



Jetson TK1 Development Kit Quick Start Guide

November 2014

DA-07498-001

Version 01

Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OR CONDITION OF TITLE, MERCHANTABILITY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE AND ON-INFRINGEMENT, ARE HEREBY EXCLUDED TO THE MAXIMUM EXTENT PERMITTED BY LAW.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication or otherwise under any patent or patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all information previously supplied. NVIDIA Corporation products are not authorized for use as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

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

Copyright

© 2014 NVIDIA Corporation. All rights reserved.

Contents

Introduction	3
Getting Started	3
Assumptions.....	3
Powering Up the Tegra Device	3
Login Credentials	4
Installing the Linux Driver Binary	4
Running Sample Code	4
CUDA	4
Multimedia	4
Flashing Jetson.....	5
Before you Begin.....	5
Downloading L4T Drivers	5
JetPack.....	6
Before you Begin.....	6
Installing JetPack TK1	7
Run JetPack TK1 Samples	14

Introduction

The NVIDIA® Jetson TK1 developer board is a full-featured device ideal for software development within the Linux environment. Standard connectors are used to access Tegra® features and interfaces, enabling a highly flexible and extensible development platform. Go to <http://developer.nvidia.com/jetson-tk1> for access to software updates and the developer SDK supporting the OS image and host development platform that you want to use. The developer SDK includes an OS image that you will load onto your Tegra device, supporting documentation, and code samples to help you get started.

This section familiarizes you with the Jetson platform; what to do when you get it out of the box, how to boot into a Graphical User Interface and run basic sample code with the pre-installed software.

Getting Started

Individual development efforts will vary and may result in modifications to the system configuration. It is recommended that you begin with the basic system configuration (as shipped) to ensure proper system operation prior to any further development.

Assumptions

- > You have a Jetson TK1 Tegra Developer System, equipped with the NVIDIA® Tegra® K1 processor.
- > Your developer system should be cabled as follows:
 - Serial cable plugged into the serial port J1A2 UART4 on the target connected to your Linux host directly or through a serial-to-USB converter. (To setup serial console on the Linux host.)
 - (Not included in the developer kit) To connect USB peripherals such as keyboard, mouse, and [optional] USB/Ethernet adapter (for network connection), a USB hub should be connected to the working USB port (J1C2 USB2) on the Jetson TK1 system.
 - An HDMI cable plugged into "J1C1 HDMI" on the target which is connected to an external HDMI display.
 - An Ethernet cable plugged into the J1D1 on board Ethernet port.

Powering Up the Tegra Device

1. Connect the AC adapter supplied in your kit to the power connector of your device.
2. Plug the power adapter into an appropriately rated electrical outlet.
3. The system should power on, if not, press and release the power button on the device. You should see a green POWER LED light illuminate and the fan will spin up

Login Credentials

- > Username: ubuntu
- > Password: ubuntu

Installing the Linux Driver Binary

Instructions on how to install the NVIDIA Linux driver binary release on your target located in:

`${HOME}/NVIDIA-INSTALLER`

1. Change directories into the NVIDIA installation directory:

```
cd ${HOME}/NVIDIA-INSTALLER
```

2. Run the installer script to extract and install the Linux driver binary release:

```
sudo ./installer.sh
```

3. Reboot the system

NOTE: the `installer.sh` should only be run once.

Running Sample Code

CUDA

Download the CUDA SDK Demo Samples from:

http://developer.download.nvidia.com/embedded/jetson/TK1/2014-03-24/CUDA_SDK_Samples.tgz

Multimedia

Download the following H.264 video sample:

<http://www.bigbuckbunny.org/index.php/download/>

Play video in `nvgstplayer` by:

```
Nvgstplayer -i <full path to video>
```

Flashing Jetson

This section steps you through process to re-flash your Jetson TK1 Development Kit.

Before you Begin

- > You have a Jetson TK1 Tegra Developer Kit equipped with the NVIDIA Tegra K1 processor.
- > You have a host machine that is running Linux.
- > Your developer system is cabled as follows:
 - Serial cable plugged into the serial port J1A2 UART4 on the target connected to your Linux host directly or through a serial-to-USB converter. (To setup serial console on the Linux host.)
 - USB Micro-B cable connecting Jetson TK1 (J1E1 USB0) to your Linux host for flashing.
 - (Not included in the developer kit) To connect USB peripherals such as keyboard, mouse, and [optional] USB/Ethernet adapter (for network connection), a USB hub should be connected to the working USB port (J1C2 USB2) on the Jetson TK1 system.
 - An HDMI cable plugged into "J1C1 HDMI1" on the target which is connected to an external HDMI display.
 - An Ethernet cable plugged into the J1D1 on board Ethernet port.

Downloading L4T Drivers

`${RELEASE_NAME}` in the following directions refers to the respective package name as follows:

- > Jetson TK1: Tegra124_Linux_Rxx.x.x_armhf.tbz2

Download the latest L4T release package for your developer system and the sample file system from <https://developer.nvidia.com/linux-tegra>

1. Untar the release package and assemble the sample file system:

```
sudo tar xpf ${RELEASE_NAME}
cd Linux_for_Tegra/rootfs/
sudo tar xpf ../../Tegra_Linux_Sample-Root-
Filesystem_Rxx.x.x_armhf.tbz2
cd ../
sudo ./apply_binaries.sh
```

2. Flash the sample file system onto the system's internal eMMC:

- a) Put your system into "reset recovery mode" by holding down the RECOVERY button while pressing and releasing the RESET button once on the main board.

- b) Ensure your Linux host system is connected to the target device through the USB cable for flashing.

```
sudo ./flash.sh -S 8GiB ${Jetson-tk1} mmcblk0p1
```

The target will automatically reboot upon completion of the flash. The command prompt will show up over the display that you have attached to the target. All actions are completed unless you wish to configure the graphical desktop on your setup. You now have Linux running on your developer system.

Installing the graphical desktop on your target board (if not already installed):

1. [OPTIONAL] Connect Ethernet to target via a USB-to-Ethernet adapter
2. Use `eth0` for the built-in Ethernet port (or find your USB Ethernet port with `'ifconfig -a'`):

```
sudo dhclient eth0
```
3. Check to see if Ethernet is up and running. You should see an IP address associated with `eth0`.

```
ifconfig
sudo apt-get update
sudo apt-get install ubuntu-desktop
```

4. Reboot: system will boot to the graphical desktop.

NOTE: the above steps can be used to install other packages with "`sudo apt-get install`".

Refer to the release notes provided with your software for up-to-date information on platform features and use.

JetPack

This section is intended to help you get familiar with installing Jetson TK1 Development Pack (JetPack TK1), using the tools and running sample code.

Before you Begin

- > You have a Jetson TK1 Tegra Developer Kit equipped with the NVIDIA Tegra K1 processor.
- > You have a host machine that is running Linux.
- > Your developer system is cabled as follows:
 - Serial cable plugged into the serial port J1A2 UART4 on the target connected to your Linux host directly or through a serial-to-USB converter. (To setup serial console on the Linux host.)
 - USB Micro-B cable connecting Jetson TK1 (J1E1 USB0) to your Linux host for flashing.
 - (Not included in the developer kit) To connect USB peripherals such as keyboard, mouse, and [optional] USB/Ethernet adapter (for network connection), a USB hub should be connected to the working USB port (J1C2 USB2) on the Jetson TK1 system.

- An HDMI cable plugged into "J1C1 HDMI1" on the target which is connected to an external HDMI display.
- An Ethernet cable plugged into the J1D1 on board Ethernet port.
- > Download latest JetPack TK1
 - The latest version of Jetson TK1 Development Pack (JetPack TK1) is available at [NVIDIA Developer](#).
 - All available JetPack TK1 downloads [here](#).

Installing JetPack TK1

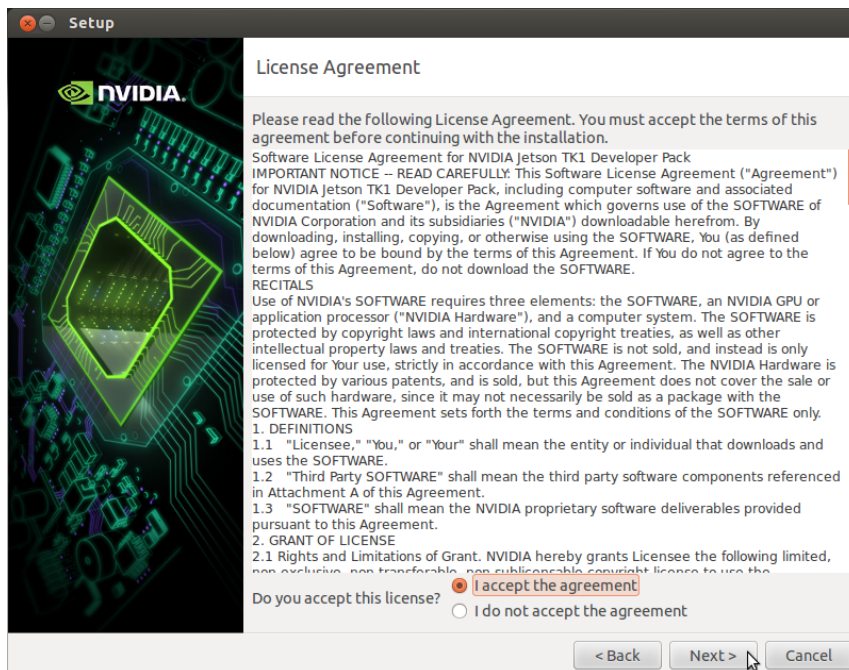
Assumes you downloaded the latest JetPack TK1 version, `jetpack-${VERSION}.run`. `${VERSION}` refers to the version string for the installer you got.

1. Add execute permissions for the `jetpack-${VERSION}.run`

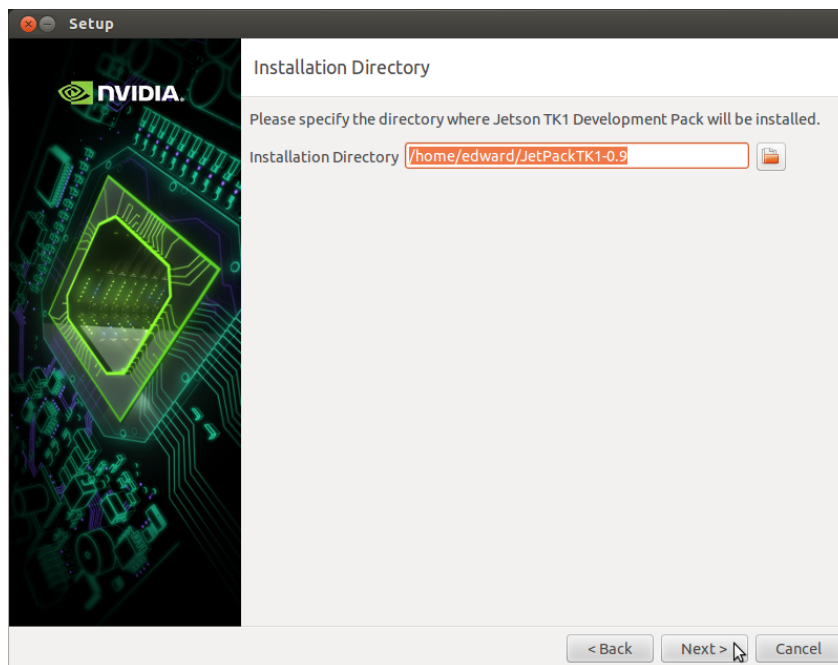
```
chmod +x jetpack-${VERSION}.run
```
2. Run the `jetpack-${VERSION}.run` by double-clicking on it in the file browser.



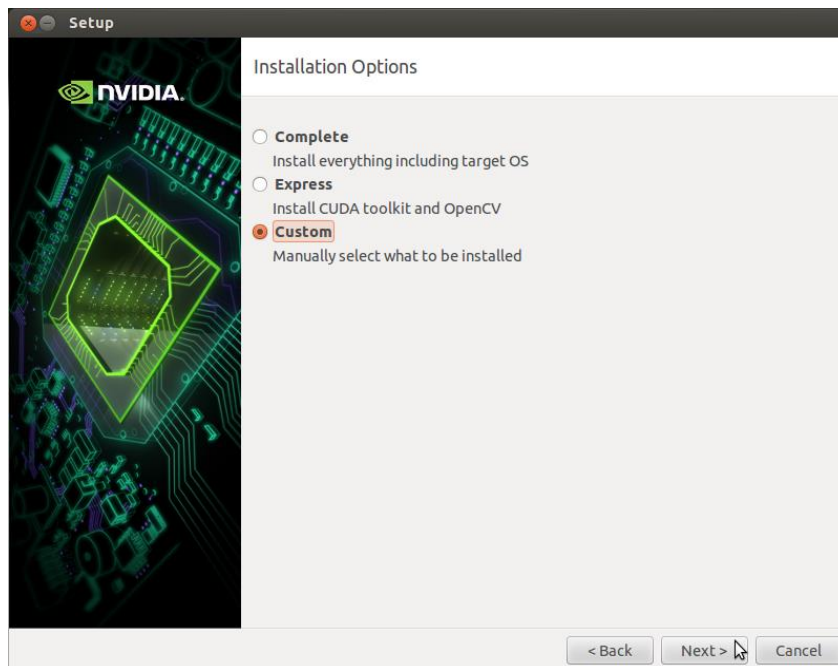
3. Accept the license agreement to continue



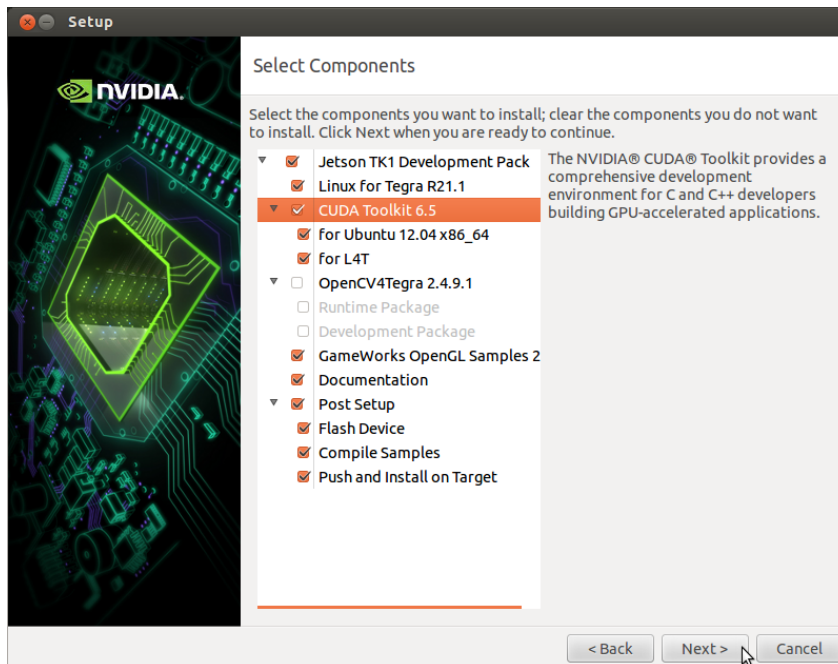
4. Enter or navigate to desired installation directory.



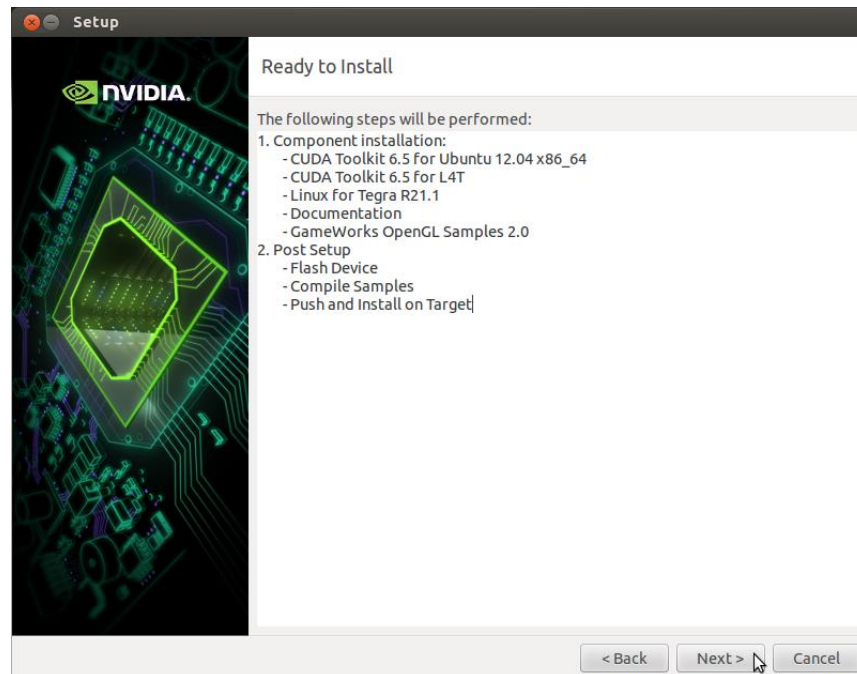
5. Select installation type.



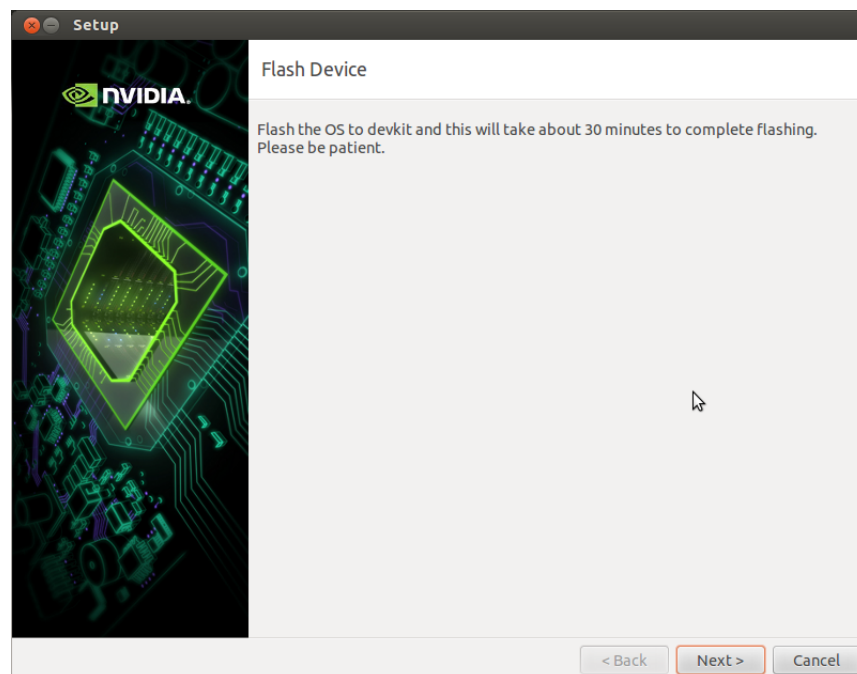
6. Select components to install.



7. Review Setup and click the Next button proceed with setup or click Back button to return to the Select Components screen.




8. Click next and follow instructions on popup window to put device into recovery mode, then begin flashing device.



Sample output when flashing OS on Jetson.

```
Flash OS
creating partition: MBP
creating partition: GP1
creating partition: APP
creating partition: DTB
creating partition: EFI
creating partition: USP
creating partition: TP1
creating partition: TP2
creating partition: TP3
creating partition: WBO
creating partition: UDA
creating partition: GPT
sending file: ppt.img
\ 2097152/2097152 bytes sent
ppt.img sent successfully
padded 5 bytes to bootloader
sending file: fastboot.bin
- 594368/594368 bytes sent
fastboot.bin sent successfully
sending file: boot.img
\ 6090752/6090752 bytes sent
boot.img sent successfully
sending file: system.img
\ 205730615/2275602980 bytes sent
```

9. After successfully flashing device, enter your device information in the dialog that comes up.



The Setup dialog box is titled "Setup" and features the NVIDIA logo on the left. The main content area is titled "Input Device Information" and contains the following text:

In order to push CUDA toolkit to Jetson TK1, please input its IP address, username, and password.

IP address can be queried by running "ifconfig" command on target.

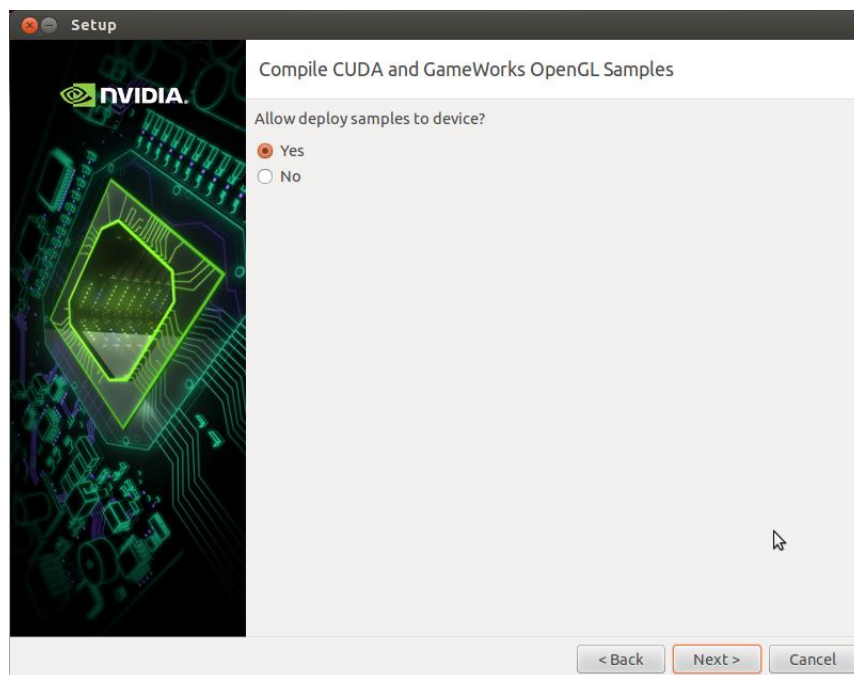
The default username/password are "ubuntu/ubuntu" without quotation.

Below the text are three input fields:

- IP Address:
- Username:
- Password:

At the bottom right of the dialog are three buttons: "< Back", "Next >", and "Cancel".

10. After installing CUDA on the device, the installer will compile CUDA and GameWorks OpenGL samples. You can choose whether to deploy the compiled binaries to device.



Sample output when compiling CUDA samples.

```

Compile CUDA Samples
make[1]: Nothing to be done for `all'.
make[1]: Leaving directory `/home/edward/JetPackTK1-0.9/NVIDIA_CUDA-6.5_Samples/1_Utillities/bandwidthTest'
make[1]: Entering directory `/home/edward/JetPackTK1-0.9/NVIDIA_CUDA-6.5_Samples/0_Simple/simpleTextureDrv'
make[1]: Nothing to be done for `all'.
make[1]: Leaving directory `/home/edward/JetPackTK1-0.9/NVIDIA_CUDA-6.5_Samples/0_Simple/simpleTextureDrv'
make[1]: Entering directory `/home/edward/JetPackTK1-0.9/NVIDIA_CUDA-6.5_Samples/0_Simple/matrixMulDrv'
make[1]: Nothing to be done for `all'.
make[1]: Leaving directory `/home/edward/JetPackTK1-0.9/NVIDIA_CUDA-6.5_Samples/0_Simple/matrixMulDrv'
make[1]: Entering directory `/home/edward/JetPackTK1-0.9/NVIDIA_CUDA-6.5_Samples/0_Simple/vectorAddDrv'
make[1]: Nothing to be done for `all'.
make[1]: Leaving directory `/home/edward/JetPackTK1-0.9/NVIDIA_CUDA-6.5_Samples/0_Simple/vectorAddDrv'
make[1]: Entering directory `/home/edward/JetPackTK1-0.9/NVIDIA_CUDA-6.5_Samples/1_Utillities/deviceQueryDrv'
make[1]: Nothing to be done for `all'.
make[1]: Leaving directory `/home/edward/JetPackTK1-0.9/NVIDIA_CUDA-6.5_Samples/1_Utillities/deviceQueryDrv'

```

11. Copy CUDA samples to device. Run `matrixMul` sample remotely on device.

```

Deploy Samples
stopwatch_linux.cpp          100% 1041    1.0KB/s  00:00
rendercheck_d3d11.cpp        100% 3784    3.7KB/s  00:00
rendercheck_d3d9.cpp         100% 5001    4.9KB/s  00:00
stopwatch.cpp                100% 3228    3.2KB/s  00:00
cuda_runtime_dynlink.cpp     100% 9834    9.6KB/s  00:00
rendercheck_d3d10.cpp        100% 3577    3.5KB/s  00:00
libGLEW.a                   100% 313KB  313.0KB/s 00:00
libGLEW.a                   100% 369KB  368.7KB/s 00:00
EULA.txt                    100%  90KB   90.2KB/s  00:00
Now running matrixMul sample....
[Matrix Multiply Using CUDA] - Starting...
GPU Device 0: "GK20A" with compute capability 3.2

MatrixA(320,320), MatrixB(640,320)
Computing result using CUDA Kernel...
done
Performance= 18.78 GFlop/s, Time= 6.980 msec, Size= 131072000 Ops, WorkgroupSize
= 1024 threads/block
Checking computed result for correctness: Result = PASS

Note: For peak performance, please refer to the matrixMulCUBLAS example.

Please press Enter key to continue

```

Compiling

JetPack TK1 automatically compiles all samples if “Compile Samples” was checked during components selection. If you selected CUDA components, CUDA samples will be found in the following directory:

`<JetPack_Install_Dir>/NVIDIA_CUDA-<version>_Samples` directory.

You could recompile the samples by running:

```
OPENMPDIR=/usr/arm-linux-gnueabi/lib EXTRA_LDFLAGS=--unresolved-
symbols=ignore-in-shared-libs ARMv7=1 GCC=arm-linux-gnueabi-g++ make
```

If you selected GameWorks OpenGL samples, GameWorks OpenGL samples will be available in the following directory:

`< JetPack_Install_Dir >/GameWorksOpenGLSamples`


You could cross compile them by running the “make” command under the “samples/build/linux-arm32” subfolder.

Run JetPack TK1 Samples

The CUDA samples directory is copied to home directory on the device by JetPack TK1. The built binaries can be found in the following directory:

```
/home/ubuntu/NVIDIA_CUDA-6.5_Samples/bin/armv7l/linux/release/gnueabi/f/
```

Run them by calling them in terminal or double clicking on them in the file browser. For example, when you run the `oceanFFT` sample, the following screen will be displayed.



GameWorks OpenGL samples will be copied to home directory on device. The built binaries can be found in the following directory:

```
/home/ubuntu/GameWorksOpenGLSamples/samples/bin/linux-arm32
```

Run them by calling them in terminal or double clicking on them in the file browser. For example, when you run `FeedbackParticlesApp`, you will see smoke particles rising.

