

CUDA Roll: Users Guide



Version 5.0 Edition



CUDA Roll: Users Guide :

Version 5.0 Edition

Published Oct 2008

Copyright © 2008 University of California and NVIDIA Corporation

Table of Contents

Preface	v
1. Requirements and Compatibility	1
1.1. Rocks Version.....	1
1.2. Hardware	1
1.3. Compatibility.....	1
2. Installing the CUDA Roll	2
2.1. Installation on a New System.....	2
2.2. Adding the Roll to an Existing Installation	6
2.3. Optional installation of the driver on the Frontend	8
3. Using the CUDA Roll	10
3.1. Summary of Installed Components	10
3.2. Testing the Installation	10
3.3. Troubleshooting.....	11
A. Rocks Copyright	13
B. Third Party Copyrights and Licenses	15
B.1. NVIDIA CUDA License	15
B.2. NVIDIA Driver License	17

List of Figures

2-1. Adding a roll.....	2
2-2. Loading the roll	2
2-3. List of available rolls	3
2-4. Selecting the CUDA Roll	4
2-5. CUDA Roll selected for installation.....	5

Preface

The purpose of the NVIDIA® CUDA™ Roll is to install and configure the device driver and full toolchain necessary to run and develop CUDA programs on a Rocks™ cluster. As of this release, this includes: CUDA Toolkit Version 2.0, CUDA SDK Version 2.0, and the NVIDIA Display Driver Version 177.70.18.

Please visit NVIDIA's site¹ to learn more about CUDA.

Notes

1. <http://www.nvidia.com/cuda>

Chapter 1. Requirements and Compatibility

1.1. Rocks Version

The CUDA Roll is for use with Rocks version 5.0 ("V"). At this time only the x86_64 architecture has been tested.

1.2. Hardware

Using CUDA requires CUDA capable devices be installed in the compute nodes. CUDA capable devices include the graphics cards in the NVIDIA G8x series and beyond and all NVIDIA® Tesla™ products. Please visit [CUDA-Enabled GPU Products¹](#) for more details.

1.3. Compatibility

Compatibility has been verified with the following optional Rolls:

- area51
- ganglia
- grid
- java
- sge
- pbs
- intel

Other Rolls may work but have not been tested.



The use of the "Jumbo Roll" which includes the boot, core, and OS roll is **not** recommended because it lacks the rpms for freeglut.



At this time the CUDA Roll is **not** compatible with the viz roll due to the different ways in which they install the device driver in the nvidia-driver rpm. We hope to have a common driver rpm in a future release.

Notes

1. http://www.nvidia.com/object/cuda_learn_products.html

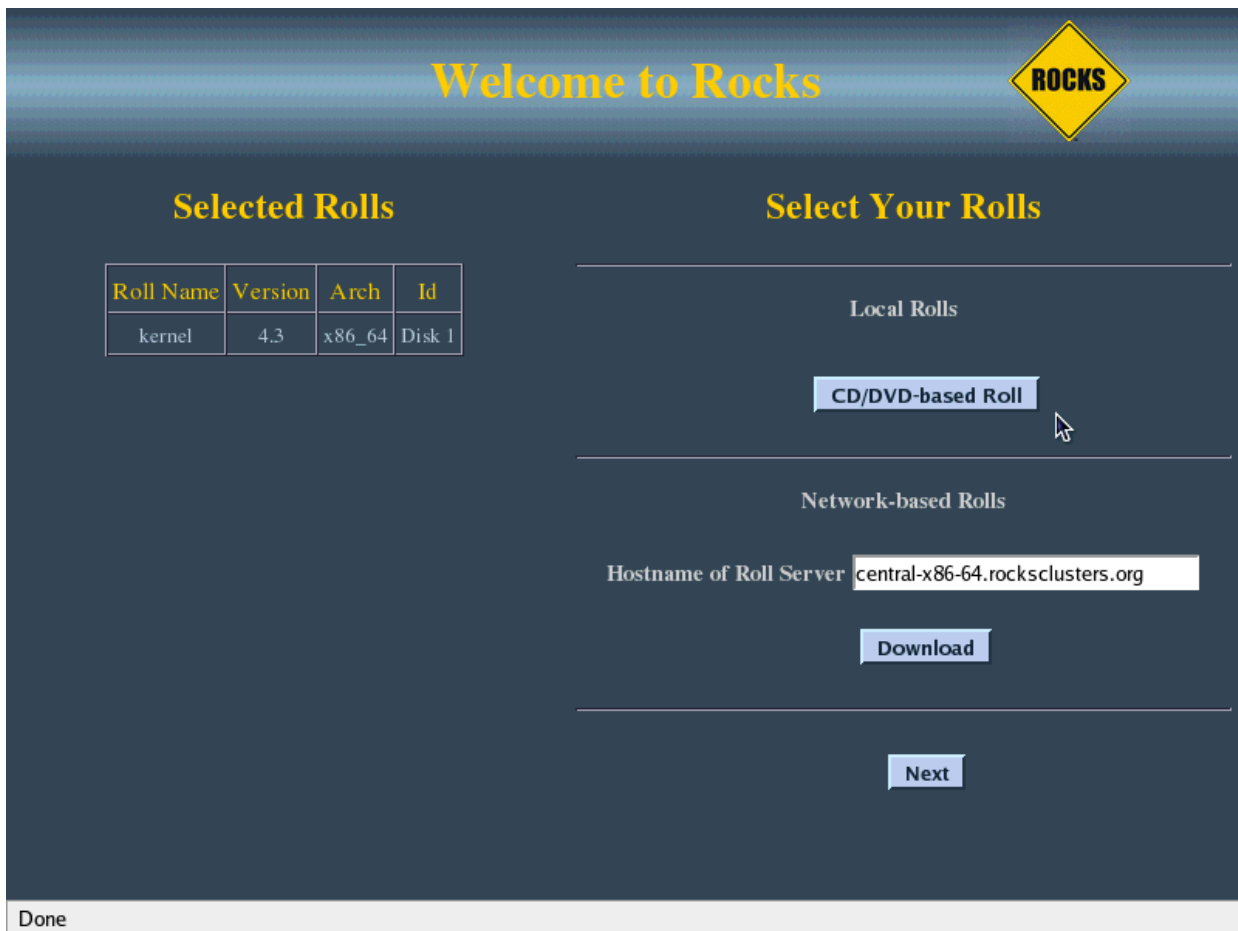
Chapter 2. Installing the CUDA Roll

The CUDA Roll can be installed during the Frontend installation step of your cluster or you can add the CUDA Roll to an existing system.

2.1. Installation on a New System

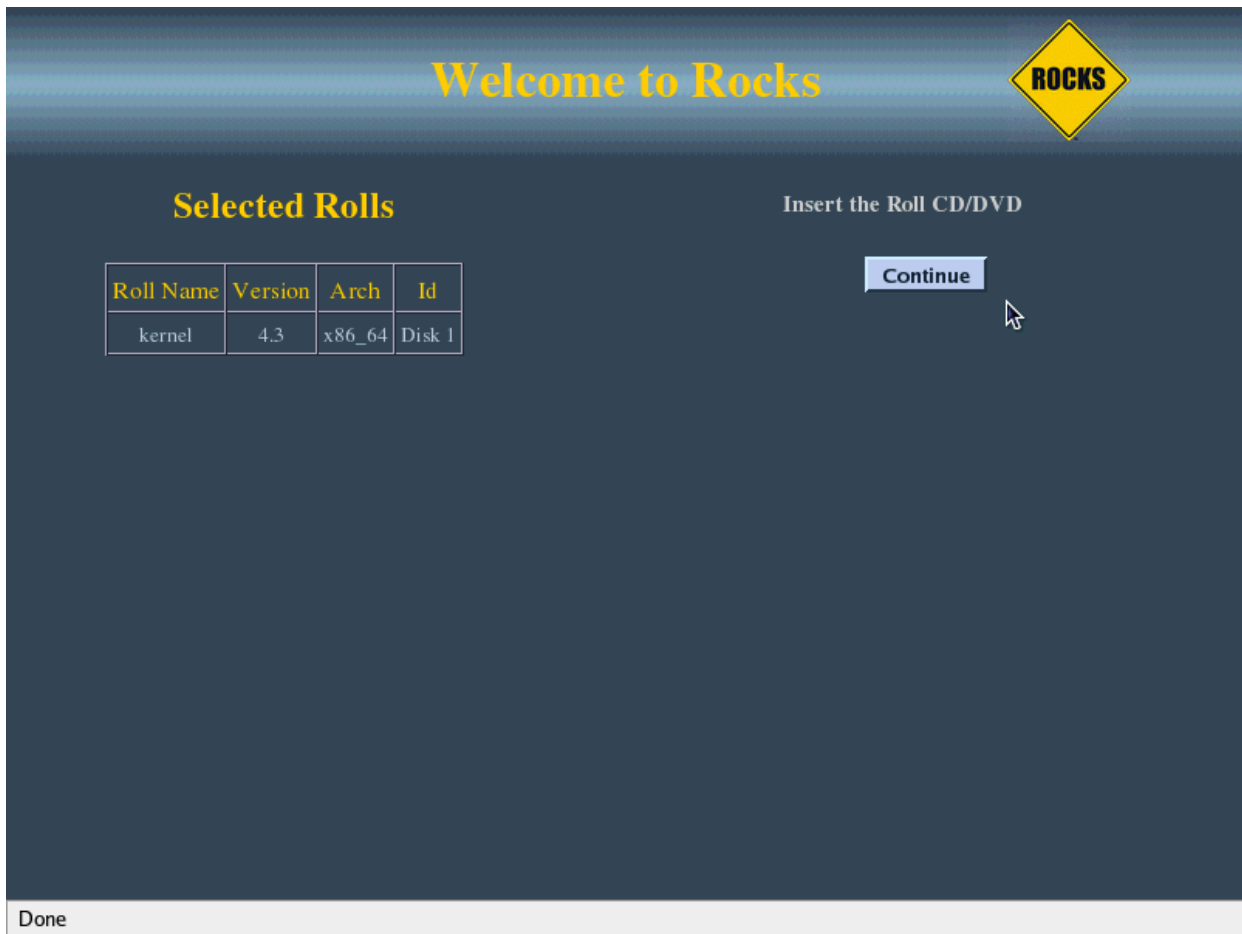
The CUDA Roll is added to a Frontend installation in exactly the same manner as other Rolls. Refer also to Section 2.2 of the Rocks Base Roll: Users Guide for more information. Note that the screenshots in the following figures are from a Rocks 4.3 install as seen in the column "Version". The install process for Rocks 5 is the same.

Figure 2-1. Adding a roll



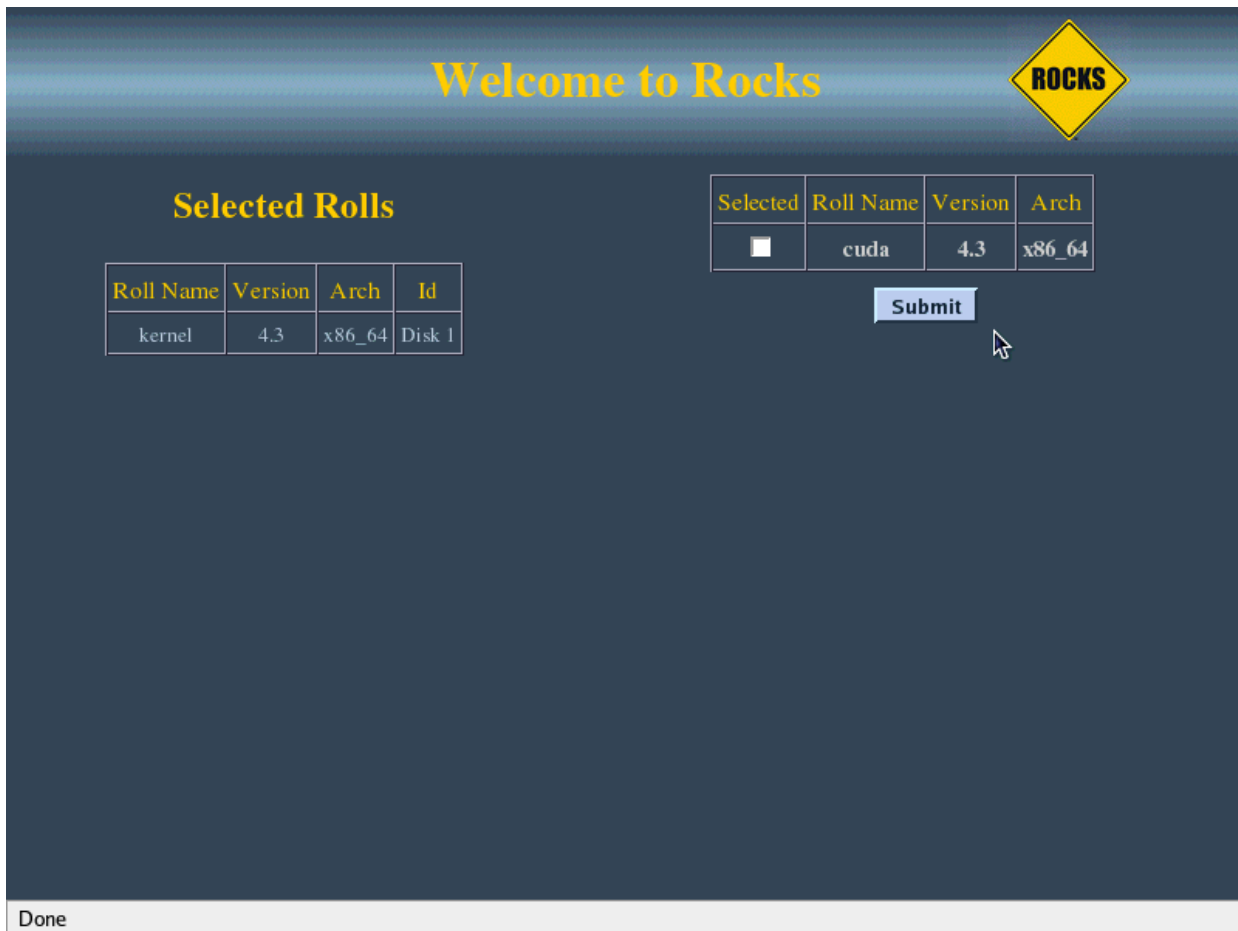
Click on the "CD/DVD-based Roll" button to indicate you have a Roll to add.

Figure 2-2. Loading the roll



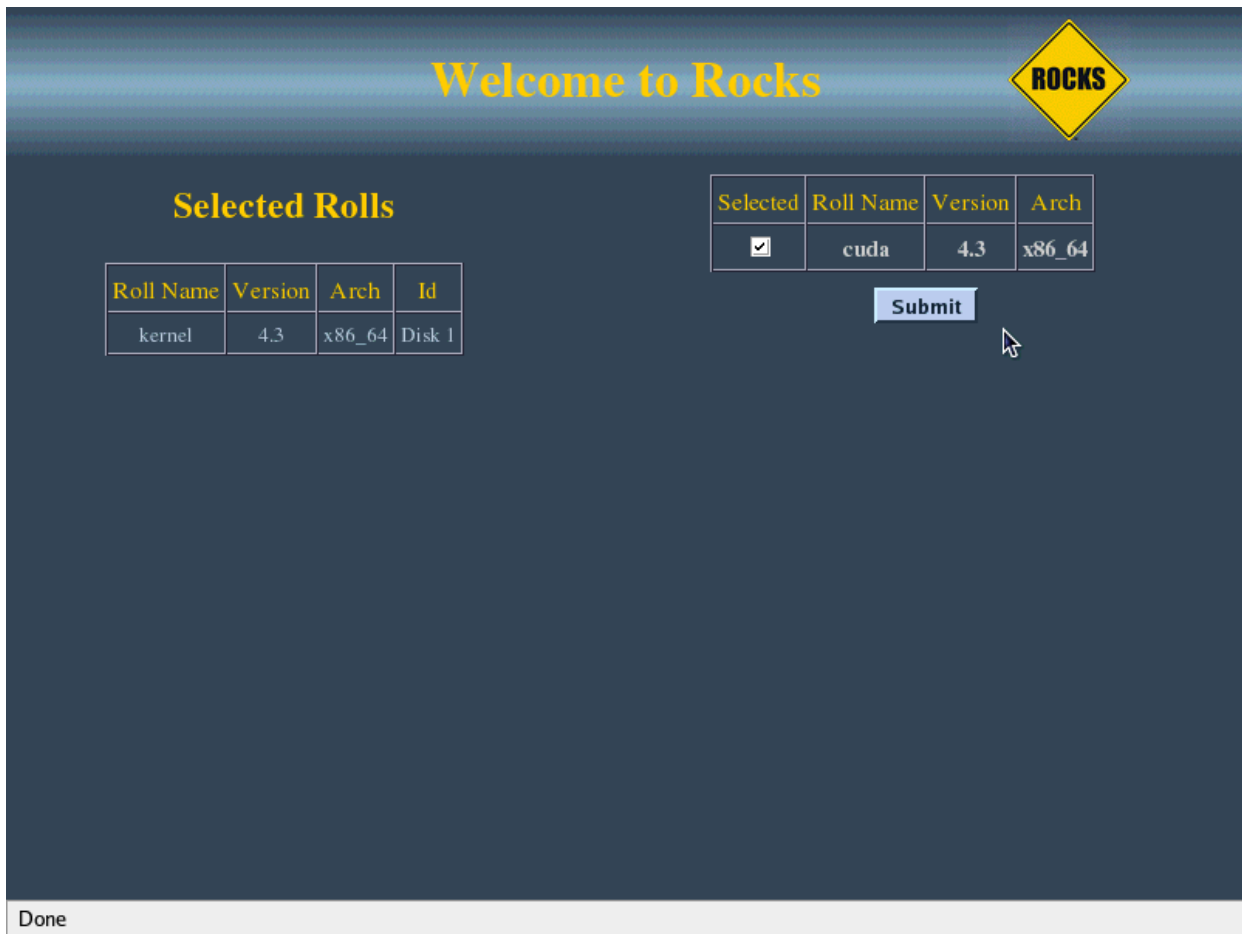
After inserting the CD, click on the "Continue" button.

Figure 2-3. List of available rolls




The available rolls will be listed on the right.

Figure 2-4. Selecting the CUDA Roll



Click on the box for CUDA and then click "Submit".

Figure 2-5. CUDA Roll selected for installation



The CUDA Roll will be added to the selected rolls on the left. You can now continue to add and select other rolls or start the installation as necessary.

Note that the device driver will not be installed on the Frontend by default. If you have an CUDA capable device in the Frontend and want to manually install the driver for it you can do so after the machine is up and running. Refer to the commands for manually installing the driver on the Frontend in the last section of this chapter.

2.2. Adding the Roll to an Existing Installation

The CUDA Roll can also be added to an existing system. Mount the CD at /mnt/cdrom or use the iso file directly:

```
# mount -o loop cuda-5.0-0.x86_64.disk1.iso /mnt/cdrom
```

Add the roll to the cluster:

```
# rocks add roll
cuda-5.0-0.x86_64.disk1.iso on /mnt/cdrom type iso9660 (rw,loop=/dev/loop0)
Copying cuda to Rolls.....185328 blocks
```

Make sure the roll is enabled by listing the rolls:

```
# rocks list roll
NAME                                VERSION ARCH  ENABLED
kernel:                             5.0      x86_64 yes
base:                                5.0      x86_64 yes
ganglia:                             5.0      x86_64 yes
hpc:                                 5.0      x86_64 yes
web-server:                          5.0      x86_64 yes
Red_Hat_Enterprise_Linux_Server_5:  5.0      x86_64 yes
cuda:                                 5.0      x86_64 yes
```

If the CUDA roll is not shown as enabled, enable the roll with the command:

```
# rocks enable roll cuda
```

Then rebuild the distribution:

```
# cd /home/install/
# rocks-dist dist
Cleaning distribution
Resolving versions (base files)
  including "kernel" (5.0,x86_64) roll...
  including "Red_Hat_Enterprise_Linux_Server_5" (5.0,x86_64) roll...
  including "hpc" (5.0,x86_64) roll...
  including "base" (5.0,x86_64) roll...
  including "cuda" (5.0,x86_64) roll...
  including "web-server" (5.0,x86_64) roll...
  including "ganglia" (5.0,x86_64) roll...
Including critical RPMS
Resolving versions (RPMs)
  including "kernel" (5.0,x86_64) roll...
  including "Red_Hat_Enterprise_Linux_Server_5" (5.0,x86_64) roll...
  including "hpc" (5.0,x86_64) roll...
  including "base" (5.0,x86_64) roll...
  including "cuda" (5.0,x86_64) roll...
  including "web-server" (5.0,x86_64) roll...
  including "ganglia" (5.0,x86_64) roll...
Resolving versions (SRPMS)
  including "kernel" (5.0,x86_64) roll...
  including "Red_Hat_Enterprise_Linux_Server_5" (5.0,x86_64) roll...
  including "hpc" (5.0,x86_64) roll...
  including "base" (5.0,x86_64) roll...
  including "cuda" (5.0,x86_64) roll...
  including "web-server" (5.0,x86_64) roll...
  including "ganglia" (5.0,x86_64) roll...
Creating files (symbolic links - fast)
Applying stage2.img
Applying updates.img
Installing XML Kickstart profiles
```

```

installing "cuda" profiles...
installing "hpc" profiles...
installing "web-server" profiles...
installing "base" profiles...
installing "kernel" profiles...
installing "ganglia" profiles...
installing "site" profiles...
Creating repository
making "torrent" files for RPMS
Cleaning distribution
Resolving versions (base files)
    including "kernel" (5.0,x86_64) roll...
    including "base" (5.0,x86_64) roll...
Including critical RPMS
Resolving versions (RPMS)
    including "kernel" (5.0,x86_64) roll...
    including "base" (5.0,x86_64) roll...
Resolving versions (SRPMS)
    including "kernel" (5.0,x86_64) roll...
    including "base" (5.0,x86_64) roll...
Creating files (symbolic links - fast)
Applying stage2.img
Applying updates.img
Installing XML Kickstart profiles
    installing "kernel" profiles...
    installing "base" profiles...
Creating repository
Linking boot stages from lan
Building Roll Links

```

Now you can try reinstalling a compute node:

```
# shoot-node compute-0-0
```

After verifying the proper operation of the compute node (see Section 3.2: Testing the Installation for more details), reinstall the rest of the nodes in the cluster.

When adding the CUDA Roll on a running system the CUDA Toolkit, CUDA SDK, and the roll users guide will not be automatically installed on the Frontend. Install them with the commands:

```
# cd /home/install/rolls/cuda/5.0/x86_64/RedHat/RPMS
# rpm -i cuda-toolkit*
# rpm -i cuda-sdk*
# rpm -i roll-cuda-usersguide*
```

The device driver is not installed on the Frontend by default. If you have an CUDA capable device in the Frontend and want to install the driver for it refer to the commands for manually installing the driver on the Frontend in the last section of this chapter.

2.3. Optional installation of the driver on the Frontend

By default the device driver will not be installed on the Frontend. If you have an CUDA capable device in the Frontend and wish to install the device driver, you can manually install the rpm and load the driver:

```
# cd /home/install/rolls/cuda/5.0/x86_64/RedHat/RPMS  
# rpm -i nvidia-driver*  
# /etc/rc.d/init.d/nvidia start
```

On subsequent reboots the device driver will be automatically loaded.

Chapter 3. Using the CUDA Roll

The CUDA Roll provides the device driver, toolkit, and SDK for developing and running CUDA enabled programs.

3.1. Summary of Installed Components

The NVIDIA Display Driver is loaded automatically at boot time by the `/etc/rc.d/init.d/nvidia` script. It attempts to load the driver, and if it fails (which happens most often because the driver hasn't been installed for the currently running kernel) will reinstall the driver. The installation of the driver will always happen upon first boot after installation and will add approximately one minute to the boot time.

The CUDA Toolkit is installed in `/usr/local/cuda` and the files `/etc/profile.d/cuda.*` automatically add the proper location to the `PATH` and `LD_LIBRARY_PATH` environment variables.

A copy of the CUDA SDK is placed in `/usr/local/NVIDIA_CUDA_SDK` with precompiled debug and release versions of the SDK examples in `/usr/local/NVIDIA_CUDA_SDK/linux/bin/debug` and `/usr/local/NVIDIA_CUDA_SDK/linux/bin/release` respectively.

3.2. Testing the Installation

The pre-compiled SDK samples can be used to verify proper setup of your cluster. For example, the `deviceQuery` example can be run on a compute node to list the CUDA capable devices:

```
[tux@compute-0-0]$ cd /usr/local/NVIDIA_CUDA_SDK/bin/linux/release
[tux@compute-0-0]$ ./deviceQuery -noprompt
There are 2 devices supporting CUDA
```

Device 0: "Tesla T10 Processor"

Major revision number:	1
Minor revision number:	3
Total amount of global memory:	4294705152 bytes
Number of multiprocessors:	30
Number of cores:	240
Total amount of constant memory:	65536 bytes
Total amount of shared memory per block:	16384 bytes
Total number of registers available per block:	16384
Warp size:	32
Maximum number of threads per block:	512
Maximum sizes of each dimension of a block:	512 x 512 x 64
Maximum sizes of each dimension of a grid:	65535 x 65535 x 1
Maximum memory pitch:	262144 bytes
Texture alignment:	256 bytes
Clock rate:	1.44 GHz
Concurrent copy and execution:	Yes

Device 1: "Tesla T10 Processor"

Major revision number:	1
Minor revision number:	3

```

Total amount of global memory:          4294705152 bytes
Number of multiprocessors:              30
Number of cores:                        240
Total amount of constant memory:        65536 bytes
Total amount of shared memory per block: 16384 bytes
Total number of registers available per block: 16384
Warp size:                              32
Maximum number of threads per block:    512
Maximum sizes of each dimension of a block: 512 x 512 x 64
Maximum sizes of each dimension of a grid: 65535 x 65535 x 1
Maximum memory pitch:                   262144 bytes
Texture alignment:                      256 bytes
Clock rate:                             1.44 GHz
Concurrent copy and execution:          Yes

```

Test PASSED

3.3. Troubleshooting

If the installation does not work, please make sure that the devices are among the list of PCI devices on the system:

```

[tux@compute-0-0 ~]$ /sbin/lspci | grep -i nvidia
01:00.0 PCI bridge: nVidia Corporation: Unknown device 05be (rev a2)
02:00.0 PCI bridge: nVidia Corporation: Unknown device 05be (rev a2)
02:01.0 PCI bridge: nVidia Corporation: Unknown device 05be (rev a2)
02:02.0 PCI bridge: nVidia Corporation: Unknown device 05be (rev a2)
02:03.0 PCI bridge: nVidia Corporation: Unknown device 05be (rev a2)
05:00.0 PCI bridge: nVidia Corporation: Unknown device 05be (rev a3)
06:00.0 PCI bridge: nVidia Corporation: Unknown device 05be (rev a3)
06:01.0 PCI bridge: nVidia Corporation: Unknown device 05be (rev a3)
06:02.0 PCI bridge: nVidia Corporation: Unknown device 05be (rev a3)
06:03.0 PCI bridge: nVidia Corporation: Unknown device 05be (rev a3)
07:00.0 3D controller: nVidia Corporation: Unknown device 05e7 (rev a1)
09:00.0 3D controller: nVidia Corporation: Unknown device 05e7 (rev a1)

```

In the output from `lspci` the entries will often show up as "Unknown device" because the hardware is too new to be in the `/usr/share/hwdata/pci.ids` file. For the example above the PCI bridge entries correspond to the host adapter card for a Tesla 1U computing system. The 3D controller entries are GPUs from a Tesla S1070.

Make sure that the proper driver version is loaded and that the device node files are present:

```

[tux@compute-0-0 ~]$ /sbin/lsmmod | grep nvidia
nvidia          7793880  32
i2c_core        56129   3 nvidia,i2c_ec,i2c_i801
[tux@compute-0-0 ~]$ cat /proc/driver/nvidia/version
NVRM version: NVIDIA UNIX x86_64 Kernel Module 177.70.18 Fri Oct 10 16:46:46 PDT 2008
GCC version: gcc version 4.1.2 20071124 (Red Hat 4.1.2-42)
[tux@compute-0-0 ~]$ ls /dev/nv*
/dev/nvidia0 /dev/nvidia1 /dev/nvidiactl

```

There should be an entry for `nvidiactl` and one entry for each NVIDIA graphics card, starting with `nvidia0`.

Please include the `nvidia-bug-report.log` file created by running `nvidia-bug-report.sh` when reporting any problems with your installation.

Appendix A. Rocks Copyright

Rocks(r)
www.rocksclusters.org
version 5.0 (V)

Copyright (c) 2000 - 2008 The Regents of the University of California.
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice unmodified and in its entirety, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising and press materials, printed or electronic, mentioning features or use of this software must display the following acknowledgement:

"This product includes software developed by the Rocks(r)
Cluster Group at the San Diego Supercomputer Center at the
University of California, San Diego and its contributors."

4. Except as permitted for the purposes of acknowledgment in paragraph 3, neither the name or logo of this software nor the names of its authors may be used to endorse or promote products derived from this software without specific prior written permission. The name of the software includes the following terms, and any derivatives thereof: "Rocks", "Rocks Clusters", and "Avalanche Installer". For licensing of the associated name, interested parties should contact Technology Transfer & Intellectual Property Services, University of California, San Diego, 9500 Gilman Drive, Mail Code 0910, La Jolla, CA 92093-0910, Ph: (858) 534-5815, FAX: (858) 534-7345, E-MAIL:invent@ucsd.edu

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Appendix B. Third Party Copyrights and Licenses

B.1. NVIDIA CUDA License

License For Customer Use of NVIDIA Software

IMPORTANT NOTICE -- READ CAREFULLY: This License For Customer Use of NVIDIA Software ("LICENSE") is the agreement which governs use of the software of NVIDIA Corporation and its subsidiaries ("NVIDIA") downloadable herefrom, including computer software and associated printed materials ("SOFTWARE"). By downloading, installing, copying, or otherwise using the SOFTWARE, you agree to be bound by the terms of this LICENSE. If you do not agree to the terms of this LICENSE, do not download the SOFTWARE.

RECITALS

Use of NVIDIA's products requires three elements: the SOFTWARE, the hardware on a graphics controller board, and a personal computer. The SOFTWARE is protected by copyright laws and international copyright treaties, as well as other intellectual property laws and treaties. The SOFTWARE is not sold, and instead is only licensed for use, strictly in accordance with this document. The hardware is protected by various patents, and is sold, but this LICENSE does not cover that sale, since it may not necessarily be sold as a package with the SOFTWARE. This LICENSE sets forth the terms and conditions of the SOFTWARE LICENSE only.

1. DEFINITIONS

1.1 Customer. Customer means the entity or individual that downloads the SOFTWARE.

2. GRANT OF LICENSE

2.1 Rights and Limitations of Grant. NVIDIA hereby grants Customer the following non-exclusive, non-transferable right to use the SOFTWARE, with the following limitations:

2.1.1 Rights. Customer may install and use one copy of the SOFTWARE on a single computer, and except for making one back-up copy of the Software, may not otherwise copy the SOFTWARE. This LICENSE of SOFTWARE may not be shared or used concurrently on different computers.

2.1.2 Linux/FreeBSD Exception. Notwithstanding the foregoing terms of Section 2.1.1, SOFTWARE designed exclusively for use on the Linux or FreeBSD operating systems, or other operating systems derived from the source code to these operating systems, may be copied and redistributed, provided that the binary files thereof are not modified in any way (except for unzipping of compressed

files).

2.1.3 Limitations.

No Reverse Engineering. Customer may not reverse engineer, decompile, or disassemble the SOFTWARE, nor attempt in any other manner to obtain the source code.

No Separation of Components. The SOFTWARE is licensed as a single product. Its component parts may not be separated for use on more than one computer, nor otherwise used separately from the other parts.

No Rental. Customer may not rent or lease the SOFTWARE to someone else.

3. TERMINATION

This LICENSE will automatically terminate if Customer fails to comply with any of the terms and conditions hereof. In such event, Customer must destroy all copies of the SOFTWARE and all of its component parts.

Defensive Suspension. If Customer commences or participates in any legal proceeding against NVIDIA, then NVIDIA may, in its sole discretion, suspend or terminate all license grants and any other rights provided under this LICENSE during the pendency of such legal proceedings.

4. COPYRIGHT

All title and copyrights in and to the SOFTWARE (including but not limited to all images, photographs, animations, video, audio, music, text, and other information incorporated into the SOFTWARE), the accompanying printed materials, and any copies of the SOFTWARE, are owned by NVIDIA, or its suppliers. The SOFTWARE is protected by copyright laws and international treaty provisions. Accordingly, Customer is required to treat the SOFTWARE like any other copyrighted material, except as otherwise allowed pursuant to this LICENSE and that it may make one copy of the SOFTWARE solely for backup or archive purposes.

5. APPLICABLE LAW

This LICENSE shall be deemed to have been made in, and shall be construed pursuant to, the laws of the State of California. The United Nations Convention on Contracts for the International Sale of Goods is specifically disclaimed.

6. DISCLAIMER OF WARRANTIES AND LIMITATION ON LIABILITY

6.1 No Warranties. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, THE SOFTWARE IS PROVIDED "AS IS" AND NVIDIA AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

6.2 No Liability for Consequential Damages. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO EVENT SHALL NVIDIA OR ITS SUPPLIERS BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT, OR CONSEQUENTIAL DAMAGES WHATSOEVER (INCLUDING,

WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION, LOSS OF BUSINESS INFORMATION, OR ANY OTHER PECUNIARY LOSS) ARISING OUT OF THE USE OF OR INABILITY TO USE THE SOFTWARE, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. MISCELLANEOUS

If any provision of this LICENSE is inconsistent with, or cannot be fully enforced under, the law, such provision will be construed as limited to the extent necessary to be consistent with and fully enforceable under the law. This LICENSE is the final, complete and exclusive agreement between the parties relating to the subject matter hereof, and supersedes all prior or contemporaneous understandings and agreements relating to such subject matter, whether oral or written. This LICENSE may only be modified in writing signed by an authorized officer of NVIDIA. Customer agrees that it will not ship, transfer or export the SOFTWARE into any country, or use the SOFTWARE in any manner, prohibited by the United States Bureau of Export Administration or any export laws, restrictions or regulations.

B.2. NVIDIA Driver License

License For Customer Use of NVIDIA Software

IMPORTANT NOTICE -- READ CAREFULLY: This License For Customer Use of NVIDIA Software ("LICENSE") is the agreement which governs use of the software of NVIDIA Corporation and its subsidiaries ("NVIDIA") downloadable herefrom, including computer software and associated printed materials ("SOFTWARE"). By downloading, installing, copying, or otherwise using the SOFTWARE, you agree to be bound by the terms of this LICENSE. If you do not agree to the terms of this LICENSE, do not download the SOFTWARE.

RECITALS

Use of NVIDIA's products requires three elements: the SOFTWARE, the hardware on a graphics controller board, and a personal computer. The SOFTWARE is protected by copyright laws and international copyright treaties, as well as other intellectual property laws and treaties. The SOFTWARE is not sold, and instead is only licensed for use, strictly in accordance with this document. The hardware is protected by various patents, and is sold, but this agreement does not cover that sale, since it may not necessarily be sold as a package with the SOFTWARE. This agreement sets forth the terms and conditions of the SOFTWARE LICENSE only.

1. DEFINITIONS

1.1 Customer. Customer means the entity or individual that downloads the SOFTWARE.

2. GRANT OF LICENSE

2.1 Rights and Limitations of Grant. NVIDIA hereby grants Customer the following non-exclusive, non-transferable right to use the SOFTWARE, with the following limitations:

2.1.1 Rights. Customer may install and use one copy of the SOFTWARE on a single computer, and except for making one back-up copy of the Software, may not otherwise copy the SOFTWARE. This LICENSE of SOFTWARE may not be shared or used concurrently on different computers.

2.1.2 Linux/FreeBSD Exception. Notwithstanding the foregoing terms of Section 2.1.1, SOFTWARE designed exclusively for use on the Linux or FreeBSD operating systems, or other operating systems derived from the source code to these operating systems, may be copied and redistributed, provided that the binary files thereof are not modified in any way (except for unzipping of compressed files).

2.1.3 Limitations.

No Reverse Engineering. Customer may not reverse engineer, decompile, or disassemble the SOFTWARE, nor attempt in any other manner to obtain the source code.

No Separation of Components. The SOFTWARE is licensed as a single product. Its component parts may not be separated for use on more than one computer, nor otherwise used separately from the other parts.

No Rental. Customer may not rent or lease the SOFTWARE to someone else.

3. TERMINATION

This LICENSE will automatically terminate if Customer fails to comply with any of the terms and conditions hereof. In such event, Customer must destroy all copies of the SOFTWARE and all of its component parts.

Defensive Suspension. If Customer commences or participates in any legal proceeding against NVIDIA, then NVIDIA may, in its sole discretion, suspend or terminate all license grants and any other rights provided under this LICENSE during the pendency of such legal proceedings.

4. COPYRIGHT

All title and copyrights in and to the SOFTWARE (including but not limited to all images, photographs, animations, video, audio,

music, text, and other information incorporated into the SOFTWARE), the accompanying printed materials, and any copies of the SOFTWARE, are owned by NVIDIA, or its suppliers. The SOFTWARE is protected by copyright laws and international treaty provisions. Accordingly, Customer is required to treat the SOFTWARE like any other copyrighted material, except as otherwise allowed pursuant to this LICENSE and that it may make one copy of the SOFTWARE solely for backup or archive purposes.

5. APPLICABLE LAW

This agreement shall be deemed to have been made in, and shall be construed pursuant to, the laws of the State of California.

6. DISCLAIMER OF WARRANTIES AND LIMITATION ON LIABILITY

6.1 No Warranties. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, THE SOFTWARE IS PROVIDED "AS IS" AND NVIDIA AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

6.2 No Liability for Consequential Damages. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO EVENT SHALL NVIDIA OR ITS SUPPLIERS BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT, OR CONSEQUENTIAL DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION, LOSS OF BUSINESS INFORMATION, OR ANY OTHER PECUNIARY LOSS) ARISING OUT OF THE USE OF OR INABILITY TO USE THE SOFTWARE, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. MISCELLANEOUS

The United Nations Convention on Contracts for the International Sale of Goods is specifically disclaimed. If any provision of this LICENSE is inconsistent with, or cannot be fully enforced under, the law, such provision will be construed as limited to the extent necessary to be consistent with and fully enforceable under the law. This agreement is the final, complete and exclusive agreement between the parties relating to the subject matter hereof, and supersedes all prior or contemporaneous understandings and agreements relating to such subject matter, whether oral or written. Customer agrees that it will not ship, transfer or export the SOFTWARE into any country, or use the SOFTWARE in any manner, prohibited by the United States Bureau of Export Administration or any export laws, restrictions or regulations. This LICENSE may only be modified in writing signed by an authorized officer of NVIDIA.