



NVIDIA CUDA Installation Guide for Linux

Installation and Verification on Linux Systems

Table of Contents

Chapter 1. Introduction.....	1
1.1. System Requirements.....	1
1.2. About This Document.....	3
Chapter 2. Pre-installation Actions.....	4
2.1. Verify You Have a CUDA-Capable GPU.....	4
2.2. Verify You Have a Supported Version of Linux.....	5
2.3. Verify the System Has gcc Installed.....	5
2.4. Verify the System has the Correct Kernel Headers and Development Packages Installed.....	5
2.5. Install MLNX_OFED.....	7
2.6. Choose an Installation Method.....	7
2.7. Download the NVIDIA CUDA Toolkit.....	8
2.8. Address Custom xorg.conf, If Applicable.....	8
2.9. Handle Conflicting Installation Methods.....	8
Chapter 3. Package Manager Installation.....	10
3.1. Overview.....	10
3.2. RHEL 7 / CentOS 7.....	11
3.2.1. Prepare RHEL 7 / CentOS 7.....	11
3.2.2. Local Repo Installation for RHEL 7 / CentOS 7.....	11
3.2.3. Network Repo Installation for RHEL 7 / CentOS 7.....	12
3.2.4. Common Installation Instructions for RHEL 7 / CentOS 7.....	12
3.2.5. Installing a Previous NVIDIA Driver Branch on RHEL 7.....	13
3.3. RHEL 8 / Rocky 8.....	13
3.3.1. Prepare RHEL 8 / Rocky 8.....	13
3.3.2. Local Repo Installation for RHEL 8 / Rocky 8.....	14
3.3.3. Network Repo Installation for RHEL 8 / Rocky 8.....	14
3.3.4. Common Instructions for RHEL 8 / Rocky 8.....	15
3.4. Fedora.....	15
3.4.1. Prepare Fedora.....	15
3.4.2. Local Repo Installation for Fedora.....	15
3.4.3. Network Repo Installation for Fedora.....	15
3.4.4. Common Installation Instructions for Fedora.....	16
3.5. SLES.....	17
3.5.1. Prepare SLES.....	17
3.5.2. Local Repo Installation for SLES.....	17

3.5.3. Network Repo Installation for SLES.....	17
3.5.4. Common Installation Instructions for SLES.....	18
3.6. OpenSUSE.....	18
3.6.1. Prepare OpenSUSE.....	18
3.6.2. Local Repo Installation for OpenSUSE.....	18
3.6.3. Network Repo Installation for OpenSUSE.....	18
3.6.4. Common Installation Instructions for OpenSUSE.....	19
3.7. WSL.....	19
3.7.1. Prepare WSL.....	19
3.7.2. Local Repo Installation for WSL.....	19
3.7.3. Network Repo Installation for WSL.....	20
3.7.4. Common Installation Instructions for WSL.....	20
3.8. Ubuntu.....	20
3.8.1. Prepare Ubuntu.....	21
3.8.2. Local Repo Installation for Ubuntu.....	21
3.8.3. Network Repo Installation for Ubuntu.....	21
3.8.4. Common Installation Instructions for Ubuntu.....	22
3.9. Debian.....	22
3.9.1. Prepare Debian.....	22
3.9.2. Local Repo Installation for Debian.....	23
3.9.3. Network Repo Installation for Debian.....	23
3.9.4. Common Installation Instructions for Debian.....	23
3.10. Additional Package Manager Capabilities.....	24
3.10.1. Available Packages.....	24
3.10.2. Optional 32-bit Packages for Linux x86_64 .deb/.rpm.....	25
3.10.3. Package Upgrades.....	25
3.10.4. Meta Packages.....	26
Chapter 4. Driver Installation.....	28
Chapter 5. NVIDIA Open GPU Kernel Modules.....	29
5.1. CUDA Runfile.....	29
5.2. Debian.....	30
5.3. Fedora.....	30
5.4. RHEL 8 / Rocky 8.....	30
5.5. RHEL 7 / CentOS 7.....	30
5.6. OpenSUSE / SLES.....	30
5.7. Ubuntu.....	31
Chapter 6. Precompiled Streams.....	32
6.1. Precompiled Streams Support Matrix.....	33

6.2. Modularity Profiles.....	33
Chapter 7. Kickstart Installation.....	35
7.1. RHEL8 / Rocky Linux 8.....	35
Chapter 8. Runfile Installation.....	36
8.1. Overview.....	36
8.2. Installation.....	36
8.3. Disabling Nouveau.....	38
8.3.1. Fedora.....	38
8.3.2. RHEL/CentOS.....	38
8.3.3. OpenSUSE.....	38
8.3.4. SLES.....	39
8.3.5. WSL.....	39
8.3.6. Ubuntu.....	39
8.3.7. Debian.....	39
8.4. Device Node Verification.....	39
8.5. Advanced Options.....	40
8.6. Uninstallation.....	41
Chapter 9. Conda Installation.....	42
9.1. Conda Overview.....	42
9.2. Installation.....	42
9.3. Uninstallation.....	42
9.4. Installing Previous CUDA Releases.....	42
Chapter 10. Pip Wheels.....	44
Chapter 11. Tarball and Zip Archive Deliverables.....	46
11.1. Parsing Redistrib JSON.....	47
11.2. Importing Tarballs into CMake.....	48
11.3. Importing Tarballs into Bazel.....	48
Chapter 12. CUDA Cross-Platform Environment.....	49
12.1. CUDA Cross-Platform Installation.....	49
12.2. CUDA Cross-Platform Samples.....	50
Chapter 13. Post-installation Actions.....	51
13.1. Mandatory Actions.....	51
13.1.1. Environment Setup.....	51
13.1.2. POWER9 Setup.....	51
13.2. Recommended Actions.....	53
13.2.1. Install Persistence Daemon.....	53
13.2.2. Install Writable Samples.....	53

13.2.3. Verify the Installation.....	53
13.2.3.1. Verify the Driver Version.....	53
13.2.3.2. Running the Binaries.....	54
13.2.4. Install Nsight Eclipse Plugins.....	55
13.3. Optional Actions.....	55
13.3.1. Install Third-party Libraries.....	56
13.3.2. Install the Source Code for cuda-gdb.....	56
13.3.3. Select the Active Version of CUDA.....	57
Chapter 14. Advanced Setup.....	58
Chapter 15. Frequently Asked Questions.....	62
How do I install the Toolkit in a different location?.....	62
Why do I see "nvcc: No such file or directory" when I try to build a CUDA application?.....	62
Why do I see "error while loading shared libraries: <lib name>: cannot open shared object file: No such file or directory" when I try to run a CUDA application that uses a CUDA library?.....	63
Why do I see multiple "404 Not Found" errors when updating my repository meta-data on Ubuntu?.....	63
How can I tell X to ignore a GPU for compute-only use?.....	63
Why doesn't the cuda-repo package install the CUDA Toolkit and Drivers?.....	63
How do I get CUDA to work on a laptop with an iGPU and a dGPU running Ubuntu14.04?.....	64
What do I do if the display does not load, or CUDA does not work, after performing a system update?.....	64
How do I install a CUDA driver with a version less than 367 using a network repo?.....	65
How do I install an older CUDA version using a network repo?.....	65
Why does the installation on SUSE install the Mesa-dri-nouveau dependency?.....	65
How do I handle "Errors were encountered while processing: glx-diversions"?.....	66
Chapter 16. Additional Considerations.....	67
Chapter 17. Removing CUDA Toolkit and Driver.....	68

Chapter 1. Introduction

CUDA[®] is a parallel computing platform and programming model invented by NVIDIA[®]. It enables dramatic increases in computing performance by harnessing the power of the graphics processing unit (GPU).

CUDA was developed with several design goals in mind:

- ▶ Provide a small set of extensions to standard programming languages, like C, that enable a straightforward implementation of parallel algorithms. With CUDA C/C++, programmers can focus on the task of parallelization of the algorithms rather than spending time on their implementation.
- ▶ Support heterogeneous computation where applications use both the CPU and GPU. Serial portions of applications are run on the CPU, and parallel portions are offloaded to the GPU. As such, CUDA can be incrementally applied to existing applications. The CPU and GPU are treated as separate devices that have their own memory spaces. This configuration also allows simultaneous computation on the CPU and GPU without contention for memory resources.

CUDA-capable GPUs have hundreds of cores that can collectively run thousands of computing threads. These cores have shared resources including a register file and a shared memory. The on-chip shared memory allows parallel tasks running on these cores to share data without sending it over the system memory bus.

This guide will show you how to install and check the correct operation of the CUDA development tools.

1.1. System Requirements

To use CUDA on your system, you will need the following installed:

- ▶ CUDA-capable GPU
- ▶ A supported version of Linux with a gcc compiler and toolchain
- ▶ NVIDIA CUDA Toolkit (available at <https://developer.nvidia.com/cuda-downloads>)

The CUDA development environment relies on tight integration with the host development environment, including the host compiler and C runtime libraries, and is therefore only supported on distribution versions that have been qualified for this CUDA Toolkit release.

The following table lists the supported Linux distributions. Please review the footnotes associated with the table.

Table 1. Native Linux Distribution Support in CUDA 11.7

Distribution	Kernel ¹	Default GCC	GLIBC	GCC ^{2,3}	ICC ³	NVHPC	XLC ³	CLANG	Arm C/C++
x86_64									
RHEL 8.y (y <= 5)	4.18.0-348	8.5.0	2.28	11	2021	22.02	NO	13	NO
RHEL 7.y (y <= 9)	3.10.0-1160	6.x	2.17						
CentOS 7.y (y <= 9)	3.10.0-1160	6.x	2.17						
OpenSUSE Leap 15.y (y <= 3)	5.3.18-57	7.5.0	2.31						
Rocky Linux 8.5	4.18.0-348	8.5.0	2.28						
SUSE SLES 15.y (y <= 3)	5.3.18-57	7.5.0	2.31						
Ubuntu 22.04	5.15.0-25	11.2.0	2.35						
Ubuntu 20.04.z (z <= 4)	5.13.0-30	9.3.0	2.31						
Ubuntu 18.04.z (z <= 6)	5.4.0-89	7.5.0	2.27						
Debian 11.3	5.10.0-13	10.2.1	2.31						
Fedora 35	5.14.10	11.2.1	2.34						
Arm64 sbsa									
RHEL 8.y (y <= 5)	4.18.0-348	8.5.0	2.28	11	NO	22.02	NO	13	21.1
SUSE SLES 15.y (y <= 3)	5.3.18-57	7.5.0	2.31						
Ubuntu 22.04	5.15.0-25	11.2.0	2.35						
Ubuntu 20.04.z (z <= 4)	5.4.0-100	9.3.0	2.31						
Arm64 Jetson									
Ubuntu 20.04.z (z <= 4)	5.10.65-tegra	9.3.0	2.31	10.2	NO	22.02	16.1.x	13	NO
POWER 9									
RHEL 8.y (y <= 5)	4.18.0-348	8.5.0	2.28	11	NO	22.02	16.1.x	13	NO

(1) The following notes apply to the kernel versions supported by CUDA:

- For specific kernel versions supported on Red Hat Enterprise Linux (RHEL), visit <https://access.redhat.com/articles/3078>.

- ▶ For a list of kernel versions including the release dates for SUSE Linux Enterprise Server (SLES) is available at <https://www.suse.com/support/kb/doc/?id=000019587>.
- ▶ For Ubuntu LTS on x86-64, the Server LTS kernel (e.g. 4.15.x for 18.04) is supported in CUDA 11.7. Visit <https://wiki.ubuntu.com/Kernel/Support> for more information.

(2) Note that starting with CUDA 11.0, the minimum recommended GCC compiler is at least GCC 6 due to C++11 requirements in CUDA libraries e.g. cuFFT and CUB. On distributions such as RHEL 7 or CentOS 7 that may use an older GCC toolchain by default, it is recommended to use a newer GCC toolchain with CUDA 11.0. Newer GCC toolchains are available with the [Red Hat Developer Toolset](#). For platforms that ship a compiler version older than GCC 6 by default, linking to static cuBLAS and cuDNN using the default compiler is not supported.

(3) Minor versions of the following compilers listed: of GCC, ICC, NVHPC and XLC, as host compilers for `nvcc` are supported.

1.2. About This Document

This document is intended for readers familiar with the Linux environment and the compilation of C programs from the command line. You do not need previous experience with CUDA or experience with parallel computation. Note: This guide covers installation only on systems with X Windows installed.



Note: Many commands in this document might require *superuser* privileges. On most distributions of Linux, this will require you to log in as root. For systems that have enabled the `sudo` package, use the `sudo` prefix for all necessary commands.

Chapter 2. Pre-installation Actions

Some actions must be taken before the CUDA Toolkit and Driver can be installed on Linux:

- ▶ Verify the system has a CUDA-capable GPU.
- ▶ Verify the system is running a supported version of Linux.
- ▶ Verify the system has gcc installed.
- ▶ Verify the system has the correct kernel headers and development packages installed.
- ▶ Download the NVIDIA CUDA Toolkit.
- ▶ Handle conflicting installation methods.



Note: You can override the install-time prerequisite checks by running the installer with the `-override` flag. Remember that the prerequisites will still be required to use the NVIDIA CUDA Toolkit.

2.1. Verify You Have a CUDA-Capable GPU

To verify that your GPU is CUDA-capable, go to your distribution's equivalent of System Properties, or, from the command line, enter:

```
lspci | grep -i nvidia
```

If you do not see any settings, update the PCI hardware database that Linux maintains by entering `update-pciids` (generally found in `/sbin`) at the command line and rerun the previous `lspci` command.

If your graphics card is from NVIDIA and it is listed in <https://developer.nvidia.com/cuda-gpus>, your GPU is CUDA-capable.

The Release Notes for the CUDA Toolkit also contain a list of supported products.

2.2. Verify You Have a Supported Version of Linux

The CUDA Development Tools are only supported on some specific distributions of Linux. These are listed in the CUDA Toolkit release notes.

To determine which distribution and release number you're running, type the following at the command line:

```
uname -m && cat /etc/*release
```

You should see output similar to the following, modified for your particular system:

```
x86_64  
Red Hat Enterprise Linux Workstation release 6.0 (Santiago)
```

The `x86_64` line indicates you are running on a 64-bit system. The remainder gives information about your distribution.

2.3. Verify the System Has gcc Installed

The `gcc` compiler is required for development using the CUDA Toolkit. It is not required for running CUDA applications. It is generally installed as part of the Linux installation, and in most cases the version of `gcc` installed with a supported version of Linux will work correctly.

To verify the version of `gcc` installed on your system, type the following on the command line:

```
gcc --version
```

If an error message displays, you need to install the *development tools* from your Linux distribution or obtain a version of `gcc` and its accompanying toolchain from the Web.

2.4. Verify the System has the Correct Kernel Headers and Development Packages Installed

The CUDA Driver requires that the kernel headers and development packages for the running version of the kernel be installed at the time of the driver installation, as well whenever the driver is rebuilt. For example, if your system is running kernel version 3.17.4-301, the 3.17.4-301 kernel headers and development packages must also be installed.

While the Runfile installation performs no package validation, the RPM and Deb installations of the driver will make an attempt to install the kernel header and development packages if no version of these packages is currently installed. However, it will install the latest version of these packages, which may or may not match the version of the kernel your system is

using. **Therefore, it is best to manually ensure the correct version of the kernel headers and development packages are installed prior to installing the CUDA Drivers, as well as whenever you change the kernel version.**

The version of the kernel your system is running can be found by running the following command:

```
uname -r
```

This is the version of the kernel headers and development packages that must be installed prior to installing the CUDA Drivers. This command will be used multiple times below to specify the version of the packages to install. Note that below are the common-case scenarios for kernel usage. More advanced cases, such as custom kernel branches, should ensure that their kernel headers and sources match the kernel build they are running.



Note: If you perform a system update which changes the version of the linux kernel being used, make sure to rerun the commands below to ensure you have the correct kernel headers and kernel development packages installed. Otherwise, the CUDA Driver will fail to work with the new kernel.

RHEL 7 / CentOS 7

The kernel headers and development packages for the currently running kernel can be installed with:

```
sudo yum install kernel-devel-$(uname -r) kernel-headers-$(uname -r)
```

Fedora / RHEL 8 / Rocky Linux 8

The kernel headers and development packages for the currently running kernel can be installed with:

```
sudo dnf install kernel-devel-$(uname -r) kernel-headers-$(uname -r)
```

If matching kernel-headers and kernel-devel packages are not available for the currently running kernel version, you may need to use the previously shipped version of these packages. See https://bugzilla.redhat.com/show_bug.cgi?id=1986132 for more information.

OpenSUSE / SLES

The kernel development packages for the currently running kernel can be installed with:

```
sudo zypper install -y kernel-<variant>-devel=<version>
```

To run the above command, you will need the variant and version of the currently running kernel. Use the output of the `uname` command to determine the currently running kernel's variant and version:

```
$ uname -r
3.16.6-2-default
```

In the above example, the variant is `default` and version is `3.16.6-2`.

The kernel development packages for the default kernel variant can be installed with:

```
sudo zypper install -y kernel-default-devel=$(uname -r | sed 's/\-default//')
```

WSL

This section does not need to be performed for WSL.

Ubuntu

The kernel headers and development packages for the currently running kernel can be installed with:

```
sudo apt-get install linux-headers-$(uname -r)
```

2.5. Install MLNX_OFED

If you intend to use GPUDirectStorage (GDS), you must install the CUDA package and MLNX_OFED package.

GDS packages can be installed using the CUDA packaging guide. Follow the instructions in [MLNX_OFED Requirements and Installation](#).

GDS is supported in two different modes: GDS (default/full perf mode) and Compatibility mode. Installation instructions for them differ slightly. Compatibility mode is the only mode that is supported on certain distributions due to software dependency limitations.

Full GDS support is restricted to the following Linux distros:

- ▶ Ubuntu 18.04, 20.04
- ▶ RHEL 8.3, RHEL 8.4

2.6. Choose an Installation Method

The CUDA Toolkit can be installed using either of two different installation mechanisms: distribution-specific packages (RPM and Deb packages), or a distribution-independent package (runfile packages).

The distribution-independent package has the advantage of working across a wider set of Linux distributions, but does not update the distribution's native package management system. The distribution-specific packages interface with the distribution's native package management system. It is recommended to use the distribution-specific packages, where possible.



Note: Standalone installers are not provided for architectures other than the `x86_64` release. For both native as well as cross development, the toolkit must be installed using the

distribution-specific installer. See the [CUDA Cross-Platform Installation](#) section for more details.

2.7. Download the NVIDIA CUDA Toolkit

The NVIDIA CUDA Toolkit is available at <https://developer.nvidia.com/cuda-downloads>.

Choose the platform you are using and download the NVIDIA CUDA Toolkit.

The CUDA Toolkit contains the CUDA driver and tools needed to create, build and run a CUDA application as well as libraries, header files, and other resources.

Download Verification

The download can be verified by comparing the MD5 checksum posted at <https://developer.download.nvidia.com/compute/cuda/11.7.0/docs/sidebar/md5sum.txt> with that of the downloaded file. If either of the checksums differ, the downloaded file is corrupt and needs to be downloaded again.

To calculate the MD5 checksum of the downloaded file, run the following:

```
md5sum <file>
```

2.8. Address Custom xorg.conf, If Applicable

The driver relies on an automatically generated `xorg.conf` file at `/etc/X11/xorg.conf`. If a custom-built `xorg.conf` file is present, this functionality will be disabled and the driver may not work. You can try removing the existing `xorg.conf` file, or adding the contents of `/etc/X11/xorg.conf.d/00-nvidia.conf` to the `xorg.conf` file. The `xorg.conf` file will most likely need manual tweaking for systems with a non-trivial GPU configuration.

2.9. Handle Conflicting Installation Methods

Before installing CUDA, any previous installations that could conflict should be uninstalled. This will not affect systems which have not had CUDA installed previously, or systems where the installation method has been preserved (RPM/Deb vs. Runfile). See the following charts for specifics.

Table 2. CUDA Toolkit Installation Compatibility Matrix

	Installed Toolkit Version == X.Y		Installed Toolkit Version != X.Y	
	RPM/Deb	run	RPM/Deb	run

Installing Toolkit Version X.Y	RPM/Deb	No Action	Uninstall Run	No Action	No Action
	run	Uninstall RPM/Deb	Uninstall Run	No Action	No Action

Table 3. NVIDIA Driver Installation Compatibility Matrix

		Installed Driver Version == X.Y		Installed Driver Version != X.Y	
		RPM/Deb	run	RPM/Deb	run
Installing Driver Version X.Y	RPM/Deb	No Action	Uninstall Run	No Action	Uninstall Run
	run	Uninstall RPM/Deb	No Action	Uninstall RPM/Deb	No Action

Use the following command to uninstall a Toolkit runfile installation:

```
sudo /usr/local/cuda-X.Y/bin/cuda-uninstaller
```

Use the following command to uninstall a Driver runfile installation:

```
sudo /usr/bin/nvidia-uninstall
```

Use the following commands to uninstall an RPM/Deb installation:

```
sudo dnf remove <package_name> # RHEL 8 / Rocky Linux 8
```

```
sudo yum remove <package_name> # RHEL7 / CentOS 7
```

```
sudo dnf remove <package_name> # Fedora
```

```
sudo zypper remove <package_name> # OpenSUSE / SLES
```

```
sudo apt-get --purge remove <package_name> # Ubuntu
```

Chapter 3. Package Manager Installation

Basic instructions can be found in the [Quick Start Guide](#). Read on for more detailed instructions.

3.1. Overview

Installation using RPM or Debian packages interfaces with your system's package management system. When using RPM or Debian local repo installers, the downloaded package contains a repository snapshot stored on the local filesystem in `/var/`. Such a package only informs the package manager where to find the actual installation packages, but will not install them.

If the online network repository is enabled, RPM or Debian packages will be automatically downloaded at installation time using the package manager: `apt-get`, `dnf`, `yum`, or `zypper`.

Distribution-specific instructions detail how to install CUDA:

- ▶ [RHEL 7/ CentOS 7](#)
- ▶ [RHEL 8 / Rocky Linux 8](#)
- ▶ [Fedora](#)
- ▶ [SLES](#)
- ▶ [OpenSUSE](#)
- ▶ [WSL](#)
- ▶ [Ubuntu](#)
- ▶ [Debian](#)

Finally, some helpful [package manager capabilities](#) are detailed.

These instructions are for native development only. For cross-platform development, see the [CUDA Cross-Platform Environment](#) section.



Note: Optional components such as `nvidia-fs`, `libnvidia_nscq`, and `fabricmanager` are not installed by default and will have to be installed separately as needed.

3.2. RHEL 7 / CentOS 7

3.2.1. Prepare RHEL 7 / CentOS 7

1. Perform the [pre-installation actions](#).
2. **Satisfy third-party package dependency:**
 - ▶ **Satisfy DKMS dependency:** The NVIDIA driver RPM packages depend on other external packages, such as DKMS and `libvdpa`. Those packages are only available on third-party repositories, such as [EPEL](#). Any such third-party repositories must be added to the package manager repository database before installing the NVIDIA driver RPM packages, or missing dependencies will prevent the installation from proceeding.

To enable EPEL:

```
sudo yum install https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm
```

- ▶ **Enable optional repos:**

On **RHEL 7 Linux** only, execute the following steps to enable optional repositories.

- ▶ **On x86_64 workstation:**

```
subscription-manager repos --enable=rhel-7-workstation-optional-rpms
```

- ▶ **On POWER9 system:**

```
subscription-manager repos --enable=rhel-7-for-power-9-optional-rpms
```

- ▶ **On x86_64 server:**

```
subscription-manager repos --enable=rhel-7-server-optional-rpms
```

3. **Optional – Remove Outdated Signing Key:**

```
sudo rpm --erase gpg-pubkey-7fa2af80*
```

4. Choose an installation method: [local repo](#) or [network repo](#).

3.2.2. Local Repo Installation for RHEL 7 / CentOS 7

1. **Install local repository onto file system:**

```
sudo rpm --install cuda-repo-<distro>-X-Y-local-<version>*.<arch>.rpm
```

3.2.3. Network Repo Installation for RHEL 7 / CentOS 7

1. Enable the network repo:

```
sudo yum-config-manager --add-repo https://developer.download.nvidia.com/compute/cuda/repos/$distro/$arch/cuda-$distro.repo
```

where `$distro/$arch` should be replaced by one of the following:

- ▶ `rhel7/x86_64`
- ▶ `rhel7/ppc64le`

2. Install the new CUDA public GPG key:

The new GPG public key for the CUDA repository (RPM-based distros) is [d42d0685](https://developer.download.nvidia.com/compute/cuda/repos/$distro/$arch/cuda-$distro.repo).

On a fresh installation of RHEL, the yum package manager will prompt the user to accept new keys when installing packages the first time. Indicate you accept the change when prompted.

For upgrades, you must also also fetch an updated `.repo` entry:

```
sudo yum-config-manager --add-repo https://developer.download.nvidia.com/compute/cuda/repos/$distro/$arch/cuda-$distro.repo
```

3. Clean Yum repository cache:

```
sudo yum clean expire-cache
```

3.2.4. Common Installation Instructions for RHEL 7 / CentOS 7

These instructions apply to both local and network installation.

1. Install CUDA SDK:

```
sudo yum install nvidia-driver-latest-dkms
sudo yum install cuda
sudo yum install cuda-drivers
```

2. Add `libcuda.so` symbolic link, if necessary:

The `libcuda.so` library is installed in the `/usr/lib{,64}/nvidia` directory. For pre-existing projects which use `libcuda.so`, it may be useful to add a symbolic link from `libcuda.so` in the `/usr/lib{,64}` directory.

3. Reboot the system:

```
sudo reboot
```

4. Perform the [post-installation actions](#).

3.2.5. Installing a Previous NVIDIA Driver Branch on RHEL 7

To perform a network install of a previous NVIDIA driver branch on RHEL 7, use the commands below:

```
version=DRIVER_VERSION;
stream="latest-dkms";
list=("kmod-nvidia-$stream-$version")
list+=("nvidia-driver-$stream-cuda-$version")
list+=("nvidia-driver-$stream-cuda-libs-$version")
list+=("nvidia-driver-$stream-devel-$version")
list+=("nvidia-driver-$stream-$version")
list+=("nvidia-driver-$stream-NVML-$version")
list+=("nvidia-driver-$stream-NvFBOpenGL-$version")
list+=("nvidia-driver-$stream-libs-$version")
list+=("nvidia-fabric-manager-$version")
list+=("nvidia-libXNVCtrl-$version")
list+=("nvidia-libXNVCtrl-devel-$version")
list+=("nvidia-modprobe-$stream-$version")
list+=("nvidia-persistenced-$stream-$version")
list+=("nvidia-settings-$version")
list+=("nvidia-xconfig-$stream-$version")
sudo yum --setopt=obsoletes=0 install ${list[@]}
sudo yum install cuda-drivers-fabricmanager-DRIVER_BRANCH
```

Where *DRIVER_VERSION* is the full version, for example, 470.82.01, and *DRIVER_BRANCH* is, for example, 470.

3.3. RHEL 8 / Rocky 8

3.3.1. Prepare RHEL 8 / Rocky 8

1. Perform the [pre-installation actions](#).
2. **Satisfy third-party package dependency:**
 - ▶ **Satisfy DKMS dependency:** The NVIDIA driver RPM packages depend on other external packages, such as DKMS and libvdpau. Those packages are only available on third-party repositories, such as [EPEL](#). Any such third-party repositories must be added to the package manager repository database before installing the NVIDIA driver RPM packages, or missing dependencies will prevent the installation from proceeding.

To enable EPEL:

```
sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
```

- ▶ **Enable optional repos:**

On **RHEL 8 Linux** only, execute the following steps to enable optional repositories.

- ▶ **On x86_64 systems:**

```
subscription-manager repos --enable=rhel-8-for-x86_64-appstream-rpms
```

```
subscription-manager repos --enable=rhel-8-for-x86_64-baseos-rpms
subscription-manager repos --enable=codeready-builder-for-rhel-8-x86_64-rpms
```

► **On POWER9 systems:**

```
subscription-manager repos --enable=rhel-8-for-ppc64le-appstream-rpms
subscription-manager repos --enable=rhel-8-for-ppc64le-baseos-rpms
subscription-manager repos --enable=codeready-builder-for-rhel-8-ppc64le-rpms
```

3. **Remove Outdated Signing Key:**

```
sudo rpm --erase gpg-pubkey-7fa2af80*
```

4. Choose an installation method: [local repo](#) or [network repo](#).

3.3.2. Local Repo Installation for RHEL 8 / Rocky 8

1. **Install local repository on file system:**

```
sudo rpm --install cuda-repo-<distro>-X-Y-local-<version>*.<arch>.rpm
```

3.3.3. Network Repo Installation for RHEL 8 / Rocky 8

1. **Enable the the network repo:**

```
sudo dnf config-manager --add-repo https://developer.download.nvidia.com/compute/cuda/repos/$distro/$arch/cuda-$distro.repo
```

where `$distro/$arch` should be replaced by one of the following:

- `rhel8/cross-linux-sbsa`
- `rhel8/ppc64le`
- `rhel8/sbsa`
- `rhel8/x86_64`

2. **Install the new CUDA public GPG key:**

The new GPG public key for the CUDA repository (RPM-based distros) is [d42d0685](#).

On a fresh installation of RHEL, the dnf package manager will prompt the user to accept new keys when installing packages the first time. Indicate you accept the change when prompted.

For upgrades, you must also also fetch an updated `.repo` entry:

```
sudo dnf config-manager --add-repo https://developer.download.nvidia.com/compute/cuda/repos/$distro/$arch/cuda-$distro.repo
```

3. **Clean Yum repository cache:**

```
sudo dnf clean expire-cache
```

3.3.4. Common Instructions for RHEL 8 / Rocky 8

These instructions apply to both local and network installation.

1. **Install CUDA SDK:**

```
sudo dnf module install nvidia-driver:latest-dkms
sudo dnf install cuda
```

2. **Install GPUDirect Filesystem:**

```
sudo dnf install nvidia-gds
```

3. **Add libcuda.so symbolic link, if necessary**

The `libcuda.so` library is installed in the `/usr/lib{,64}/nvidia` directory. For pre-existing projects which use `libcuda.so`, it may be useful to add a symbolic link from `libcuda.so` in the `/usr/lib{,64}` directory.

4. **Reboot the system:**

```
sudo reboot
```

5. Perform the [post-installation actions](#).

3.4. Fedora

3.4.1. Prepare Fedora

1. Perform the [pre-installation actions](#).

2. **Remove Outdated Signing Key:**

```
sudo rpm --erase gpg-pubkey-7fa2af80*
```

3. Choose an installation method: [local repo](#) or [network repo](#).

3.4.2. Local Repo Installation for Fedora

1. **Install local repository on file system:**

```
sudo rpm --install cuda-repo-<distro>-X-Y-local-<version>*.x86_64.rpm
```

where `$distro` is `fedora33` or `fedora35`, for example.

3.4.3. Network Repo Installation for Fedora

1. **Enable the network repo:**

```
sudo dnf config-manager --add-repo https://developer.download.nvidia.com/compute/cuda/repos/$distro/x86_64/cuda-$distro.repo
```

where `$distro` should be replaced by one of the following:

- ▶ fedora32
- ▶ fedora33
- ▶ fedora34
- ▶ fedora35

2. Install the new CUDA public GPG key:

The new GPG public key for the CUDA repository (RPM-based distros) is [d42d0685](https://developer.download.nvidia.com/compute/cuda/repos/$distro/x86_64/cuda-$distro.repo).

On a fresh installation of Fedora, the dnf package manager will prompt the user to accept new keys when installing packages the first time. Indicate you accept the change when prompted.

For upgrades, you must also fetch an updated `.repo` entry:

```
sudo dnf config-manager --add-repo https://developer.download.nvidia.com/compute/cuda/repos/$distro/x86_64/cuda-$distro.repo
```

3. Clean DNF repository cache:

```
sudo dnf clean expire-cache
```

3.4.4. Common Installation Instructions for Fedora

These instructions apply to both local and network installation for Fedora.

1. Install CUDA SDK:

```
sudo dnf module install nvidia-driver:latest-dkms
sudo dnf install cuda
```

NOTE: The CUDA driver installation may fail if the RPMFusion non-free repository is enabled. In this case, CUDA installations should temporarily disable the RPMFusion non-free repository:

```
sudo dnf --disablerepo="rpmfusion-nonfree*" install cuda
```

It may be necessary to rebuild the grub configuration files, particularly if you use a non-default partition scheme. If so, then run this below command, and reboot the system:

```
sudo grub2-mkconfig -o /boot/grub2/grub.cfg
```

2. Reboot the system:

```
sudo reboot
```

3. Add `libcuda.so` symbolic link, if necessary:

The `libcuda.so` library is installed in the `/usr/lib{,64}/nvidia` directory. For pre-existing projects which use `libcuda.so`, it may be useful to add a symbolic link from `libcuda.so` in the `/usr/lib{,64}` directory.

4. Perform the [post-installation actions](#).

3.5. SLES

3.5.1. Prepare SLES

1. Perform the [pre-installation actions](#).
2. On SLES12 SP4, install the Mesa-libgl-devel Linux packages before proceeding. See [Mesa-libGL-devel](#).
3. **Add the user to the video group:**

```
sudo usermod -a -G video <username>
```

4. **Remove Outdated Signing Key:**

```
sudo rpm --erase gpg-pubkey-7fa2af80*
```

5. Choose an installation method: [local repo](#) or [network repo](#).

3.5.2. Local Repo Installation for SLES

1. **Install local repository on file system:**

```
sudo rpm --install cuda-repo-sles15-X-Y-local-<version>*.x86_64.rpm
```

3.5.3. Network Repo Installation for SLES

1. Enable the network repo:

```
sudo zypper addrepo https://developer.download.nvidia.com/compute/cuda/repos/
$distro/$arch/cuda-$distro.repo
```

where \$distro/\$arch should be replaced by one of the following:

- ▶ sles15/cross-linux-sbsa
- ▶ sles15/sbsa
- ▶ sles15/x86_64

2. **Install the new CUDA public GPG key:**

The new GPG public key for the CUDA repository (RPM-based distros) is [d42d0685](#).

On a fresh installation of SLES, the zypper package manager will prompt the user to accept new keys when installing packages the first time. Indicate you accept the change when prompted.

For upgrades, you must also also fetch an updated .repo entry:

```
sudo zypper removerepo cuda-$distro-$arch
sudo zypper addrepo https://developer.download.nvidia.com/compute/cuda/repos/
$distro/$arch/cuda-$distro.repo
```

3. Refresh Zypper repository cache:

```
sudo SUSEConnect --product PackageHub/15/<architecture>
sudo zypper refresh
```

3.5.4. Common Installation Instructions for SLES

These instructions apply to both local and network installation for SLES.

1. Install CUDA SDK:

```
sudo zypper install cuda
```

2. Install CUDA Samples GL dependencies:

Refer to [CUDA Cross-Platform Samples](#).

3. Reboot the system:

```
sudo reboot
```

4. Perform the [post-installation actions](#).

3.6. OpenSUSE

3.6.1. Prepare OpenSUSE

1. Perform the [pre-installation actions](#).

2. Add the user to the video group:

```
sudo usermod -a -G video <username>
```

3. Remove Outdated Signing Key:

```
sudo rpm --erase gpg-pubkey-7fa2af80*
```

4. Choose an installation method: [local repo](#) or [network repo](#).

3.6.2. Local Repo Installation for OpenSUSE

1. Install local repository on file system:

```
sudo rpm --install cuda-repo-opensuse15-<version>.x86_64.rpm
```

3.6.3. Network Repo Installation for OpenSUSE

1. Enable the network repo:

```
sudo zypper addrepo https://developer.download.nvidia.com/compute/cuda/repos/
opensuse15/x86_64/cuda-opensuse15.repo
```

2. Install the new CUDA public GPG key:

The new GPG public key for the CUDA repository (RPM-based distros) is [d42d0685](#). On fresh installation of openSUSE, the zypper package manager will prompt the user to accept new keys when installing packages the first time. Indicate you accept the change when prompted.

For upgrades, you must also also fetch an updated .repo entry:

```
sudo zypper removerepo cuda-opensuse15-x86_64
sudo zypper addrepo https://developer.download.nvidia.com/compute/cuda/repos/
opensuse15/x86_64/cuda-opensuse15.repo
```

3. Refresh Zypper repository cache:

```
sudo zypper refresh
```

3.6.4. Common Installation Instructions for OpenSUSE

These instructions apply to both local and network installation for OpenSUSE.

1. Install CUDA SDK:

```
sudo zypper install cuda
```

2. Reboot the system:

```
sudo reboot
```

3. Perform the [post-installation actions](#).

3.7. WSL

These instructions must be used if you are installing in a WSL environment. Do not use the Ubuntu instructions in this case; it is important to not install the `cuda-drivers` packages within the WSL environment.

3.7.1. Prepare WSL

1. Perform the [pre-installation actions](#).

2. Remove Outdated Signing Key:

```
sudo apt-key del 7fa2af80
```

3. Choose an installation method: local repo or network repo.

3.7.2. Local Repo Installation for WSL

1. Install local repository on file system:

```
sudo dpkg -i cuda-repo-wsl-ubuntu-X-Y-local_<version>*_x86_64.deb
```

2. Enroll ephemeral public GPG key:

```
sudo cp /var/cuda-repo-wsl-ubuntu-X-Y-local/cuda-*-keyring.gpg /usr/share/
keyrings/
```

3.7.3. Network Repo Installation for WSL

The new GPG public key for the CUDA repository (Debian-based distros) is [3bf863cc](#). This must be enrolled on the system, either using the `cuda-keyring` package or manually; the `apt-key` command is deprecated and not recommended.

1. Install the `newcuda-keyring` package:

```
wget https://developer.download.nvidia.com/compute/cuda/repos/wsl-ubuntu/x86_64/
cuda-keyring_1.0-1_all.deb
sudo dpkg -i cuda-keyring_1.0-1_all.deb
```

Or if you are unable to install the `cuda-keyring` package, you can optionally:

a). Enroll the new signing key manually:

```
wget https://developer.download.nvidia.com/compute/cuda/repos/wsl-ubuntu/
x86_64/cuda-wsl-ubuntu-keyring.gpg
sudo mv cuda-wsl-ubuntu-keyring.gpg /usr/share/keyrings/cuda-archive-
keyring.gpg
```

b). Enable the network repository:

```
echo "deb [signed-by=/usr/share/keyrings/cuda-archive-keyring.gpg] https://
developer.download.nvidia.com/compute/cuda/repos/wsl-ubuntu/x86_64/ /" | sudo
tee /etc/apt/sources.list.d/cuda-wsl-ubuntu-x86_64.list
```

c). Add pin file to prioritize CUDA repository:

```
wget https://developer.download.nvidia.com/compute/cuda/repos/wsl-ubuntu/
x86_64/cuda-wsl-ubuntu.pin
sudo mv cuda-wsl-ubuntu.pin /etc/apt/preferences.d/cuda-repository-pin-600
```

3.7.4. Common Installation Instructions for WSL

These instructions apply to both local and network installation for WSL.

1. Update the Apt repository cache:

```
sudo apt-get update
```

2. Install CUDA SDK:

```
sudo apt-get install cuda
```

3. Perform the [post-installation actions](#).

3.8. Ubuntu

3.8.1. Prepare Ubuntu

1. Perform the [pre-installation actions](#).
2. **Remove Outdated Signing Key:**

```
sudo apt-key del 7fa2af80
```

3. Choose an installation method: local repo or network repo.

3.8.2. Local Repo Installation for Ubuntu

1. **Install local repository on file system:**

```
sudo dpkg -i cuda-repo-<distro>_<version>_<architecture>.deb
```

2. **Enroll ephemeral public GPG key:**

```
sudo cp /var/cuda-repo-<distro>-X-Y-local/cuda-*-keyring.gpg /usr/share/keyrings/
```

3.8.3. Network Repo Installation for Ubuntu

The new GPG public key for the CUDA repository is [3bf863cc](#). This must be enrolled on the system, either using the `cuda-keyring` package or manually; the `apt-key` command is deprecated and not recommended.

1. **Install the new `cuda-keyring` package:**

```
wget https://developer.download.nvidia.com/compute/cuda/repos/$distro/$arch/cuda-keyring_1.0-1_all.deb
sudo dpkg -i cuda-keyring_1.0-1_all.deb
```

where `$distro/$arch` should be replaced by one of the following:

- ▶ `ubuntu1604/x86_64`
- ▶ `ubuntu1804/cross-linux-sbsa`
- ▶ `ubuntu1804/ppc64el`
- ▶ `ubuntu1804/sbsa`
- ▶ `ubuntu1804/x86_64`
- ▶ `ubuntu2004/cross-linux-sbsa`
- ▶ `ubuntu2004/sbsa`
- ▶ `ubuntu2004/x86_64`
- ▶ `ubuntu2204/sbsa`
- ▶ `ubuntu2204/x86_64`

```
sudo dpkg -i cuda-keyring_1.0-1_all.deb
```

2. Or if you are unable to install the `cuda-keyring` package, you can optionally:

a). **Enroll the new signing key manually:**

```
wget https://developer.download.nvidia.com/compute/cuda/repos/<distro>/<arch>/
cuda-<distro>-keyring.gpg
sudo mv cuda-<distro>-keyring.gpg /usr/share/keyrings/cuda-archive-keyring.gpg
```

b). **Enable the network repository:**

```
echo "deb [signed-by=/usr/share/keyrings/cuda-archive-keyring.gpg] https://
developer.download.nvidia.com/compute/cuda/repos/<distro>/<arch>/ /" | sudo
tee /etc/apt/sources.list.d/cuda-<distro>-<arch>.list
```

c). **Add pin file to prioritize CUDA repository:**

```
wget https://developer.download.nvidia.com/compute/cuda/repos/<distro>/<arch>/
cuda-<distro>.pin
sudo mv cuda-<distro>.pin /etc/apt/preferences.d/cuda-repository-pin-600
```

3.8.4. Common Installation Instructions for Ubuntu

These instructions apply to both local and network installation for Ubuntu.

1. **Update the Apt repository cache:**

```
sudo apt-get update
```

2. **Install CUDA SDK:**



Note: These two commands must be executed separately.

```
sudo apt-get install cuda
```

To include all GDS packages:

```
sudo apt-get install nvidia-gds
```

3. **Reboot the system**

```
sudo reboot
```

4. Perform the [Post-installation Actions](#)

3.9. Debian

3.9.1. Prepare Debian

1. Perform the [pre-installation actions](#).

2. **Enable the contrib repository:**

```
sudo add-apt-repository contrib
```

3. Remove Outdated Signing Key:

```
sudo apt-key del 7fa2af80
```

4. Choose an installation method: [local repo](#) or [network repo](#).

3.9.2. Local Repo Installation for Debian

1. Install local repository on file system:

```
sudo dpkg -i cuda-repo-<distro>-X-Y-local_<version>*_x86_64.deb
```

2. Enroll ephemeral public GPG key:

```
sudo cp /var/cuda-repo-<distro>-X-Y-local/cuda-*-keyring.gpg /usr/share/keyrings/
```

3.9.3. Network Repo Installation for Debian

The new GPG public key for the CUDA repository (Debian-based distros) is [3bf863cc](#). This must be enrolled on the system, either using the `cuda-keyring` package or manually; the `apt-key` command is deprecated and not recommended.

1. Install the new cuda-keyring package:

```
wget https://developer.download.nvidia.com/compute/cuda/repos/<distro>/<arch>/  
cuda-keyring_1.0-1_all.deb
```

where `$distro/$arch` should be replaced by one of the following:

- ▶ `debian10/x86_64`
- ▶ `debian11/x86_64`

```
sudo dpkg -i cuda-keyring_1.0-1_all.deb
```

2. Or if you are unable to install the `cuda-keyring` package, you can optionally:

a). Enroll the new signing key manually:

```
wget https://developer.download.nvidia.com/compute/cuda/repos/<distro>/x86_64/  
cuda-<distro>-keyring.gpg  
sudo mv cuda-<distro>-keyring.gpg /usr/share/keyrings/cuda-archive-keyring.gpg
```

b). Enable the network repository:

```
echo "deb [signed-by=/usr/share/keyrings/cuda-archive-keyring.gpg] https://  
developer.download.nvidia.com/compute/cuda/repos/<distro>/x86_64/ /" | sudo  
tee /etc/apt/sources.list.d/cuda-<distro>-x86_64.list
```

c). Add pin file to prioritize CUDA repository:

```
wget https://developer.download.nvidia.com/compute/cuda/repos/<distro>/x86_64/  
cuda-<distro>.pin  
sudo mv cuda-<distro>.pin /etc/apt/preferences.d/cuda-repository-pin-600
```

3.9.4. Common Installation Instructions for Debian

These instructions apply to both local and network installation for Debian.

1. Update the Apt repository cache:

```
sudo apt-get update
```



Note: If you are using Debian 10, you may instead need to run:

```
sudo apt-get --allow-releaseinfo-change update
```

2. Install CUDA SDK:

```
sudo apt-get -y install cuda
```

3. Reboot the system:

```
sudo reboot
```

4. Perform the [post-installation actions](#).

3.10. Additional Package Manager Capabilities

Below are some additional capabilities of the package manager that users can take advantage of.

3.10.1. Available Packages

The recommended installation package is the `cuda` package. This package will install the full set of other CUDA packages required for native development and should cover most scenarios.

The `cuda` package installs all the available packages for native developments. That includes the compiler, the debugger, the profiler, the math libraries, and so on. For `x86_64` platforms, this also includes Nsight Eclipse Edition and the visual profilers. It also includes the NVIDIA driver package.

On supported platforms, the `cuda-cross-aarch64` and `cuda-cross-ppc64e1` packages install all the packages required for cross-platform development to ARMv8 and POWER8, respectively. The libraries and header files of the target architecture's display driver package are also installed to enable the cross compilation of driver applications. The `cuda-cross-<arch>` packages do not install the native display driver.

The packages installed by the packages above can also be installed individually by specifying their names explicitly. The list of available packages be can obtained with:

```
yum --disablerepo="*" --enablerepo="cuda*" list available # RedHat
```

```
dnf --disablerepo="*" --enablerepo="cuda*" list available # Fedora
```

```
zypper packages -r cuda # OpenSUSE & SLES
```

```
cat /var/lib/apt/lists/*cuda*Packages | grep "Package:" # Ubuntu
```

3.10.2. Optional 32-bit Packages for Linux x86_64 .deb/.rpm

These packages provide 32-bit driver libraries needed for things such as Steam (popular game app store/launcher), older video games, and some compute applications.

For Debian 11:

```
sudo dpkg --add-architecture i386
sudo apt-get update
sudo apt-get install libcuda1-i386 nvidia-driver-libs-i386
```

For Ubuntu:

```
sudo dpkg --add-architecture i386
sudo apt-get update
sudo apt-get install libnvidia-compute-<branch>:i386 libnvidia-decode-<branch>:i386 \
  libnvidia-encode-<branch>:i386 libnvidia-extra-<branch>:i386 libnvidia-fbc1-
<branch>:i386 \
  libnvidia-gl-<branch>:i386
```

Where <branch> is the driver version, for example 495.

For Fedora and RHEL8:

```
sudo dnf install nvidia-driver-cuda-libs.i686 nvidia-driver-devel.i686 \
  nvidia-driver-libs.i686 nvidia-driver-NvFBCOpenGL.i686 nvidia-driver-NVML.i686
```



Note: There is no modularity profile support.

For openSUSE/SLES:

No extra installation is required, the `nvidia-g1G05` package already contains the 32-bit libraries.

3.10.3. Package Upgrades

The `cuda` package points to the latest stable release of the CUDA Toolkit. When a new version is available, use the following commands to upgrade the toolkit and driver:

```
sudo yum install cuda # RHEL7
sudo dnf upgrade cuda # Fedora/RHEL8
sudo zypper install cuda # OpenSUSE & SLES
sudo apt-get install cuda # Ubuntu
```

The `cuda-cross-<arch>` packages can also be upgraded in the same manner.

The `cuda-drivers` package points to the latest driver release available in the CUDA repository. When a new version is available, use the following commands to upgrade the driver:

```
sudo yum install nvidia-driver-latest-dkms # RHEL7
sudo yum install cuda-drivers # RHEL7
sudo dnf module update nvidia-driver:latest-dkms # RHEL8/Fedora
sudo zypper install cuda-drivers nvidia-gfxG04-kmp-default # OpenSUSE & SLES
sudo apt-get install cuda-drivers # Ubuntu
```

Some desktop environments, such as GNOME or KDE, will display a notification alert when new packages are available.

To avoid any automatic upgrade, and lock down the toolkit installation to the X.Y release, install the `cuda-X-Y` or `cuda-cross-<arch>-X-Y` package.

Side-by-side installations are supported. For instance, to install both the X.Y CUDA Toolkit and the X.Y+1 CUDA Toolkit, install the `cuda-X.Y` and `cuda-X.Y+1` packages.

3.10.4. Meta Packages

Meta packages are RPM/Deb/Conda packages which contain no (or few) files but have multiple dependencies. They are used to install many CUDA packages when you may not know the details of the packages you want. Below is the list of meta packages.

Table 4. Meta Packages Available for CUDA 11.7

Meta Package	Purpose
<code>cuda</code>	Installs all CUDA Toolkit and Driver packages. Handles upgrading to the next version of the <code>cuda</code> package when it's released.
<code>cuda-11-7</code>	Installs all CUDA Toolkit and Driver packages. Remains at version 11.7 until an additional version of CUDA is installed.
<code>cuda-toolkit-11-7</code>	Installs all CUDA Toolkit packages required to develop CUDA applications. Does not include the driver.
<code>cuda-tools-11-7</code>	Installs all CUDA command line and visual tools.
<code>cuda-runtime-11-7</code>	Installs all CUDA Toolkit packages required to run CUDA applications, as well as the Driver packages.
<code>cuda-compiler-11-7</code>	Installs all CUDA compiler packages.
<code>cuda-libraries-11-7</code>	Installs all runtime CUDA Library packages.
<code>cuda-libraries-dev-11-7</code>	Installs all development CUDA Library packages.

Meta Package	Purpose
cuda-drivers	Installs all Driver packages. Handles upgrading to the next version of the Driver packages when they're released.

Chapter 4. Driver Installation

This section is for users who want to install a specific driver version.

For Debian and Ubuntu:

```
sudo apt-get install cuda-drivers-<branch>
```

For example:

```
sudo apt-get install cuda-drivers-418
```

For OpenSUSE and SLES:

```
sudo zypper install cuda-drivers-<branch>
```

For example:

```
sudo zypper install cuda-drivers-450
```

This allows you to get the highest version in the specified branch.

For Fedora and RHEL8:

```
sudo dnf module install nvidia-driver:<stream>/<profile>
```

where profile by default is "default" and does not need to be specified.

- ▶ Example dkms streams: 450-dkms or latest-dkms
- ▶ Example precompiled streams: 450 or latest



Note: Precompiled streams are only supported on RHEL8 x86_64.

To uninstall or change streams on Fedora and RHEL8:

```
sudo dnf module remove --all nvidia-driver  
sudo dnf module reset nvidia-driver
```

Chapter 5. NVIDIA Open GPU Kernel Modules

The NVIDIA Linux GPU Driver contains several kernel modules:

- ▶ `nvidia.ko`
- ▶ `nvidia-modeset.ko`
- ▶ `nvidia-uvm.ko`
- ▶ `nvidia-drm.ko`
- ▶ `nvidia-peermem.ko`

Starting in the 515 driver release series, two "flavors" of these kernel modules are provided:

- ▶ **Proprietary**- this is the flavor that NVIDIA has historically shipped.
- ▶ **Open-source** - published kernel modules that are dual licensed MIT/GPLv2. These are new starting in release 515. With every driver release, the source code to the open kernel modules will be published on <https://github.com/NVIDIA/open-gpu-kernel-modules> and a tarball will be provided on <https://download.nvidia.com/XFree86/>.

Verify that your NVIDIA GPU is at least Turing or newer generation.

```
lspci | grep VGA
```

Experimental support for GeForce and Quadro SKUs can be enabled with:

```
echo "options nvidia NVreg_OpenRmEnableUnsupportedGpus=1" | sudo tee /etc/modprobe.d/nvidia-gsp.conf
```

To install NVIDIA Open GPU Kernel Modules, follow the instructions below.

5.1. CUDA Runfile

Pass the CLI argument to the CUDA runfile to opt in to NVIDIA Open GPU Kernel Modules:

```
sh cuda_<release>_<version>_linux.run -m=kernel-open
```

5.2. Debian

1. Install the NVIDIA Open GPU Kernel Modules package:

```
sudo apt-get install nvidia-kernel-open-dkms
```

2. Install the rest of the NVIDIA driver packages:

```
sudo apt-get install cuda-drivers
```

5.3. Fedora

1. Install the NVIDIA Open GPU Kernel Modules package and the rest of the NVIDIA driver packages:

```
sudo dnf module install nvidia-driver:open-dkms
```

5.4. RHEL 8 / Rocky 8

1. Install the NVIDIA Open GPU Kernel Modules package and the rest of the NVIDIA driver packages:

```
sudo dnf module install nvidia-driver:open-dkms
```

5.5. RHEL 7 / CentOS 7

1. Install the NVIDIA Open GPU Kernel Modules package:

```
sudo yum install kmod-nvidia-open-dkms
```

2. Install the rest of the NVIDIA driver packages (except nvidia-settings):

```
sudo yum install nvidia-driver-latest-dkms
```

3. Install nvidia-settings:

```
sudo yum install cuda-drivers
```

5.6. OpenSUSE / SLES

1. Install the NVIDIA Open GPU Kernel Modules package:

```
sudo zypper install nvidia-open-gfxG05-kmp-default
```

2. Install the rest of the NVIDIA driver packages:

```
sudo zypper install cuda-drivers
```

5.7. Ubuntu

1. Install the NVIDIA Open GPU Kernel Modules package:

```
sudo apt-get install nvidia-kernel-open-515
```

2. Install the rest of the NVIDIA driver packages:

```
sudo apt-get install cuda-drivers-515
```

Chapter 6. Precompiled Streams

Precompiled streams offer an optional method of streamlining the installation process.

The advantages of precompiled streams:

- ▶ Precompiled: faster boot up after driver and/or kernel updates
- ▶ Pre-tested: kernel and driver combination has been validated
- ▶ Removes gcc dependency: no compiler installation required
- ▶ Removes dkms dependency: enabling EPEL repository not required
- ▶ Removes kernel-devel and kernel-headers dependencies: no black screen if matching packages are missing

When using precompiled drivers, a plugin for the dnf package manager is enabled that cleans up stale .ko files. To prevent system breakages, the NVIDIA dnf plugin also prevents upgrading to a kernel for which no precompiled driver yet exists. This can delay the application of security fixes but ensures that a tested kernel and driver combination is always used. A warning is displayed by dnf during that upgrade situation:

```
NOTE: Skipping kernel installation since no NVIDIA driver kernel module package
kmod-nvidia-${driver}-${kernel} ... could be found
```

Packaging templates and instructions are provided on GitHub to allow you to maintain your own precompiled kernel module packages for custom kernels and derivative Linux distros: [NVIDIA/yum-packaging-precompiled-kmod](https://github.com/NVIDIA/yum-packaging-precompiled-kmod)

To use the new driver packages on RHEL 8:

1. First, ensure that the Red Hat repositories are enabled:

```
subscription-manager repos --enable=rhel-8-for-x86_64-appstream-rpms
subscription-manager repos --enable=rhel-8-for-x86_64-baseos-rpms
```

2. Choose **one** of the four options below depending on the desired driver:

- ▶ `latest` always updates to the highest versioned driver (precompiled):

```
sudo dnf module install nvidia-driver:latest
```

- ▶ `<id>` locks the driver updates to the specified driver branch (precompiled):

```
sudo dnf module install nvidia-driver:<id>
```



Note: Replace `<id>` with the appropriate driver branch streams, for example 455, 450, 440, or 418.

- ▶ `latest-dkms` always updates to the highest versioned driver (non-precompiled):

```
sudo dnf module install nvidia-driver:latest-dkms
```



Note: This is the default stream.

- ▶ `<id>-dkms` locks the driver updates to the specified driver branch (non-precompiled):

```
sudo dnf module install nvidia-driver:<id>-dkms
```



Note: Valid streams include 455-dkms, 450-dkms, 440-dkms, and 418-dkms.

6.1. Precompiled Streams Support Matrix

This table shows the supported precompiled and legacy DKMS streams for each driver.

NVIDIA Driver	Precompiled Stream	Legacy DKMS Stream
Highest version	latest	latest-dkms
Locked at 455.x	455	455-dkms
Locked at 450.x	450	450-dkms
Locked at 440.x	440	440-dkms
Locked at 418.x	418	418-dkms

6.2. Modularity Profiles

Modularity profiles work with any supported modularity stream and allow for additional use cases. These modularity profiles are available on RHEL8 and Fedora.

Table 5. List of nvidia-driver Module Profiles

Stream	Profile	Use Case
Default	<code>/default</code>	Installs all the driver packages in a stream.
Kickstart	<code>/ks</code>	Performs unattended Linux OS installation using a config file.

Stream	Profile	Use Case
NVSwitch Fabric	/fm	Installs all the driver packages plus components required for bootstrapping an NVSwitch system (including the Fabric Manager and NSCQ telemetry).
Source	/src	Source headers for compilation (precompiled streams only).

For example:

```
sudo dnf module nvidia-driver:<stream>/default
sudo dnf module nvidia-driver:<stream>/ks
sudo dnf module nvidia-driver:<stream>/fm
sudo dnf module nvidia-driver:<stream>/src
```

You can install multiple modularity profiles using BASH curly brace expansion, for example:

```
sudo dnf module install nvidia-driver:latest/{default,src}
```

See <https://developer.nvidia.com/blog/streamlining-nvidia-driver-deployment-on-rhel-8-with-modularity-streams> in the Developer Blog and https://developer.download.nvidia.com/compute/cuda/repos/rhel8/x86_64/precompiled/ for more information.

Chapter 7. Kickstart Installation

7.1. RHEL8 / Rocky Linux 8

1. Enable the EPEL repository:

```
repo --name=epel --baseurl=http://download.fedoraproject.org/pub/epel/8/Everything/x86_64/
```

2. Enable the CUDA repository:

```
repo --name=cuda-rhel8 --baseurl=https://developer.download.nvidia.com/compute/cuda/repos/rhel8/x86_64/
```

3. In the packages section of the `ks.cfg` file, make sure you are using the `/ks` profile and `:latest-dkms` stream:

```
@nvidia-driver:latest-dkms/ks
```

4. Perform the [post-installation actions](#).

Chapter 8. Runfile Installation

Basic instructions can be found in the [Quick Start Guide](#). Read on for more detailed instructions.

This section describes the installation and configuration of CUDA when using the standalone installer. The standalone installer is a ".run" file and is completely self-contained.

8.1. Overview

The Runfile installation installs the NVIDIA Driver and CUDA Toolkit via an interactive ncurses-based interface.

The [installation steps](#) are listed below. Distribution-specific instructions on [disabling the Nouveau drivers](#) as well as steps for [verifying device node creation](#) are also provided.

Finally, [advanced options](#) for the installer and [uninstallation steps](#) are detailed below.

The Runfile installation does not include support for cross-platform development. For cross-platform development, see the [CUDA Cross-Platform Environment](#) section.

8.2. Installation

1. Perform the [pre-installation actions](#).
2. [Disable the Nouveau drivers](#).
3. Reboot into text mode (runlevel 3).

This can usually be accomplished by adding the number "3" to the end of the system's kernel boot parameters.

Since the NVIDIA drivers are not yet installed, the text terminals may not display correctly. Temporarily adding "nomodeset" to the system's kernel boot parameters may fix this issue.

Consult your system's bootloader documentation for information on how to make the above boot parameter changes.

The reboot is required to completely unload the Nouveau drivers and prevent the graphical interface from loading. The CUDA driver cannot be installed while the Nouveau drivers are loaded or while the graphical interface is active.

4. Verify that the Nouveau drivers are not loaded. If the Nouveau drivers are still loaded, consult your distribution's documentation to see if further steps are needed to disable Nouveau.
5. Run the installer and follow the on-screen prompts:

```
sudo sh cuda_<version>_linux.run
```

The installer will prompt for the following:

- ▶ EULA Acceptance
- ▶ CUDA Driver installation
- ▶ CUDA Toolkit installation, location, and `/usr/local/cuda` symbolic link

The default installation location for the toolkit is `/usr/local/cuda-11.7`:

The `/usr/local/cuda` symbolic link points to the location where the CUDA Toolkit was installed. This link allows projects to use the latest CUDA Toolkit without any configuration file update.

The installer must be executed with sufficient privileges to perform some actions. When the current privileges are insufficient to perform an action, the installer will ask for the user's password to attempt to install with root privileges. Actions that cause the installer to attempt to install with root privileges are:

- ▶ installing the CUDA Driver
- ▶ installing the CUDA Toolkit to a location the user does not have permission to write to
- ▶ creating the `/usr/local/cuda` symbolic link

Running the installer with **sudo**, as shown above, will give permission to install to directories that require root permissions. Directories and files created while running the installer with **sudo** will have root ownership.

If installing the driver, the installer will also ask if the OpenGL libraries should be installed. If the GPU used for display is not an NVIDIA GPU, the NVIDIA OpenGL libraries should not be installed. Otherwise, the OpenGL libraries used by the graphics driver of the non-NVIDIA GPU will be overwritten and the GUI will not work. If performing a silent installation, the `--no-opengl-libs` option should be used to prevent the OpenGL libraries from being installed. See the [Advanced Options](#) section for more details.

If the GPU used for display is an NVIDIA GPU, the X server configuration file, `/etc/x11/xorg.conf`, may need to be modified. In some cases, `nvidia-xconfig` can be used to automatically generate an `xorg.conf` file that works for the system. For non-standard systems, such as those with more than one GPU, it is recommended to manually edit the `xorg.conf` file. Consult the `xorg.conf` documentation for more information.



Note: Installing Mesa may overwrite the `/usr/lib/libGL.so` that was previously installed by the NVIDIA driver, so a reinstallation of the NVIDIA driver might be required after installing these libraries.

6. Reboot the system to reload the graphical interface:

```
sudo reboot
```

7. Verify the [device nodes](#) are created properly.
8. Perform the [post-installation actions](#).

8.3. Disabling Nouveau

To install the Display Driver, the Nouveau drivers must first be disabled. Each distribution of Linux has a different method for disabling Nouveau.

The Nouveau drivers are loaded if the following command prints anything:

```
lsmod | grep nouveau
```

8.3.1. Fedora

1. Create a file at `/usr/lib/modprobe.d/blacklist-nouveau.conf` with the following contents:

```
blacklist nouveau
options nouveau modeset=0
```

2. Regenerate the kernel initramfs:

```
sudo dracut --force
```

3. Run the following command:

```
sudo grub2-mkconfig -o /boot/grub2/grub.cfg
```

4. Reboot the system.

8.3.2. RHEL/CentOS

1. Create a file at `/etc/modprobe.d/blacklist-nouveau.conf` with the following contents:

```
blacklist nouveau
options nouveau modeset=0
```

2. Regenerate the kernel initramfs:

```
sudo dracut --force
```

8.3.3. OpenSUSE

1. Create a file at `/etc/modprobe.d/blacklist-nouveau.conf` with the following contents:

```
blacklist nouveau
options nouveau modeset=0
```

2. Regenerate the kernel initrd:

```
sudo /sbin/mkinitrd
```

8.3.4. SLES

No actions to disable Nouveau are required as Nouveau is not installed on SLES.

8.3.5. WSL

No actions to disable Nouveau are required as Nouveau is not installed on WSL.

8.3.6. Ubuntu

1. Create a file at `/etc/modprobe.d/blacklist-nouveau.conf` with the following contents:

```
blacklist nouveau
options nouveau modeset=0
```

2. Regenerate the kernel initramfs:

```
sudo update-initramfs -u
```

8.3.7. Debian

1. Create a file at `/etc/modprobe.d/blacklist-nouveau.conf` with the following contents:

```
blacklist nouveau
options nouveau modeset=0
```

2. Regenerate the kernel initramfs:

```
sudo update-initramfs -u
```

8.4. Device Node Verification

Check that the device files `/dev/nvidia*` exist and have the correct (0666) file permissions. These files are used by the CUDA Driver to communicate with the kernel-mode portion of the NVIDIA Driver. Applications that use the NVIDIA driver, such as a CUDA application or the X server (if any), will normally automatically create these files if they are missing using the `setuid nvidia-modprobe` tool that is bundled with the NVIDIA Driver. However, some systems disallow `setuid` binaries, so if these files do not exist, you can create them manually by using a startup script such as the one below:

```
#!/bin/bash

/sbin/modprobe nvidia

if [ "$?" -eq 0 ]; then
  # Count the number of NVIDIA controllers found.
  NVDEVS=`lspci | grep -i NVIDIA`
  N3D=`echo "$NVDEVS" | grep "3D controller" | wc -l`
  NVGA=`echo "$NVDEVS" | grep "VGA compatible controller" | wc -l`

  N=`expr $N3D + $NVGA - 1`
  for i in `seq 0 $N`; do
    mknod -m 666 /dev/nvidia$i c 195 $i
  done
fi
```

```

done

mknod -m 666 /dev/nvidiactl c 195 255

else
  exit 1
fi

/sbin/modprobe nvidia-uvvm

if [ "$?" -eq 0 ]; then
  # Find out the major device number used by the nvidia-uvvm driver
  D=`grep nvidia-uvvm /proc/devices | awk '{print $1}'`

  mknod -m 666 /dev/nvidia-uvvm c $D 0
else
  exit 1
fi

```

8.5. Advanced Options

Action	Options Used	Explanation
Silent Installation	<code>--silent</code>	Required for any silent installation. Performs an installation with no further user-input and minimal command-line output based on the options provided below. Silent installations are useful for scripting the installation of CUDA. Using this option implies acceptance of the EULA. The following flags can be used to customize the actions taken during installation. At least one of <code>--driver</code> , <code>--uninstall</code> , and <code>--toolkit</code> must be passed if running with non-root permissions.
	<code>--driver</code>	Install the CUDA Driver.
	<code>--toolkit</code>	Install the CUDA Toolkit.
	<code>--toolkitpath=<path></code>	Install the CUDA Toolkit to the <code><path></code> directory. If not provided, the default path of <code>/usr/local/cuda-11.7</code> is used.
	<code>--defaultroot=<path></code>	Install libraries to the <code><path></code> directory. If the <code><path></code> is not provided, then the default path of your distribution is used. <i>This only applies to the libraries installed outside of the CUDA Toolkit path.</i>
Extraction	<code>--extract=<path></code>	Extracts to the <code><path></code> the following: the driver runfile, the raw files of the toolkit to <code><path></code> . This is especially useful when one wants to install the driver using one or more of the command-line options provided by the driver installer which are not exposed in this installer.
Overriding Installation Checks	<code>--override</code>	Ignores compiler, third-party library, and toolkit detection checks which would prevent the CUDA Toolkit from installing.

Action	Options Used	Explanation
No OpenGL Libraries	<code>--no-opengl-libs</code>	Prevents the driver installation from installing NVIDIA's GL libraries. Useful for systems where the display is driven by a non-NVIDIA GPU. In such systems, NVIDIA's GL libraries could prevent X from loading properly.
No man pages	<code>--no-man-page</code>	Do not install the man pages under <code>/usr/share/man</code> .
Overriding Kernel Source	<code>--kernel-source-path=<path></code>	Tells the driver installation to use <code><path></code> as the kernel source directory when building the NVIDIA kernel module. Required for systems where the kernel source is installed to a non-standard location.
Running nvidia-xconfig	<code>--run-nvidia-xconfig</code>	Tells the driver installation to run <code>nvidia-xconfig</code> to update the system X configuration file so that the NVIDIA X driver is used. The pre-existing X configuration file will be backed up.
No nvidia-drm kernel module	<code>--no-drm</code>	Do not install the <code>nvidia-drm</code> kernel module. This option should only be used to work around failures to build or install the <code>nvidia-drm</code> kernel module on systems that do not need the provided features.
Custom Temporary Directory Selection	<code>--tmpdir=<path></code>	Performs any temporary actions within <code><path></code> instead of <code>/tmp</code> . Useful in cases where <code>/tmp</code> cannot be used (doesn't exist, is full, is mounted with 'noexec', etc.).
Show Installer Options	<code>--help</code>	Prints the list of command-line options to stdout.

8.6. Uninstallation

To uninstall the CUDA Toolkit, run the uninstallation script provided in the bin directory of the toolkit. By default, it is located in `/usr/local/cuda-11.7/bin`:

```
sudo /usr/local/cuda-11.7/bin/cuda-uninstaller
```

To uninstall the NVIDIA Driver, run `nvidia-uninstall`:

```
sudo /usr/bin/nvidia-uninstall
```

To enable the Nouveau drivers, remove the blacklist file created in the [Disabling Nouveau](#) section, and regenerate the kernel `initramfs/initrd` again as described in that section.

Chapter 9. Conda Installation

This section describes the installation and configuration of CUDA when using the Conda installer. The Conda packages are available at <https://anaconda.org/nvidia>.

9.1. Conda Overview

The Conda installation installs the CUDA Toolkit. The installation steps are listed below.

9.2. Installation

To perform a basic install of all CUDA Toolkit components using Conda, run the following command:

```
conda install cuda -c nvidia
```

9.3. Uninstallation

To uninstall the CUDA Toolkit using Conda, run the following command:

```
conda remove cuda
```

9.4. Installing Previous CUDA Releases

All Conda packages released under a specific CUDA version are labeled with that release version. To install a previous version, include that label in the `install` command such as:

```
conda install cuda -c nvidia/label/cuda-11.3.0
```

**Note:**

Some CUDA releases do not move to new versions of all installable components. When this is the case these components will be moved to the new label, and you may need to modify the install command to include both labels such as:

```
conda install cuda -c nvidia/label/cuda-11.3.0 -c nvidia/label/cuda-11.3.1
```


This example will install all packages released as part of CUDA 11.3.0.

Chapter 10. Pip Wheels

NVIDIA provides Python Wheels for installing CUDA through pip, primarily for using CUDA with Python. These packages are intended for runtime use and do not currently include developer tools (these can be installed separately).

Please note that with this installation method, CUDA installation environment is managed via pip and additional care must be taken to set up your host environment to use CUDA outside the pip environment.

Prerequisites

To install Wheels, you must first install the `nvidia-pyindex` package, which is required in order to set up your pip installation to fetch additional Python modules from the NVIDIA NGC PyPI repo. If your pip and setuptools Python modules are not up-to-date, then use the following command to upgrade these Python modules. If these Python modules are out-of-date then the commands which follow later in this section may fail.

```
python3 -m pip install --upgrade setuptools pip wheel
```

You should now be able to install the `nvidia-pyindex` module.

```
python3 -m pip install nvidia-pyindex
```

If your project is using a `requirements.txt` file, then you can add the following line to your `requirements.txt` file as an alternative to installing the `nvidia-pyindex` package:

```
--extra-index-url https://pypi.ngc.nvidia.com
```

Procedure

Install the CUDA runtime package:

```
python3 -m pip install nvidia-cuda-runtime-cu11
```

Optionally, install additional packages as listed below using the following command:

```
python3 -m pip install nvidia-<library>
```

Metapackages

The following metapackages will install the latest version of the named component on Linux for the indicated CUDA version. "cu11" should be read as "cuda11".

- ▶ `nvidia-cuda-runtime-cu11`

- ▶ `nvidia-cuda-cupti-cu11`
- ▶ `nvidia-cuda-nvcc-cu11`
- ▶ `nvidia-nvml-dev-cu11`
- ▶ `nvidia-cuda-nvrtc-cu11`
- ▶ `nvidia-nvtx-cu11`
- ▶ `nvidia-cuda-sanitizer-api-cu11`
- ▶ `nvidia-cublas-cu11`
- ▶ `nvidia-cufft-cu11`
- ▶ `nvidia-curand-cu11`
- ▶ `nvidia-cusolver-cu11`
- ▶ `nvidia-cuspars-cu11`
- ▶ `nvidia-npp-cu11`
- ▶ `nvidia-nvjpeg-cu11`


These metapackages install the following packages:

- ▶ `nvidia-nvml-dev-cu114`
- ▶ `nvidia-cuda-nvcc-cu114`
- ▶ `nvidia-cuda-runtime-cu114`
- ▶ `nvidia-cuda-cupti-cu114`
- ▶ `nvidia-cublas-cu114`
- ▶ `nvidia-cuda-sanitizer-api-cu114`
- ▶ `nvidia-nvtx-cu114`
- ▶ `nvidia-cuda-nvrtc-cu114`
- ▶ `nvidia-npp-cu114`
- ▶ `nvidia-cuspars-cu114`
- ▶ `nvidia-cusolver-cu114`
- ▶ `nvidia-curand-cu114`
- ▶ `nvidia-cufft-cu114`
- ▶ `nvidia-nvjpeg-cu114`

Chapter 11. Tarball and Zip Archive Deliverables

In an effort to meet the needs of a growing customer base requiring alternative installer packaging formats, as well as a means of input into community CI/CD systems, tarball and zip archives for each component.

These tarball and zip archives are provided at <https://developer.download.nvidia.com/compute/cuda/redist/>.



These .tar.xz and .zip archives do not replace existing packages such as .deb, .rpm, runfile, conda, etc. and are not meant for general consumption, as they are not installers. However this standardized approach will replace existing .txz archives.

For each release, a JSON manifest is provided such as **redistrib_11.4.2.json**, which corresponds to the CUDA 11.4.2 release label (CUDA 11.4 update 2) which includes the release date, the name of each component, license name, relative URL for each platform and checksums.

Package maintainers are advised to check the provided LICENSE for each component prior to redistribution. Instructions for developers using CMake and Bazel build systems are provided in the next sections.

11.1. Parsing Redistrib JSON

The following example of a JSON manifest contains keys for each component: name, license, version, and a platform array which includes relative_path, sha256, md5, and size (bytes) for each archive.

```
{
  "release_date": "2021-09-07",
  "cuda_cudart": {
    "name": "CUDA Runtime (cudart)",
    "license": "CUDA Toolkit",
    "version": "11.4.108",
    "linux-x86_64": {
      "relative_path": "cuda_cudart/linux-x86_64/cuda_cudart-linux-
x86_64-11.4.108-archive.tar.xz",
      "sha256":
"d08a1b731e5175aa3ae06a6d1c6b3059dd9ea13836d947018ea5e3ec2ca3d62b",
      "md5": "da198656b27a3559004c3b7f20e5d074",
      "size": "828300"
    },
    "linux-ppc64le": {
      "relative_path": "cuda_cudart/linux-ppc64le/cuda_cudart-linux-
ppc64le-11.4.108-archive.tar.xz",
      "sha256":
"831dffe062ae3ebda3d3c4010d0ee4e40a01fd5e6358098a87bb318ea7c79e0c",
      "md5": "ca73328e3f8e2bb5b1f2184c98c3a510",
      "size": "776840"
    },
    "linux-sbsa": {
      "relative_path": "cuda_cudart/linux-sbsa/cuda_cudart-linux-
sbsa-11.4.108-archive.tar.xz",
      "sha256":
"2ab9599bbaebdcf59add73d1f1a352ae619f8cb5ccec254093c98efd4c14553c",
      "md5": "aeb5c19661f06b6398741015ba368102",
      "size": "782372"
    },
    "windows-x86_64": {
      "relative_path": "cuda_cudart/windows-x86_64/cuda_cudart-windows-
x86_64-11.4.108-archive.zip",
      "sha256":
"b59756c27658d1ea87a17c06d064d1336576431cd64da5d1790d909e455d06d3",
      "md5": "7f6837a46b78198402429a3760ab28fc",
      "size": "2897751"
    }
  }
}
```

A JSON schema is provided at <https://developer.download.nvidia.com/compute/redist/redistrib-v2.schema.json>.

A sample script that parses these JSON manifests is available on [GitHub](#):

- ▶ Downloads each archive
- ▶ Validates SHA256 checksums
- ▶ Extracts archives
- ▶ Flattens into a collapsed directory structure

11.2. Importing Tarballs into CMake

The recommended module for importing these tarballs into the CMake build system is via [FindCUDAToolkit](#) (3.17 and newer).



Note: The FindCUDA module is deprecated.

The path to the extraction location can be specified with the `CUDAToolkit_ROOT` environmental variable. For example `CMakeLists.txt` and commands, see [cmake/1_FindCUDAToolkit/](#).

For older versions of CMake, the [ExternalProject_Add](#) module is an alternative method. For example `CMakeLists.txt` file and commands, see [cmake/2_ExternalProject/](#).

11.3. Importing Tarballs into Bazel

The recommended method of importing these tarballs into the Bazel build system is using [http_archive](#) and [pkg_tar](#).

For an example, see [bazel/1_pkg_tar/](#).

Chapter 12. CUDA Cross-Platform Environment

Cross-platform development is only supported on Ubuntu systems, and is only provided via the Package Manager installation process.

We recommend selecting Ubuntu 18.04 as your cross-platform development environment. This selection helps prevent host/target incompatibilities, such as GCC or GLIBC version mismatches.

12.1. CUDA Cross-Platform Installation

Some of the following steps may have already been performed as part of the [native Ubuntu installation](#). Such steps can safely be skipped.

These steps should be performed on the x86_64 host system, rather than the target system. To install the native CUDA Toolkit on the target system, refer to the native [Ubuntu](#) installation section.

1. Perform the [pre-installation actions](#).
2. Install repository meta-data package with:

```
sudo dpkg -i cuda-repo-cross-<identifier>_all.deb
```

where `<identifier>` indicates the operating system, architecture, and/or the version of the package.

3. Update the Apt repository cache:
4. Install the appropriate cross-platform CUDA Toolkit:

```
sudo apt-get update
```

- a). For aarch64:

```
sudo apt-get install cuda-cross-aarch64
```

- b). For QNX:

```
sudo apt-get install cuda-cross-qnx
```

5. Perform the [post-installation actions](#).

12.2. CUDA Cross-Platform Samples

CUDA Samples are now located in <https://github.com/nvidia/cuda-samples>, which includes instructions for obtaining, building, and running the samples.

Chapter 13. Post-installation Actions

The post-installation actions must be manually performed. These actions are split into mandatory, recommended, and optional sections.

13.1. Mandatory Actions

Some actions must be taken after the installation before the CUDA Toolkit and Driver can be used.

13.1.1. Environment Setup

The `PATH` variable needs to include `export PATH=/usr/local/cuda-11.7/bin${PATH:+:${PATH}}`. Nsight Compute has moved to `/opt/nvidia/nsight-compute/` only in rpm/deb installation method. When using `.run` installer it is still located under `/usr/local/cuda-11.7/`.

To add this path to the `PATH` variable:

```
export PATH=/usr/local/cuda-11.7/bin${PATH:+:${PATH}}
```

In addition, when using the runfile installation method, the `LD_LIBRARY_PATH` variable needs to contain `/usr/local/cuda-11.7/lib64` on a 64-bit system, or `/usr/local/cuda-11.7/lib` on a 32-bit system

- ▶ To change the environment variables for 64-bit operating systems:

```
export LD_LIBRARY_PATH=/usr/local/cuda-11.7/lib64\  
${LD_LIBRARY_PATH:+:${LD_LIBRARY_PATH}}
```

- ▶ To change the environment variables for 32-bit operating systems:

```
export LD_LIBRARY_PATH=/usr/local/cuda-11.7/lib\  
${LD_LIBRARY_PATH:+:${LD_LIBRARY_PATH}}
```

Note that the above paths change when using a custom install path with the runfile installation method.

13.1.2. POWER9 Setup

Because of the addition of new features specific to the NVIDIA POWER9 CUDA driver, there are some additional setup requirements in order for the driver to function properly. These

additional steps are not handled by the installation of CUDA packages, and failure to ensure these extra requirements are met will result in a non-functional CUDA driver installation.

There are two changes that need to be made manually after installing the NVIDIA CUDA driver to ensure proper operation:

1. The NVIDIA Persistence Daemon should be automatically started for POWER9 installations. Check that it is running with the following command:

```
systemctl status nvidia-persistenced
```

If it is not active, run the following command:

```
sudo systemctl enable nvidia-persistenced
```

2. Disable a udev rule installed by default in some Linux distributions that cause hot-pluggable memory to be automatically onlined when it is physically probed. This behavior prevents NVIDIA software from bringing NVIDIA device memory online with non-default settings. This udev rule must be disabled in order for the NVIDIA CUDA driver to function properly on POWER9 systems.

On RedHat Enterprise Linux 8.1, this rule can be found in:

```
/lib/udev/rules.d/40-redhat.rules
```

On Ubuntu 18.04, this rule can be found in:

```
/lib/udev/rules.d/40-vm-hotadd.rules
```

The rule generally takes a form where it detects the addition of a memory block and changes the 'state' attribute to online. For example, in RHEL8, the rule looks like this:

```
SUBSYSTEM=="memory", ACTION=="add", PROGRAM="/bin/uname -p", RESULT!="s390*",
ATTR{state}=="offline", ATTR{state}="online"
```

This rule must be disabled by copying the file to `/etc/udev/rules.d` and commenting out, removing, or changing the hot-pluggable memory rule in the `/etc` copy so that it does not apply to POWER9 NVIDIA systems. For example, on RHEL 7.5 and earlier:

```
sudo cp /lib/udev/rules.d/40-redhat.rules /etc/udev/rules.d
sudo sed -i 's/SUBSYSTEM=="memory", ACTION=="add"/d' /etc/udev/rules.d/40-
redhat.rules
```

On RHEL 7.6 and later versions:

```
sudo cp /lib/udev/rules.d/40-redhat.rules /etc/udev/rules.d
sudo sed -i 's/SUBSYSTEM!="memory",.*GOTO="memory_hotplug_end"/SUBSYSTEM=="*",
GOTO="memory_hotplug_end"/' /etc/udev/rules.d/40-redhat.rules
```

You will need to reboot the system to initialize the above changes.



Note: For NUMA best practices on IBM Newell POWER9, see [NUMA Best Practices](#).

13.2. Recommended Actions

Other actions are recommended to verify the integrity of the installation.

13.2.1. Install Persistence Daemon

NVIDIA is providing a user-space daemon on Linux to support persistence of driver state across CUDA job runs. The daemon approach provides a more elegant and robust solution to this problem than persistence mode. For more details on the NVIDIA Persistence Daemon, see the documentation [here](#).

The NVIDIA Persistence Daemon can be started as the root user by running:

```
/usr/bin/nvidia-persistenced --verbose
```

This command should be run on boot. Consult your Linux distribution's init documentation for details on how to automate this.

13.2.2. Install Writable Samples

CUDA Samples are now located in <https://github.com/nvidia/cuda-samples>, which includes instructions for obtaining, building, and running the samples.

13.2.3. Verify the Installation

Before continuing, it is important to verify that the CUDA toolkit can find and communicate correctly with the CUDA-capable hardware. To do this, you need to compile and run some of the sample programs, located in <https://github.com/nvidia/cuda-samples>.



Note: Ensure the PATH and, if using the runfile installation method, LD_LIBRARY_PATH variables are [set correctly](#).

13.2.3.1. Verify the Driver Version

If you installed the driver, verify that the correct version of it is loaded. If you did not install the driver, or are using an operating system where the driver is not loaded via a kernel module, such as L4T, skip this step.

When the driver is loaded, the driver version can be found by executing the command

```
cat /proc/driver/nvidia/version
```

Note that this command will not work on an iGPU/dGPU system.

13.2.3.2. Running the Binaries

After compilation, find and run `deviceQuery` from <https://github.com/nvidia/cuda-samples>. If the CUDA software is installed and configured correctly, the output for `deviceQuery` should look similar to that shown in [Figure 1](#).

Figure 1. Valid Results from `deviceQuery` CUDA Sample

```

./deviceQuery Starting...

CUDA Device Query (Runtime API) version (CUDA static linking)

Detected 1 CUDA Capable device(s)

Device 0: "Tesla K20c"
  CUDA Driver Version / Runtime Version      6.0 / 6.0
  CUDA Capability Major/Minor version number: 3.5
  Total amount of global memory:             4800 MBytes (5032706048 bytes)
  (13) Multiprocessors, (192) CUDA Cores/MP: 2496 CUDA Cores
  GPU Clock rate:                           706 MHz (0.71 GHz)
  Memory Clock rate:                         2600 Mhz
  Memory Bus Width:                          320-bit
  L2 Cache Size:                             1310720 bytes
  Maximum Texture Dimension Size (x,y,z)    1D=(65536), 2D=(65536, 65536), 3D=(4096, 4096, 4096)
  Maximum Layered 1D Texture Size, (num) layers 1D=(16384), 2048 Layers
  Maximum Layered 2D Texture Size, (num) layers 2D=(16384, 16384), 2048 Layers
  Total amount of constant memory:          65536 bytes
  Total amount of shared memory per block:   49152 bytes
  Total number of registers available per block: 65536
  Warp size:                                 32
  Maximum number of threads per multiprocessor: 2048
  Maximum number of threads per block:      1024
  Max dimension size of a thread block (x,y,z): (1024, 1024, 64)
  Max dimension size of a grid size (x,y,z): (2147483647, 65535, 65535)
  Maximum memory pitch:                     2147483647 bytes
  Texture alignment:                         512 bytes
  Concurrent copy and kernel execution:     Yes with 2 copy engine(s)
  Run time limit on kernels:                 No
  Integrated GPU sharing Host Memory:        No
  Support host page-locked memory mapping:  Yes
  Alignment requirement for Surfaces:        Yes
  Device has ECC support:                    Enabled
  Device supports Unified Addressing (UVA):  Yes
  Device PCI Bus ID / PCI location ID:      2 / 0
  Compute Mode:
    < Default (multiple host threads can use ::cudaSetDevice() with device simultaneously) >

deviceQuery, CUDA Driver = CUDART, CUDA Driver Version = 6.0, CUDA Runtime Version = 6.0, NumDevs = 1, Device0 = Tesla K20c
Result = PASS

```

The exact appearance and the output lines might be different on your system. The important outcomes are that a device was found (the first highlighted line), that the device matches the one on your system (the second highlighted line), and that the test passed (the final highlighted line).

If a CUDA-capable device and the CUDA Driver are installed but `deviceQuery` reports that no CUDA-capable devices are present, this likely means that the `/dev/nvidia*` files are missing or have the wrong permissions.


On systems where `SELinux` is enabled, you might need to temporarily disable this security feature to run `deviceQuery`. To do this, type:

```
setenforce 0
```

from the command line as the *superuser*.

Running the `bandwidthTest` program ensures that the system and the CUDA-capable device are able to communicate correctly. Its output is shown in [Figure 2](#).

Figure 2. Valid Results from bandwidthTest CUDA Sample



```
[CUDA Bandwidth Test] - Starting...
Running on...

Device 0: Quadro K5000
Quick Mode

Host to Device Bandwidth, 1 Device(s)
PINNED Memory Transfers
  Transfer Size (Bytes)      Bandwidth(MB/s)
  33554432                   5798.4

Device to Host Bandwidth, 1 Device(s)
PINNED Memory Transfers
  Transfer Size (Bytes)      Bandwidth(MB/s)
  33554432                   6378.4

Device to Device Bandwidth, 1 Device(s)
PINNED Memory Transfers
  Transfer Size (Bytes)      Bandwidth(MB/s)
  33554432                   133606.8

Result = PASS
```

Note that the measurements for your CUDA-capable device description will vary from system to system. The important point is that you obtain measurements, and that the second-to-last line (in [Figure 2](#)) confirms that all necessary tests passed.

Should the tests not pass, make sure you have a CUDA-capable NVIDIA GPU on your system and make sure it is properly installed.

If you run into difficulties with the link step (such as libraries not being found), consult the *Linux Release Notes* found in <https://github.com/nvidia/cuda-samples>.

13.2.4. Install Nsight Eclipse Plugins

To install Nsight Eclipse plugins, an installation script is provided:

```
/usr/local/cuda-11.7/bin/nsight_ee_plugins_manage.sh install <eclipse-dir>
```

Refer to [Nsight Eclipse Plugins Installation Guide](#) for more details.

13.3. Optional Actions

Other options are not necessary to use the CUDA Toolkit, but are available to provide additional features.

13.3.1. Install Third-party Libraries

Some CUDA samples use third-party libraries which may not be installed by default on your system. These samples attempt to detect any required libraries when building.

If a library is not detected, it waives itself and warns you which library is missing. To build and run these samples, you must install the missing libraries. These dependencies may be installed if the RPM or Deb `cuda-samples-11-7` package is used. In cases where these dependencies are not installed, follow the instructions below.

RHEL 7 / CentOS 7

```
sudo yum install freeglut-devel libX11-devel libXi-devel libXmu-devel \
make mesa-libGLU-devel freeimage-devel
```

RHEL 8 / Rocky Linux 8

```
sudo dnf install freeglut-devel libX11-devel libXi-devel libXmu-devel \
make mesa-libGLU-devel freeimage-devel
```

Fedora

```
sudo dnf install freeglut-devel libX11-devel libXi-devel libXmu-devel \
make mesa-libGLU-devel freeimage-devel
```

SLES

```
sudo zypper install libglut3 libX11-devel libXi6 libXmu6 libGLU1 make
```

OpenSUSE

```
sudo zypper install freeglut-devel libX11-devel libXi-devel libXmu-devel \
make Mesa-libGL-devel freeimage-devel
```

Ubuntu

```
sudo apt-get install g++ freeglut3-dev build-essential libx11-dev \
libxmu-dev libxi-dev libglul-mesa libglul-mesa-dev libfreeimage-dev
```

Debian

```
sudo apt-get install g++ freeglut3-dev build-essential libx11-dev \
libxmu-dev libxi-dev libglul-mesa libglul-mesa-dev libfreeimage-dev
```

13.3.2. Install the Source Code for `cuda-gdb`

The `cuda-gdb` source must be explicitly selected for installation with the runfile installation method. During the installation, in the component selection page, expand the component "CUDA Tools 11.0" and select `cuda-gdb-src` for installation. It is unchecked by default.

To obtain a copy of the source code for `cuda-gdb` using the RPM and Debian installation methods, the `cuda-gdb-src` package must be installed.

The source code is installed as a tarball in the `/usr/local/cuda-11.7/extras` directory.

13.3.3. Select the Active Version of CUDA

For applications that rely on the symlinks `/usr/local/cuda` and `/usr/local/cuda-MAJOR`, you may wish to change to a different installed version of CUDA using the provided alternatives.

To show the active version of CUDA and all available versions:

```
update-alternatives --display cuda
```

To show the active minor version of a given major CUDA release:

```
update-alternatives --display cuda-11
```

To update the active version of CUDA:

```
sudo update-alternatives --config cuda
```

Chapter 14. Advanced Setup

Below is information on some advanced setup scenarios which are not covered in the basic instructions above.

Table 6. Advanced Setup Scenarios when Installing CUDA

Scenario	Instructions
Install CUDA using the Package Manager installation method without installing the NVIDIA GL libraries.	<p>Fedora</p> <p>Install CUDA using the following command:</p> <pre>sudo dnf install cuda-toolkit-11-7 \ nvidia-driver-cuda akmod-nvidia</pre> <p>Follow the instructions here to ensure that Nouveau is disabled.</p> <p>If performing an upgrade over a previous installation, the NVIDIA kernel module may need to be rebuilt by following the instructions here.</p> <p>OpenSUSE/SLES</p> <p>On some system configurations the NVIDIA GL libraries may need to be locked before installation using:</p> <pre>sudo zypper addlock nvidia-glg04</pre> <p>Install CUDA using the following command:</p> <pre>sudo zypper install --no-recommends cuda-toolkit-11-7 \ nvidia-computeG04 \ nvidia-gfxG04-kmp-default</pre> <p>Follow the instructions here to ensure that Nouveau is disabled.</p> <p>Ubuntu</p> <p>This functionality isn't supported on Ubuntu. Instead, the driver packages integrate with the Bumblebee framework to provide a solution for users who wish to control what applications the NVIDIA drivers are used for. See Ubuntu's Bumblebee wiki for more information.</p>
Upgrade from a RPM/ Deb driver installation which includes the diagnostic driver packages to a driver	<p>RHEL/CentOS</p> <p>Remove diagnostic packages using the following command:</p> <pre>sudo yum remove cuda-drivers-diagnostic \ xorg-x11-drv-nvidia-diagnostic</pre>

Scenario	Instructions
<p>installation which does not include the diagnostic driver packages.</p>	<p>Follow the instructions here to continue installation as normal.</p> <p>Fedora</p> <p>Remove diagnostic packages using the following command:</p> <pre>sudo dnf remove cuda-drivers-diagnostic \ xorg-x11-drv-nvidia-diagnostic</pre> <p>Follow the instructions here to continue installation as normal.</p> <p>OpenSUSE/SLES</p> <p>Remove diagnostic packages using the following command:</p> <pre>sudo zypper remove cuda-drivers-diagnostic \ nvidia-diagnosticG04</pre> <p>Follow the instructions here to continue installation as normal.</p> <p>Ubuntu</p> <p>Remove diagnostic packages using the following command:</p> <pre>sudo apt-get purge cuda-drivers-diagnostic \ nvidia-384-diagnostic</pre> <p>Follow the instructions here to continue installation as normal.</p>
<p>Use a specific GPU for rendering the display.</p>	<p>Add or replace a Device entry in your <code>xorg.conf</code> file, located at <code>/etc/X11/xorg.conf</code>. The Device entry should resemble the following:</p> <pre>Section "Device" Identifier "Device0" Driver "driver_name" VendorName "vendor_name" BusID "bus_id" EndSection</pre> <p>The details you will need to add differ on a case-by-case basis. For example, if you have two NVIDIA GPUs and you want the first GPU to be used for display, you would replace <code>"driver_name"</code> with <code>"nvidia"</code>, <code>"vendor_name"</code> with <code>"NVIDIA Corporation"</code> and <code>"bus_id"</code> with the Bus ID of the GPU.</p> <p>The Bus ID will resemble <code>"PCI:00:02.0"</code> and can be found by running <code>lspci</code>.</p>
<p>Install CUDA to a specific directory using the Package Manager installation method.</p>	<p>RPM</p> <p>The RPM packages don't support custom install locations through the package managers (Yum and Zypper), but it is possible to install the RPM packages to a custom location using rpm's <code>--relocate</code> parameter:</p> <pre>sudo rpm --install --relocate /usr/local/cuda-11.7=/new/ toolkit package.rpm</pre> <p>You will need to install the packages in the correct dependency order; this task is normally taken care of by the package managers. For example, if package <code>"foo"</code> has a dependency on package <code>"bar"</code>, you should install</p>

Scenario	Instructions
	<p>package "bar" first, and package "foo" second. You can check the dependencies of a RPM package as follows:</p> <pre>rpm -qRp package.rpm</pre> <p>Note that the driver packages cannot be relocated.</p> <p>Deb</p> <p>The Deb packages do not support custom install locations. It is however possible to extract the contents of the Deb packages and move the files to the desired install location. See the next scenario for more details on extracting Deb packages.</p>
Extract the contents of the installers.	<p>Runfile</p> <p>The Runfile can be extracted into the standalone Toolkit and Driver Runfiles by using the <code>--extract</code> parameter. The Toolkit standalone Runfiles can be further extracted by running:</p> <pre>./runfile.run --tar mxvf</pre> <p>The Driver Runfile can be extracted by running:</p> <pre>./runfile.run -x</pre> <p>RPM</p> <p>The RPM packages can be extracted by running:</p> <pre>rpm2cpio package.rpm cpio -idmv</pre> <p>Deb</p> <p>The Deb packages can be extracted by running:</p> <pre>dpkg-deb -x package.deb output_dir</pre>
<p>Modify Ubuntu's apt package manager to query specific architectures for specific repositories.</p> <p>This is useful when a foreign architecture has been added, causing "404 Not Found" errors to appear when the repository meta-data is updated.</p>	<p>Each repository you wish to restrict to specific architectures must have its <code>sources.list</code> entry modified. This is done by modifying the <code>/etc/apt/sources.list</code> file and any files containing repositories you wish to restrict under the <code>/etc/apt/sources.list.d/</code> directory. Normally, it is sufficient to modify only the entries in <code>/etc/apt/sources.list</code></p> <p>An architecture-restricted repository entry looks like:</p> <pre>deb [arch=<arch1>,<arch2>] <url></pre> <p>For example, if you wanted to restrict a repository to only the amd64 and i386 architectures, it would look like:</p> <pre>deb [arch=amd64,i386] <url></pre> <p>It is not necessary to restrict the <code>deb-src</code> repositories, as these repositories don't provide architecture-specific packages.</p> <p>For more details, see the <code>sources.list</code> manpage.</p>

Scenario	Instructions
<p>The nvidia.ko kernel module fails to load, saying some symbols are unknown.</p> <p>For example:</p> <pre>nvidia: Unknown symbol drm_open (err 0)</pre>	<p>Check to see if there are any optionally installable modules that might provide these symbols which are not currently installed.</p> <p>For the example of the <code>drm_open</code> symbol, check to see if there are any packages which provide <code>drm_open</code> and are not already installed. For instance, on Ubuntu 14.04, the <code>linux-image-extra</code> package provides the DRM kernel module (which provides <code>drm_open</code>). This package is optional even though the kernel headers reflect the availability of DRM regardless of whether this package is installed or not.</p>
<p>The runfile installer fails to extract due to limited space in the TMP directory.</p>	<p>This can occur on systems with limited storage in the TMP directory (usually <code>/tmp</code>), or on systems which use a <code>tmpfs</code> in memory to handle temporary storage. In this case, the <code>--tmpdir</code> command-line option should be used to instruct the runfile to use a directory with sufficient space to extract into. More information on this option can be found here.</p>
<p>Re-enable Wayland after installing the RPM driver on Fedora.</p>	<p>Wayland is disabled during installation of the Fedora driver RPM due to compatibility issues. To re-enable Wayland, comment out this line in <code>/etc/gdm/custom.conf</code>:</p> <pre>WaylandEnable=false</pre>
<p>In case of the error:</p> <pre>E: Failed to fetch file:/var/cuda-repo File not found</pre>	<p>Debian and Ubuntu</p> <p>This can occur when installing CUDA after uninstalling a different version. Use the following command before installation:</p> <pre>sudo rm -v /var/lib/apt/lists/*cuda* /var/lib/apt/lists/*nvidia*</pre>

Chapter 15. Frequently Asked Questions

How do I install the Toolkit in a different location?

The Runfile installation asks where you wish to install the Toolkit during an interactive install. If installing using a non-interactive install, you can use the `--toolkitpath` parameter to change the install location:

```
./runfile.run --silent \  
              --toolkit --toolkitpath=/my/new/toolkit
```

The RPM and Deb packages cannot be installed to a custom install location directly using the package managers. See the "Install CUDA to a specific directory using the Package Manager installation method" scenario in the [Advanced Setup](#) section for more information.

Why do I see "nvcc: No such file or directory" when I try to build a CUDA application?

Your PATH environment variable is not set up correctly. Ensure that your PATH includes the bin directory where you installed the Toolkit, usually `/usr/local/cuda-11.7/bin`.

```
export PATH=/usr/local/cuda-11.7/bin${PATH:+:${PATH}}
```

Why do I see "error while loading shared libraries: <lib name>: cannot open shared object file: No such file or directory" when I try to run a CUDA application that uses a CUDA library?

Your `LD_LIBRARY_PATH` environment variable is not set up correctly. Ensure that your `LD_LIBRARY_PATH` includes the `lib` and/or `lib64` directory where you installed the Toolkit, usually `/usr/local/cuda-11.7/lib{,64}`:

```
export LD_LIBRARY_PATH=/usr/local/cuda-11.7/lib\
${LD_LIBRARY_PATH:+:${LD_LIBRARY_PATH}}
```

Why do I see multiple "404 Not Found" errors when updating my repository meta-data on Ubuntu?

These errors occur after adding a foreign architecture because `apt` is attempting to query for each architecture within each repository listed in the system's `sources.list` file. Repositories that do not host packages for the newly added architecture will present this error. While noisy, the error itself does no harm. Please see the [Advanced Setup](#) section for details on how to modify your `sources.list` file to prevent these errors.

How can I tell X to ignore a GPU for compute-only use?

To make sure `X` doesn't use a certain GPU for display, you need to specify which **other** GPU to use for display. For more information, please refer to the "Use a specific GPU for rendering the display" scenario in the [Advanced Setup](#) section.

Why doesn't the `cuda-repo` package install the CUDA Toolkit and Drivers?

When using `RPM` or `Deb`, the downloaded package is a repository package. Such a package only informs the package manager where to find the actual installation packages, but will not install them.

See the [Package Manager Installation](#) section for more details.

How do I get CUDA to work on a laptop with an iGPU and a dGPU running Ubuntu14.04?

After installing CUDA, set the driver value for the `intel` device in `/etc/X11/xorg.conf` to `'modesetting'` as shown below:

```
Section "Device"
    Identifier "intel"
    Driver "modesetting"
    ...
EndSection
```

To prevent Ubuntu from reverting the change in `xorg.conf`, edit `/etc/default/grub` to add `"nougpu manager"` to `GRUB_CMDLINE_LINUX_DEFAULT`.

Run the following command to update grub before rebooting:

```
sudo update-grub
```

What do I do if the display does not load, or CUDA does not work, after performing a system update?

System updates may include an updated Linux kernel. In many cases, a new Linux kernel will be installed without properly updating the required Linux kernel headers and development packages. To ensure the CUDA driver continues to work when performing a system update, rerun the commands in the [Kernel Headers and Development Packages](#) section.

Additionally, on Fedora, the Akmods framework will sometimes fail to correctly rebuild the NVIDIA kernel module packages when a new Linux kernel is installed. When this happens, it is usually sufficient to invoke Akmods manually and regenerate the module mapping files by running the following commands in a virtual console, and then rebooting:

```
sudo akmods --force
sudo depmod
```

You can reach a virtual console by hitting `ctrl+alt+f2` at the same time.

How do I install a CUDA driver with a version less than 367 using a network repo?

To install a CUDA driver at a version earlier than 367 using a network repo, the required packages will need to be explicitly installed at the desired version. For example, to install 352.99, instead of installing the `cuda-drivers` metapackage at version 352.99, you will need to install all required packages of `cuda-drivers` at version 352.99.

How do I install an older CUDA version using a network repo?

Depending on your system configuration, you may not be able to install old versions of CUDA using the `cuda` metapackage. In order to install a specific version of CUDA, you may need to specify all of the packages that would normally be installed by the `cuda` metapackage at the version you want to install.

If you are using `yum` to install certain packages at an older version, the dependencies may not resolve as expected. In this case you may need to pass `--setopt=obsoletes=0` to `yum` to allow an install of packages which are obsoleted at a later version than you are trying to install.

Why does the installation on SUSE install the Mesa-dri-nouveau dependency?

This dependency comes from the SUSE repositories and shouldn't affect the use of the NVIDIA driver or the CUDA Toolkit. To disable this dependency, you can lock that package with the following command:

```
sudo zypper al Mesa-dri-nouveau
```

How do I handle "Errors were encountered while processing: glx-diversions"?

This sometimes occurs when trying to uninstall CUDA after a clean .deb installation. Run the following commands:

```
sudo apt-get install glx-diversions --reinstall  
sudo apt-get remove nvidia-alternative
```

Then re-run the commands from [Removing CUDA Toolkit and Driver](#).

Chapter 16. Additional Considerations

Now that you have CUDA-capable hardware and the NVIDIA CUDA Toolkit installed, you can examine and enjoy the numerous included programs. To begin using CUDA to accelerate the performance of your own applications, consult the *CUDA C++ Programming Guide*, located in `/usr/local/cuda-11.7/doc`.

A number of helpful development tools are included in the CUDA Toolkit to assist you as you develop your CUDA programs, such as NVIDIA® Nsight™ Eclipse Edition, NVIDIA Visual Profiler, CUDA-GDB, and CUDA-MEMCHECK.

For technical support on programming questions, consult and participate in the developer forums at <https://developer.nvidia.com/cuda/>.

Chapter 17. Removing CUDA Toolkit and Driver

Follow the below steps to properly uninstall the CUDA Toolkit and NVIDIA Drivers from your system. These steps will ensure that the uninstallation will be clean.

RHEL 8 / Rocky Linux 8

To remove CUDA Toolkit:

```
sudo dnf remove "cuda*" "*cublas*" "*cufft*" "*cufile*" "*curand*" \  
"*cusolver*" "*cuspars*" "*gds-tools*" "*npp*" "*nvjpeg*" "nsight"
```

To remove NVIDIA Drivers:

```
sudo dnf module remove --all nvidia-driver
```

To reset the module stream:

```
sudo dnf module reset nvidia-driver
```

RHEL 7 / CentOS 7

To remove CUDA Toolkit:

```
sudo yum remove "cuda*" "*cublas*" "*cufft*" "*cufile*" "*curand*" \  
"*cusolver*" "*cuspars*" "*gds-tools*" "*npp*" "*nvjpeg*" "nsight"
```

To remove NVIDIA Drivers:

```
sudo yum remove "*nvidia*"
```

Fedora

To remove CUDA Toolkit:

```
sudo dnf remove "cuda*" "*cublas*" "*cufft*" "*cufile*" "*curand*" \  
"*cusolver*" "*cuspars*" "*gds-tools*" "*npp*" "*nvjpeg*" "nsight"
```

To remove 3rd party NVIDIA Drivers:

```
sudo dnf remove "*nvidia*"
```

To remove NVIDIA Drivers:

```
sudo dnf module remove --all nvidia-driver
```

To reset the module stream:

```
sudo dnf module reset nvidia-driver
```

OpenSUSE / SLES

To remove CUDA Toolkit:

```
sudo zypper remove "cuda*" "*cublas*" "*cufft*" "*cufile*" "*curand*" \  
"*cusolver*" "*cusparse*" "*gds-tools*" "*npp*" "*nvjpeg*" "nsight*"
```

To remove NVIDIA Drivers:

```
sudo zypper remove "*nvidia*"
```

Ubuntu and Debian

To remove CUDA Toolkit:

```
sudo apt-get --purge remove "cuda*" "*cublas*" "*cufft*" "*cufile*" "*curand*" \  
"*cusolver*" "*cusparse*" "*gds-tools*" "*npp*" "*nvjpeg*" "nsight*"
```

To remove NVIDIA Drivers:

```
sudo apt-get --purge remove "*nvidia*"
```

To clean up the uninstall:

```
sudo apt-get autoremove
```

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

Trademarks

NVIDIA and the NVIDIA logo are trademarks or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2009-2022 NVIDIA Corporation & affiliates. All rights reserved.