NVIDIA ACCELERATED LINUX DRIVER SET
Release 20 Notes

Driver Version 1.0-2313

Table of Contents

List of Tables

1

INTRODUCTION

This chapter contains the following major topics:

Note: A text version of this document is available in:

```
usr/share/doc/NVIDIA_GLS-1.0/README
```

About Release Notes

These Release Notes provides information about the current Release 20 NVIDIA® Accelerated Linux Driver Set. NVIDIA provides these notes to enable add-in-card (AIC) producers and original equipment manufacturers (OEMs) to monitor performance improvements and software problem (bug) resolutions in each documented version of the driver.

This book explains how to install, configure, and use the NVIDIA Accelerated Linux Driver Set. It also describes current and historic software problem resolutions and software enhancements *and* contains Troubleshooting, Frequently Asked Questions, and other contact and support information.

About the NVIDIA Accelerated Linux Driver Set

The NVIDIA Accelerated Linux Driver Set brings both accelerated 2D functionality and high performance OpenGL support to Linux x86 with the use of NVIDIA products, which are **graphics processing units (GPUs)**. The NVIDIA nForce product is an **integrated graphics processor (IGP)**.

NVIDIA has a **Unified Driver Architecture (UDA)** model, which means that one driver set can be used with all supported NVIDIA hardware. This driver set provides optimized hardware acceleration of OpenGL applications through a direct-rendering X Server.

Supported NVIDIA products are listed in Table 1.2 and supported features are discussed in "Supported NVIDIA Features" on page 5.

Minimum Operating System Requirements

This release includes drivers for the Linux operating systems listed in Table 1.1.

Table 1.1 Minimum Operating System Requirements

Module	Version	Determining the Version
Linux Kernel	2.2.12	`# cat /proc/version`
XFree86	4.0.1	`# XFree86 -version`
Kernel modutils	2.1.121	`# insmod -V`
If you need to build the NVdriver kernel module, use the following components:		
`binutils`	2.9.5	`# size --version`
`GNU make`	3.77	`# make --version`
`gcc`	2.7.2.3	`# gcc --version`
If you build the NVdriver kernel module from source RPMs (RedHat Package Manager utility), use the following component:		
`spec-helper` RPM		`# rpm -qi spec-helper`

Notes and Tips on Operating Systems

• XFree86 can be retrieved from `www.xfree86.org`. Software packages may also be available through your Linux distributor.

• All official stable kernel releases from version 2.2.12 and higher are supported.

  Prerelease versions such as "2.4.3-pre2" *and* "development" kernels such as 2.3.*x* or 2.5.*x* are *not supported*.

The Linux kernel sources can be retrieved from `www.kernel.org` or one of its mirrors.

- `binutils` and `gcc` are required *only* if you install the `NVIDIA_kernel` package by **SRPM (Source-RPM)** or tar file and can be retrieved from `www.gnu.org` or one of its mirrors.

  Note: binutils and gcc are *not* required by binary RPM installations.

- If you are using XFree86, but do not have a file `/var/log/XFree86.0.log`, then you probably have a version 3.*x* of XFree86 and *must* upgrade.

- If you are setting up `XFree86 4.x` for the first time, it is often easier to begin with one of the open source drivers that ships with XFree86 (either "nv", "vga", or "vesa"). Once XFree86 is operating properly with the open source driver, it is easier to switch to the NVIDIA driver.

- NVIDIA products may not work with older versions of the "nv" driver that shipped with XFree86. For example, the "nv" driver that shipped with XFree86 (version 4.0.1) did not support graphics cards based on the NVIDIA GeForce2 family and the Quadro2 MXR products. However, this was fixed in XFree86 (version 4.0.2).

- GeForce2 Go (product for laptop computers) is supported by the NVIDIA Accelerated Linux Driver Set but *not* by the open source "nv" driver for XFree86.

Supported NVIDIA Products

Table 1.2 lists the NVIDIA products (chips) supported by the current version of the NVIDIA Accelerated Linux Driver Set.

Table 1.2   Supported NVIDIA Products

NVIDIA Desktop Product	NVIDIA Workstation Product	Device PCI ID
nForce		
	Quadro DCC™	0x0203
GeForce3™ Ti 500		0x0202
GeForce3™ Ti 200		0x0201
GeForce3™		0x0200
	Quadro2 Pro™	0x0153
GeForce2™ Ultra™		0x0152
GeForce2 Ti		0x0151
GeForce2 GTS™		0x0150
GeForce2 Pro™		0x0150
	Quadro2 MXR™	0x0113

Table 1.2   Supported NVIDIA Products (continued)

NVIDIA Desktop Product	NVIDIA Workstation Product	Device PCI ID
	Quadro2 Go™	0x0113
	Quadro2 EX™	0x0113
GeForce2 Go™		0x0112
GeForce2 MX 100		0x0111
GeForce2 MX 200		0x0111
GeForce2 MX 400		0x0110
GeForce2 MX™		0x0110
	Quadro™	0x0103
GeForce DDR™		0x0101
GeForce 256™		0x0100
RIVA™ TNT2™ Family		
RIVA TNT2		0x0028
RIVA TNT2 (Ultra)		0x0029
RIVA TNT2 (Vanta)		0x002C
RIVA TNT2 (M64)		0x002D
RIVA TNT2 (Integrated)		0x00A0
RIVA TNT™		0x0020

Notes and Tips on Supported NVIDIA Products

- GeForce2 Go (product for laptop computers) is supported by the NVIDIA Accelerated Linux Driver Set but *not* by the open source "nv" driver for XFree86.

- RIVA 128/128ZX products are supported by the open source "nv" driver for XFree86 but not by the NVIDIA Accelerated Linux Driver Set.

- Some Release 20 features support only certain NVIDIA products, which is indicated, where applicable. (*See* Supported NVIDIA Features.)

- If you want to check the VGA device IDs for comparison with the values shown in Table 1.2, use one of these methods:

  - "cat /proc/pci" *or*

  - "lspci -n" to dump hex information. With this method, you need to retrieve the line with the NVIDIA vendor ID of "Class 0300: 10de", as in the following example:

    ```
    0x10de: 01:00.0 Class 0300: 10de:0028 (rev 15)
    ```

    The non-NVIDIA information presented in the above line will vary depending of type of card you are using.

Supported NVIDIA Features

Also see relevant sections in Chapter 2: "NVIDIA Linux Driver History" on page 10 for recent and historic lists of software issue resolutions and enhancements.

Note: NVIDIA Linux Release 20 driver features are *not* supported by the Release 10 driver. However, Release 10 features *are* supported by the Release 20 driver.

Release 20

Starting with Release 20 (version 1.0-2313), the NVIDIA Accelerated Linux Driver Set supports the following features:

• Unified Back Buffer (UBB)

• Page Flipping

• Window Flipping

• OpenGL 1.3 (the current ARB-approved version of OpenGL)

• `pbuffer` and `fbconfig` extensions provide support for accelerated off-screen rendering.

Note: The first three features can provide performance gains in certain situations.

Release 10

GeForce 3

Release 10 is the first driver release to fully support the GeForce3 nfiniteFX™ engine in OpenGL using Texture and Vertex programs. Developers can now expose the full functionality of GeForce3 under Linux.

GeForce2 Go and Quadro2 Go

The Release 10 driver also added support for mobile (laptop) platforms with the GeForce2 Go and Quadro2 Go products, hot key switching, and improved driver stability on mobile platforms.

TV Output Feature

Graphics cards based on the NVIDIA chipset with a TV-Out (S-Video) connector can be used to send the display to a television as another display device, such as a CRT (monitor) or a digital flat panel (DFP) display. The TV

can be used by itself, or (on appropriate graphics cards) in conjunction with another display device in a TwinView configuration.

If a TV is the only display device connected to your graphics card, it will be used as the primary display when you start-up your system; that is, the display will come up on the TV just as if it was a CRT.

Release 6

TwinView

The TwinView feature is only supported on NVIDIA products that support dual-display functionality, i.e., GeForce2 MX, GeForce2 Go, and Quadro2 MXR.

TwinView is a mode of operation where two display devices can display the contents of a single X screen in any arbitrary configuration. TwinView supports a variety of display options, such as digital flat panels, red-green-blue (RGB) monitors, TVs, and analog flat panels. This method of using multiple monitors has several distinct advantages over other techniques (such as Xinerama), as outlined here:

• A single X screen is used. The NVIDIA driver conceals all information about multiple display devices from the X Server, which only acknowledges one screen.

• Both display devices share one frame buffer. Thus, all the functionality present on a single display (e.g. accelerated OpenGL) is available in TwinView.

• No additional overhead is needed to emulate having a single desktop.

Known Product Limitations

Software Issues

TwinView and Hot Key Switching (Toshiba Satellite 2800 Series Laptops)

TwinView and hot key switching are currently not supported on Toshiba Satellite 2800 series laptops. (*See* "Known Laptop Issues" on page 47.)

X Server and Changing AGP Drivers

Under Linux32, after starting X twice with different AGP support, the X Server may crash causing the system to become unusable. To work around this problem, if you change your AGP driver, be sure to restart the system before restarting X.

OpenGL + Xinerama

Currently, OpenGL is not functional with Xinerama.

OpenGL and dlopen()

There are some issues with the older version of the glibc dynamic loader (e.g., the version that shipped with RedHat 7.2) and applications such as Quake3 and Radiant that use the dlopen() function.

For reference, see the explanation for "If your OpenGL applications crash and print out the following warning:" on page 52 in Chapter 5.

DPMS and TwinView

DPMS (Display Power Management System) modes "suspend" and "standby" do not work correctly on a second CRT when using TwinView. The screen becomes blank instead of the monitor being set to the requested DPMS state.

DPMS and Flat Panel

The DPMS modes "suspend" and "standby" do not work correctly on a flat panel display. The screen becomes blank instead of the flat panel being set to the requested DPMS state.

Multicard, Multimonitor

In certain cases under a multicard/multimonitor configuration, the secondary card may not be initialized correctly by the NVdriver kernel module. You can work around this by enabling the "UseInt10Module" option in the XFree86 configuration file to soft-boot all secondary cards.

For further information, *see* "Option "UseInt10Module" "boolean"" on page 29.

Laptops

If you are using a laptop, refer to "Known Laptop Issues" on page 47.

Hardware Issues

This section describes problems that will not be fixed. Usually, the source of the problem is beyond the control of NVIDIA.

- "Gigabyte GA-6BX Motherboard" on page 8
- "VIA KX133 and 694X Chipsets With AGP 2X" on page 8
- "Irongate Chipsets With AGP 1X" on page 8
- "ALi Chipsets: ALi1541 and ALi1647" on page 9

Gigabyte GA-6BX Motherboard

This motherboard uses a LinFinity regulator on the 3.3-V rail that is rated to only 5 A — less than the AGP specification, which requires 6 A. When diagnostics or applications are running, the temperature of the regulator rises, causing the voltage to the NVIDIA chip to drop as low as 2.2 V. Under these circumstances, the regulator cannot supply the current on the 3.3-V rail that the NVIDIA chip requires.

This problem does not occur when the graphics card has a switching regulator or when an external power supply is connected to the 3.3-V rail.

VIA KX133 and 694X Chipsets With AGP 2X

On Athlon motherboards with the VIA KX133 or 694X chipset, such the ASUS K7V motherboard, NVIDIA drivers default to AGP 2X mode to work around insufficient drive strength on one of the signals. You can force AGP 4x by enabling `NVreg_EnableVia4x` in `os-registry.c` and rebuilding the NVIDIA kernel module.

Note: Forcing AGP *on* may cause the driver to become unstable on these chipsets.

Irongate Chipsets With AGP 1X

AGP 1X transfers are used on Athlon motherboards with the Irongate chipset to work around a problem with the signal integrity of the chipset.

ALi Chipsets: ALi1541 and ALi1647

On Acer Laboratories Inc.ALi1541 and ALi1647 chipsets, NVIDIA drivers disable AGP to work around timing and signal integrity issues. You can force AGP to be enabled on these chipsets by enabling `NVreg_EnableALiAGP` in `os-registry.c` and rebuilding the NVIDIA kernel module.

Note: Forcing AGP *on* may cause the driver to become unstable on these chipsets.

2

NVIDIA LINUX DRIVER HISTORY

This chapter contains the following major topics:

- "NVIDIA Linux Driver Versions" on page 10
- "Overview of Resolved Issues and Enhancements" on page 10

NVIDIA Linux Driver Versions

Release 20 is the latest release of the NVIDIA Accelerated Linux Driver Set. Table 2.1 contains a summary of driver releases and the versions associated with them. Some versions listed may not have been released outside of NVIDIA.

Table 2.1   NVIDIA Linux Driver Versions

Driver	Versions	Comments
Release 20	1.0-2313 --	Releases ongoing
Release 10	1.0-1101 through 1.0-1541	
Release 6	0.9-1 through 0. 9-769	

Overview of Resolved Issues and Enhancements

This section contains key software enhancements and resolved issues in the following releases of the NVIDIA Linux drivers:

- "Release 20: Resolved Issues and Enhancements" on page 11
- "Release 10: Resolved Issues and Enhancements" on page 12
- "Release 6: Resolved Issues and Enhancements" on page 14

Release 20: Resolved Issues and Enhancements

Version 1.0-2313

- Added support for NVIDIA nForce (Crush11/12) chip.

- Added support for OpenGL 1.3.

- Added SGIX_pbuffer and SGIX_fbconfig extensions.

- Added support for Unified Back Buffer, Page Flipping, and Window Flipping features.

- Improved performance of SPECviewperf, Quake, and immediate mode applications.

- Improved driver stability on AMD platforms.

- Fixed TwinView problem that caused garbage to appear on the screen when starting X with a null option on head 0.

- Added workaround for XAA bug that caused systems to hang when Lisa Screensaver is run. The workaround is to add the following line to the "Device" section of the XF86Config file: Option "XaaNoSolidFillTrap". This option prevents XAA from breaking wide lines (and polygons) into trapezoids and avoids an XAA clipping problem.

- Fixed a system hang that occurred on some chipsets when taking the X server down (after having run once successfully) and restarting it while using the DVI-I connected to a flat panel.

- Fixed problem where the performer application Perfly would hang when run in forked-draw mode.

- Fixed problem where the X driver would segfault when given an invalid MetaMode.

- Fixed problem where the console was not restored properly after entering X and returning to the console.

- Fixed problem where X did not redraw completely after a screen blank on GeForce2 Go and Quadro2 Go.

- Fixed problem with initializing the secondary card. In most cases, the secondary card is posted correctly. In the cases where it isn't, a workaround was created to initialize the card via X. (See "Multicard, Multimonitor" on page 7.)

- Added workarounds to enable AGP on ALi chipsets and enable 4X AGP on VIA chipsets. (*See* "VIA KX133 and 694X Chipsets With AGP 2X" on page 8 and "ALi Chipsets: ALi1541 and ALi1647" on page 9.)

Release 10: Resolved Issues and Enhancements

Resolved issues and enhancements are explained in these sections:

Version 1.0-1541

- Fixed problem where starting X on GeForce3 caused screen corruption (e.g., red vertical lines).

Version 1.0-1512

- Fixed problem where garbage appears on the screen and the LCD blooms when X is started on the Toshiba 3000 series laptops.

- Changed behavior of the X server so that the NVIDIA splash screen only appears on the first run of X. The splash screen can also be disabled by setting an option in the XF86Config file; *see* "Option "NoLogo" "boolean"" on page 33 for details.

- Fixed problem where OpenGL applications would sometimes leave portions of their rendering behind when the window was closed using the "x" button on the window banner.

- Fixed problem on mobile where X would respond to the wrong hot key event under certain conditions.

- Fixed several more problems with indirect rendering.

- Fixed problem on SMP machines that occurred when VT switching while running gloss and gears with indirect rendering.

- Fixed problem where /proc/nv/card0 did not report Quadro DCC correctly.

Version 1.0-1450

- Fixed problem on TNT2, where the driver would only support up to four threads per process.

- Fixed X Server crash that occurred when running two X Servers with AGPGART.

- Fixed some problems in GLX that occurred when running multi-threaded applications.

- Fixed problems with window borders picking up color values when moved across active OpenGL applications.

- Fixed problems so that the X Server detects a Quadro DCC-based card and properly initializes.

- Fixed problem so that Redhat 7.1 SMP RPMs can correctly uninstall with the "`rpm -e`" command while X Server is running.

- Corrected default OpenGL state when indirect rendering.

Version 1.0-1420

- Added `xf86XVOffscreenImage` support so the V4L module can use the hardware scaler on YUV surfaces.

- Added support for hot key switching on mobile platforms, i.e., laptop computers.

- Fixed a hang on mobile that occurred after starting, stopping, and then restarting X.

- Fixed a problem on mobile platforms that prevented DVDs from displaying.

- Fixed a problem that caused OpenGL programs to segfault when using a graphical login with xdm/kdm, and doing the following sequence: login, mode switch, logout, log in and run an OpenGL application.

- Fixed some indirect rendering problems.

- Fixed a problem that caused Xconfigure to fail on NVIDIA drivers.

- Fixed a crash that occurred when X forwarding over SSH.

- Fixed OpenGL front buffer clipping bug.

- Improved X-Render acceleration.

- Fixed a problem that prevented X-Render acceleration on GeForce3.

- The NvAGP option now defaults to "3", which causes the driver to use AGP GART, if it is available, and NVidia AGP, otherwise.

- Fixed issue with GeForce 256/DDR where Linux PCI adapter did not correctly initialize with AGP card installed.

- Fixed typo in error format string which caused error messages to report "%" when it should have printed one of several error messages.

Version 1.0-1251

- Added preliminary GeForce2 Go support.

- Added support for GeForce3 OpenGL and GLX extensions.

- Fixed many SMP issues.

- Added TV-Out support.

- Fixed DGA depth change problem.

- Rewrote 2D off screen memory allocation.

- Fixed X-Video in TwinView.

- Added acceleration for X-Render extension.

- Fixed `GLXPixmap` rendering.

- Fixed problem with `glXMakeCurrent()` to same drawable but different display.

- Fixed problem in which OpenGL caused a segfault when reading X atoms.

- Fixed issues so that X now gets the **dots per inch (dpi)** from the monitor's **EDID (Extended Display Identification Data)** instead of defaulting to 75 dpi.

- All DPMS modes are now supported. Some DPMS issues remain for flat panels and the second display in a TwinView configuration.

- Fixed support for AGP on systems with 1 GB or more of memory.

Release 6: Resolved Issues and Enhancements

Resolved issues and enhancements are categorized by the following versions:

Version 0.9-769

- Fixed problem where an old version of the release documentation was being installed instead of the current one.

- Fixed problem where direct rendering applications were allowed to continue rendering after "`xkill`" was called.

- Fixed problem where Tribes 2 crashed when compressed (s3tc) textures were used.

- Some drawable leaks were fixed in X and GLX.

- Fixed problem where the application would hang when calling "`glXMakeCurrent()`" while holding the X Server grab.

- BIOS posting problems with GeForce2 GTS and GeForce Ultra were fixed. These problem caused a significant performance loss.

- Added support for the X Render extension.

- TwinView functionality was enhanced for each display to pan independently.

- Fixed problem on TNT and TNT2 where "`Xv(Shm)PutImage`" returned "`BadAlloc`" in high resolutions when there was not enough video bandwidth to correctly display the YUV video overlay. This works now but the resulting display has artifacts.

- Fixed problem with cursor hangs in X.

- Fixed problem with X console not restoring on some monitors.

- Fixed problem with `fork()` and OpenGL rendering

- Fixed problem with X driver module, `nvidia_drv.o`, being stripped when RPM was rebuilt.

- Added missing PCI device IDs for some TNT2 variants and GeForce3.

- Fixed problem where the kernel would often hang during X and/or OpenGL operation when on an SMP machine and using the version 2.4 kernel.

- Fixed `SYNC_TO_VBLANK` hang with 2.4 kernels.

- Fixed DPMS so that it is possible to set the "off" option. DPMS options "suspend" and "standby" are not fully supported; these options simply blank the screen. (Be sure to include `Option "DPMS"` in your XF86Config file. Refer to the XF86Config man page for detailed information.)

Version 0.96

- Fixed many SMP problems.
- Fixed memory management problems that arose with large RAM systems (500 MB plus).
- Added multi-monitor OpenGL support.
- Added TwinView support.
- Fixed more mode-line handling issues and added double-scan support.
- Fixed BIOS-posting problems with TNT2 M64 and GeForce2 MX.
- Added dynamic run-time selection between NVAGP and AGPGART.
- Fixed TNT2 OpenGL slowdowns, which were noticeable in UT.

Version 0.95

- Improved XFree86 version 4.0.1 support.
- Re-fixed console switch lockup.
- Fixed some AGP regressions resulting in better detection/support of AGPGART.
- Fixed color palette problems (xgamma, direct color visuals).
- Added BIOS-posting override to help with NVIDIA products such as TNT2 M64.
- Update included version 2.4 support to newest test kernels.

Version 0.94

- Added support for XFree86 version 4.0.1.
- Fixed mode-setting problem.
- Added AGPGART support (NVAGPGART version 0.5-5).
- Added GeForce2 MX support.
- Fixed various hangs.
- Added full-scene antialiasing (FSAA) support.
- Fixed problem where an OpenGL application malfunctioning during a console switch would crash the X Server.

Version 0.9-3

- Allowed mode-line directives in the XF86Config file to override NVIDIA auto-detection of monitor resolutions and refresh rates.

- Implemented "correct" fix for TNT memory-type problems.

- Fixed VT switch lockups.

- Fixed general ALi chipset lockups.

- Added and documented some registry keys Check `os-registry.c` in the kernel source directory for more details and options.

- Implemented workaround for Quake3 mode switch problem that caused system to crash. Note that this was a problem in the `dlopen()` function.

- Implemented major improvement in multi-threading behavior.

- Display list sharing with `glXCreateContext` now works.

- Added faster implementation of `glTexImage/glTexSubImage` and `glCopyTexImage/glCopytexSubImage` calls.

- Fixed kernel memory leak, which was related to threaded OpenGL This problem was most noticeable with XMMS.

- Fixed build problems with older version 2.2.*x* kernels (RedHat 6.0)

Version 0.9-2

- Fixed problem initializing TNT with SGRAM.

- Added better logging and messages for tracking problems.

- Added dynamic, rather than static, allocation of client data in kernel.

- Incorporated *unsupported* version 2.3 kernel changes for completeness.

- Makefile updates add "`-D_LOOSE_KERNEL_NAMES`" and default to "`make install`".

- Improved mode switching in Quake3.

- Changed installation name of libraries. Added revision .1.0.1 in the libraries.

- Temporarily forced disabling of AGP fast writes for all chips.

- Fixed monitor issues and allowed overriding of synchronization polarities.

Version 0.9-1

Initial Release

3

INSTALLING THE NVIDIA LINUX DRIVERS

This chapter contains the following major topics:

Choosing the NVIDIA Packages for Your System

The NVIDIA Accelerated Linux Driver Set consists of two packages that you need to download and install:

- `NVIDIA_GLX` package contains the OpenGL libraries and the XFree86 driver.

- `NVIDIA_kernel` package contains the NVdriver kernel module required by the X driver and OpenGL libraries in the `NVIDIA_GLX` package.

For detailed description of the components of each package, *see* "Installed Components" on page 24.

Note: You must install both packages with matching version numbers; i.e., `NVIDIA_GLX-0.9-6` should only be used with `NVIDIA_kernel-0.9-6` and *not* `NVIDIA_kernel-0.9-3`.

The packages are available in these three formats:

- **RPM**
- **SRPM**
- **Tar file**

Installation of each package type is described in the sections that follow.

The package type is largely a matter of personal preference, though note that the binary RPMs are for use only with the kernel shipped with a particular distribution; i.e., `NVIDIA_kernel-0.9-6.rh62.i386.rpm` should only be used with the uni-processor kernel shipped with RedHat version 6.2.

Where appropriate, NVIDIA has provided separate RPMs for the distinct SMP and uni-processor kernels of each distribution. If you have upgraded your kernel (either manually, or through a distribution upgrade) or a specific NVIDIA_kernel RPM is not available for your distribution, then use either the NVIDIA_kernel SRPM or tar file.

In the case where distributors ship multiple kernels (as is often the case with uni-processor and SMP systems), multiple RPMs are available, i.e., `NVIDIA_kernel-0.9-7.rh62.i386.rpm` and `NVIDIA_kernel-0.9-7.rh62.smp.i386.rpm`.

The `NVIDIA_GLX` RPM, however, is not dependent upon the kernel version, and therefore an SRPM is not required. Install the `NVIDIA_GLX` package either by RPM or tar file.
If you have upgraded

Installing the NVIDIA_kernel and NVIDIA_GLX Packages

This section contains the following sections:

Before You Begin Driver Installation

1 Exit the X Server.

2 Set your default run level so you will boot to console and not start up X.

Note: If you are unsure about how to perform this step, refer to the documentation that came with your Linux distribution.

3 The package revision numbers have been omitted in the installation instructions to make them as general as possible.

For example, if the directions are `NVIDIA_kernel.tar.gz`, replace that with the name of the driver version you are installing; for example: `NVIDIA_kernel.0.9-6.tar.gz`.

Installing by RPM

Commands

```
$ rpm -ivh NVIDIA_kernel.i386.rpm
$ rpm -ivh NVIDIA_GLX.i386.rpm
```

Commands Explained

1 Before installing from RPM, be sue you have downloaded the `NVIDIA_kernel` RPM that is appropriate for your kernel.

2 Once you have verified that you have the correct RPM, install `NVIDIA_kernel` with this command:

```
$ rpm -ivh NVIDIA_kernel.i386.rpm
```

3 Next, install the `NVIDIA_GLX` RPM with this command:

```
$ rpm -ivh NVIDIA_GLX.i386.rpm
```

Upgrading by RPM

Commands

```
$ rpm -Uvh NVIDIA_kernel.i386.rpm
$ rpm -e NVIDIA_GLX
$ rpm -ivh NVIDIA_GLX.i386.rpm
```

Commands Explained

Before upgrading from RPM, be sure that you have downloaded the NVIDIA_kernel RPM that is appropriate for your kernel. Once you have verified that you have the correct RPM, upgrade the NVIDIA_kernel package using this command:

```
$ rpm -Uvh NVIDIA_kernel.i386.rpm
```

Note: You should not use the "-U" option to upgrade the NVIDIA_GLX RPM because a problem in the uninstall section of older NVIDIA RPMs will cause some files to be incorrectly removed. Instead, use "-e" to remove the old NVIDIA_GLX RPM and then install the new one using these commands:

```
$ rpm -e NVIDIA_GLX
```
```
$ rpm -ivh NVIDIA_GLX.i386.rpm
```

Installing & Upgrading by SRPM

Commands

```
$ rpm --rebuild NVIDIA_kernel.src.rpm
$ rpm -ivh /path/to/rpms/RPMS/i386/NVIDIA_kernel.i386.rpm
$ rpm -ivh NVIDIA_GLX.i386.rpm
```

Commands Explained

To build a custom NVIDIA_kernel for your system, pass RPM the "--rebuild" flag:

```
$ rpm --rebuild NVIDIA_kernel.src.rpm
```

Watch for the line that looks something like (the path may be different):

```
Wrote: /usr/src/redhat/RPMS/i386/NVIDIA_kernel.i386.rpm
```

and use that as input for RPM to install:

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

or upgrade:

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

To install the NVIDIA_GLX package, follow the instructions above for either installing or upgrading NVIDIA_GLX from RPM.

Installing & Upgrading by Tar File

Commands

```
$ tar xvzf NVIDIA_kernel.tar.gz
$ tar xvzf NVIDIA_GLX.tar.gz
$ cd NVIDIA_kernel
$ make install
$ cd ../NVIDIA_GLX
$ make install
```

Commands Explained

To install from tar file, follow these steps:

1 Unpack each file with these commands:

```
$ tar xvzf NVIDIA_kernel.tar.gz
$ tar xvzf NVIDIA_GLX.tar.gz
```

2 Change to the NVIDIA_kernel directory. Then use the "make install" command to compile the kernel interface to the NVdriver; link the NVdriver; copy the NVdriver into place; and attempt to insert the NVdriver into the running kernel.

```
$ cd NVIDIA_kernel
$ make install
```

3 Move into the NVIDIA_GLX directory by using the "make install" command to copy required OpenGL and XFree86 files into place.

```
$ cd ../NVIDIA_GLX
$ make install
```

Editing Your XF86Config File

Note: This documentation uses "XF86Config" to refer to your configuration file; you may have a different name for the file.

When XFree86 4.0 was released, it used a slightly different XF86Config file syntax than the 3.*x* series used. Therefore, to allow both 3.*x* and 4.*x* versions of XFree86 to co-exist on the same system, it was decided that XFree86 4.x will use the configuration file /etc/X11/XF86Config-4, if this file exists. If this file does not exit, XFree86 4.x will use /etc/X11/XF86Config.

Note: X searches a large path to find the config files. For a complete description of the search path, it is strongly recommended that you refer to the XF86Config man page.

Verify the configuration file that XFree86 is using. To do so, you can locate a line beginning with "`==) Using config file:`" in your XFree86 log file (`/var/log/XFree86.0.log`.)

If you do not have a working XF86Config file, there are several ways to begin:

• Refer to the sample configuration file is included with XFree86.

• Refer to the sample configuration file is included in the `NVIDIA_GLX` package and installed in `/usr/share/doc/NVIDIA_GLX-1.0`.

• You can also use a program such as `xf86config`. Some distributions provide their own tool for generating an XF86Config file. (For details on the XF86Config file syntax, refer to the man page.)

If you already have an XF86Config file working with a different driver, such as the `nv` driver, then follow these steps:

1 Find the relevant Device section and replace the line:

```
Driver "nv"
```

with

```
Driver "nvidia"
```

2 In the Module section, verify that you have:

```
Load   "glx"
```

3 *Remove* the following lines, if they exist:

```
Load   "dri"
Load   "GLcore"
```

4 There are many options that can be added to the XF86Config file to fine-tune the NVIDIA XFree86 driver. (*See* "XF86Config Options: Release 10" on page 30 for a complete list of these options.)

5 Once you have configured your XF86Config file, you are ready to restart X and begin using the accelerated OpenGL libraries.

After you restart X, you can run any OpenGL application, which will automatically use the new NVIDIA libraries. If you encounter any problems, *see* "Troubleshooting, Questions, & Resources" on page 48.

Installed Components

The NVIDIA Accelerated Linux Driver Set consists of several components:

- The **XFree86 driver**, **GLX module**, **libGL**, and **libGLcore** components are included in the NVIDIA_GLX package.

- The **NVdriver kernel module** component is included in the NVIDIA_kernel package.

- **Documentation and the OpenGL and GLX header files** are part of the NVIDIA_GLX package and are installed in:

    /usr/share/doc/NVIDIA_GLX-1.0

Note: The file shown in parenthesis is the full name of the component after installation. x.y.z denotes the current version — appropriate symlinks are created during installation.

- **An XFree86 driver**

    (/usr/X11R6/lib/modules/drivers/nvidia_drv.o)

    is required by XFree86 to make use of your NVIDIA hardware. The nvidia_drv.o driver is binary compatible with XFree86 version 4.0.1 and higher.

- **A GLX extension module**

    (/usr/X11R6/lib/modules/extensions/libglx.so.x.y.z)

    is used by XFree86 to provide server-side GLX support.

- **An OpenGL library**

    (/usr/lib/libGL.so.x.y.z)

    provides the API entry points for all OpenGL and GLX function calls. At run-time, OpenGL applications link to this library.

- **An OpenGL core library**

    (/usr/lib/libGLcore.so.x.y.z)

    is implicitly used by libGL and libglx and contains the core accelerated 3D functionality.

    Note: *Do not* explicitly load the OpenGL core library in your XF86Config file; this task is performed by libglx.

- **A kernel module**

    (/lib/modules/`uname -r`/video/NVdriver) <u>or</u>
    (/lib/modules/`uname -r`/kernel/drivers/video/NVdriver)

provides low-level access to your NVIDIA hardware for all of the above components. It is generally loaded into the kernel when the X Server is started and is used by the XFree86 driver and OpenGL.

`NVdriver` has two components:

- the **binary-only core** *and*

- a **kernel interface** that must be compiled specifically for your kernel version.

Note: The Linux kernel does not have a consistent binary interface such as XFree86, so it is important that this kernel interface is matched to the version of the kernel that you are using. You can accomplish this by compiling yourself or using precompiled binaries provided for the kernels shipped with some of the more common Linux distributions.

- **OpenGL and GLX header files**

  `(/usr/share/doc/NVIDIA_GLX-1.0/usr/include/GL/gl.h)`

  `(/usr/share/doc/NVIDIA_GLX-1.0/usr/include/GL/glx.h)`

  In most cases, the system-supplied headers in `/usr/include/GL` will suffice for OpenGL development. But NVIDIA provided these headers because they contain the most up-to-date versions of the NVIDIA OpenGL extensions. If you want to use these headers, it is recommended that you copy them to `/usr/include/GL`.

Installing Libraries

Note: Problems will arise if applications use the wrong version of a library, which can occur if either old libGL libraries or obsolete symlinks exist. If you have reason to believe that your installation may encounter problems, check that the following files are in place; these files are part of the NVIDIA Accelerated Linux Driver Set and include their symlinks:

- `/usr/X11R6/lib/modules/drivers/nvidia_drv.o`
- `/usr/X11R6/lib/modules/extensions/libglx.so.x.y.z`
- `/usr/X11R6/lib/modules/extensions/libglx.so -> libglx.so.x.y.z`
- `/usr/lib/libGL.so.x.y.z`
- `/usr/lib/libGL.so.x -> libGL.so.x.y.z`
- `/usr/lib/libGL.so -> libGL.so.x`
- `/usr/lib/libGLcore.so.x.y.z`
- `/usr/lib/libGLcore.so.x -> libGLcore.so.x.y.z`
- `/lib/modules/`uname -r`/video/NVdriver`, or
- `/lib/modules/`uname -r`/kernel/drivers/video/NVdriver`

Installing the `NVIDIA_kernel` package also creates the `/dev` files:

`crw-rw-rw- 1 root root     195,   0 Feb 15 17:21 nvidia0`

```
crw-rw-rw- 1 root root      195,   1 Feb 15 17:21 nvidia1
crw-rw-rw- 1 root root      195,   2 Feb 15 17:21 nvidia2
crw-rw-rw- 1 root root      195,   3 Feb 15 17:21 nvidia3
crw-rw-rw- 1 root root      195, 255 Feb 15 17:21 nvidiactl
```

If there are other libraries with a "so* name" that conflicts with that of the NVIDIA libraries, "ldconfig" may create the wrong symlinks. In this case, it is recommended that you follow these steps:

1 Manually remove or *rename* conflicting libraries. (*See* the **Note** at the end of these steps.)

Note: Be sure to rename clashing libraries to a name that "ldconfig" will not identify; for example, you can prepend "xxx" to a library name.)

2 Rerun "ldconfig" and check that the correct symlinks were made. Some libraries that often create conflicts are

/usr/X11R6/lib/libGL.so* and
/usr/X11R6/lib/libGLcore.so*.

3 Once you've verified the libraries, then verify that the application is using the correct libraries.

For example, to check that the application /usr/X11R6/bin/gears is using the NVIDIA libraries, you would issue the following command:

```
$ ldd /usr/X11R6/bin/gears
  libglut.so.3 => /usr/lib/libglut.so.3 (0x40014000)
  libGLU.so.1 => /usr/lib/libGLU.so.1 (0x40046000)
  libGL.so.1 => /usr/lib/libGL.so.1 (0x40062000)
  libc.so.6 => /lib/libc.so.6 (0x4009f000)
  libSM.so.6 => /usr/X11R6/lib/libSM.so.6 (0x4018d000)
  libICE.so.6 => /usr/X11R6/lib/libICE.so.6 (0x40196000)
  libXmu.so.6 => /usr/X11R6/lib/libXmu.so.6 (0x401ac000)
  libXext.so.6 => /usr/X11R6/lib/libXext.so.6 (0x401c0000)
  libXi.so.6 => /usr/X11R6/lib/libXi.so.6 (0x401cd000)
  libX11.so.6 => /usr/X11R6/lib/libX11.so.6 (0x401d6000)
  libGLcore.so.1 => /usr/lib/libGLcore.so.1 (0x402ab000)
  libm.so.6 => /lib/libm.so.6 (0x4048d000)
  libdl.so.2 => /lib/libdl.so.2 (0x404a9000)
  /lib/ld-linux.so.2 => /lib/ld-linux.so.2 (0x40000000)
  libXt.so.6 => /usr/X11R6/lib/libXt.so.6 (0x404ac000)
```

Note: If the files being used for libGL and libGLcore are not NVIDIA libraries, you need to either remove the libraries that are getting in the way, or adjust your "ld" search path. If you are not familiar with this process, you may want to read the man pages for "ldconfig" and "ldd" for pointers.

For information on troubleshooting the NVIDIA Accelerated Linux Driver Set, *see* .

4

CONFIGURING NVIDIA LINUX DRIVERS

This chapter contains the following major topics:

Note: This documentation uses "XF86Config" to refer to your configuration file; you may be using a different file name.

XF86Config Options: New in Release 20

Note: Release 20 options are *not* supported by Release 10 or older NVIDIA Linux drivers.

The following options are supported by the NVIDIA XFree86 driver for Release 20 and need to be included in your XF86Config configuration file:

Unified Back Buffer (UBB), Page Flipping, and Window Flipping features provide performance gains under certain conditions.

Option "PageFlip" "*boolean*"

Note: The Page Flipping feature is available on all NVIDIA GeForce or newer hardware, which excludes TNT/TNT2 products.

Default: Option "PageFlip" "on" enables Page Flipping.

To disable the option, specify "off" in the above statement.

Page Flipping is enabled (by default) in the case of a single full-screen unobscured OpenGL application when synchronizing to VBLANK. Buffer swapping is done by changing the buffer rather than copying the back buffer contents to the front buffer. This mechanism allows for a much higher performance and non-tearing swapping during the retrace (when __GL_SYNC_TO_VBLANK is set).

Option "UBB" "*boolean*"

Note: UBB is available *only* on Quadro and Quadro2 family products.

Default: Option "UBB" "on" enables UBB.

To disable UBB, specify "off" in the above statement.

UBB is enabled by default when there is sufficient video memory available so that all windows share the same back, stencil, and depth buffer. When there are many windows, the back, stencil, and depth buffer usage will never exceed the size of that used by a full-screen window.

Note: For a single small window, the back, stencil, and depth buffer usage equals that of a full-screen window. In this case, the video RAM may be used more efficiently when UBB is disabled.

Option "UseInt10Module" "*boolean*"

Default: Option "UseInt10Module" "off"

In this "default" condition, POSTing the cards is done through the NVdriver kernel module.

To enable use of the XFree86 Int10 module to soft-boot all secondary cards rather than POSTing the cards through the NVdriver kernel module, specify "on" in the above statement.

Option "WindowFlip" "*boolean*"

Note: Window Flipping is a feature that requires UBB and therefore only available on Quadro and Quadro2 family products.

Default: Option "WindowFlip" "off" disables Window Flipping.

To enable Window Flipping:

1 Confirm that UBB is enabled. (*See* Option "PageFlip" "boolean" in the previous section.)

2 Include this line in your configuration file:

Option "WindowFlip" "on".

Window Flipping only works when there is a single OpenGL window. This OpenGL windows's buffers can be swapped by changing the buffer rather than copying the contents of the back buffer to the front buffer. This is similar to the Page Flipping functionality but removes the restriction that the window must be unobscured and full-screen in order for the feature to work.

XF86Config Options: Release 10

Release 10 options *are* supported by Release 20 and any newer NVIDIA Linux drivers.

The following options are supported by the NVIDIA XFree86 driver and should be specified in your XF86Config file.

• "Option "UseEdidFreqs" "boolean"" on page 34

Unified Back Buffer (UBB), Page Flipping, and Window Flipping features provide performance gains under certain conditions.

Option "ConnectedMonitor" "*string*"

Default: String is NULL.

Examples:

```
Option "ConnectedMonitor" "TV"
Option "ConnectedMonitor" "CRT"
Option "ConnectedMonitor" "CRT,DFP"
Option "ConnectedMonitor" "CRT,TV"
```

This option allows you to override the device that the NVIDIA kernel module detects as connected to your graphics card.

This feature may be useful under these circumstances:

• If any of your display devices does not support detection using **Display Data Channel (DDC)** protocols

• If you use a **KVM (keyboard/video/mouse)** switch and you are switched away when X is started. In such a situation, the NVIDIA kernel module cannot detect the display devices that are connected and the NVIDIA X driver assumes you have a single CRT connected.

• If you use a digital flat panel instead of a CRT, use this option to explicitly communicate to the NVIDIA X driver the device that is connected.

Valid values for this option are:

 • `"CRT"` (cathode ray tube / analog monitor)
 • `"DFP"` (digital flat panel)
 • `"TV"` (television)

If using TwinView, this option may be a comma-separated list of display devices; for example:

 • `"CRT,CRT"`
 • `"CRT,DFP"`
 • `"CRT,TV"`

As in all XF86Config entries, spaces are ignored and all entries are case insensitive.

Option "CursorShadow" "*boolean*"

This option is only available on GeForce2 MX or later hardware, which includes GeForce2 MX (200/400), GeForce2 Pro, GeForce 2 Ultra, GeForce2 Go, GeForce3, Quadro, Quadro DDC, Quadro2 (MXR/EX), Quadro2 Pro and Quadro2 Go.

Default: `Option "CursorShadow" "off"`

This option enables (`"on"`) or disables (`"off"`) use of a shadow with the hardware accelerated cursor. The shadow is a black translucent replica of your cursor shape at a given offset from the real cursor.

Option "CursorShadowAlpha" "*integer*"

Default: `Option "CursorShadowAlpha" "64"`

Integer must be in the range [0, 255].

  0:   Completely transparent
255:   Completely opaque

Note: This option defines the alpha value to use for the cursor shadow and is applicable *only* if `CursorShadow` is enabled.

Option "CursorShadowXOffset" "*integer*"

Default: `Option "CursorShadowXOffset" "4"`

Integer must be in the range [0, 32].

This option defines the offset, in pixels, that the shadow image will be shifted to the right from the real cursor image. It is only applicable if `CursorShadow` is enabled.

Option "FlatPanelScalingMode" "*string*"

Examples:

```
Option "FlatPanelScalingMode" "scaled"
Option "FlatPanelScalingMode" "centered"
```

This option requests a specific flat panel scaling mode where "*string*" can be one of the following:

• **Centered**: Any connected flat panels will be configured to center the screen contents. For example, setting an 800x600 mode on a flat panel with native

resolution 1400x1050 will result in the center 800x600 pixels of the flat panel to be used with a black surrounding border.

• **Scaled**: The screen contents will be scaled to use the entire flat panel screen.

Option "HWCursor" "*boolean*"

Default: Option "HWCursor" "on"

This option enables or disables (with the "off" setting) hardware rendering of the X cursor.

Option "IgnoreEDID" "*boolean*"

Default: Option "IgnoreEDID" "off" means EDIDs *are used* to validate modes.

This option disables ("on") or enables ("off") probing of EDIDs from your monitor. Requested modes are compared against values obtained from your monitor EDIDs (if any) during mode validation.

Caution: Some monitors may not provide accurate information about its physical capabilities. However, ignoring the values that the monitor provides (i.e., setting the option to "on") may help to validate a certain mode but *could result in serious problems* and, therefore, should be used with extreme care.

Option "NoDCC" "*boolean*"

Synonym for "Option "IgnoreEDID" "boolean"" on page 33.

Option "NoLogo" "*boolean*"

Default: Option "NoLogo" "off" *enables* the NVIDIA logo splash screen when the X server is started.

To disable the splash screen during X start-up, specify the "on" setting.

Option "NoRenderAccel" "*boolean*"

Default: Option "NoRenderAccel" "off" *enables* hardware acceleration of the RENDER extension.

To disable acceleration of the RENDER extension, specify "on" in the above statement.

Option "NvAGP" "*integer*"

Default: Option "NvAGP" "3"

Note: The default was **1** through NVIDIA Linux driver version 1.0-1251.

This option configures AGP support; *integer* value can be one of:

   0: Disable AGP.

   1: Use the NVIDIA internal AGP support, if possible.

   2: Use AGPGART, if possible.

   3: Use any AGP support (try AGPGART, then the NVIDIA AGP).

Note: The NVIDIA internal AGP support cannot work if AGPGART is either statically compiled into your kernel or is built as a module and loaded into your kernel. (Some distributions load AGPGART into the kernel during system start-up.)

Option "SWCursor" "*boolean*"

Default: Option "SWCursor" "off" disables software rendering of the X cursor.

To enable software rendering of the X cursor, specify "on" in the above statement.

Option "UseEdidFreqs" "*boolean*"

Default: Option "UseEdidFreqs" "off" means that the frequencies provided by the display device's EDID will *not* be used.

When "on" is specified in the above command, the X server uses the HorizSync and VertRefresh ranges given in a display device's EDID, if any. EDID-supplied range information will override the HorizSync and VertRefresh ranges specified in the "Monitor" section of the XFree86 configuration file.

If a display device does not provide an EDID, or the EDID does not specify an HSync or VRefresh range, then the X server will default to the HorizSync and VertRefresh ranges specified in the "Monitor" section of the XFree86 configuration file.

OpenGL Environment Variable Settings

Full-Scene Antialiasing (FSAA)

Antialiasing is a technique used to smooth the edges of objects in a scene to reduce the jagged "stairstep" effect that sometimes appears. Full-scene antialiasing is supported on the following NVIDIA products:

GeForce3	GeForce2 GTS	Quadro2 MXR/EX
GeForce2 Go	GeForce2 MX	Quadro2 Go
GeForce2 Ultra	GeForce 256	Quadro2 Pro
GeForce2 Pro	Quadro DCC	Quadro

By setting the `__GL_FSAA_MODE` environment variable you can enable full-scene antialiasing in any OpenGL application when using one of the above NVIDIA products.

Several antialiasing methods are available and you can select among them by setting the `__GL_FSAA_MODE` environment variable using one of the values in Table 4.1. The values correspond to antialiasing methods.

Note: Increasing the number of samples taken during FSAA rendering may decrease performance.

Table 4.1   Values for the `__GL_FSAA_MODE` Environment Variable

Value	GeForce/Geforce2/Quadro Description	GeForce3 Description
0	FSAA is disabled.	FSAA is disabled.
1	FSAA is disabled.	2 x 2 oversampling with texture LOD bias
2	FSAA is disabled.	2 x 2 Quincunx
3	1.5 x 1.5 oversampling	FSAA is disabled.
4	2 x 2 oversampling with no texture LOD bias	4 x 4 Bilinear
5	FSAA is disabled.	4 x 4 Gaussian

VBLANK Synchronizing

Setting the environment variable `__GL_SYNC_TO_VBLANK` to a non-zero value forces `glXSwapBuffers` to synchronize to your monitor's vertical refresh rate. This implies that the code performs a swap only during the vertical blanking period on GeForce or newer hardware (i.e., any NVIDIA product except any in the TNT and TNT2 families).

Note: This option is available only on NVIDIA "3D" products, such as those in the GeForce/GeForce2/GeForce3 and equivalent workstation family.

Configuring AGP

The `"NvAgp"` option in your XF86Config file provides several ways to configure the use of AGP by the NVdriver kernel module.

1 You can use either the NVIDIA AGP module (NVAGP) or the AGP module that comes with the Linux kernel (AGPGART).

`Option "NvAgp" "0"`   ... disables AGP support

`Option "NvAgp" "1"`   ... use NVAGP, if possible

`Option "NvAgp" "2"`   ... use AGPGART, if possible

`Option "NvAgp" "3"`   ... try AGPGART; if that fails, try NVAGP

Default: `"3"`

(The default was `"1"` through NVIDIA Linux driver version 1.0-1251.)

2 It is recommended that you use the AGP module that works best with your AGP chipset. If you are experiencing problems with stability, you may want to start by disabling AGP and observing if that solves the problems. You can then experiment with either of the other AGP modules.

3 You can check your AGP status with this command:

`cat /proc/nv/card0`

4 To use the Linux AGPGART module, you must compile it with your kernel, either statically linked in or built as a module.

> **Note:** NVIDIA AGP support cannot be used if AGPGART is loaded in the kernel. It is recommended that you compile AGPGART as a module and verify that it is not loaded when trying to use NVIDIA AGP.
>
> Changing AGP drivers generally requires a reboot before the changes actually take effect.

5 Rebuild and reinstall the new driver using "`make`", which forces the driver to ignore the BIOS of the NVIDIA product and use your values.

AGP Chipsets Supported by NVIDIA AGP

The following AGP chipsets are supported by the NVIDIA AGP module; for all other chipsets, it is recommended that you use the AGPGART module.

• Intel 440LX

• Intel 440BX

• Intel 440GX

- Intel 815 ("Solano")

- Intel 820 ("Camino")

- Intel 840 ("Carmel")

- Intel 845 ("Brookdale")

- Intel 850 ("Tehama")

- Intel 860 ("Colusa")

- AMD 751 ("Irongate")

- AMD 761 ("IGD4")

- AMD 762 ("IGD4 MP")

- VIA 8371

- VIA 82C694X

- VIA KT133

- RCC 6585HE

- Micron SAMDDR ("Samurai")

- Micron SCIDDR ("Scimitar")

- nForce AGP

Enabling and Configuring TwinView

For a description of the TwinView feature, *see* "Release 6" on page 6.

This sections contains the following topics:

To enable TwinView, you must specify the following options in the "Screen" section of your XF86Config file.

```
Option "TwinView"
Option "SecondMonitorHorizSync"     "<hsync range(s)>"
Option "SecondMonitorVertRefresh"   "<vrefresh range(s)>"
Option "MetaModes"                  "<list of metamodes>"
```

You may also use any of the following options, though they are *not required*.

```
Option                    "<relationship of head (display)
"TwinViewOrientation"     1 to head (display) 0>"
Option "ConnectedMonitor" "<list of connected display
                           devices>"
```

Descriptions of the options are provided in the sections that follow.

Option "TwinView" "*boolean*"

Default: Option "TwinView" "off" (TwinView is disabled.)

To enable TwinView, you must enable this option with the "on" setting as shown below; otherwise all other TwinView related options are ignored.

Option "TwinView" "on"

Option "SecondMonitorHorizSync" "*range(s)*"

Default: None

This option is similar to the "HorizSync" entry in the "Monitor" section of your XF86Config file, but applies to the second monitor when using TwinView.

According to the XF86Config man page, the "ranges" may be a comma-separated list of distinct values and/or ranges of values, where a range is given by two distinct values separated by a dash. The HorizSync value is given in KHz.

If you trust the EDID of your display device, you may want to use the "UseEdidFreqs" option instead of this option. (For details, *see* "Option "UseEdidFreqs" "boolean"" on page 34.)

Option "SecondMonitorVertRefresh" "*range(s)*"

Default: None

This option is similar to the "VertRefresh" entry in the "Monitor" section of your XF86Config file, but applies to the second monitor when using TwinView.

According to the XF86Config man page, the "ranges" may be a comma-separated list of distinct values and/or ranges of values, where a range is given by two distinct values separated by a dash. The VertRefresh value is given in Hz.

If you trust the EDID of your display device, you may want to use the "UseEdidFreqs" option instead of this option. (For details, *see* "Option "UseEdidFreqs" "boolean"" on page 34.)

Option "MetaModes" "*string*"

Default: None

A single MetaMode describes the mode that should be used on each display device at a given time. Multiple MetaModes list the combinations of modes and the sequence in which they should be used. When the NVIDIA driver communicates the available modes to X, it is really the minimal bounding box of the MetaMode that is communicated, while the "per display device" mode is kept internal to the NVIDIA driver. In MetaMode syntax, modes within a MetaMode are separated by a comma and multiple MetaModes are separated by semicolons. For example:

```
"<mode name 0>, <mode name 1>; <mode name 2>, <mode name 3>; ..."
```

Where <mode name 0> is the name of the mode to be used on display device 0 concurrently with <mode name 1> used on display device 1. A mode switch will then cause <mode name 2> to be used on display device 0 and <mode name 3> to be used on display device 1.

- **An actual MetaMode entry** from the XF86Config TwinView sample configuration file is as follows:

```
Option "MetaModes" "1280x1024,1280x1024;
1024x768,1024x768"
```

- I**f you do not want a display device to be active for a certain MetaMode**, you can use the mode name "NULL", or simply omit the mode name entirely, as follows:

```
"1600x1200, NULL; NULL, 1024x768"
```

  or

```
"1600x1200; , 1024x768"
```

- **Optionally, mode names can be followed by "offset" information** to control the positioning of the display devices within the virtual screen space; for example:

```
"1600x1200 +0+0, 1024x768 +1600+0; ..."
```

Offset descriptions follow the conventions used in the X `"-geometry"` command line option; i.e. both positive and negative offsets are valid, though negative offsets are only allowed when a virtual screen size is explicitly given in the XF86Config file.

When no offsets are given for a MetaMode, the offsets are computed following the value of the `"TwinViewOrientation"` option. (*See* "Option "TwinViewOrientation" "string"" on page 41.)

Note: If offsets are given for any one of the modes in a single MetaMode, then offsets are expected for all modes within that single MetaMode; in such a case, offsets are assumed to be +0+0 when not given.

When not explicitly given, the virtual screen size is computed as the bounding box of all MetaMode bounding boxes.

MetaModes with a bounding box larger than an explicitly given virtual screen size are discarded.

- **A MetaMode string can be further modified with a** "**panning domain" specification**; for example:

```
"1024x768 @1600x1200, 800x600 @1600x1200"
```

A panning domain is the area in which a display device's viewport is panned to follow the mouse.

Panning takes place on two levels under TwinView:

- *First*, an individual display device's viewport is panned within its panning domain, as long as the viewport is contained by the bounding box of the MetaMode.

- *Second***,** once the mouse leaves the bounding box of the MetaMode, the entire MetaMode (i.e., all display devices) is panned to follow the mouse within the virtual screen.

Note: The panning domains of the display devices default to being clamped to the position of the display devices' viewports. Therefore, the default functionality is that viewports remain "locked" together and only perform the second type of panning.

The most beneficial use of panning domains is to eliminate dead areas, which are regions of the virtual screen that are inaccessible due to display devices with different resolutions. For example:

- Specifying `"1600x1200, 1024x768"` produces an inaccessible region below the 1024x768 display.

- Specifying `"1600x1200, 1024x768 @1024x1200"` as a panning domain for the second display device: provides access to that dead area by allowing you to pan the 1024x768 viewport up and down in the 1024x1200 panning domain.

- **Offsets can be used in conjunction with panning domains** to position the panning domains in the virtual screen space. The offset describes the panning domain and only affects the viewport in that the viewport must be contained within the panning domain. For example, the following line describes two modes, each with a panning domain width of 1900 pixels where the second display is positioned below the first:

    `"1600x1200 @1900x1200 +0+0, 1024x768 @1900x768 +0+1200"`

Note: If a MetaMode string is not specified, then the X driver uses the modes listed in the relevant "`Display`" subsection of the XF86Config file, attempting to place matching modes on each display device.

Option "TwinViewOrientation" "*string*"

Default: `Option "TwinViewOrientation" "RightOf"`

Valid values are:

- `"RightOf"`
- `"LeftOf"`
- `"Above"`
- `"Below"`
- `"Clone"`

This option controls the positioning of the second display device relative to the first within the virtual X screen, when offsets are not explicitly given in the MetaModes.

Note: Under the TwinView Clone option, both display devices are assigned offsets of 0,0.

Option "ConnectedMonitor" "*string*"

See "Option "ConnectedMonitor" "string"" on page 31 for application to TwinView.

Enabling and Configuring TV-Out

This sections contains the following topics:

A TV monitor can be connected to an NVIDIA product-based graphics card with a TV-Out (S-Video) connector so that the TV functions as any other display device, such as a CRT or digital flat panel (DFP). The TV can be used by itself, or (on appropriate graphics cards) in conjunction with another display device in a TwinView configuration. (For TwinView features, *see* the previous section "Enabling and Configuring TwinView" on page 37).

If a TV is the only display device connected to your graphics card, the TV will be used as the primary display when you start-up your system; that is, the console will come up on the TV just as if it were a CRT.

XF86Config File Settings

To use your TV with X, note the following settings in your XF86Config file:

- `VertRefresh` and `HorizSync` values in the "`Monitor`" section of the XF86Config file; confirm that the values are appropriate for your television. Values are generally:
  - `HorizSync 30-50`
  - `VertRefresh 60`
- Valid modes for TV in the "`Screen`" section of the XF86Config file are:
  - `640x480`
  - `800x600`
  - `1024x768` (if the TV encoder on your graphics card is a BrookTree 871)

  Note: Your XFree86 log file should indicate the encoder that you have; locate the line: `(--) NVIDIA(0):  TV Encoder detected as`

- "`TV Standard`" is an option that you need to *add* to the "`Screen`" section of your XF86Config file. *See*

- "`ConnectedMonitor`" option (). You can use this option to tell X to use the TV for display. This option is required *only* if your TV is not detected by the graphics card, or you normally use a CRT or digital flat panel as your start-up display but want to redirect X to use the TV.

  Example command line in your XF86Config file:

  ```
  Option "ConnectedMonitor" "TV"
  ```

- "`TVOutFormat`" is an option that you can use to force S-Video or Composite output.

Option "TVOutFormat" "*string*"

Add the "`TVOutFormat`" option to force S-Video or Composite output. Without this option, the driver auto-detects the output format but may not do so correctly. The output format can be forced with *one* of these options:

```
Option "TVOutFormat" "SVIDEO"
```
or
```
Option "TVOutFormat" "COMPOSITE"
```

Option "TVStandard" "*string*"

Add the "`TVStandard`" option to the "`Screen`" section of your XF86Config file. Replace "*string*" with a valid TV output format, as listed in Table 4.2. A sample line in the XF86Config file is:

```
Option "TVStandard" "NTSC-M"
```

If you don't specify a TVStandard, or you specify an invalid value, the *default* "`NTSC-M`" value is used.

Default: "`NTSC-M`"

Note: If your country is not in the list of countries, select the country closest to your location.

Table 4.2   TV Output Formats by Country

TV Output Format	Country Where Used
PAL-B	Belgium, Denmark, Finland, Germany, Guinea, Hong Kong, India, Indonesia, Italy, Malaysia, The Netherlands, Norway, Portugal, Singapore, Spain, Sweden, and Switzerland
PAL-D	China and North Korea

Table 4.2   TV Output Formats by Country (continued)

TV Output Format	Country Where Used
PAL-G	Denmark, Finland, Germany, Italy, Malaysia, The Netherlands, Norway, Portugal, Spain, Sweden, and Switzerland
PAL-H	Belgium
PAL-I	Hong Kong and United Kingdom
PAL-K1	Guinea
PAL-M	Brazil
PAL-N	France, Paraguay, and Uruguay
PAL-NC	Argentina
NTSC-J	Japan
NTSC-M	Canada, Chile, Colombia, Costa Rica, Ecuador, Haiti, Honduras, Mexico, Panama, Puerto Rico, South Korea, Taiwan, United States of America, and Venezuela

Configuring a Laptop

This section contains the following topics:

Standard Functionality

Configuring NVIDIA Linux drivers on a laptop computer is similar to the process used in a desktop environment.

To enable drivers for a laptop under NVIDIA Accelerated Linux Driver Set Release 10 (version 1.0-1251), users were required to pass the option NVreg_Mobile to the NVdriver kernel module using either the modprobe command:

```
modprobe NVdriver NVreg_Mobile=X
```

or by adding the following to the kernel module configuration file (usually either /etc/conf.modules or /etc/modules.conf):

```
options NVdriver NVreg_Mobile=X
```

where X is:

- **1** if using a GeForce2 Go or a Quadro2 Go in a Dell laptop

- **2** if using a GeForce2 Go or a Quadro2 Go in a Satellite 2800 series Toshiba laptop

- **4** if using a GeForce2 Go or a Quadro2 Go in a Satellite 3000 series Toshiba laptop

- **5** if using a GeForce2 Go or Quadro2 Go in a Gateway laptop

For example, for a Toshiba Satellite laptop with GeForce2 Go, you would add the following line to your kernel module configuration file:

```
options NVdriver NVreg_Mobile=2
```

Note: Versions after 1.0-1251, including the current version of the drivers, no longer require the `NVreg_Mobile` option, though it can be used to override what is detected.

TwinView Functionality

Both GeForce2 Go and Quadro2 Go support TwinView. TwinView on a laptop can be configured in the same way as on a desktop computer. (*See* "Enabling and Configuring TwinView" on page 37.)

In a TwinView configuration, when using the laptop's internal flat panel and an external CRT:

- The CRT is the primary display device. Specify the CRT's `HorizSync` and `VertRefresh` values in the "`Monitor`" section of your XF86Config file.

- The flat panel is the secondary display device. Specify the flat panel's `HorizSync` and `VertRefresh` values using the "`SecondMonitorHorizSync`" and "`SecondMonitorVertRefresh`" options. (*See* "Option "SecondMonitorHorizSync" "range(s)"" on page 38 and "Option "SecondMonitorVertRefresh" "range(s)"" on page 39.)

   You can also use the "`UseEdidFreqs`" option to obtain the `HorizSync` and `VertRefresh` values from the EDID of each display device, in which case, you don't have to set these values in your XF86Config file.

   Note: Use this method of obtaining `HorizSync` and `VertRefresh` values *only* if you if you trust the reported EDIDs of your display device. For details, *see* "Option "UseEdidFreqs" "boolean"" on page 34.

Using Hot Keys to Switch Display Devices

Laptop computers using GeForce2 Go can react to an LCD/CRT hot key event, toggling between each of the connected display devices and each possible combination of the connected display devices

Note: Only two display devices may be active at a time.

Note: TwinView, as configured in your XF86Config file, and hot key functionality are mutually exclusive; that is, if you enable TwinView in your XF86Config file, then the NVIDIA X driver ignores hot key events.

The hot key functionality lets you dynamically connect and remove display devices to/from your laptop and hot key to them without restarting X.

When using this feature, it is a good idea to use the "UseEdidFreqs" option so that the HorizSync and VertRefresh values for each display device can be retrieved from its EDID. Otherwise, the "Monitor" section of the XF86Config file will be interpreted differently with each hot key event.

When X is started, or when a change is detected in the list of connected display devices, a new hot key sequence list is constructed. This list contains the display devices that will be used with each hot key event.

When a hot key event occurs, the next hot key state in the sequence is chosen. Each mode requested in the XF86Config file is validated against the constraint of each display device and the resulting modes are made available for that display device.

If multiple display devices are to be active at once, then the modes from each display device are paired together. If an exact match (same resolution) can't be found, then the closest match is used and the display device with the smaller resolution is panned within the resolution of the other display device.

When VT-switching away from X, the VGA console will always be restored on the display device on which it was present when X was started. Similarly, when VT-switching back into X, the same display device configuration will be used as when you VT-switched away from X, regardless of what LCD/CRT hot key activity occurred while VT-switched away.

Non-Standard Modes on LCD Displays

Some users have had difficulty programming a 1400x1050 mode (the native resolution of some laptop LCDs). In version 4.0.3, XFree86 added several 1400x1050 modes to its database of default modes, but if you're using an older version of XFree86, here is a modeline that you can use.

```
# -- 1400x1050 --

# 1400x1050 @ 60Hz, 65.8 kHz hsync

Modeline  "1400x1050" 129   1400  1464  1656  1960

                                  1050  1051  1054  1100  +HSync    +VSync
```

Known Laptop Issues

- Power Management is currently not supported on laptop computers.
- LCD/CRT hot key switching is currently not supported on Toshiba Satellite 2800 series laptops.
- TwinView is currently not supported on Toshiba Satellite 2800 series laptops.
- The video overlay only works on the first display device on which you started X.

  For example, if you perform these steps, the video will not appear on the second display device:

  a Start X on the internal LCD.

  b Run a video application that uses the video overlay (i.e., the "Video Overlay" adaptor advertised through the XV extension).

  c Hot key switch to add a second display device.

  Note: To work around this problem, you can either configure the video application to use the "Video Blitter" adaptor advertised through the XV extension (this is always available) or hot key switch to the display device on which you want to use the video overlay *before* starting X.

5

TROUBLESHOOTING, QUESTIONS, & RESOURCES

This chapter contains the following major topics:

Troubleshooting: General

Here are a few guidelines to keep in mind when troubleshooting problems with the NVIDIA Accelerated Linux Driver Set:

- One of the most useful tools for diagnosing problems is the `XFree86` log file in `/var/log`; the file is named `/var/log/XFree86.<#>.log`.

  - Lines that begin with "`(II)`" are informational.

  - Lines that begin with "`(WW)`" are warnings.

  - Lines that begin with "`(EE)`"are errors.

- Verify that the correct configuration file (i.e., the configuration file you are editing) is being used; find the line that begins with:

  ```
  (==) Using config file:
  ```

- **Important:** Verify that the NVIDIA driver is being used instead of the "nv" driver by locating the line:

  ```
  (II) LoadModule: "nvidia"
  ```

  Lines from the driver should begin with: `(II) NVIDIA(0)`

• By default, the NVIDIA X driver prints relatively few messages to `stderr` and the `XFree86` log file.

If you need to troubleshoot, it may be helpful to enable more verbose output by using the XFree86 command line options "`-verbose`" and "`-logverbose`", which can be used to set the verbosity level for the `stderr` and log file messages, respectively.

The NVIDIA X driver outputs more messages when the verbosity level is at or above 5. (XFree86 defaults to verbosity level 1 for `stderr` and level 3 for the log file.) Therefore, to enable verbose messaging from the NVIDIA X driver to both the log file and `stderr`, you can start X by using the following command:

```
startx -- -verbose 5 -logverbose 5
```

• Nothing will work if the NVdriver kernel module doesn't function properly.

If you see a line in the X log file such as:

```
(EE) NVIDIA(0): Failed to initialize the NVdriver kernel
module!
```

then, most likely, a problem exists with the NVdriver kernel module. Follow these troubleshooting guidelines:

a Verify that, if you installed from RPM, that the RPM was built specifically for the kernel you are using.

b Check that the module "`/sbin/lsmod`" is loaded; if it is not loaded, try loading it explicitly with "`insmod`" or "`modprobe`". (Be sure to exit the X Server before installing a new kernel module.)

If you receive errors about unresolved symbols, the kernel module has most likely been built using header files for a different kernel revision than what you are running. You can explicitly control the kernel header files that are used by building NVdriver from the `NVIDIA_kernel tar` file by using the command:

```
make install SYSINCLUDE=/path/to/kernel/headers
```

Note: The convention for the location of kernel header files is in a state of transition, as is the location of kernel modules. If the kernel module

fails to load properly, `modprobe`/`insmod` may be attempting to load an older kernel module, assuming you've upgraded. Changing to the directory with the new kernel module and using the command "`insmod ./NVdriver`" may help.

Note: The NVdriver may print error messages to indicate a problem. To view these messages, check `/var/log/messages` or the location where syslog is directed to place kernel messages.

- **If X starts, but OpenGL causes problems**, most likely a problem exists with other libraries that are in the way or the presence of obsolete symlinks. (*See* "Installed Components" on page 24.)

  Note: You may be able to fix this problem by rerunning "`ldconfig`".

  a Use "`xdpyinfo`" to verify that the following (correct) extensions are present:

  GLX

  NV-GLX

  NVIDIA-GLX

  If these three extensions are not present, then there is most likely a problem with the loading of the `glx` module or its inability to implicitly load `GLcore`.

  b Check your XF86Config file to verify that you are loading `glx`. (*See* "Editing Your XF86Config File" on page 22.)

  If your XF86Config file is correct, then check the XFree86 log file for warnings and errors pertaining to GLX. Also check that all of the necessary symlinks are in place. (*See* "Installed Components" on page 24.).

- **If you are trying to install/upgrade by SRPM and the command**

  `rpm --rebuild NVIDIA_kernel-__DRIVER_VERSION__.src.rpm`

  only prints out a list of RPM command line options, then you probably do not have the RPM Development packages installed.

  In most cases you can resolve this problem by installing the RPM Development package for your distribution. Alternatively, you can install/upgrade by tar file since the tar files do not require RPM.

- **If you are trying to install/upgrade by SRPM and the command:**

  `rpm --rebuild NVIDIA_kernel-__DRIVER_VERSION__.src.rpm`

  only prints out a list of RPM command line options, then you probably do not have the RPM development packages installed. In most situations you can fix this problem by installing the RPM development package for your

distribution. Alternatively, you can install/upgrade by tar file as the tar files don't require rpm.

- **If you are trying to install/upgrade by SRPM and the command:**

  ```
  rpm --rebuild NVIDIA_kernel-__DRIVER_VERSION__.src.rpm'
  ```

  reports and error of:

  ```
  NVIDIA_kernel-.src.rpm:no such file or directory
  ```

  Then you need to install the rpm-build package for your distribution. Alternatively, you can install/upgrade by tar file as the tar files don't require rpm.

- **If installing the NVIDIA_kernel module results in an error message** such as the following:

  ```
  #error Modules should never use kernel-headers system headers
  #error but headers from an appropriate kernel-source
  ```

  then you need to install the source for the Linux kernel. In most cases, your can fix this problem by installing the kernel-source package for your distribution.

- **If your OpenGL applications exit with the following error message**:

  ```
  Error: Could not open /dev/nvidiactl because the permissions
  are too restrictive. Please see the TROUBLESHOOTING section of
  /usr/share/doc/NVIDIA_GLX-1.0/README for steps to correct.
  ```

  then it is likely that a security module for the Linux-PAM (Pluggable Authentication Modules for Linux) system may be changing the permissions on the NVIDIA device files.

  To correct the problem, it is recommended that you disable this security feature by following the steps below.

  a Different Linux distributions have different files to control.

     If your system has the file /etc/security/console.perms, you need to edit the file by removing the line that starts with "<dri>".

     If your system has the file /etc/logindevperms, you need to edit the file by removing the line that lists /dev/nvidiactl.

     (The above steps will prevent the PAM security system from modifying the permissions on the NVIDIA device files.)

  b Next, you need to reset the permissions on the device files to their original settings and owner by using the following commands:

     ```
     chmod 0666 /dev/nvidia* chown root /dev/nvidia*
     ```

- **If your OpenGL applications crash and print out the following warning:**

  ```
  WARNING: Your system is running with a buggy dynamic loader.
  This may cause crashes in certain applications. If you
  experience crashes you can try setting the environment variable
  __GL_SINGLE_THREADED. For more information please consult
  (sec-04) TROUBLESHOOTING in the file /usr/share/doc/NVIDIA_GLX-
  1.0/README.
  ```

  then the dynamic loader on your system has a software problem, which will cause applications to crash under these conditions:

  - The application is linked with the `pthreads` library *and*

  - The application calls the `dlopen()` function to open the library `libGL.so` multiple times during its execution *and*

  - The system has a version of the Linux loader with the software problem.

  This problem is present in older versions of the dynamic loader. Distributions that shipped with this loader include but are not limited to RedHat Linux 6.2 and Mandrake Linux 7.1.

  If the crashing application is single-threaded then setting the environment variable `__GL_SINGLE_THREADED` to any value will prevent the crash.

  - In the `bash` shell, use the `export__GL_SINGLE_THREADED` command**.**

  - In `csh` and derivatives, use the `setenv __GL_SINGLE_THREADED` command**.**

  Previous releases of the **NVIDIA Accelerated Linux Driver Set** attempted to work around this problem; however, the workaround caused problems with other applications and was removed after version 1.0-1541.

- **If Quake3 crashes when changing video modes**, then you are probably experiencing the problem described above. Please check the text output for the "WARNING" message shown previously.

  To resolve the problem, before running Quake3, set `__GL_SINGLE_THREADED` as described above.

Troubleshooting: ALi Chipset Users

The NVIDIA Driver for Linux, Version 0.9-3, fixed the majority of problems with ALi chipsets, though some problems still exist. The following tips help stabilize problematic systems.

Tips:
- Disable TURBO AGP MODE in the BIOS.
- When using a P5A, upgrade to BIOS Revision 1002 BETA 2.
- When using 1007, 1007A, or 1009, adjust the IO Recovery Time to 4 cycles.
- AGP is disabled by default on some ALi chipsets (e.g, ALi1541 and ALi1647) to work around severe system stability problems (e.g., timing and signal integrity issues) that occur when using these chipsets. You can force AGP *on* for these chipsets by enabling NVreg_EnableALiAGP in os-registry.c and rebuilding the NVIDIA kernel module. (*See* the comments for NVreg_EnableALiAGP in os-registry.c to force AGP on.).

  Note: Forcing AGP *on* may cause the driver to become unstable on these chipsets.

Troubleshooting: NVIDIA TNT Users

The NVIDIA Linux Driver (Version 0.9-3) corrected the problem with SGRAM/SDRAM TNT products.

Note: In the rare case that your NVIDIA TNT-based card has the wrong BIOS installed, the driver will fail.

If the driver fails, follow these steps:

1 Watch your monitor as the system starts up. The very first, brief screen will identify the type of video memory for your card, which will be either SGRAM or SDRAM.

2 Obtain the most recent NVIDIA_kernel tar file.

3 Edit the file os-registry.c from the kernel module sources.

4 Locate the variable NVreg_VideoMemoryTypeOverride and set the value of this variable to the type of memory you have. (Numerically, see the line just above it).

5 Since this variable is not normally used, change the "#if 0" (about 10 lines above the variable) to "#if 1".

6 Rebuild and reinstall the new driver using the "`make`" command.

Frequently Asked Questions - General

Question: When I start X, it fails and my XFree86 log file contains:

```
II) LoadModule: "nvidia"
II) Loading /usr/X11R6/lib/modules/drivers/nvidia_drv.o
No symbols found in this module
EE) Failed to load /usr/X11R6/lib/modules/drivers/nvidia_drv.o
(II) UnloadModule: "nvidia"
EE) Failed to load module "nvidia" (loader failed, 256) ...
(EE) No drivers available.
```

Answer: The `nvidia_drv.o` X driver has been stripped of needed symbols; some versions of RPM (incorrectly) strip object files while installing. You probably need to upgrade your version of RPM. Or, you can install the `NVIDIA_GLX` package from tar file.

———

Question: Why doesn't the NVdriver work with DevFS?

Answer: DevFS will be supported in a future NVIDIA release. In the meantime, you will need to recreate the NVIDIA device nodes after each reboot. Several patches have been suggested by users to make NVdriver DevFS-aware. You may want to try one of these patches; a Web search can provide you several options.

———

Question: My system runs but seems unstable. What's wrong?

Answer: You might be using the wrong AGP module. (*See* "Configuring AGP" on page 36 for details.)

———

Question: The kernel module doesn't get loaded dynamically when X starts; I always have to do "`modprobe NVdriver`" first. What's wrong?

Answer: Verify that the line "`alias char-major-195 NVdriver`" appears in your module configuration file, which usually is <u>one</u> of the following:

- `/etc/conf.modules`
- `/etc/modules.conf`
- `/etc/modutils/alias`

For details, consult the documentation that came with your distribution.

Question: I can't build the NVdriver kernel module, or I can build the NVdriver kernel module but `modprobe/insmod` **fails to load the module into my kernel. What's wrong?**

Answer: These problems are generally caused by the build using the wrong kernel header files; i.e., header files for a different kernel version than the one you are running.

The previous convention was that kernel header files are stored in:

`/usr/include/linux/`

but that is being deprecated in favor of:

`/lib/modules/'uname -r'/build/include.`

The `NVIDIA_kernel Makefile` can determine the location on your system. However, if you encounter a problem, you can force the build to use certain header files by using this command:

`make SYSINCLUDE=/path/to/kernel/headers`

Of course, for this process to work, you need the appropriate kernel header files installed on your system.

Note: *Consult the documentation that came with your distribution;* some distributions do not install the kernel header files by default, or they install headers that do not work properly with the kernel you are running.

Question: Why do OpenGL applications run so slowly?

Answer: The application is probably using a different library (that still remains on your system) instead of the NVIDIA-supplied OpenGL library.

For details, *see* "Installed Components" on page 24.

Question: Why are there are problems running Quake2.

Answer: Quake2 requires minor setup to get it started.

a First, in the Quake2 directory, the installation program creates a symlink called "`libGL.so`" that points to "`libMesaGL.so.`" *Remove* or *rename* this symlink.

b Then, to run Quake2 in OpenGL mode, use this command:

`quake2 +set vid_ref glx +set gl_driver libGL.so`

Quake2 does not seem to support any kind of full-screen mode but you can run your X Server at the resolution on which Quake2 runs in order to emulate full-screen mode.

Question: There are problems running Heretic II. What can I do?

Answer: Heretic II also installs, by default, a symlink called libGL.so in the application directory. You can remove or rename this symlink, since the system will then find the default `libGL.so`, which the NVIDIA drivers install in `/usr/lib`. From within Heretic II, you can then set your render mode to OpenGL in the video menu.

A patch is also available for Heretic II from lokigames at:

```
http://www.lokigames.com/products/heretic2/updates.php3
```

_____

Question: Where can I get `gl.h` or `glx.h` so I can compile OpenGL programs?

Answer: Most systems come with these headers installed. However, NVIDIA has provided its own `gl.h` and `glx.h` files in case your system does not have them or in case you want to develop OpenGL applications that use the new NVIDIA OpenGL extensions. These files have been installed in:

```
/usr/share/doc/NVIDIA_GLX-1.0/usr/include/GL
```

to avoid conflicting with the system-installed versions. To use these headers, copy them to `/usr/include/GL`.

_____

Question: Can I receive E-mail notification of new NVIDIA Accelerated Linux Driver Set releases?

Answer: Yes. Complete the form at:

```
http://www.nvidia.com/view.asp?FO=driver_update
```

_____

Question: Why does my system hang when VT-switching if I have `rivafb` enabled?

Answer: Currently, you cannot use both `rivafb` and the NVdriver kernel module simultaneously due to software limitations.

Note: In general, using two independent software drivers to drive the same piece of hardware _is not recommended_.

_____

Question: Why do I receive the following error message when compiling the NVdriver kernel module?

> You appear to be compiling the NVdriver kernel module with a
> compiler different from the one that was used to compile the
> running kernel. This may be perfectly fine, but there are cases
> where this can lead to unexpected behavior and system crashes.
>
> If you know what you are doing and want to override this check,
> you can do so by setting IGNORE_CC_MISMATCH.
>
> In any other case, set the CC environment variable to the name
> of the compiler that was used to compile the kernel.

Answer: You should compile the NVdriver kernel module with the same
compiler version that was used to compile your kernel. Some Linux kernel data
structures are dependent on the version of gcc used to compile them; for
example, in `include/linux/spinlock.h`:

```
.....
* Most gcc versions have a nasty bug with empty initializers.
*/
#if (__GNUC__ > 2)
typedef struct { } rwlock_t;
#define RW_LOCK_UNLOCKED (rwlock_t) { }
#else
typedef struct { int gcc_is_buggy; } rwlock_t;
#define RW_LOCK_UNLOCKED (rwlock_t) { 0 }
#endif
```

If the kernel is compiled with gcc 2.*x*, but gcc 3.x is used when the open files in
NVdriver are built (or vice versa), the size of `rwlock_t` will vary and internal
function calls such as `ioremap` will fail.

To verify the version of gcc that was used to compile your kernel, you can
examine the output of:

```
cat /proc/version
```

To verify the version of gcc that is currently in your `$PATH`, you can examine the
output of:

```
gcc -v
```

_____

Question: Why does X fail with error `"Failed to allocate LUT context
DMA"`?

Answer: This is one of the possible consequences of compiling NVdriver with a
different gcc version than used to compile the Linux kernel (see above).

Question: NVIDIA_kernel RPMs are not available for release N from <insert your favorite distro here>. I tried installing the RPM for version N-1, but that didn't work. What should I do?

Answer: As explained in , if a specific NVIDIA_kernel RPM is not available for your distribution, then use either the NVIDIA_kernel SRPM or tar file.

_____

Question: The following messages appear when I install the NVIDIA_GLX package.

```
The above file(s) possibly belong to a conflicting MESA rpm.
--- They have been renamed to xxx.<originalFile>.RPMSAVE to
--- avoid conflicting with the files contained within this
--- package.
--- Please see /usr/share/doc/NVIDIA_GLX-1.0/README, section
--- "(sec-05) FREQUENTLY ASKED QUESTIONS" for more details.
```

What's wrong?

Answer: This is an informational message stating that conflicting files have been moved aside to ensure that your applications find the newly installed OpenGL libraries. If you uninstall the NVIDIA_GLX package, the original files will be automatically restored.

Frequently Asked Questions: TwinView

Question: Nothing gets displayed on my second monitor; what's wrong?

Answer: Monitors (including most older monitors) that do not support monitor detection using DDC protocols cannot be detected by your NVIDIA product. You need to explicitly tell the NVIDIA XFree86 driver the type of device that is connected by using the "ConnectedMonitor" option; for example:

```
Option "ConnectedMonitor" "CRT, CRT"
```

For additional details on using this option, *see* .

_____

Question: Will window managers be able to appropriately place windows (i.e., avoid placing windows across both display devices or in inaccessible regions of the virtual desktop)?

Answer: Yes. The NVIDIA X driver provides a Xinerama extension that allows X clients (such as window managers) to call XineramaQueryScreens() to

detect the current TwinView configuration. Note that the Xinerama protocol doesn't provide a way to inform clients when a configuration change occurs. So, if you modeswitch to a different MetaMode, your window manager will continue to detect the previous configuration. By using the Xinerama extension in conjunction with the XF86VidMode extension to get modeswitch events, window managers should be able to determine the TwinView configuration at any given time.

Another solution is to use panning domains to eliminate inaccessible regions of the virtual screen. *See* "Option "MetaModes" "string"" on page 39.

––

Question: Why can't I get a resolution of 1600x1200 on the second display device?

Answer: Since the second display device was designed to be a digital flat panel, the Pixel Clock for the second display device is only 150 MHz. This effectively limits the resolution on the second display device to somewhere around 1280x1024.

Note: For a description of how Pixel Clock frequencies limit the programmable modes, see the *XFree86 Video Timings How To* documentation.

––

Question: Do video overlays work across both display devices?

Answer: Hardware video overlays only work on the first display device. The current solution is that blitted video is used on TwinView configuration.

––

Question: How are virtual screen dimensions determined in TwinView?

Answer: After all requested modes have been validated, and the offsets for the viewport for each MetaMode have been computed, the NVIDIA driver computes the bounding box of the viewports for each MetaMode. The maximum bounding box width and height is then figured.

Note: A side effect of this process is that the virtual width and virtual height may come from different MetaModes. Given the following MetaMode string:

"1600x1200,NULL; 1024x768+0+0, 1024x768+0+768",

the resulting virtual screen size will be 1600 x 1536.

––

Question: Can I play full-screen games across both display devices?

Answer: Yes. While the details of configuration will vary among games, the basic idea is that a MetaMode presents X with a mode whose resolution is the

bounding box of the viewports for that MetaMode. For example, the following lines:

```
Option "MetaModes" "1024x768,1024x768; 800x600,800x600"
Option "TwinViewOrientation" "RightOf"
```

produce two modes: one with a resolution of 2048x768 and another with a resolution of 1600x600. Games such as Quake 3 Arena use the VidMode extension to discover the resolutions of the modes currently available.

To configure Quake 3 Arena to use the above MetaMode string, add the following lines to your q3config.cfg file:

```
seta r_customaspect "1"
seta r_customheight "600"
seta r_customwidth  "1600"
seta r_fullscreen   "1"
seta r_mode         "-1"
```

Note: Given the above configuration, there is no mode with a resolution of 800x600 (remember that the MetaMode 800x600, 800x600 has a resolution of 1600x600), so if you change Quake 3 Arena to use a resolution of 800x600, it will display in the lower left corner of your screen with the rest of the screen grayed out.

To have single-display modes also available, an appropriate MetaMode string could be:

```
"800x600,800x600; 1024x768,NULL; 800x600,NULL; 640x480,NULL"
```

More precise configuration information for specific games is beyond the scope of this document. However, the above examples coupled with numerous online sources can point you in the right direction.

Contacting Us

- The URL for the new Web forum for NVIDIA Linux driver issues is:

  ```
  http://www.nvnews.net/cgi-bin/ultimatebb.cgi?ubb=forum&f=29
  ```

  This is the preferable tool for seeking help; users can post questions, answer other users' questions, and search the archives of previous postings.

- If, all else fails, you can contact NVIDIA for support at:

  ```
  linux-bugs@nvidia.com
  ```

  Note: Please only send email to this address after you've followed the troubleshooting sections in this chapter and requested help on the nvnew.net web forum.

Additional Resources

- **Linux OpenGL ABI**

  `http://oss.sgi.com/projects/ogl-sample/ABI/`

- **NVIDIA Linux HowTo**

  `http://www.linuxdoc.org/HOWTO/mini/Nvidia-OpenGL-Configuration/index.html`

- **OpenGL:** `www.opengl.org`

- **The XFree86 Project:** `www.xfree86.org`

- **#nvidia** `(irc.openprojects.net)`

A

PROGRAMMING MODES

This appendix contains the following major topics:

Introduction

The NVIDIA Accelerated Linux Driver Set supports all standard VGA and VESA modes, most user-written custom mode lines, as well as double-scan and interlaced modes.

Note: Double-scan modes are supported on all hardware, and interlaced modes are supported on GeForce 256, GeForce DDR, Quadro, GeForce2 GTS/ GeForce2 Pro, GeForce2 Ti, GeForce2 Ultra, Quadro2 Pro, and all TNT products.

In general, your display device (such as a CRT, DFP, or TV) can be a greater constraint on usable modes than either your NVIDIA product-based graphics card *or* the NVIDIA Accelerated Linux Driver Set.

To request one or more standard modes for use in X, you can add a "`Modes`" line, as shown below, in the appropriate "`Display`" subsection of your XF86Config file:

```
Modes "1600x1200" "1024x768" "640x480"
```

Refer to the XF86Config(4/5) man page for further details.

Note: The documentation that follows is primarily of interest if you create your own custom mode lines, experiment with `xvidtune(1)`, or are simply interested in further knowledge. This document is neither an explanation nor a guide to crafting custom mode lines for XFree86, which is offered in documents such as the *XFree86 Video Timings HowTo*; refer to the `www.linuxdoc.org` web site.

Depth, Bits Per Pixel, and Pitch

The bits used per pixel is an important issue when considering the maximum programmable resolution, though not a direct concern when programming modes. A discussion of the terms "depth" vs. "bits per pixel" follows.

Depth is the number of bits of data are stored per pixel. Supported depths are 8, 15, 16, and 24. Most video hardware, however, stores **pixel data** in sizes of 8, 16, or 32 bits; this is the amount of memory allocated per pixel. When you specify the depth, X selects the **bits per pixel (bpp)** size in which to store the data. Table A.1 lists the bits per pixel used for each supported depth..

Table A.1   Bits Per Pixel Used for Depth

Depth	Bits Per Pixel
8	8
15	16
16	16
24	32

The **pitch** is the number of bytes in the linear frame buffer between the data of one pixel and the data of the pixel immediately below that one. Pitch can be represented by the formula:

Pitch = HR * (bpp/8)

> **Pitch =** Horizontal Resolution (HR) multiplied by the Bytes Per Pixel
> **Bytes Per Pixel =** Bits Per Pixel (bpp) divided by 8

Note: In practice, the pitch may be greater than this product because video hardware often requires the pitch to be a multiple of a certain value.

Maximum Resolutions

The NVIDIA Accelerated Linux Driver Set and NVIDIA product-based graphics cards support resolutions up to 2048x1536, though the maximum resolution that your system can support is also limited by the amount of video memory (referenced in Useful Formulas in the next section) and the maximum supported resolution of your display device (i.e., CRT, DFP, or TV).

Note: While use of a video overlay does not limit the maximum resolution or refresh rate, video memory bandwidth used by a programmed mode does affect the quality of the overlay.

Useful Formulas

Video Memory Used

The maximum resolution is a function of the amount of video memory *and* the bits per pixel (bpp) that you want to use, as represented by the formula:

Amount of Video Memory Used = HR * VR * (bpp/8)

Amount of Video Memory Used =	Horizontal Resolution (HR) multiplied by the Vertical Resolution (VR) multiplied by the Bytes Per Pixel
Bytes Per Pixel =	Bits Per Pixel (bpp) divided by 8

Note: Technically, the Video Memory Used = (Pitch * Vertical Resolution) where the Pitch may be slightly greater than "HR * (bpp/8)" to accommodate hardware requirements that specify the pitch to be a multiple of a certain value.

This discussion *only* refers to memory usage for the frame buffer; video memory is also used by other operations, such as OpenGL or pixmap caching.

Resolution, Pixel Clock, and Vertical Refresh Rate

The relationship between resolution, pixel clock (i.e., dot clock), and vertical refresh rate can be described by this formula:

RR = PCLK / (HFL * VFL)

Refresh Rate (RR) =	Pixel Clock (PCLK) divided by the Total Number of Pixels
Total Number of Pixels =	Horizontal Frame Length (HFL) multiplied by the Vertical Frame Length (VFL).

Note: The frame length in VFL refers to the actual frame length and not simply the visible resolution.

As described in the *XFree86 Video Timings HowTo* documentation, the above formula may be rewritten as:

PCLK = RR * HFL * VFL

Given a maximum pixel clock, you can adjust the **RR**, **HFL** and **VFL** as needed.

The pixel clock is reported in the log file when you run X with verbose logging, as shown below:

```
startx -- -logverbose 5
```

Your XFree86.0.log should contain several lines (as shown below) to indicate the maximum pixel clock at each bit per pixel size.

```
(--) NVIDIA(0): Display Device 0: maximum pixel clock at  8 bpp:
350 MHz
(--) NVIDIA(0): Display Device 0: maximum pixel clock at 16 bpp:
350 MHz
(--) NVIDIA(0): Display Device 0: maximum pixel clock at 32 bpp:
300 MHz
```

Mode Validation

During the pre-initialization phase of the X Server, the NVIDIA X driver validates all requested modes using these steps:

1 Takes the intersection of the HorizSync and VertRefresh ranges provided by the user in the XF86Config file with the ranges reported by the monitor in the EDID. This function can be disabled by using the "IgnoreEDID" option, in which case the X driver will accept the HorizSync and VertRefresh provided by the user.

2 Calls the xf86ValidateModes() helper function, which finds modes with the user-specified names in the XF86Config file, pruning out modes with:

- invalid horizontal sync frequencies or vertical refresh rates,

- pixel clocks larger than the maximum pixel clock for the graphics card, and

- resolutions larger than the virtual screen size, if a virtual screen size was specified in the XF86Config file).

Several other constraints are applied to this mode search. (*See* xc/programs/Xserver/hw/xfree86/common/xf86Mode.c:xf86ValidateModes().)

3 All modes returned from `xf86ValidateModes()` are then examined to ensure that their resolutions are not greater than the highest mode reported by the monitor's EDID. (As explained earlier, this function may be disabled with the `"IgnoreEDID"` option.)

If the display is a TV, each mode is checked to ensure that it has a resolution that is supported by the TV encoder. Usually, only resolutions of 800x600 and 640x480 are supported by the encoder.

4 All remaining modes are then checked to ensure sure they pass the constraints described in "Additional Mode Constraints" on page 66.

Note: In order to catch potentially invalid modes submitted by `XF86VidModeExtension` (e.g., `xvidtune(1)`), the last two steps are also performed when each mode is programmed.

Note: Under TwinView configuration, the above validation is performed for the requested modes for each display device.

Additional Mode Constraints

Following is a list of additional constraints on the mode parameters:

• The Horizontal Resolution (HR) must be a multiple of 4 *and* less than or equal to 2048.

• The Horizontal Blanking Width (HBW) must be a multiple of 4 *and* less than or equal to 1024. The formula is:

HBW = max(HFL,HSE) – min(HR,HSS)

max =	maximum of	
HFL, =	Horizontal Frame Length *and*	
HSE =	Horizontal Sync End	
– =	*minus*	
min =	minimum of	
HR, =	Horizontal Resolution *and*	
HSS =	Horizontal Sync Start	

• The Horizontal Sync Start (HSS) value must be a multiple of 4 *and* be less than or equal to 4088.

• The Horizontal Sync Width (HSW) must be a multiple of 4 *and* less than or equal to 256. The formula is:

HSW = HSE – HSS

where:

HSE	=	Horizontal Sync End
–	=	*minus*
HSS	=	Horizontal Sync Start

- The **Horizontal Frame Length** (**HFL**) must be a multiple of 4 *and* less than or equal to 4128 *and* greater than or equal to 40.

- The **Vertical Resolution (VR)** must be less than or equal to 2048.

- The Vertical Blanking Width (VBW) must be less than or equal to 128. The formula is:

  VBW = max(VFL,VSE) – min(VR,VSS)

max	=	maximum of
VFL,	=	Vertical Frame Length *and*
VSE	=	Vertical Sync End
–	=	*minus*
min	=	minimum of
VR,	=	Vertical Resolution *and*
VSS	=	Vertical Sync Start

- The Vertical Sync Start (VSS) value must be less than or equal to 2047.

- The Vertical Sync Width (VSW) must be less than or equal to 16. The formula is:

  VBW = VSE - VSS

VSE	=	Vertical Sync End
–	=	*minus*
VSS	=	Vertical Sync Start

- The **Vertical Frame Length (VFL)** must be less than or equal to 2049 *and* greater than or equal to 2.

Example Mode Line

The following is an example mode line that contains each abbreviation that is used in the constraints listed above:

```
# Custom Mode line for the SGI 1600SW Flatpanel
#       name        PCLK HR   HSS  HSE  HFL  VR   VSS  VSE  VFL
Modeline "sgi1600x1024" 106.9 1600 1632 1656 1672 1024 1027 1030
1067
```

Note: An XFree86 modeline generator conforming to the GTF Standard has been posted to the XFree86 Xpert mailing list:

http://www.xfree86.org/pipermail/xpert/2001-October/
012070.html

For additional modeline generators, search for "Modeline" on
freshmeat.net.