

NVIDIA 550.67 Display Driver and CUDA 12.4 Installation Guide

April 12, 2024



This page intentionally left blank

1. Introduction

This document assists the user in installing the NVIDIA 550.67 display driver, the RedHawk pagepool driver, and the CUDA 12.4 toolkit.

This installation media may be installed on RedHawk versions 7.5 and newer.

2. Pre-Installation

The system must be properly prepared to install the NVIDIA drivers.

2.1. Ensure that the RedHawk kernel source RPM is installed.

Check for the following package:

```
# sudo rpm -qa ccur-kernel-source
```

Or

```
# sudo apt -qq list ccur-kernel-source
```

2.2. Blacklist the Nouveau drivers:

```
# sudo blscfg --kopt-add=modprobe.blacklist=nouveau
```

Note, older systems and Ubuntu based installations utilize the command `ccur-grub2` instead of `blscfg`.

2.3. Reboot the system.

2.4. Place the system in multi-user mode.

```
# sudo init 3
```

3. Installation

The `install` script will prepare the kernel source and automate the installation of NVIDIA drivers and the CUDA toolkit for RedHawk systems.

Please use one of the following commands to install the appropriate subset of NVIDIA drivers and CUDA.

3.1. Install CUDA toolkit and NVIDIA drivers for all suitable RedHawk kernels.

```
# sudo bash install --cuda --all-kernels
```

3.2. Install CUDA toolkit and NVIDIA driver for the currently booted RedHawk kernel.

```
# sudo bash install --cuda
```

3.3. Install NVIDIA driver for the currently booted RedHawk kernel.

```
# sudo bash install
```

3.4. Install CUDA toolkit.

```
# sudo bash install --cuda --skip-driver
```

NOTES

Installing the NVIDIA driver on 7.x systems may result in CentOS/RHEL kernels failing to work with the nouveau driver. This may be resolved by updating Gnome to the latest version.

Installing the NVIDIA driver on systems with limited space in /tmp will fail to install. Rerun the above commands with the following:

```
# export TMPDIR=/some/place/with/more/storage
```

4. Pagepool Driver

For each NVIDIA driver installation, the RedHawk pagepool driver is also installed. This driver is responsible for pre-allocating memory for usage by the NVIDIA driver. This prevents inter-processor interrupts (IPIs) and TLB flushes from impacting real-time operations.

The RedHawk pagepool driver installation includes two utility commands: `ccur-pagepool` and `pagepool-memory`.

4.1. Pagepool-memory

This is the primary method for monitoring the current pagepool memory usage and to modify the size of the pagepool. The pagepool will grow dynamically with sufficient GPU usage. Please refer to `pagepool-memory -h` for more information on this command.

4.2. Ccur-pagepool

This command modifies NVIDIA kernel module (.ko) files to use the RedHawk pagepool driver. Concurrent Real-Time validates the shipped NVIDIA drivers against the pagepool driver for 100% compatibility.