# Nvidia Binary Driver User Guide for Linux

## Table of Contents

## Introduction

During OS installation, the installer will attempt to use the open source driver (nv) that is included in the X Window package. This is a high quality 2D driver suitable for most desktop use. However, there are several cases where the Nvidia binary driver is required:

- The display adapter is unsupported by the current X Window driver

- More than one display is connected to the adapter

- The card is driving an LCD display with a resolution of 1600x1200 or higher

In these cases, the Nvidia binary driver will be pre-installed at the factory. If the operating system is re-installed, the Nvidia driver will need to be installed afterward. The Nvidia driver must also be installed if 3D acceleration is needed. For convenience, the Nvidia binary driver is included on the ASL Driver CD. As of this writing, the latest version is 260.19.29.

⚠ **Caution**
Whenever the Linux kernel is updated, the Nvidia driver must be re-installed following the instructions in this guide. Failure to do this will result in a non-functional X server.

## Installation

To install, upgrade or re-install the Nvidia binary driver, use the following steps:

1. Log in as the root user, or use the **su** command to become root.

2. If the X Server is running, terminate it by running the following command from an X terminal or console:

```
# init 3
```

3. Verify that the kernel source or development package has been installed. If necessary, install the appropriate package before proceeding to step 5. Note that the kernel source or development package must match the running kernel. If the package is missing or has the incorrect version, the Nvidia driver installation will fail.

    To verify that the correct package is present, run the following command:

```
# rpm -qa | grep kernel
```

and search for one of the following packages:

- kernel-source-*kernel version*

- kernel-sourcecode-*kernel version*

- kernel-smp-devel-*kernel version*

- kernel-devel-*kernel version*

4. Verify that the C/C++ Compiler and Tools package has been installed. This package is required to compile the Nvidia module. If necessary, install the compiler package before proceeding to step 5.

5. Mount the ASL Driver CD and install the Nvidia binary display driver.

   For 32-bit Linux (use ASL Driver CD):

   ```
   # mount /dev/cdrom /mnt
   # cd /mnt/ASL-packages
   # sh NVIDIA-Linux-x86-260.19.29.run
   ```

   For 64-bit Linux (use ASL Driver CD):

   ```
   # mount /dev/cdrom /mnt
   # cd /mnt/ASL-packages
   # sh NVIDIA-Linux-x86_64-260.19.29.run
   ```

   Under SuSE 10.0, the driver disk is automatically mounted when inserted into the drive. In addition, the normal installer interface does not work properly. To address these differences, the installer should be started as shown below:

   32-bit:

   ```
   # cd /media/ASL_DRIVER/ASL-packages
   # sh NVIDIA-Linux-x86-260.19.29-pkg1.run --ui=none
   ```

   64-bit:

   ```
   # cd /media/ASL_DRIVER/ASL-packages
   # sh NVIDIA-Linux-x86_64-260.19.29-pkg2.run --ui=none
   ```

   During installation, the software based OpenGL libraries are replaced with a 3D accelerated version. In addition, the installer compiles a custom driver module (nvidia) used for accelerated operation. In the 64-bit version, the package includes 32-bit compatible OpenGL libraries. When prompted, select 'Yes' to install them.

   Starting with release 1.0-8174, the installer can update the existing X configuration file to use the Nvidia binary driver. However, this step should be skipped. Instead, use the ASL predefined X configuration file as described in step 6.

6. Copy a predefined X configuration file from the ASL Driver CD.

**Note**

If the driver is being upgraded or re-installed due to a kernel upgrade, it is not necessary to perform this step.

For the following distributions, use the ASL Driver CD and copy the configuration file from `APPS/Xorg/config` to `/etc/X11/xorg.conf`:

### Table 1. Xorg configuration files

| Distribution | Filename |
|---|---|
| Fedora Core 2-5, RHEL4, CentOS 4.X | xorg.nvidia.conf |
| | xorg.nvidia.conf.SM243T (Samsung SM243T LCD) |
| Fedora Core 6, RHEL 5.X, CentOS 5.X | xorg.nvidia.conf.fc6 |
| Fedora 7 | xorg.nvidia.conf.f7 |
| Fedora 8 | xorg.nvidia.conf.f8 |
| Fedora 9 | xorg.nvidia.conf.f9 |
| Fedora 10 | xorg.nvidia.conf.f10 |
| SuSE 10.0 | xorg.nvidia.conf.suse10 |
| SuSE 10.1 | xorg.nvidia.conf.suse10.1 |

For the following distributions, use the ASL Driver CD and copy the configuration file from `APPS/XFree86/config` to `/etc/X11/XF86Config`:

### Table 2. XFree86 configuration files

| Distribution | Filename |
|---|---|
| RHEL 3, CentOS 3 | XF86Config-4.nvidia.twinview |
| | XF86Config-4.nvidia.twinview-1600_1200 |
| | XF86Config-4.nvidia |

# Configuration

**Caution**

Because the Nvidia binary driver is not supported or included in any Linux distributions, do not use system configuration tools to setup the display. To modify the display such as changing the screen resolution or color depth, edit the X configuration file manually.

Using DDC-enabled monitors, the Nvidia binary driver can configure the screen at the optimal refresh rate automatically. For older monitors that do not have DDC, it is necessary to enter the monitor horizontal sync rate (Khz) and vertical refresh rate (Hz) manually.

To use DDC, comment out the lines below in the X configuration file:

```
#    HorizSync  31.5-64.3
#    VertRefresh 50-160
```

```
#     Option        "SecondMonitorHorizSync" "31.5-64"
#     Option        "SecondMonitorVertRefresh" "50-160"
```

To specify the monitor frequency range manually, remove the # symbol. For example:

```
    HorizSync  31.5-64.3
    VertRefresh 50-160
    Option        "SecondMonitorHorizSync" "31.5-64"
    Option        "SecondMonitorVertRefresh" "50-160"
```

HorizSync and VertRefresh apply to the primary display (first monitor). The secondary display (second monitor) uses the values of SecondMonitorHorizSync and SecondMonitorVertRefresh.

Details about the customization of the display driver can be found at the link below:

Nvidia README [ftp://download.nvidia.com/XFree86/Linux-x86/260.19.29/README/README.txt]

After the Nvidia binary driver has been installed successfully and the X configuration file has been created, test the X server by running **startx**. While loading the X Server, the Nvidia logo will be displayed briefly. That is a quick way to verify that the Nvidia binary driver is being i used.

# Special Notes

- Whenever a new kernel is installed, it is necessary to re-install the Nvidia binary driver. Otherwise X will not work. After installing a new kernel, set the system runlevel to 3 before rebooting the system. After the system has booted with the new kernel, then re-install the Nvidia binary driver. Afterward, set the system runlevel back to 5 if preferred. To change the runlevel, edit /etc/inittab and change the value of id.

- If the display is not rendered correctly, check the X log file (/var/log/XFree86.0.log or /var/log/Xorg.0.log). It contains useful information for debugging purposes.

- If the monitor has two connectors (DVI and DB-15), only use one connector when hooking up to the video card.

- Changes in the X configuration file do not take affect until the X Server is restarted. To restart the X Server, press **Ctrl**+**Alt**+**Backspace** simultaneously. Restarting the X Server will terminate all existing X applications. Thus it is important to save their data first.

- By default, the predefined X configuration files disable 3D acceleration to improve system stability. To enable 3D acceleration, edit the X configuration file (/etc/X11/XF86Config or /etc/X11/xorg.conf) and change the value of NvAGP from 0 to 1.

- When using the Nvidia binary driver, do not upgrade or re-install the XFree86/Xorg Mesa development package(s) afterward. That will overwrite the Nvidia OpenGL libraries causing X applications to crash. During the online update, always skip the package(s) below:

### Table 3. Mesa packages

| Distribution | Package name(s) |
|---|---|
| RHEL 3, CentOS 3 | XFree86-Mesa-libGL |
| | XFree86-Mesa-libGLU |
| SuSE 10 | Xorg-x11-Mesa-devel |

| Distribution | Package name(s) |
|---|---|
| | Xorg-x11-Mesa-devel-32-bit |
| SuSE 10.1 | Mesa-devel |
| | Mesa-devel-32-bit |
| Fedora Core 2-4 | Xorg-x11-Mesa-libGL |
| Fedora Core 5-10, RHEL 5.X, CentOS 5.X | mesa-libGL |
| | mesa-libGL-devel |
| | mesa-libGLU |
| | mesa-libGLU-devel |
| RHEL 4, CentOS 4.x | Xorg-X11-Mesa-libGL |
| | Xorg-X11-Mesa-libGLU |

# Linux Distribution Notes

• Always skip the Nvidia package when running online update through **Yast** or **Yast2** under SuSE Linux. The SuSE version is always older and should be skipped.

• Red Hat Enterprise Linux 3 requires kernel 2.4.21-15EL or higher for systems running with 8GB or more of memory. Using an older kernel will cause the system to lock up when starting the X Server.

# Technical Support

Prior to contacting technical support, please visit the ASL online FAQ for quick answers to common hardware and software issues.The answers may be searched by keyword or browsed by category. These documents are maintained with up to date information.

FAQ [http://faq.aslab.com]

To request replacement of a defective part (except monitor) on systems under warranty, please send an E-mail to `<techsupport@aslab.com>` with the following information:

• Serial number or invoice number of the system

• Detailed description of the problem

• Shipping address

If the transaction involves multiple systems, ASL highly recommends creating one ticket per system. This will avoid any mixup when the RMA is processed.

Defective monitors will be handled directly by the manufacturer. Here is the contact information:

• Viewsonic: 800 888 8583 (US) or 866-262-1967 (Canada)

• IIyama: 800 594 7480