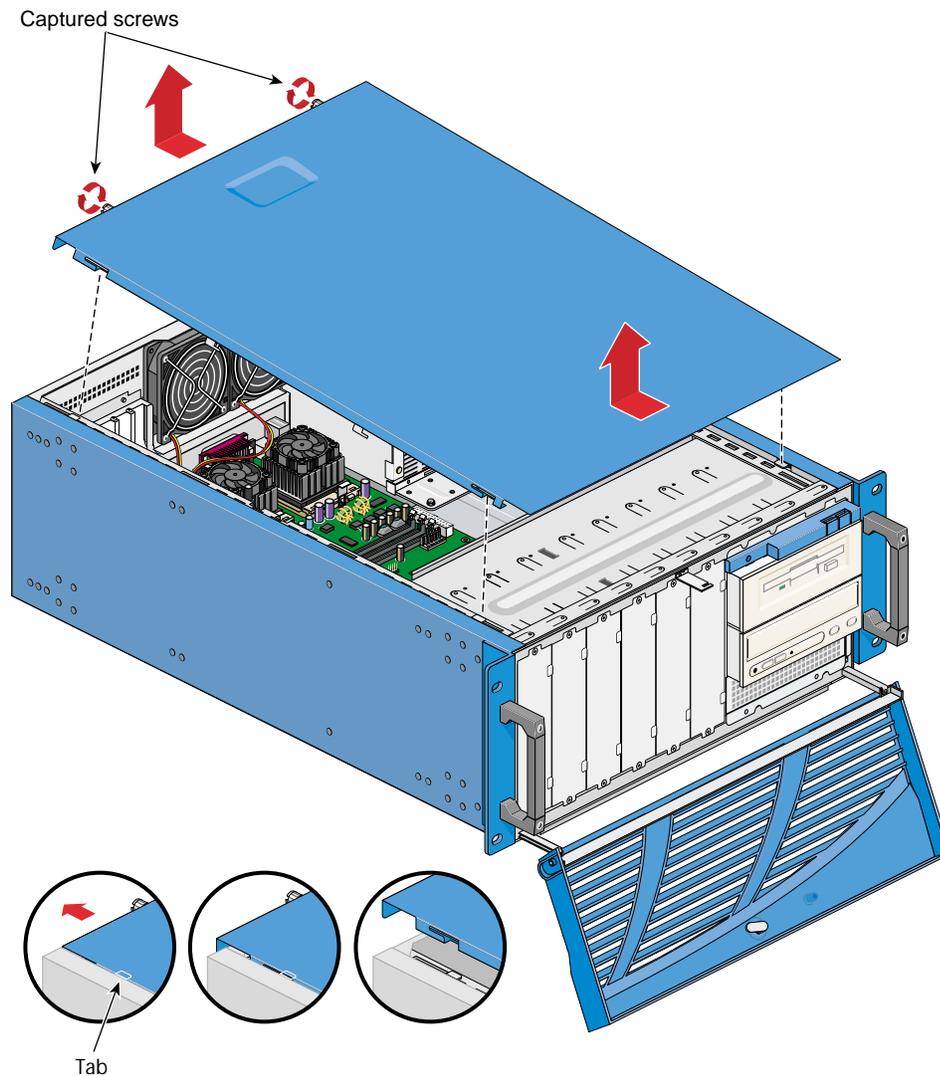


Figure 9 Removing a Node Cover

2. Pull the node cover a centimeter or two (about half an inch) toward the back of the node; refer to Figure 9.
3. Lifting the back of the node cover, pull it up and away from the node; refer to Figure 9.

Note: To replace the node cover, reverse the steps in this subsection.

Replacing a Graphics Card

Note: For Linux systems, you must install the new NVIDIA graphics driver (1.0 1551) for every channel node, including the master node, **before** you replace the old graphics cards with the new cards.

To replace the graphics card in each node, follow these steps:

1. Access the node's internal components. Refer to "Preparing a Node for the Upgrade".
2. For Series 12 nodes, detach the twisted-pair cable from its connector on the graphics card; refer to Figure 10.

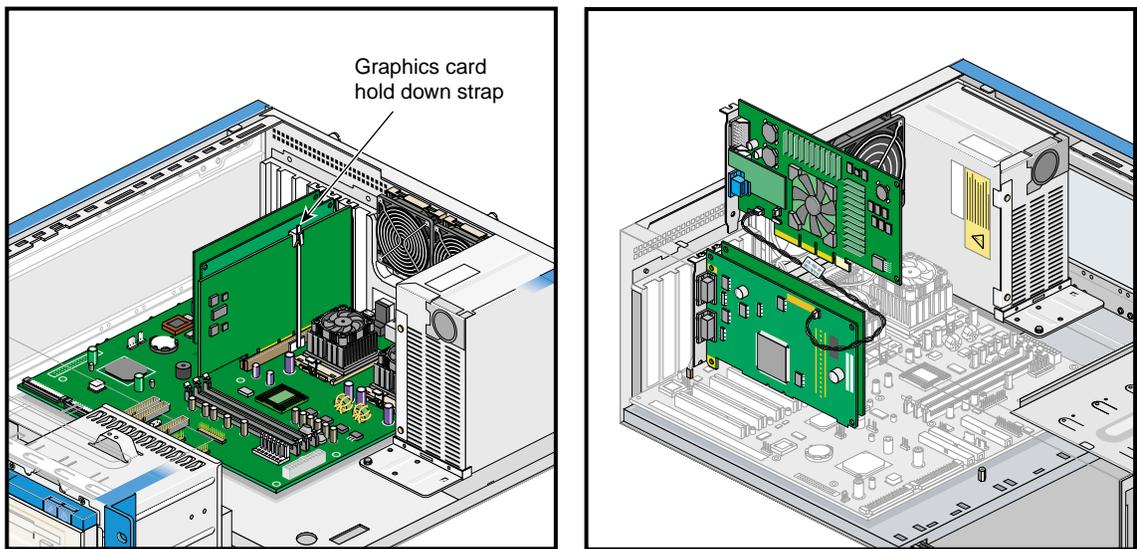


Figure 10 Graphics Cluster Series 12 Graphics Card Cabling

3. Note the position of the AGP retention mechanism; then release the mechanism.
4. For all nodes, carefully pull out the card.
5. Remove the replacement graphics card from the shipping container. Save the shipping materials for return to SGI.
6. If applicable, put the tie wrap from the removed card on the replacement graphics card.
7. For Series 12 nodes, connect the free end of the twisted-pair cable to the connector on the replacement graphics card before you insert the card in the AGP slot.
8. For all nodes, firmly and carefully insert the replacement graphics card in the AGP slot.

Expansion cards in the SGI Graphics Cluster reside in specific slots, as outlined in Table 1.

Table 1 Expansion Slots and Cards in the SGI Graphics Cluster

Slot	Card
AGP	Graphics card
PCI slot 1 (closest to the AGP slot)	Optional gigabit Ethernet card
PCI slot 2	SGI ImageSync card; SGI Graphics Cluster Series 12 only
PCI slot 3	Network interface card (secondary Ethernet), master node only; included with SGI Graphics Cluster Series 12
PCI slot 4	Empty; available for customer option
PCI slot 5 (closest to the chassis wall)	Commercial audio card: included with master node of SGI Graphics Cluster Series 12

Inserting a Node in a Rack

1. Replace the node cover and secure it with the two captive screws. (Refer again to “Removing a Node Cover”.)
2. If necessary, install the node in the slide rails—with an assistant’s help. (Refer again to “Removing a Node From a Rack”.)
3. Refer again to “Removing a Node From a Rack” as you complete the following steps:
 - a. Slide the node into the chassis and secure it with screws.
 - b. Connect all cables and the power cord.

Preparing the System for Operation

After you replace all graphics cards and install all node boards in the rack(s), restore the system to operation by using the following procedure:

1. Push in each rack’s anti-tip shelf.
2. If applicable, power on each rack PDU with its circuit breaker.
3. Power on the monitor connected to the master node, at the bottom of the rack. Power on other peripherals as desired.

4. Turn on the master node. The green power indicator LED on the front panel illuminates. Figure 11 shows the front controls.

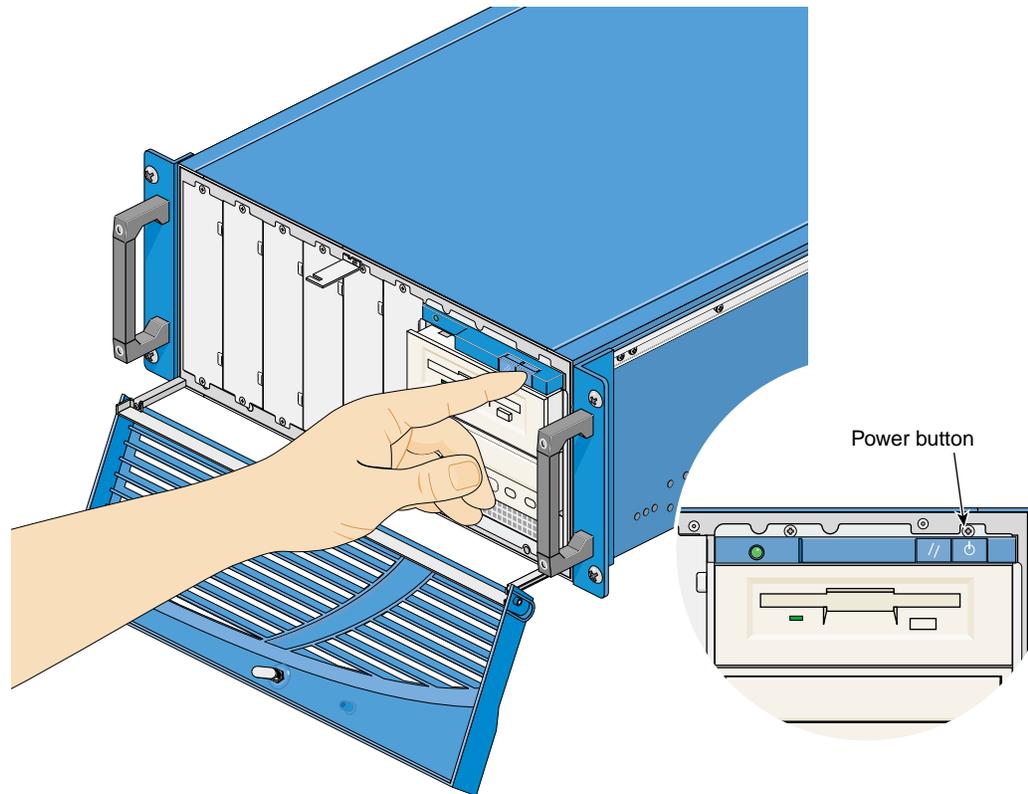


Figure 11 Node Front Panel Controls

The Num Lock, Scroll Lock, and Caps Lock LEDs on the master node's keyboard flash briefly. The system starts and displays a boot message on the monitor connected to the master node.

Note: If the power indicator LED illuminates but the SGI splash screen and boot information do not appear on the master node monitor, press the reset button. If the system does not turn on or boot after you press the power button, see the next section for the possible causes of the boot failure.

5. Power on each channel node with the button at the far right above the floppy disk drive.
6. Close each nodes bezel and lock it, if desired.
7. Close the rack doors.
8. Return the keys to the appropriate location.
9. Install the graphics drivers as described in the following procedure, "Installing the Graphics Drivers".

Installing the Graphics Drivers

This section describes how to install the new graphics driver in SGI Graphics Cluster Series 11 and Series 12 systems running the Windows NT 4.0 or Red Hat Linux 7.1 operating system.

Windows NT Systems

1. Install the new VR9 (NV20) graphics cards in the system nodes. Refer to “Replacing a Graphics Card” for details.

2. Boot the system and log into Windows NT 4 as Administrator.

The system comes up in low-resolution mode (VGA) at this time.

3. Insert the graphics driver CD.

A Web browser appears, displaying links to the graphics drivers and to a README file that contains this procedure.

4. In the Web browser, select the **Drivers** link; then click on the **Windows Workstation NT 4.0** link.

A list of NV20 (VR9) windows drivers appears.

5. Locate the `setup.exe` file in the list and double click it.

The driver installation process starts.

6. Choose **NVIDIA QUADRO DCC (SGI)** when presented with a menu selection.

Note: If you cannot install the driver via the browser, you can install it via the DOS command line or by accessing the `setup.exe` file on the CD (`D:\Drivers\WinNT4\setup.exe`).

7. When the driver installation completes, restart the system.

8. Set the display resolution as required.

To access the list of available screen resolutions and refresh rates, perform the following steps:

- a. From the operating system Start menu, go to **Settings/Control Panel** and double-click **Display**.
- b. From the Display Properties dialog, click the **Settings** tab.
- c. Click “List All Modes...” to access the table of supported screen resolutions and refresh rates.

The upgrade is now complete.

Linux SGI Graphics Cluster Systems

Note: You must install the new NVIDIA graphics driver (1.0 1551) for every channel node, including the master node, **before** you replace the graphics cards.

Note: After you install the 1.0-1551 graphics driver, you must install the new OpenGL Performer eoe, dev, and SGI DataSync modules.

If you do not have a valid OpenGL Performer licence, install the OpenGL Performer demo license, which is valid until July 1, 2002. After that time, download the `performer_demo_license rpm` from the following URL:

`http://www.sgi.com/software/performer/`

If you have a permanent SGI DataSync license, create a backup copy of the `license.dat` file before you install the SGI DataSync module. After you install the module, replace the new `license.dat` file with the backup copy.

1. Boot the machine.
2. Log in as `root`.
3. Insert the graphics driver CD.
4. Open a Linux shell and mount the graphics driver CD:

```
mount /dev/cdrom -t iso9660 /mnt/cdrom
```
5. Access the directory that contains the graphics driver for the SGI kernel:

```
cd /mnt/cdrom/SGI/RPMS
```
6. Install the NVIDIA kernel module and GLX module:
 - For Series 11 systems:

```
rpm -Uvh NVIDIA_kernel-1.0-1551.i386up.rpm NVIDIA_GLX-1.0-1551.i386.rpm
```
 - For Series 12 systems:

```
rpm -Uvh NVIDIA_kernel-1.0-1551.i386smp.rpm NVIDIA_GLX-1.0-1551.i386.rpm
```
7. Ensure that the `rpm` installation completed successfully.

Note: If you did not install the graphics driver successfully, refer to “Sources of Graphics Driver Installation Failures”.
8. Install the new OpenGL Performer `eoe` and `dev rpm` modules:
 - a.

```
cd /mnt/cdrom/SGI/RPMS
```
 - b.

```
rpm -Uvh performer_eoe-2.4.2-4.i386.rpm performer_dev-2.4.2-4.i386.rpm
```
9. If you do not have a valid OpenGL Performer license, install the OpenGL Performer demo license:

```
rpm -Uvh performer_demo_license-2.4.2-4.i386.rpm
```

10. If the machine has a permanent SGI DataSync license, create a backup copy of the `license.dat` file:
 - a. View the `/usr/share/DataSync/license.dat` file:


```
less /usr/share/DataSync/license.dat
```
 - b. Check the second line of the `license.dat` file:
 - `HOSTID=DEMO` indicates that you have a demo license; proceed to Step 11.
 - `HOSTID=<your machine's hostid>` indicates that you have a permanent license, which you must keep; complete Step c.
 - c. If you have a permanent license, copy the `license.dat` file to `license.dat.perm`:


```
cp license.dat license.dat.perm
```
11. Install the new SGI DataSync module:
 - a.

```
cd /mnt/cdrom/SGI/RPMS
```
 - b.

```
rpm -Uvh DataSync-1.0.1.i386.rpm
```
12. If the machine has a permanent SGI DataSync license, restore the `license.dat` file:


```
cp license.dat.perm license.dat
```
13. Reboot the machine.
14. Install the new VR9 (NV20) graphics cards in the system nodes. Refer to “Replacing a Graphics Card” for details.

The upgrade is now complete.

Sources of Graphics Driver Installation Failures

If you could not install the graphics driver successfully, one of the following problems may have caused the failure:

- The kernel source and kernel headers may not have been installed when the system was rebuilt.

To correct this problem, reinstall the dev environment from the SGI Graphics Cluster Resource CD 1.0 or VWE-3.1 (GC11) CD.

The Development portion, including the C compiler and associated modules, of the Red Hat 7.1 distribution is required. (The system ships with the Development portion installed.)
- You may need to recompile the NVIDIA kernel.

To recompile the NVIDIA kernel, refer to the following procedure, “Recompiling the NVIDIA Kernel on Linux Systems.”

Recompiling the NVIDIA Kernel on Linux Systems

If you could not install the graphics driver successfully, then the Linux kernel may not be an SGI specific kernel or you may need to recompile the NVIDIA kernel or modules.

1. Boot the machine.
2. Log in as `root`.
3. Insert the graphics driver CD.
4. Open a Linux shell and mount the graphics driver CD:

```
mount /dev/cdrom -t iso9660 /mnt/cdrom
```

5. Access the directory that contains the NVIDIA 1.0-1551 graphics driver source code:

```
cd /mnt/cdrom/SGI/SRPMS
```

6. Recompile the graphics kernel:

```
rpm --rebuild NVIDIA_kernel-1.0-1551.src.rpm
```

7. Install the rebuilt graphics kernel module:

```
rpm -Uvh /usr/src/redhat/RPMS/i386/NVIDIA_kernel-1.0-1551.rpm
```

If the `rpm` command does not complete successfully and a message notes that the NVIDIA kernel is already installed, execute the following commands to install the new kernel:

```
rpm -e --nodeps NVIDIA_kernel
```

```
rpm -ivh /usr/src/redhat/RPMS/i386/NVIDIA_kernel-1.0-1551.rpm
```

Note: On the original NVIDIA kernel that shipped on the SGI Graphics Cluster Linux Resource CD 1.0 and VWE-3.1 CD, the `rpm -e` command fails due to a bug in the NVIDIA script. You must upgrade via the `rpm -Uvh` command to enable the `rpm -e` command to function.

8. If you did not install it earlier, install the NVIDIA GLX module:

```
a. cd /mnt/cdrom/SGI/RPMS
```

```
b. rpm -Uvh NVIDIA_GLX-1.0-1551.i386.rpm
```

9. Install the new OpenGL Performer `eo` and `dev` rpm modules:

```
a. cd /mnt/cdrom/SGI/RPMS
```

```
b. rpm -Uvh performer_eoe-2.4.2-4.i386.rpm  
performer_dev-2.4.2-4.i386.rpm
```

10. If you do not have a valid OpenGL Performer license, install the OpenGL Performer demo license:

```
rpm -Uvh performer_demo_license-2.4.2-4.i386.rpm
```

11. If the machine has a permanent SGI DataSync license, create a backup copy of the `license.dat` file:

```
a. View the /usr/share/DataSync/license.dat file:
```

```
less /usr/share/DataSync/license.dat
```

```
b. Check the second line of the license.dat file:
```

- HOSTID=DEMO indicates that you have a demo license; proceed to Step 12.
- HOSTID=<your machine's hostid> indicates that you have a permanent license, which you must keep; complete Step c.

- c. If you have a permanent license, copy the `license.dat` file to `license.dat.perm`:

```
cp license.dat license.dat.perm
```

12. Install the new SGI DataSync module:

- a. `cd /mnt/cdrom/SGI/RPMS`
- b. `rpm -Uvh DataSync-1.0.1.i386.rpm`

13. If the machine has a permanent SGI DataSync license, restore the `license.dat` file:

```
cp license.dat.perm license.dat
```

14. Reboot the machine. You should be able to boot the channel into the X Window System environment.

The upgrade is now complete.

Installing the Driver in Command-line Mode -- Linux Systems

If the X server does not run and you cannot enter the X Window System environment (the screen blinks), use the following process to accomplish the driver upgrade.

1. Enter the following keyboard command:

```
Ctrl+Alt+F2
```

2. Login as `root`.
3. Put the system into run level 3:

```
init 3
```

4. Press `Enter` to access command-line mode.
5. Upgrade the graphics driver according to the instructions in "Linux SGI Graphics Cluster Systems".

Returning NV16 Graphics Cards

Package the old NV16 graphics cards in the NV20 shipping containers. Place the return address label (P/N 024-1707-001) on the shipping container(s). Return the old cards to SGI.