



TEGRA LINUX DRIVER PACKAGE DEVELOPMENT GUIDE

PG_06706-R27 | March 3, 2017 | R27.1 Release

Development Guide



TABLE OF CONTENTS

Quick Start Guide.....	
Architecture.....	
Package Manifest.....	
Kernel.....	
Kernel Supplements TBZ2.....	
Boot Loader.....	
NV Tegra.....	
Nvgstapps TBZ2.....	
Config TBZ2.....	
NVIDIA Drivers TBZ2.....	
NVIDIA Drivers TBZ2.....	
Advanced Setup and Customization.....	
Reference Board Preparation.....	
Boot Options.....	
Linux Host System Prerequisites.....	
Extracting Tegra Linux Driver Package.....	
Setting Up Your File System.....	
Sample Root File System.....	
Setting Up the Root File System.....	
Step 1: Set Up the Root File System.....	
Step 2: Copy the rootfs to the Device.....	
Flashing the Boot Loader and Kernel.....	
Flash Procedure.....	
Flash Script Usage.....	
Increasing the Internal Memory Partition for the Root File System.....	
Determining the Success of a Driver Update.....	
Installing Additional Packages.....	
Installing Additional NVIDIA Packages.....	
Installing Additional Ubuntu Packages.....	
Configuring NFS Root on the Linux Host.....	
Getting the Kernel Sources.....	
Sync with git.....	
Manually Download and Expand.....	
Building the NVIDIA Kernel.....	
Building External Kernel Modules.....	
OpenGL/EGL Gears Test Application.....	

- GStreamer-based Multimedia Playback (NvGstPlayer).....
- Installing GStreamer.....
- Using NvGstPlayer.....
- Gstreamer-based Camera Capture (NvGstCapture).....
- Determining Version and Platform Information.....
- NVIDIA Bug Reporting Script.....

Software Features.....

- Boot Loaders.....
- Toolchain.....
- Kernel.....
- I/O.....
- CUDA.....
- Graphics.....
- EGL and OpenGL ES Support.....
- Video Decoders.....
- Video Encoders.....
- Display Outputs.....
- Conversion, Scaling, and Rotation Formats.....
- CSI and USB Camera Formats.....

U-Boot Customization.....

- Requirements.....
- Downloading and Building U-Boot.....
- Before Flashing.....
- Flashing U-Boot Full L4T Image.....
- Flashing U-Boot Only.....
- Changing the eMMC Partition Layout.....
- Building the Device Tree Compiler.....
- Replacing the Kernel.....
- Copying the Root Filesystem.....
- Boot Sequence and Sysboot Configuration Files.....
- Optimizing U-Boot Boot Time.....
- Disabling PCIe Support.....
- Disabling USB Support.....
- Environment Configuration.....
- Setting Environment Variables.....
- Compile-Time.....
- Manufacturing and Flashing Time.....
- extlinux.conf Modifications.....
- Debugging U-Boot Environment.....
- Interrupting U-Boot.....
- Getting Help.....
- Listing a Directory Structure.....
- Listing the Contents of a Directory.....
- Printing the U-Boot Environment.....
- Printing/Setting Environment Variables.....

Mass Storage Partitions Configuration.....

- Partition Configuration File.....
- Root Element.....
- Device Element.....
- Partition Element.....
- Default Partition Overview.....
- Example CFGs.....

Kernel Boot Time Optimization.....

- Device Tree Nodes.....
- Real-time Clock.....

- Environment Configuration.....
 - Disable Console over UART.....
- Secondary Bootloader.....
- Compile-Time Configuration.....
 - Asynchronous Probe.....
 - File System.....
 - Sound.....
 - Modularize Kernel Driver.....
 - Disable Debugging.....

Camera Software Development Solution.....

- Camera Architecture Stack.....
- Camera API Matrix.....
- Approaches for Validating and Testing the V4L2 Driver.....
 - Applications Using libargus Low-level APIs.....
 - Applications Using GStreamer with the nvcamerasrc Plugin.....
 - Applications Using GStreamer with V4L2 Source Plugin.....
 - Applications Using V4L2 IOCTL Directly.....
- ISP Support.....
- Document Change History.....

V4L2 Sensor Driver Programming Guide.....

Sensor Driver Programming Guide.....

- Camera Modules.....
- Individual Imaging Device.....
- V4L2 Kernel Driver.....
 - Macro Definitions.....
 - Sensor-Private Data.....
 - Configuring Regmap.....
 - Configuring Controls.....
 - Setting Up Control Registers.....
 - Read-Write Wrapper in the Register.....
 - Power Functions.....
 - Setting Up V4L2 Subdevice and Camera Common.....
 - Control Handlers.....
 - Set Control.....
 - Get-Volatile Control.....
 - Other Control-Related Functions.....
 - Boot-Time Initialization.....
 - Removing Sensor Drivers.....
- Device Registration.....
 - Using Plugin Manager.....
 - Using Main Platform Device Tree File.....
- How to Verify the V4L2 Sensor Driver.....
- Debugging Tips.....
- Mode Tables.....
- Open Source Software Licenses.....
 - GNU General Public License 2.0.....
 - mtd-utils.....
 - mkfs.....
 - ubinize.....
 - FreeType 6 version 2.4.12.....
 - ddccontrol version 0.2.....
 - FreeImage 3 version 3.15.1.....
 - License Information.....
 - GNU Lesser General Public License 2.1.....
 - GNU Lesser General Public License.....

Tegra ASoC Driver.....

- ALSA.....
- Tegra ASoC Driver Overview.....

- DAPM.....
- Device Tree.....
- Audio Driver.....
- Tegra Audio Hub.....
 - Tegra Audio Hub Architecture.....
 - Software Architecture.....
- Tegra Platform Driver.....
 - ADMAIF.....
 - Playback Hardware Devices in the Tegra ASoC Driver.....
 - Capture Hardware Devices in the Tegra ASoC Driver.....
- Tegra Codec Driver.....
 - XBAR.....
 - AMX.....
 - ADX.....
 - I2S.....
 - Mixer.....
 - SFC.....
 - SPDIF.....
 - DMIC.....
 - MVC.....
 - OPE.....
- Tegra Machine Driver.....
 - Tegra X2.....
 - Machine Specific DAI links.....
 - Audio Path.....
 - Tegra X2 Audio Path.....
 - XBAR Route Setting for Tegra X2.....
- Dynamic Audio Routing.....
 - Case 1: Internal AHUB TDM Path.....
 - Modify Case 1 to Record on I2S3 (I2S Mode) And Output On I2S4 (TDM Mode).....
- Codec Driver Instantiation via Device Tree.....
- TDM Slot Mapping.....
- Clocking and Power Management.....
- Audio Playback/Record Examples.....
- Troubleshooting.....
- Miscellaneous Examples.....
 - Simple Internal Audio Path.....
 - Routing Commands.....
 - Testing Commands.....
 - I2S-x and I2S-y Under Same Clock Domain.....
 - Routing Commands.....
 - Testing Commands.....

Building the AARCH 64 Toolchain and glibc.....

- Toolchain Information.....
- Downloading the Pre-Built Toolchain.....
- Building the Toolchain.....
- Troubleshooting.....

MB1 Platform Configuration.....

- Location of Configuration Files.....
- Pinmux, and GPIO Configuration.....
- Prod Configuration.....
- Pad Voltage Configuration.....
- PMIC Configuration.....
 - Rail-Specific Parameters.....
 - Commands.....
 - PWM specific commands.....
 - Generic Format.....
 - Secure PMC Scratch Register Configurations for BootROM.....
 - AO Block Parameters.....
 - I2C type-specific parameters.....

- Commands.....
- Security Configuration Register Configurations.....
- Miscellaneous Configurations.....
 - Debug.....
 - AOTAG.....
 - Clock.....
 - CPU Parameters.....
 - AST settings.....
 - SW Carveout.....
 - I2C Settings.....
 - Dev Parameters.....
 - Watchdog timer controller settings.....
- Generating MB1 Configuration Files.....
- Pinmux, GPIO, and Pad.....
- PMIC.....
- BootROM.....
- Flashing.....
- Configuring Pinmux, GPIO and PAD.....
 - Pinmux Configuration Process.....
 - Generating DT File.....
 - Generating Configuration files using dts2cfg.....
 - To convert pinmux and GPIO files to CFG format.....
 - To convert PAD file to CFG format.....

Boot Flow.....

- Responsibilities of the Boot Software.....
- High Level Tegra Software Boot.....
- BootROM.....
- Boot Loader Components.....
 - Common Driver Framework.....
 - Microboot1.....
- TegraBoot.....
 - TegraBoot-BPMP(tboot-BPMP).....
 - TegraBoot-CPU.....
 - CBoot.....
- U-Boot.....
- TegraFlash.....
- Security.....
 - Validation of Binaries.....
 - Boot Binaries and Associated Signing.....
- Secure Boot.....
 - Secure Boot Overview.....
- Secure Boot Flow.....
 - Build System.....

Display Configuration and Bringup.....

- Setting HDMI or DP Screen Resolution.....
- Mirroring or Extending Displays.....
- Determining Display Timing Values.....
- Configuring the Display Serial Interface.....
- Configuring the Display Port.....
- Configuring Pulse Width Modulation (PWM).....
- Panel Driver and Device Tree.....

PDF Documents.....

Appendix.....

- Tegra Stats Utility.....
 - Reported Statistics.....
 - Options.....

Running Tegra Stats.....
Re-Deploying Tegra Stats.....

FAQ.....

Linux FAQs.....
-----------------	-------

Glossary.....

Licenses.....

NVIDIA Software.....
License For Customer Use of NVIDIA Software.....
RECITALS.....
Sample File System.....
GST OpenMAX.....
GNU LESSER GENERAL PUBLIC LICENSE.....
How to Apply These Terms to Your New Libraries.....
GST EGL.....
GStreamer EGL/GLES Sink.....
Linux Kernel.....
GNU GENERAL PUBLIC LICENSE.....
mkbooting and mkubootscript.....
Apache License.....
Copyright: WIDE Project.....
GNU General Public License.....
Copyright: Regents of the University of California.....
U-Boot and mkimage.....
GNU GENERAL PUBLIC LICENSE, Version 2, June 1991.....
mkbctpart.....
brcm_patchram_plus.....
libnvcam_imageencoder.so.....
libscf.....
License Agreement for Protocol Buffers.....
License Agreement for Open Source Computer Vision Library.....
License Agreement for OpenCV Tutorial Library.....
License Agreement for Open Source Computer Vision Library.....
License Agreement for OpenCV Tutorial Library.....
JasPer License Version 2.0.....
JasPer License Version 2.0.....
JasPer License Version 2.0.....
The Independent JPEG Group's JPEG software.....
GNU GENERAL PUBLIC LICENSE.....
Threading Building Blocks.....
ZLIB DATA COMPRESSION LIBRARY.....
gstvideocuda.....
bpmp and tos-img.....

Legal Information.....

Quick Start Guide

This topic provides the steps to quickly get started using the L4T package with the Jetson TX 2 Developer Kit.

Note:

The most current Quick Start Guide for the Linux for Tegra Board Support Package is available at:

`<local_path>/t186ref_release_aarch64/l4t_quick_start_guide.txt`

Prerequisites

- A Jetson TX2 Developer Kit
- Your Jetson TX2 carrier board must be cabled as follows:
 - Serial cable plugged into the serial port on the target connected to your Linux host directly or through a serial-to-USB converter. This is required to setup serial console on the Linux host.
 - An Ethernet cable is plugged into the on-board Ethernet port.
 - An HDMI cable connects the carrier board to an external HDMI display.
 - (Not included in the developer kit) To connect USB peripherals such as keyboard, mouse, and USB/Ethernet adapter (for network connection), a USB hub must be connected to the working to the working USB port on the system.

To download the latest release package

1. Download the latest version of the L4T release package and sample filesystem for the Jetson Developer Kit.

The package is available from the NVIDIA Embedded Developer Zone at:

<https://developer.nvidia.com/linux-tegra>

If NVIDIA does not yet provide a public release for your Jetson Developer Kit, contact your NVIDIA Support Representative to obtain the latest release.

2. Untar the files and assemble the rootfs with the following commands:

```
sudo tar xpf $<L4T_RELEASE_PACKAGE>
cd Linux_for_Tegra/rootfs/
sudo tar xpf ../../$<SAMPLE_FS_PACKAGE>
cd ..
sudo ./apply_binaries.sh
```

Where:

- `<L4T_RELEASE_PACKAGE>` is `Tegra186_Linux_R27.0.1_aarch64.tbz2` for the Jetson TX2 release package filename.
- `<SAMPLE_FS_PACKAGE>` is `Tegra_Linux_Sample-Root-Filesystem_R27.0.1_aarch64.tbz2` for the sample filesystem package filename.

To flash the rootfilesystem to internal eMMC

1. Ensure your Linux host system is connected to the Jetson Developer Kit micro-USB port with a USB cable.
2. With the system powered on, place the Jetson Developer Kit into Force USB Recovery Mode as follows.
 - Press and hold the RECOVERY FORCE button.
 - While depressing the RECOVERY FORCE button, press and release the RESET (RST) button.
 - Wait 2 seconds and release the RECOVERY FORCE button.
3. Flash the L4T release onto the Jetson Developer Kit by executing the following command on your Linux host system:

```
sudo ./flash.sh $<platform> mmcblk0p1
```

Where `<platform>` is `jetson-tx2`.

The flashing procedure takes approximately 10 minutes or more on slower host systems.

Upon completion, the Jetson Developer Kit automatically reboots. The command prompt window opens on all display devices that are attached to the Jetson Developer Kit, and on the serial console.

The default login credentials are:

- Username: `nvidia`
- Password: `nvidia`

To install additional optional Ubuntu packages

1. Connect the Jetson Developer Kit to your wired network using a network cable.
2. Acquire an IP address.
 - If you are not running a graphical desktop environment that automatically configures your network connection, acquire an IP address with the following command:

```
sudo dhclient <interface>
```

Where `<interface>` is the name of the wired ethernet adapter for the Jetson Developer Kit's wired. Obtain the name by executing the following command:

```
ip link
```

The above command returns `eth0`, `enx<mac_address>`, or a similar name.

3. Install the package using the following commands:

```
sudo apt update
sudo apt install <ubuntu-desktop>
```

Where `<ubuntu-desktop>` is the name of the packages you wish to install.

4. Reboot the system using the following command:

```
sudo reboot
```

Consult the *Release Notes* for up-to-date platform and feature usage information.

Architecture

The following diagram shows the architecture of L4T and related components.



The following topics are included in the *Development Guide*.

Components	Description
•Sample Code/Applications	•Sample source code for developing embedded applications for the Jetson platform.
•NVIDIA [®] TensorRT™	<ul style="list-style-type: none"> •A high-performance neural network inference engine for production deployment of deep learning applications. TensorRT can be used to rapidly optimize, validate and deploy trained neural networks for inference to the Jetson platform. The Multimedia API provides samples that use this engine. •NVIDIA[®] JetPack installs this component as part of the Low-Level Multimedia API.
•V4L2	<ul style="list-style-type: none"> •A standard Linux V4L2 application that uses direct kernel IOCTL calls to access V4L2 functionality for encode, decode, scaling, and other media functions. •Tegra V4L2 camera driver bypasses the Tegra ISP and is based on Version 2 of the Linux kernel video capture and output device API and driver framework. •JetPack installs this component as part of the Low-Level Multimedia API.
•libArgus	<ul style="list-style-type: none"> •A low-level API based on the camera core stack. •JetPack installs this component as part of the Low-Level Multimedia API.
•Nvcamerasrc	<ul style="list-style-type: none"> •An NVIDIA camera GStreamer plugin that provides many options to control ISP properties. •JetPack installs this component as part of the Low-Level Multimedia API.
•GStreamer Framework	•Provides a higher-level Multimedia API. This framework provides a parallel capability to the collection of APIs in the Low-Level Multimedia API.

	<ul style="list-style-type: none"> •The GStreamer framework is installed with the download for the Jetson Driver. It is also installed by NVIDIA JetPack. •This framework is an alternative to Multimedia API.
•Multimedia API	<ul style="list-style-type: none"> •A collection of lower-level APIs that support flexible application development by providing better control over the underlying hardware blocks. (Not shown in the architecture diagram.) •Multimedia API includes: <ul style="list-style-type: none"> • libargus for imaging applications • V4L2 API for encoding, decoding, scaling, and other media functions • NVOSD for On-Screen display • Buffer Utility for buffer allocation, management, and sharing •It includes examples that demonstrate image processing with CUDA, object detection and classification with cuDNN, TensorRT, and OpenCV4Tegra usage. •This API is an alternative to using a framework such as GStreamer.
•X11	•X11 X Window System driver.
•libjpeg	•C library for reading and writing JPEG image files.
•CUDA	•NVIDIA [®] CUDA [®] parallel computing platform and API for CUDA-enabled GPU.
•VisionWorks	•NVIDIA [®] VisionWorks™ software development package for computer vision (CV) and image processing.
•EGL	•Interface between Khronos rendering APIs such as OpenGL ES or OpenVG and the underlying native platform window system.
•OpenGL ES	•Cross-platform API for full-function 2D and 3D graphics on embedded systems
•cuDNN	•NVIDIA [®] CUDA [®] Deep Neural Network library.

Package Manifest

The NVIDIA[®] Tegra[®] Linux Driver Package is provided in the following tar file:

```
Tegra<t-arch>_Linux_R<release_num>_<release_type>.tbz2
```

Tip: In the above expression, float your cursor over a place-holder to reveal the currently-defined value.

The following table lists the directories (denoted by a trailing slash /) and top-level files that are created when you expand the tar file.

Directory or Filename	Description
•bootloader/	•Boot loader and related components. For more information on this directory, see Boot Loader in this topic.
•bootloader/<platform ver>/	•Platform-specific files.
•bootloader/<platform ver>/BCT/	•Platform-specific Boot Configuration Table (BCT) files.
•bootloader/<platform ver>/cfg/	•Configuration files for specific <platform ver>.
•bootloader/<platform ver>/<board_and_rev>/	•Boot-related DTB and configuration files for the specific <board_and_rev>.
•	
•camera-rtcpu-sce.bin	•Firmware for handling RT aspects of frame capture in the camera hardware pipeline.
•kernel/	•Kernel images and kernel modules. For more information on this directory, see Kernel in this topic.
•nv_tegra/	•NVIDIA drivers and sample applications.
•nv_tegra/nv_sample_apps/	•NVIDIA sample applications.
•rootfs/	•Staging directory for the root filesystem.
•rootfs/README.txt	•README on installing the root filesystem.
•apply_binaries.sh	•Script to apply nv_tegra components.
•elf-get-entry.py	•Python script to extract and print entry point address of an ELF-format binary.
•flash.sh	•Script to flash the boot loader and kernel from the package.

<ul style="list-style-type: none"> •<platform>.conf •<board_and_rev>.conf •<board_and_rev>.conf.common 	<ul style="list-style-type: none"> •Configuration file(s) for <code>flash.sh</code> specific to the development platform represented by <platform> or <board_and_rev>.
<ul style="list-style-type: none"> •source_sync.sh 	<ul style="list-style-type: none"> •Script to download kernel and U-Boot source.

Documentation

Tegra Linux Driver Package (L4T) also includes the following documentation:

- Tegra_Linux_Driver_Package_Release_Notes_<release>.pdf
- Tegra_Linux_Driver_Package_Documents_<release>.tar

Kernel

The `kernel` directory contains the following directories (denoted by a trailing slash /) and files.

Directory or Filename	Description
•dtb/	•SoC-specific kernel Device Tree Binary (DTB) files.
•dtb/*.dtb	•DTB files specific to various board types.
•dtc	•Device-tree-compiler binary.
•Image	•Kernel binary image.
•kernel_headers.tbz2	•Kernel header files needed for compiling kernel modules. You can download these headers and sources from the <code>nv_tegra</code> git server.
•kernel_supplements.tbz2	•Loadable kernel modules specific to the included kernel zImage that was built with the <code>defconfig</code> enabled for the device.
•LICENSE	•GNU General Public License (GPL).
•LICENSE.dtc	•GNU General Public License (GPL) for the device-tree-compiler binary.
•pinmux/	•-
•pinmux/addr_info.txt	•Contains the database of pins available on tegra. Information includes, pin name, address, pinmux functionality of pin.
•pinmux/gpio_addr_info.txt	•Contains the database of tegra GPIO port name and its base address.
•pinmux/pad_info.txt	•Contains the database of tegra pads. Information includes pad name, voltage

	type (1.2V, 1.8V, 3.3V) and bit fields information in pad register.
•pinmux/pinmux-dts2cfg.py	•Script tool used to generate CFG format file from given pinmux, gpio, pad platform dts files.
•pinmux/por_val.txt	•Contains the database of pinmux address and value pairs which should not be modified while running the tool.
•pinmux/README.txt	•Contains details on how to run the script.

Kernel Supplements TBZ2

The following table lists the contents available upon decompressing the `kernel_supplements.tbz2` archive, located at:

```
kernel/kernel_supplements.tbz2
```

Filename	Description
•lib/	•-
•lib/modules/	•-
•lib/modules/<linux_version>/*	•Kernel modules.

Boot Loader

The `bootloader` directory contains the following directories (denoted by a trailing slash /) and files.

Directory or Filename	Description
•<platform ver>/	•<platform ver>-specific boot loader directory.
•<platform ver>/BCT/	•Platform-specific BCT directory.
•<platform ver>/BCT/*.cfg	•Platform-specific Boot Configuration Table (BCT) files.
•<platform ver>/cfg/	•Platform-specific configuration directory.
•<platform ver>/cfg/ board_config_ers_e2220.xml	•Platform-specific configuration file.
•<platform ver>/cfg/ board_config_p2595.xml	•Platform-specific configuration file.
•<platform ver>/cfg/ board_config_p2597.xml	•Platform-specific configuration file.

•<platform ver>/cfg/board_config_p2597-devkit.xml	•Platform-specific configuration file.
•<platform ver>/cfg/flash_l4t_t186.xml	•Platform-specific configuration file.
•<platform ver>/cfg/gnu_linux_fastboot_emmc_full.cfg	•Platform-specific configuration file.
•<platform ver>/cfg/gnu_linux_tegraboot_emmc_full.xml	•Platform-specific configuration file.
•<platform ver>/*.dtb	•Platform-specific device tree files.
•<platform ver>/DTB/*.dtb	•Platform-specific device tree files.
•<platform ver>/<board_and_rev>/	•Boot loader <board_and_rev>-specific directory.
•<platform ver>/<board_and_rev>/extlinux.conf.emmc	•<board_and_rev>-specific U-Boot config file for booting off the internal EMMC.
•<platform ver>/<board_and_rev>/extlinux.conf.nfs	•<board_and_rev>-specific U-Boot config file for booting off the nfs root.
•<platform ver>/<board_and_rev>/extlinux.conf.sdcard	•<board_and_rev>-specific U-Boot config file for booting off the SD card.
•<platform ver>/<board_and_rev>/extlinux.conf.usb	•<board_and_rev>-specific U-Boot config file for booting off USB flash storage device.
•<platform ver>/<board_and_rev>/u-boot	•<board_and_rev>-specific U-Boot boot loader binary.
•<platform ver>/<board_and_rev>/u-boot.bin	•<board_and_rev>-specific U-Boot boot loader binary.
•<platform ver>/<board_and_rev>/u-boot.dtb	•<board_and_rev>-specific U-Boot device tree binary.
•<platform ver>/<board_and_rev>/u-boot-dtb.bin	•<board_and_rev>-specific U-Boot device tree binary.
•cboot.bin	•CPU binary to load the kernel. It also supports Fastboot, charging, and display.
•LICENSE.cboot	•LICENSE file for the cboot.bin binary.
•<platform ver>/nvtboot.bin	•Tegra boot specific boot loader binary (AVP bootloader, microboot, miniloader).
•<platform ver>/warmboot.bin	•Warm boot binary.
•bpmf.bin	•Boot and power management firmware.
•chkbinfo	•TBD

•eks.img	•Encrypted key set data file containing 8 zero-length keys.
•exec-uboot.sh	•Shell script used to load U-Boot into RAM and execute.
•gen-tboot-img.py	•Script used by the bootlaoder to add an nvtboot-specific header during the flash process.
•l4t_initrd.img	•L4T initrd image based on minimal Ubuntu environment.
•LICENSE	•Tegra software license.
•LICENSE.bpmp	•License file for bpmp.bin.
•LICENSE.camera-rtcpu	•License file for camera-rtcpu.
•LICENSE.chkbdinfo	•License file for chkbdinfo.
•LICENSE.mkbctpart	•License file for mkbctpart.
•LICENSE.mkbootimg	•License file for mkbootimg.
•LICENSE.mkgpt	•License file for the mkgpt tool.
•LICENSE.mksparse	•License file for the mksparse tool.
•LICENSE.spe.bin	•License for the spe tool.
•LICENSE.tos.img.arm-trusted-firmware	•License for the tos.img firmware.
•LICENSE.u-boot	•License for U-Boot.
•mb1_prod.bin	•First stage software boot binary that runs post Boot ROM in coldboot mode
•mb1_recovery_prod.bin	•First stage sw boot binary that runs post Boot ROM in recovery mode.
•mce_mts_d15_prod_cr.bin	•CPU microcode for jetson-tx2.
•preboot_d15_prod_cr.bin	•CPU microcode for jetson-tx2.
•mkbootimg	•Tool for img creation.
•mkgpt	•Tool that encodes both primary and secondary GPT into flashable binary image files.
•mkimage	•U-Boot tool for vmlinux.uimg creation.
•mksparse	•Sparse image flashing with the boot loader.
•mkubootscript	•Tool for flashing U-Boot.

•nvtboot_cpu.bin	•CPU part of Tegraboot for TLK hand over transition.
•nvtboot_recovery.bin	•AVP bootrom applet binary used by Tegraflash
•nvtboot_recovery_cpu.bin	•CPU part of Tegraboot used for RCM boot for MODS.
•spe.bin	•NVIDIA sensorhub microcontroller firmware.
•tegrabct_v2	•BCT operation helper binary.
•tegradevflash_v2	•Boot loader device communication library.
•tegraflash.py	•Script used to flash the board.
•tegraflash_internal.py	•Helper implementation API script for tegraflash.py.
•tegrahost_v2	•Boot loader encryption binary.
•tegraparser_v2	•Parses partition layout, common BCT configuration, fuse bypass configuration and NVIDIA Configuration Table (NCT).
•tegrarcm_v2	•Bootrom RCM communications binary.
•tegrasign_v2	•TegraSign creates signature data for PKC operating mode and hash, and encrypted data for SBK operating mode.
•tos.img	•The monitor binary running in the EL3 exception space on ARMv8 CPUs.

NV Tegra

The `nv_tegra` directory contains the following directories (denoted by a trailing slash /) and files.

Direcotry or Filename	Description
•nv_sample_apps/	•NVIDIA sample applications.
•nv_sample_apps/LICENSE.gst-openmax	•License for the <code>libgstomx.so</code> , <code>libgstnvegl-1.0.so.0</code> , and <code>libnvgstjpeg.so</code> libraries included in <code>nvgstapps.tbz2</code> .
•nv_sample_apps/LICENSE.gstvideocuda	•License for Gstreamer 1.0 CUDA post-processing plugin library.
•nv_sample_apps/nvgstapps.tbz2	•NVIDIA gstreamer components and

	applications. For details, see Nvgstapps TBZ2 table (below).
•nv_sample_apps/nvgstcapture-<version>_README.txt	•Read Me for NVIDIA Gstreamer-based camera capture application (nvgstcapture).
•nv_sample_apps/nvgstplayer-<version>_README.txt	•Read Me for NVIDIA Gstreamer-based multimedia player (nvgstplayer).
•config.tbz2	•Configuration files specific to the sample filesystem. For details, see Config TBZ2 table (below).
•LICENSE	•Tegra software license.
•LICENSE.brcm_patchram_plus	•License for brcm_patchram_plus.
•LICENSE.libargus	•License for the libargus API.
•LICENSE.libglvnd	•License file for libglvnd.
•LICENSE.libnvcam_imageencoder	•License for image encoder.
•LICENSE.libscf	•License for core camera driver.
•LICENSE.libtegrav4l2	•License for libtegrav4l2.
•nvidia_drivers.tbz2	•NVIDIA driver components.
•nv_tools.tbz2	•The <code>tegrastats</code> application, a script for calculations for loads, frequencies, RAM sizes, using existing sysfs nodes. Refer to the <i>Development Guide</i> for usage.

Nvgstapps TBZ2

The following table lists the directories (denoted by a trailing slash /) and files available upon decompressing the `nvgstapps.tbz2` archive, located at:

`nv_tegra/nv_sample_apps/nvgstapps.tbz2`

Filename	Description
•usr/bin/gst-install	•Script to build gstreamer from sources. Version can be specified with the <code>--version</code> option (1.6.0 is the default).
•usr/bin/nvgstcapture-<version>	•Multimedia capture camera application.
•usr/bin/nvgstplayer-<version>	•Multimedia video player application.
•usr/lib/< ABI directory >/gstreamer-<version>/	•Plug-ins and drivers for gstreamer.
•usr/lib/< ABI directory >/gstreamer-<version>/libgstnvcamera.so	•gst-plugin library for camera.

Filename	Description
•usr/lib/<ABI_directory>/gststreamer-1.0/libgstnveglglessink.so	•Accelerated Egl based renderer element for gststreamer-1.0.
•usr/lib/<ABI_directory>/gststreamer-1.0/libgstnveglstreamsrc.so	•EGLStream Consumer functionality library.
•usr/lib/<ABI_directory>/gststreamer-1.0/libgstnvegltransform.so	•NVM buffer conversion to EGLImage plugin library.
•usr/lib/<ABI_directory>/gststreamer-1.0/libgstnvivafilter.so	•CUDA post-processing plugin library for gststreamer-1.0.
•usr/lib/<ABI_directory>/gststreamer-<version>/libgstnvidconv.so	•NVIDIA proprietary GStreamer conversion plug-in library.
•usr/lib/<ABI_directory>/gststreamer-1.0/libgstnvideosink.so	•GStreamer 1.0 EGLProducer video sink plugin.
•usr/lib/<ABI_directory>/gststreamer-<version>/libgstomx.so	•OpenMax driver.
•usr/lib/<ABI_directory>/gststreamer-1.0/libgstvideocuda.so	•Gstreamer 1.0 CUDA post-processing plugin library.
•usr/lib/<ABI_directory>/gststreamer-<version>/libnvgstjpeg.so	•Accelerated libjpeg based jpeg decoding and encoding library.
•usr/lib/<ABI_directory>/libgstnvegl-1.0.so.0	•Gstreamer EGL API wrapper library.
•usr/lib/<ABI_directory>/libgstnvexifmeta.so	•Gstreamer buffer exif metadata library.
•usr/lib/<ABI_directory>/libgstnvivameta.so	•Interface library used to add and get gst metadata.
•usr/lib/<ABI_directory>/libsample_cudaprocess.so	•"gst-nvivafilter" sample for cuda post-processing.

Config TBZ2

The following table lists the contents available upon decompressing the `config.tbz2` archive, located at:

```
nv_tegra/config.tbz2
```

Filename	Description
•etc/init/argus-daemon.conf	•Configuration file for the argus-daemon.
•etc/init/nv.conf	•NVIDIA-specific initialization script.
•etc/init/nvcamera-daemon.conf	•nvcamera-daemon service configuration launcher.
•etc/init/nvfb.conf	•NVIDIA specific first-boot script.

Filename	Description
•etc/init/nvwifibt.conf	•NVIDIA bluetooth/wifi init script.
•etc/init/ttyS0.conf	•Initialization script for getty on ttyS0.
•etc/modprobe.d	•Configuration directory/file for modprobe.
•etc/modprobe.d/bcmdhd.conf	•NVIDIA specific modprobe configuration file for bcmdhd driver loading.
•etc/nv/nvfirstboot	•Control file used for for first boot.
•etc/nvpmode.conf	•Configuration file for nvpmode, containing clock settings for each power mode.
•etc/pulse/daemon.conf	•Configuration file for the PulseAudio daemon.
•etc/pulse/default.pa.hdmi	•PulseAudio configuration file.
•etc/pulse/default.pa.orig	•PulseAudio configuration file.
•etc/sysctl.d/10-console-messages.conf	•Contains the default loglevel for the console messages.
•etc/sysctl.d/90-tegra-settings.conf	•Control file for sysrq.
•etc/systemd/	•-
•etc/systemd/nvfb-pre.sh	•NVIDIA pre-init boot script.
•etc/systemd/nvfb.sh	•NVIDIA first boot script.
•etc/systemd/nvpulse.sh	•Service to launch etc/systemd/nvpulse.sh script.
•etc/systemd/nv.sh	•NVIDIA init script.
•etc/systemd/nvwifibt-pre.sh	•NVIDIA bluetooth/Wi-Fi pre-init script.
•etc/systemd/nvwifibt.sh	•NVIDIA bluetooth/Wi-Fi init script.
•etc/systemd/system/	•-
•etc/systemd/system/argus-daemon.service	•Libargus daemon systemd startup service.
•etc/systemd/system/nvcamera-daemon.service	•Service to launch nvcamera-daemon.
•etc/systemd/system/nvfb.service	•Service to launch etc/systemd/nvfb.sh script.
•etc/systemd/system/nvpmode.service	•Systemd service executed at boot to set the MaxEDP mode.

Filename	Description
•etc/systemd/system/nvpulse.service	•Service to launch etc/systemd/nvpulse.sh script.
•etc/systemd/system/nv.service	•Service to launch etc/systemd/nv.sh script.
•etc/systemd/system/nwifibt.service	•Service to launch etc/systemd/nwifibut-pre and nwifibt scripts.
•etc/systemd/system/serial-getty@ttyS0.service.d/	•-
•etc/systemd/system/serial-getty@ttyS0.service.d/autologin.conf	•Configuration file to enable autologin on console.
•etc/udev/rules.d/90-alsa-asound-tegra.rules	•Rules configuration for proper asound.conf selection.
•etc/udev/rules.d/91-xorg-conf-tegra.rules	•Rules configuration for proper xorg.conf selection.
•etc/udev/rules.d/92-hdmi-audio-tegra.rules	•Rules configuration for proper /etc/pulse/default.pa selection.
•etc/udev/rules.d/99-nv-wifibt.rules	•Rules configuration for Wi-Fi and Bluetooth.
•etc/udev/rules.d/99-tegra-devices.rules	•Permission setting for Tegra devices.
•etc/udev/rules.d/99-tegra-mmc-ra.rules	•SD card read_ahead_kb configuration.
•etc/X11/xorg.conf	•Configuration file for xorg.
•etc/X11/xorg.conf.*	•Configuration file for <board and rev> -specific xorg.
•etc/amixer_settings	•Audio configuration.
•etc/asound.conf.hdmi	•ALSA sound configuration for HDMI audio.
•etc/asound.conf.tegrahda	•Asound configuration file.
•etc/asound.conf.tegrasndt186ref	•Asound configuration file for Tegra code-name Parker audio supermodule.
•etc/asound.conf.tegrasnd< platform ver >	•ALSA sound configuration for onboard audio.
•etc/encntune.conf	•Default multimedia encoding parameters for NVIDIA reference platforms.
•etc/modules	•Lists bluedroid as a supporting module for Bluetooth.
•etc/motd	•One time factory output message file.

Filename	Description
•etc/wpa_supplicant.conf	•Sample WPA supplicant.
•lib/udev/rules.d/73-special-net-names.rules	•Network interface naming rule from "udev" debian package.

NVIDIA Drivers TBZ2

The following table lists the contents available upon decompressing the `nvidia_drivers.tbz2` archive, located at:

```
nv_tegra/nvidia_drivers.tbz2
```

Filename	Description
•etc/ld.so.conf.d/nvidia-tegra.conf	•Ldconf file for tegra directories.
•etc/nv_tegra_release	•Tegra driver versioning file.
•lib/firmware/brcm/	•BCM firmware directory.
•lib/firmware/brcm/fw_bcmdhd.bin	•Firmware for Jetson-tx1 on-board wifi.
•lib/firmware/brcm/nvram.txt	•File containing tuning parameters for the jetson-tx1 on-board wifi.
•lib/firmware/gp10b/	•Firmware files for Jetson TX2.
•lib/firmware/gp10b/acr_ucose.bin	•High secure mode PMU code.
•lib/firmware/gp10b/fecs.bin	•GPU FECS firmware.
•lib/firmware/gp10b/fecs_sig.bin	•Signature of FECS microcode.
•lib/firmware/gp10b/gpccs.bin	•GPU GPCCS firmware.
•lib/firmware/tegra12x/gpccs_sig.bin	•Signature of GPU GPCCS firmware.
•lib/firmware/gp10b/gpmu_ucose.bin	•GPU PMU ucode firmware
•lib/firmware/gp10b/gpmu_ucose_desc.bin	•Descriptor data for LS PMU.
•lib/firmware/gp10b/gpmu_ucose_image.bin	•Low-secure mode PMU code.
•lib/firmware/gp10b/gpu2cde.bin	•GPU shader program used for converting GPU compression metadata to be read by VIC and Display.
•lib/firmware/gp10b/NETB_img.bin	•GPU device hardware description.
•lib/firmware/gp10b/pmu_bl.bin	•Boot loader loading acr_ucose.bin.
•lib/firmware/gp10b/pmu_sig.bin	•Signature of gpmu_ucose_image.bin.

Filename	Description
•lib/firmware/tegra18x/nvhost_nvdec020.fw	•NVDEC firmware file for video decode.
•lib/firmware/tegra18x/nvhost_nvdec020_ns.fw	•NVDEC firmware that runs without boot loader.
•lib/firmware/tegra18x/nvhost_nvdec020_prod.fw	•NVDEC firmware.
•lib/firmware/tegra18x/nvhost_nvdec_bl020_prod.fw	•NVDEC boot loader firmware.
•lib/firmware/tegra18x/nvhost_nvenc061.fw	•NVENC firmware file for video decode.
•lib/firmware/tegra18x/nvhost_nvjpg011.fw	•NVJP firmware file for jpeg encode and decode.
•lib/firmware/gp10b/vic04_ucose.bin	•VIC hardware-specific ucode control firmware.
•lib/firmware/tegra18x_xusb_firmware	•USB 3.0 firmware.
•lib/firmware/bcm4354.hcd	•Bluetooth firmware for the BCM4354 chip.
•	•
•usr/bin/nvidia-bug-report-tegra.sh	•NVIDIA bug reporting script. Run for usage tips.
•usr/lib/	•-
•usr/lib/< ABI directory >/	•-
•usr/lib/< ABI directory >/libv4l/plugins/libv4l2_nvvidconv.so	•Gstreamer (nv to raw and raw to nv) conversion plugin.
•usr/lib/< ABI directory >/libv4l/plugins/libv4l2_nvvideocodec.so	•Video encode/decode libv4l2 plugin library.
•usr/lib/< ABI directory >/tegra/	•-
•usr/lib/< ABI directory >/tegra/libargus.so	•libargus camera library.
•usr/lib/< ABI directory >/tegra/libargus_socketclient.so	•A library that supports the multiprocess implementation of the libargus API.
•usr/lib/< ABI directory >/tegra/libargus_socketserver.so	•A library that supports the multiprocess implementation of the libargus API.
•usr/lib/< ABI directory >/tegra/libcuda.so.1.1	•CUDA library.
•usr/lib/< ABI directory >/tegra/libdrm.so.2	•Alternative OSS libdrm library.

Filename	Description
•usr/lib/<ABI_directory>/tegra/libGL.so.1	•GL graphics support library.
•usr/lib/<ABI_directory>/tegra/libGLdispatch.so.0	•OpenGL dispatching and TLS library.
•usr/lib/<ABI_directory>/tegra/libglx.so	•GLX extension module for X. Module is used by the X server to provide server-side GLX support.
•usr/lib/<ABI_directory>/tegra/libnvapputil.so	•Host (x86) shared object for application utilities.
•usr/lib/<ABI_directory>/tegra/libnvavp.so	•User-space interface to the AVP for audio/video acceleration via the nvavp kernel driver.
•usr/lib/<ABI_directory>/tegra/libnvbuf_utils.so.1.0.0	•libv2 helper library.
•usr/lib/<ABI_directory>/tegra/libnvcameratools.so	•Supporting library for NVIDIA camera utilities.
•usr/lib/<ABI_directory>/tegra/libnvcamerautils.so	•Supporting library for NVIDIA camera utilities.
•usr/lib/<ABI_directory>/tegra/libnvcam_imageencoder.so	•Library encodes camera YUV frames to JPEG using the NVIDIA TVMR architecture.
•usr/lib/<ABI_directory>/tegra/libnvcamlog.so	•Camera runtime tracing and logging helper library.
•usr/lib/<ABI_directory>/tegra/libnvcolorutil.so	•NvColor utility library.
•usr/lib/<ABI_directory>/tegra/libnvdc.so	•DC driver file.
•usr/lib/<ABI_directory>/tegra/libnvddk_2d_v2.so	•DDK 2D.
•usr/lib/<ABI_directory>/tegra/libnvddk_vic.so	•DDK VIC.
•usr/lib/<ABI_directory>/tegra/libnveglstream_camconsumer.so	•The libargus consumer library.
•usr/lib/<ABI_directory>/tegra/libnveglstreamproducer.so	•Library that implements EGLStream Producer functionality.
•usr/lib/<ABI_directory>/tegra/libnvidia-eglcore.so.<release_num>	•EGL core library.
•usr/lib/<ABI_directory>/tegra/libnvidia-fatbinaryloader.so.<release_num>	•A library that supports interactions between the CUDA driver and CUDA fatbinaries. Fatbinary is a container format that packages different PTX

Filename	Description
	and Cubins compiled for different architectures.
•usr/lib/<ABI_directory>/tegra/libnvidia-glc core.so.<release_num>	•OpenGL core library. This library is implicitly used by libGL and by libglx, and contains the core accelerated 3D functionality.
•usr/lib/<ABI_directory>/tegra/libnvidia-glsi.so.<release_num>	•OpenGL System Interaction library.
•usr/lib/<ABI_directory>/tegra/libnvidia-ptxjitcompiler.so.<release_num>	•A library that provides a JIT compiler that compiles PTX into GPU machine code and is used by the CUDA driver.
•usr/lib/<ABI_directory>/tegra/libnvidia-rmapi-tegra.so.<release_num>	•Utility library that implements common code for using kernel-level graphics drivers on Tegra.
•usr/lib/<ABI_directory>/tegra/libnvidia-tls.so.<release_num>	•NVIDIA tls libraries.
•usr/lib/<ABI_directory>/tegra/libnvjpeg.so	•Accelerated libjpeg for Tegra.
•usr/lib/<ABI_directory>/tegra/libnvll.so	•Utility library used by libdrm.so.2 for memory management.
•usr/lib/<ABI_directory>/tegra/libnvmedia.so	•Multimedia programming API to access hardware units like encoder, decoder, and video post-processing on Tegra.
•usr/lib/<ABI_directory>/tegra/libnvmm.so	•NVIDIA Multimedia Framework.
•usr/lib/<ABI_directory>/tegra/libnvmm_contentpipe.so	•Content pipe implementation (file source abstraction).
•usr/lib/<ABI_directory>/tegra/libnvmm-lite.so	•NVIDIA Multimedia driver.
•usr/lib/<ABI_directory>/tegra/libnvmm-lite_image.so	•NVIDIA Multimedia image driver.
•usr/lib/<ABI_directory>/tegra/libnvmm-lite_utils.so	•NVIDIA Multimedia utilities.
•usr/lib/<ABI_directory>/tegra/libnvmm-lite_video.so	•NVIDIA Multimedia video driver.
•usr/lib/<ABI_directory>/tegra/libnvmm_parser.so	•NVIDIA Multimedia parser.
•usr/lib/<ABI_directory>/tegra/libnvmm_utils.so	•Multimedia Framework utilities.

Filename	Description
•usr/lib/<ABI_directory>/tegra/libnvodm_imager.so	•Tegra development platform ODM adaptation for imager.
•usr/lib/<ABI_directory>/tegra/libnvomx.so	•OpenMAX IL implementation.
•usr/lib/<ABI_directory>/tegra/libnvomxilclient.so	•OpenMAX IL client.
•usr/lib/<ABI_directory>/tegra/libnvos.so	•NVIDIA OS abstraction library.
•usr/lib/<ABI_directory>/tegra/libnvparses.so	•Parser used for NVIDIA NvMMLite.
•usr/lib/<ABI_directory>/tegra/libnvrm.so	•Resource Manager kernel interface.
•usr/lib/<ABI_directory>/tegra/libnvrm_gpu.so	•NVIDIA kernel graphics driver abstraction library.
•usr/lib/<ABI_directory>/tegra/libnvrm_graphics.so	•Resource Manager (NvRM) graphics host, AVP communication library, and graphics drivers.
•usr/lib/<ABI_directory>/tegra/libnvtestresults.so	•Test results library.
•usr/lib/<ABI_directory>/tegra/tegra/libnvtnr.so	•Temporal Noise Reduction (TNR) interface.
•usr/lib/<ABI_directory>/tegra/libnvtvmr.so	•Multimedia Tegra video mixer/renderer.
•usr/lib/<ABI_directory>/tegra/libnvwinsys.so	•Winsys library.
•usr/lib/<ABI_directory>/tegra/libOpenGL.so	•Provides symbols for OpenGL entry point library.
•usr/lib/<ABI_directory>/tegra/libscf.so	•Core camera driver.
•usr/lib/<ABI_directory>/tegra/libtegrav4l2.so	•V4L2 driver for Tegra.
•usr/lib/<ABI_directory>/tegra/nvidia_icd.json	•Vulkan ICD configuration file.
•usr/lib/<ABI_directory>/tegra-egl/	•-
•usr/lib/<ABI_directory>/tegra-egl/ld.so.conf	•Ldconf file for tegra-egl directories.
•usr/lib/<ABI_directory>/tegra-egl/libEGL.so.1	•OpenGL ES driver file.
•usr/lib/<ABI_directory>/tegra-egl/libEGL_nvidia.so.0	•OpenGL ES driver file.

Filename	Description
•usr/lib/<ABI_directory>/tegra-egl/libGLESv1_CM.so.1	•OpenGL ES driver file.
•usr/lib/<ABI_directory>/tegra-egl/libGLESv2.so.2	•OpenGL ES driver file.
•usr/lib/xorg/	•X Windows System libraries and drivers
•usr/lib/xorg/modules/	•-
•usr/lib/xorg/modules/drivers/	•-
•usr/lib/xorg/modules/drivers/nvidia_drv.so	•Tegra X driver.
•usr/lib/xorg/modules/extensions/	•-
•usr/lib/xorg/modules/extensions/libglx.so	•Symbolic link pointing to /usr/lib/<ABI_directory>/tegra/libglx.so in the rootfs.
•usr/sbin/	•-
•usr/sbin/argus_daemon	•System daemon process that is run at startup and provides libargus clients access to the libargus camera API in a multiprocess environment.
•usr/sbin/brcm_patchram_plus	•Utility for loading the broadcom bluetooth firmware.
•usr/sbin/nvcamera-daemon	•Daemon process for using multiple or simultaneous camera instances on L4T platform using core_scf library.
•usr/sbin/nvtunerd	•Support for image quality tuning tools.
•usr/share/doc/nvidia-tegra/LICENSE.brcm_patchram_plus	•License for brcm_patchram_plus.
•usr/share/doc/nvidia-tegra/LICENSE.libargus	•License for the libargus API.
•usr/share/doc/nvidia-tegra/LICENSE.libglvnd	•License file for libglvnd.
•usr/share/doc/nvidia-tegra/LICENSE.libnvcam_imageencoder	•License for image encoder.
•usr/share/doc/nvidia-tegra/LICENSE.libtegrav4l2	•License for libtegrav4l2.
•usr/share/doc/procps/copyright	•-
•usr/share/doc/udev/copyright	•-

Filename	Description
•var/nvidia/	•-
•var/nvidia/nvcam/	•-
•var/nvidia/nvcam/apps/	•-
•var/nvidia/nvcam/apps/README.txt	•Nvcam application README.txt file.
•var/nvidia/nvcam/input/	•-
•var/nvidia/nvcam/input/ model_frontal.xml	•-
•var/nvidia/nvcam/input/README.txt	•Nvcam input README.txt file.
•var/nvidia/nvcam/output/	•-
•var/nvidia/nvcam/output/README.txt	•Nvcam output README.txt file.
•var/nvidia/nvcam/settings/	•-
•var/nvidia/nvcam/settings/README.txt	•Nvcam settings README.txt file.

NVIDIA Drivers TBZ2

The following table lists the contents available upon decompressing the `nv_tools.tbz2` archive, located at:

`nv_tegra/nv_tools.tbz2`

Filename	Description
•home/nvidia/jetson_clocks.sh	•Clock frequency control script.
•home/nvidia/tegrastats	•Performance measurement utility.
•home/ubuntu/jetson_clocks.sh	•Clock frequency control script.
•home/ubuntu/tegrastats	•Performance measurement utility.

Advanced Setup and Customization

To use with NVIDIA[®] Tegra[®] Linux Driver Package (L4T), you can manually configure and setup the software drivers before use or allow JetPack to perform the standard setup.

Consult your board documentation for guidance on manually setting up and configuring your reference board.

Note:

NVIDIA JetPack performs the following:

- Install the Linux for Tegra package
- Configure or flash your Tegra device
- Run samples

If you installed using JetPack, skip these topics. For more information, see the JetPack documentation.

Reference Board Preparation

When developing systems and application software with L4T, you run and test your code on an actual reference platform, such as the NVIDIA[®] Jetson[™] TX2 developer kit. Your code targets this hardware directly, rather than a software simulator or emulator.

Accordingly, you must acquire and set up your reference board before using L4T. Consult your board documentation for guidance on setting up and configuring your board.

Although the reference board supports a variety of peripheral devices, you can start developing on L4T with a board that has the following:

- One of the storage devices specified in [Boot Options](#)
- A USB cable to plug into the board recovery port

Boot Options

Boot L4T on the Jetson TX2 reference board from a root file system (rootfs) on integrated, attached, or network-accessible storage. The boot loader must be loaded from the internal eMMC. Root filesystem options include:

- USB stick (formatted to EXT4)
- USB hard disk (formatted to EXT4)
- SD card (formatted to EXT4)
- Internal eMMC
- Network File System (NFS)

Linux Host System Prerequisites

To use L4T on a Linux host system, the following hardware and software prerequisites must be met:

- Host PC running Linux `<os_ver_host>`.
- A kernel image (Image). L4T contains a kernel image for your use. Alternatively, you can download and rebuild the kernel image from source.
- Boot loader. Flashing on a Tegra X2 series (Jetson TX2) developer board requires a boot loader, which is a combination of NVIDIA T-Boot (nvtboot) and U-Boot.
- NFS if you intend to boot L4T on the reference board from your Linux host system or a network-accessible server.
- A USB cable to plug into the recovery port.

Extracting Tegra Linux Driver Package

Use the following procedures to extract your L4T package. Commands in the examples assume you extracted the release package in `~/`.

To extract Tegra Linux Driver Package

- Extract the package manually by executing the following command:

```
$ sudo tar -vxjf Tegra<t-arch|ver>_Linux_R<release_num>_<release_type>.tbz2
```

Tip:

In the above expression, float your cursor over a placeholder to reveal the currently-defined value.

Setting Up Your File System

L4T requires a root file system. You must create a Linux host system and copy it to your reference board.

Sample Root File System

L4T comes with a pre-built sample root file system created for the Jetson TX2 developer kit. If you wish to create an Ubuntu sample root file system, see:

<https://wiki.ubuntu.com/ARM/RootfsFromScratch>

Setting Up the Root File System

Before booting the target board, configure the root file system (rootfs) to:

- Set up the rootfs
- Copy it to the rootfs on the device

Step 1: Set Up the Root File System

This procedure uses the sample file system provided by NVIDIA as the base. If you wish to use your own file system, set the `LDK_ROOTFS_DIR` environment variable to point to the location of your rootfs and skip the steps for setting the root file system.

To set up the rootfs

1. Download the following file to your home directory:

```
Tegra-Linux-Sample-Root-Filesystem_<release_type>.tbz2
```

This file contains the NVIDIA-provided sample root file system.

2. Extract the compressed file as follows:

- Navigate to the rootfs directory of the extracted NVIDIA driver package.

```
$ cd <your_L4T_root>/Linux_for_Tegra/rootfs
```

Where `<your_L4T_root>` is your L4T root directory, which is assumed to be your home directory (`~`).

For more information, see [Extracting Tegra Linux Driver Package](#).

- Extract the sample file system to the rootfs directory.

```
$ sudo tar -jxpf ../../Tegra-Linux-Sample-Root-Filesystem_<release_type>.tbz2
```

3. Run the `apply_binaries.sh` script to copy the NVIDIA user space libraries into the target file system.

```
$ cd ..
```

```
$ sudo ./apply_binaries.sh
```

4. If you are using a different rootfs, or if you have already configured your rootfs, apply the NVIDIA user space libraries by setting the `LDK_ROOTFS_DIR` environment variable to point to your rootfs. Then run the script, as shown above, to copy the binaries into your target file system.

If the `apply_binaries.sh` script installs the binaries correctly, the last message output from the script is "Success!".

You have now completed setting up the root filesystem. Proceed to flash the rootfs onto the target Tegra device.

Step 2: Copy the rootfs to the Device

Use these procedures to copy the file system to the Tegra device.

1. Pick a device to place your rootfs.
 - If you are using the internal EMMC, skip ahead to [Flashing the Bootloader and Kernel](#).
2. If you prefer to use an external storage device for the root filesystem, use the following procedure.

To copy the file system to an external storage device

1. Plug your rootfs device into the host system.
2. If your device is not formatted as Ext4, enter the following command to format it with an Ext4 file system:

```
$ sudo mkfs.ext4 /dev/sd<port><device_number>
```

Where:

- `<port>` is the port to which your device is mounted.
- `<device_number>` is the device number of the device attached to the port. You can use the `dmesg` command to determine the port.

3. If needed, mount your device with the following command:

```
$ sudo mount /dev/sdX1 <mntpoint>
```

Where <mntpoint> is the mount point on the host system for your rootfs device.

4. Copy the file system. If `LDK_ROOTFS_DIR` is set, execute these commands:

```
$ cd ${LDK_ROOTFS_DIR}
$ sudo cp -a * <mntpoint> && sync
```

5. If it is not set, copy the rootfs directory that is included in the release by executing the following commands:

```
$ cd <your_L4T_root>/Linux_for_Tegra/rootfs
$ sudo cp -a * <mntpoint> && sync
```

6. After copying the content to the external disk or device, unmount the disk and connect it to the target Tegra device.

Proceed to [Flashing the Boot Loader and Kernel](#).

Flashing the Boot Loader and Kernel

This section describes the steps to flash and boot the target Tegra device. It also provides usage information for the `flash.sh` helper script.

Flash Procedure

First, flash the board with the boot loader and kernel, and optionally, flash the rootfs to internal eMMC.

Prerequisites

The following directories must be present:

- `bootloader`—boot loader plus flashing tools (NvFlash, CFG, [BCT](#), etc.)
- `kernel`—a kernel `zImage` / `vmlinux.uimg`, DTB files, and kernel modules
- `rootfs`—the root file system that you download (This directory starts empty and you populate it with the sample file system.)
- `nv_tegra`—NVIDIA® Tegra® user space binaries and sample applications

Additionally, before running the following commands, you must have the USB cable connected to the recovery port.

To flash the boot loader and kernel

1. Put the target into reset/recovery mode.
 - Power on the carrier board and hold the RECOVERY button.
 - Then press the RESET button.
2. Run the `flash.sh` script that is in the top level directory of this release. The script must be supplied with the target board (`jetson-tx2`) for the root file system:

```
sudo ./flash.sh <platform> <rootdev>
```

Where `<rootdev>` depends on where the root file system will be:

- If the root file system will be on the Jetson TX2 internal eMMC, execute the script as follows:

```
sudo ./flash.sh jetson-tx2 mmcblk0p1
```

- If the root file system will be on a USB disk, execute the script as follows:

```
sudo ./flash.sh jetson-tx2 sda1
```

- If a SATA device is connected, that device enumerates as `sda1`.

- If the root file system will be on an SD card, execute the script as follows:

```
sudo ./flash.sh jetson-tx2 mmcblk1p1
```

The above examples are for U-Boot. For Fastboot, add the following argument:

```
-L <PATH_TO_FASTBOOT_BIN_FILE>
```

For example:

```
sudo ./flash.sh -L bootloader/<platform|ver>/fastboot.bin <platform> <rootdev>
```

This loads the boot loader and kernel.

Flash Script Usage

Locate the most up-to-date usage information by running `flash.sh -h` (using the `flash.sh` script included in the release). The basic usage is as follows.

```
sudo ./flash.sh [options] <platform> <rootdev>
```

Specify the required parameters and one or more of the options in the following table.

Parameters	Description	
•<platform>	•Is the <platform> for your release.	
•<rootdev>	•Is one of following:	
•	•mmcblk0p1	•Specifies internal eMMC.
	•mmcblk1p1	•Specifies external SDCARD.
	•sda1	•Specifies external USB device (such as, USB memory stick or HDD).
	•eth0	•Specifies nfsroot via external USB Ethernet interface.
Options	Description	
•-h	•Specifies to print this usage information.	

•-b <bct_file>	•Specifies the NvFlash Boot Configuration Table (BCT) file.
•-c <cfg_file>	•Specifies the NvFlash configuration file.
•-d <dtb_file>	•Optionally specifies a device tree file to use instead of the default.
•-e <emmc_file>	•Specifies the eMMC size of the target device.
•-f <flashapp>	•Specifies the path to flash application: nvflash or tegra-rcm.
•-i	•Specifies to pass the user kernel command line to the kernel as-is.
•-k <partition id>	•Specifies the kernel partition ID to be updated (minimum = 5).
•-n <nfs args>	•Specifies the static NFS network assignments: <Client IP>:<Server IP>:<Gateway IP>:<Netmask>
•-o <odmdata>	•Specifies the ODM data value.
•-p	•Total eMMC HW boot partition size.
•-r	•Specifies to skip building and reuse existing <code>system.img</code> .
•-s <ubootscript>	•Specifies the boot script file for U-Boot.
•-C <cmdline>	•Specifies the kernel command line. Warning: Each option in this kernel command-line gets higher precedence over the same option from fastboot. In case of NFS booting, this script adds NFS booting related arguments if the -i option is omitted.
•-F <flasher>	•Specifies the flash server, such as <code>fastboot.bin</code> .
•-I <initrd>	•Specifies <code>initrd</code> file. Null <code>initrd</code> is the default.
•-K <kernel>	•Specifies the kernel image, such as <code>zImage</code> .
•-L <bootloader>	•Specifies the full path to the boot loader, such as <code>fastboot.bin</code> or <code>u-boot.bin</code> .
•-P <end_of_PPT_plus 1>	•Specifies the sum of the primary GPT start address, the size of PPT, plus 1.
•-R <rootfs_dir>	•Specifies the sample rootfs directory.
•-N <nfsroot>	•Specifies the nfsroot, for example: •<my IP addr>:/my/exported/nfs/rootfs
•-S <size>	•Specifies the rootfs size in bytes. This is valid only for internal rootdev. KiB, MiB, GiB style shorthand is allowed. For example, 1GiB signifies 1024 * 1024 * 1024 bytes.
•-T <ITS_file>	•ITS file name. Valid only for u-boot.

Increasing the Internal Memory Partition for the Root File System

The suggested rootfs partition size for the Jetson TX2 platform is 15 gigabytes (GB). It is specified by default in the `<target_board>.conf` file used by the `flash.sh` script.

Use the `-S <size-in-bytes>` argument to `flash.sh` to change the partition size.

To flash for a larger partition

- Execute the following command:

```
$ sudo ./flash.sh -S <size> <platform> <rootdev>
```

Where:

- `<size>` is the desired size for the partition, such as 8589934592 (or 8 GiB) for 8 GB, if you want to decrease the size of the partition.
- `<rootdev>` is the rootfs partition internal memory, for example `mmcblk0p1`.

Determining the Success of a Driver Update

After updating drivers on a target board, verify that the update completed successfully. You can determine the success or failure of a driver update by using the following commands.

To determine the success of a driver update

- Execute the following command on a booted target device:

```
$ shasum -c /etc/nv_tegra_release
```

If the driver update succeeded, the output displays the word *OK* after the file name. A typical success message looks like this:

```
/usr/lib/xorg/modules/drivers/nvidia_drv.so: OK
```

The driver update fails if the file is missing. A typical error message looks like this:

```
shasum: /usr/lib/xorg/modules/drivers/nvidia_drv.so: No such file or directory
/usr/lib/xorg/modules/drivers/nvidia_drv.so: FAILED open or read
```

The driver update also fails if the new file is not the same as the existing file, producing an error such as:

```
/usr/lib/xorg/modules/drivers/nvidia_drv.so: FAILED
```

Installing Additional Packages

L4T comes with additional NVIDIA packages, including packages for Ubuntu and Google Chrome.

Installing Additional NVIDIA Packages

Additional NVIDIA packages may be posted alongside the release. To make full use of the features in the release, install these additional packages.

Directly after the `apply_binaries` step in [Setting Up the Root File System](#), you can install the package into the configured rootfs.

Installing Additional Ubuntu Packages

Install additional packages from Ubuntu, using the provided sample file system.

L4T is tested with the provided sample file system Ubuntu packages. Periodic Ubuntu package updates from Canonical are not validated.

To receive notifications

1. Locate and edit the following file:

```
/etc/apt/sources.list
```

2. Add the following line:

```
deb http://ports.ubuntu.com/ubuntu-ports <OS version (target)>-updates main universe
```

For example, for a rootfs based on the Xenial Xerus distribution of Ubuntu, add the line:

```
deb http://ports.ubuntu.com/ubuntu-ports xenial-updates main universe
```

Prerequisite

You have attached an Ethernet cable to the device through either the Ethernet port (if available) or through the USB Ethernet adapter.

To install more packages

1. Boot the target device.
2. Verify your Ethernet connection.
3. Update the package list by executing:

```
$ sudo apt-get update
```

Ensure that you run `sudo apt-get update` and not `apt-get upgrade`, which upgrades already installed packages. Do not confuse the two commands.

4. Install packages using `apt-get`. For example, to install `wget` execute this command:

```
$ sudo apt-get install wget
```

Configuring NFS Root on the Linux Host

To boot the target device from NFS, you must provide an NFS root mount point on your Linux host machine. The general steps for configuring an NFS root on the Linux host are as follows.

Prerequisites

- An Ethernet connection to install packages on the host.
- An Ethernet connection on the target.

To configure NFS root on the Linux host

1. Install the `nfs` components on your host machine:

```
$ sudo apt-get install nfs-common nfs-kernel-server
```

2. The NFS server must know which directories you want to 'export' for clients. This information is specified in the `/etc/exports` file.

- Modify `/etc/exports` to look somewhat like this:

```
$ /nfsroot *(rw,nohide,insecure,no_subtree_check,async,no_root_squash)
```

- After adding the entry, restart using the following command:

```
$ sudo /etc/init.d/nfs-kernel-server restart
```

3. Create an `/nfsroot` directory on your Linux host machine:

```
$ sudo mkdir /nfsroot
```

4. Copy the file system to the `nfsroot` directory:

```
$ cd ./rootfs
```

```
$ sudo cp -a * /nfsroot
```

5. Export the root point:

```
$ sudo exportfs -a
```

Alternatively, you can export or un-export all directories by using the `-a` and `-u` flags. The following command un-exports all directories:

```
$ sudo exportfs -au
```

6. (Optional) If the Ubuntu firewall blocks NFS root access, it must be disabled depending on your configuration. You can do so with the following command:

```
$ sudo ufw disable
```

7. If there are issues performing the NFS boot, to separately verify everything on the 'host' machine is configured properly, you can perform the following step on a booted target board through USB/SD/internal eMMC. It should be possible to mount the host NFS root point on the target device:

```
$ mkdir rootfs
```

```
$ sudo mount -v -o nfsvers=3 <IP-ADDR>:/nfsroot rootfs
```

Where `<IP-ADDR>` is the IP address of the Linux Host machine as taken from the `ifconfig` command. This proves that the host configuration is correct.

Note: Prior to executing the mount command on the target machine, you must install the `nfs-common` package using the following command:

```
$ sudo apt-get install nfs-common
```

To boot the target with the NFS root point, see [Flashing the Boot Loader and Kernel](#). Be sure to include the `-N` option for the nfs root point.

Getting the Kernel Sources

You can manually rebuild the kernel used for this package. Internet access is required.

Sync with git

Prerequisites

- You have installed Git. Install Git with the following command:

```
$ sudo apt-get install git-core
```

- Your system has the default Git port 9418 open for outbound connections.

To sync the kernel sources

- Get the kernel source by running the `source_sync.sh` script:

```
$ sudo ./source_sync.sh
```

When prompted enter a 'tag' name, as provided in the release notes.

You can sync to any Linux tag you like. However, the tag provided in the release notes syncs the sources to the same source revision the release binary was built from. To see a list of the available release tags, use:

```
$ git tag -l tegra-l4t*
```

Manually Download and Expand

As an alternative, you can download the kernel source files and then manually expand them. If possible, it is recommended to instead sync with git.

To manually download and expand

- In a browser, navigate to:

```
https://developer.nvidia.com/embedded/downloads
```

- Locate and download the L4T Sources for your release.
- Expand the TBZ2 file.

```
$ sudo tar -vxjf sources.tbz2
```

- Manually expand the TBZ2 files in the resulting `source` directory.

Building the NVIDIA Kernel

Follow the steps in this procedure to build the NVIDIA kernel.

Prerequisites

- You have downloaded the kernel source code

To build the Tegra Kernel

- Export the following environment variables:

```
$ export CROSS_COMPILE=<crossbin>
$ export TEGRA_KERNEL_OUT=<outdir>
$ export ARCH=arm64
```

Where:

- <crossbin> is the prefix applied to form the path to the toolchain for cross compilation targeting arm64, e.g., gcc.

The reference arm64 toolchain is:

```
<aarch64_toolchain_install>/bin/aarch64-unknown-linux-gnu-
```

This example requires GCC 4.8 or above. See [Building the AARCH 64 Toolchain and glibc](#) for information on how to obtain the reference toolchains.

- <outdir> is the desired destination for the compiled kernel.

- Execute the following commands to create the .config:

```
$ cd <myworkspace>/<kernel_source>
$ mkdir $TEGRA_KERNEL_OUT
```

Where <kernel_source> directory contains the kernel sources.

```
$ make O=$TEGRA_KERNEL_OUT tegra18_defconfig
```

Where <myworkspace> is the parent of the Git root.

- Execute the following commands to build the kernel:

```
$ make O=$TEGRA_KERNEL_OUT zImage
```

- Execute the following command to create the kernel device tree components:

```
$ make O=$TEGRA_KERNEL_OUT dtbs
```

- Execute the following commands to build the kernel modules (and optionally install them)

```
$ make O=$TEGRA_KERNEL_OUT modules
$ make O=$TEGRA_KERNEL_OUT modules_install INSTALL_MOD_PATH=<your_destination>
```

- Copy both the uncompressed (Image) and compressed (zImage) kernel images over the ones present in the following directory:

```
Linux_for_Tegra/kernel
```

- Replace the contents of Linux_for_Tegra/kernel/dtb/ with the contents of:


```
$TEGRA__KERNEL_OUT/arch/arm64/boot/dts/
```

- Archive the kernel modules created using the following command:

```
cd <modules_install_path>
tar --owner root --group root -cjf kernel_supplements.tbz2 lib/modules
```

Use the archive to replace the one in the kernel directory of the release:

```
Linux_for_Tegra/kernel/kernel_supplements.tbz2
```

Building External Kernel Modules

The procedures in this section describe how to build an out-of-tree kernel module against kernel headers included in the BSP.

The kernel headers are installed at: `/usr/src/linux-headers-<kernel version>`

To compile natively on the target system

- Determine user space architecture with the following command:

```
getconf LONG_BIT
```

The command returns 64 for AARCH64.

- In AARCH64 user space, prepare the kernel headers with the following commands:

```
$ cd /usr/src/linux-headers-`uname -r`
$ sudo make modules_prepare
```

- Build the kernel module.

To cross-compile on another system

- Uncompress the following kernel headers to a local directory:

```
<top>/Linux_for_Tegra/kernel/kernel_headers.tbz2
```

with the following commands:

```
cd <local src dir>
tar xpf <top>/Linux_for_Tegra/kernel/kernel_headers.tbz2
```

- Set up the cross-compile toolchain with the following commands:

```
export CROSS_COMPILE=<crossbin>
export ARCH=arm64
```

Where `<crossbin>` completes path to the tool chain for cross compilation targeting arm64, e.g., gcc.

For this reference toolchain:

```
<aarch64_toolchain_install>/bin/aarch64-unknown-linux-gnu-
```

For a Linaro toolchain, the path is similar to the following:

```
<linaro_install>/aarch64-unknown-linux-gnu/bin/aarch64-unknown-linux-gnu-
```

See [Building the AARCH64 Toolchains](#) for downloading and building the reference toolchain.

3. Prepare the kernel headers with the following commands:

```
cd <local src dir>/linux-headers-<kernel version>
export ARCH=arm64
sudo make modules_prepare
```

4. Specify `<local_source_directory>/linux-headers-<kernel version>` as the kernel source path, with the following line in the make file:

```
make -C <local src dir>/linux-headers-<kernel version> M=$(PWD) modules
```

5. Build the kernel module.

OpenGL/EGL Gears Test Application

To run a sample OpenGL/EGL test application, you can run the open-source Gears application.

To install and run Gears test application

1. Boot the target system with an Ethernet connection.
2. Enable package download from the “universe” repository by editing `/etc/apt/sources.list` as root:

```
$ sudo vi /etc/apt/sources.list
```

3. Uncomment the following line in the file by removing the leading # character:

```
# deb http://ports.ubuntu.com/ubuntu-ports/ xenial universe
```

4. Update the repository:

```
$ sudo apt-get update
```

5. Install the `mesa-utils` and `mesa-utils-extra` packages:

```
$ sudo apt-get install -y mesa-utils
$ sudo apt-get install -y mesa-utils-extra
```

6. At this point you should be able to run the application with the following steps:

```
$ export DISPLAY=:0
$ X&
$ /usr/bin/es2gears
```

GStreamer-based Multimedia Playback (NvGstPlayer)

Use the GStreamer open source multimedia framework and the NvGstPlayer utility for testing multimedia local playback and HTTP/RTSP streaming playback use cases. The NvGstPlayer can be used as a reference implementation.

For more information about the NvGstPlayer application, refer to the ReadMe file included with the release.

Installing GStreamer

Install GStreamer from the Internet directly on the target. The wrapper library, `gst-openmax`, is an interface between GStreamer and OpenMAX. It enables accelerated NVIDIA plug-ins in the GStreamer framework.

For more information about GStreamer, see the following website:

<http://gstreamer.freedesktop.org>

NvGstPlayer is a multimedia player test application.

Complete prerequisite steps in the file `nvgstcapture_README.txt` before running the NvGstPlayer and NvGstCapture applications.

Instructions for installing GStreamer are also included in that text file.

Using NvGstPlayer

NvGstPlayer is a command line media file player. It plays audio/video files encapsulated in MP4, 3GP, AVI, ASF, WMA, MKV, M2TS, WEBM, and MOV. NvGstPlayer supports local file playback and playback over RSTP, HTTP, and UDP.

For information about NvGstPlayer runtime commands, default settings, and important notes see the `nvgstplayer_README.txt` file included in the release.

Gstreamer-based Camera Capture (NvGstCapture)

The NvGstCapture application supports GStreamer version 0.10.36 by default. NvGstCapture captures audio and video data using a microphone and camera and encapsulate encoded A/V data in the container file.

For NvGstCapture installation and usage information, see the `nvgstcapture-<VERSION>_README.txt` file included with the release at `~Linux_for_Tegra/nv_tegra/nv_sample_apps`.

Determining Version and Platform Information

Use the procedures in this section to determine the flashed BSP version, the kernel version, and other platform information.

To determine the BSP version and other platform information

- Determine the BSP version and other platform information with the following command:

```
head -1 /etc/nv_tegra_release
```

Output from that command is similar to the following:

```
#R27 (release), REVISION: 1.0, GCID: 7164062, BOARD: t186ref, EABI: aarch64, DATE:
```

To determine the kernel version

- Determine the kernel version, with the following command in the kernel directory:

```
head -4 Makefile
```

Output is similar to the following:

```
VERSION = 4
PATCHLEVEL = 4
SUBLEVEL = 15
```

Or, if the system is running, determine the kernel version with the following command:

```
uname -a
```

Output is similar to the following:

```
Linux tegra-ubuntu 4.4.15-tegra #1 SMP PREEMPT Tue Mar 03 16:29:05 PDT 2017 aarc
```

To determine boot configuration

Determine the boot configuration from the following information.

- Boot option: `l4t/config/t186ref/p2771-0000/extlinux.conf.emmc`
 - `'FDT /boot/tegra186-quill-p3310-1000-a00-00-base.dtb'` indicates the board device tree blob file. The source is located at:

```
arch/arm64/boot/dts/tegra186-quill-p3310-1000-a00.dts.
```

- Default u-boot build configuration:
 - `3rdparty/u-boot/configs/p2771-0000-000_defconfig`
 - and/or
 - `3rdparty/u-boot/configs/p2771-0000-500_defconfig`
- Default kernel build configuration: `arch/arm64/configs/tegra18_defconfig`

NVIDIA Bug Reporting Script

For debugging purposes, attach the log file to communicate issues found with the release. Use the `nvidia-bug-report-tegra.sh` script to generate log files.

To generate a log file for bug reporting

- Log into the target board and enter the following command:

```
$ sudo /usr/bin/nvidia-bug-report-tegra.sh
```

To generate a log file for bug reporting with extended logging mode

- Log into the target board and enter the following command:

```
$ sudo /usr/bin/nvidia-bug-report-tegra.sh -e
```

By default, the logfile generated by both these procedures is located at:

```
$HOME/nvidia-bug-report-tegra.log
```

Attach a log file when reporting any issues to NVIDIA, whether through email or the forums.

Software Features

NVIDIA® Tegra® Linux Driver Package (L4T) supports the following software features, which provide users a complete package to bring up Linux on targeted NVIDIA® Tegra® X2 devices.

This release supports the NVIDIA® Jetson™ TX2 developer kit and module.

Note:

- Check the *Release Notes* for constraints related to these features.

Boot Loaders

Boot Loader	Feature	Notes
•nvtboot-bpmp	•Execution CPU	•BPMP
	•Storage location	•Cold boot: eMMC
		•RCM boot: Downloaded over USB recovery port
	•Next stage storage location	•Cold boot: eMMC
		•RCM boot: Downloaded over USB recovery port
	•Next stage	•cboot
	•Storage device support	•eMMC
	•Partition table support	•GPT (with protective MBR)
	•Filesystem support	•None
	•I/O bus support	•I2C
•Console UART	•	
•cboot	•Execution CPU	•CCPLEX
	•Storage location	•Cold boot: eMMC
		•RCM boot: Downloaded over USB recovery port
	•Next stage storage location	•Cold boot: eMMC
•RCM boot: Downloaded over USB recovery port		
•Next stage	•U-boot or Linux kernel	

	•Storage device support	•eMMC
	•Partition table support	•GPT (with protective MBR)
	•Filesystem support	•None
	•I/O bus support	•I2C
	•Console	•UART
•U-Boot	•Execution CPU	•CCPLEX
	•Storage location	•Cold boot: eMMC
	•Next stage storage location	•Cold boot: eMMC
	•Next stage	•Linux kernel
	•Storage device support	•eMMC, SD card
	•Partition table support	•GPT (with protective MBR), DOS MBR
	•Filesystem support	•ext2/3/4. FAT
	•I/O bus support	•I2C, PCIe

Toolchain

Feature	Tool Chains	Notes
•Aarch64	•gcc-4.8.2-glibc-2.17	•For 64-bit Kernel, Userspace, and U-Boot •For more information see Building the Kernel .

Kernel

Interface	Feature	Notes
•DSI	•DSI Display Support	•DSI0, DSI1
	•DSI Ganged Mode	•DSI1
	•PWM Backlight	•DSI0, DSI1
	•DC Continuous Mode	•DSI0
	•DC Driven Command Mode	•-
	•Host Write	•-
	•DSI One-Shot Mode	•-
	•Dual Display	•-
	•Run Time Power Management	•DSI0, DSI1

•HDMI	•EDID Support	•Yes
	•Hot-Plug Detection Mechanism	•Yes
	•HDMI 1.4	•480p, 720p, 1080p, RGB 444 4K @ 30 Hz
	•Driver Suspend/Resume for Low Power	•Yes
	•HDMI as Primary Display	•Yes
	•Multi Display	•Yes for Mirror/Extended Mode
	•HDMI: 1.4b compliance	•Yes
	•HDMI: 2.0 compliance	•Yes
	•Audio Support	•Yes
	•HDMI 2.0 support	•4K @ 60 HZ
•Audio	•ADSP Audio	•
•Display	•DP2 support	•Yes
	•Driver Suspend/Resume for low power	•Yes
	•Support eDP as Primary Display	•
	•3-head Display	•
•Display DPO	•DSI Host Read	•DSI0
	•DSI Host Read	•DSI1
	•DSI Secondary Data during v_blank	•DSI0
	•DSI Secondary Data during v_blank	•DSI1
	•DC Driven Command Mode	•DSI0
	•DC Driven Command Mode	•DSI1
	•Host Write	•DSI0
	•Host Write	•DSI1
	•DSI One-Shot Mode	•DSI0
	•DSI One-Shot Mode	•DSI1
	•DP Support	•
	•Audio Support	•
•PWM	•PWM Operations	•PWM registration to framework
	•Prod Setting	•Tegra specific controller configuration
	•Clock accuracy calculation	•Clock calculation
•I2C	•DMA Mode	•I2C Gen1
	•Bus Clear Support	•I2C Gen1

	•Multi Master Support	•I2C Gen1
	•Normal/Byte Mode	•N/A
	•General Support	•I2S0
•JTAG	•JTAG Attach	•Debugging capability
	•JTAG Halt/Step/Go	•Debugging capability
•PCIe	•Physical Port: PCI-E 0	•General Support
	•Physical Port: CPI-E 1	•General Support
	•Physical Port: CPI-E 2	•General Support
	•Host Controller Features	•Lanes Xbar config (X4_X0_X1, X2_X1_X1, X1_X1_X1) •Hot-plug (using GPIO)
	•PCI Features	•Message Signalled Interrupts
	•PCIe Link Capabilities	•ASPM Support (L0s and L1) •L1 Clock Power Management •ASPM Support (L1.1 and L1.2)
	•Root Control	•PME Interrupt Enable
	•Extended Capabilities	•Advanced Error Reporting (AER)
	•Miscellaneous Features	•Dynamic Voltage Frequency (DVFS) •Tegra Low Power Mode (LP0) •Runtime PM
	•L1 PM Substates	•Rest All Capabilities
•Bluetooth	•Bluetooth 4.0	•BCM4354
	•BLE 4.0	•No BCM4354 (BlueZ limitation)
•Camera support •(CSI input support)	•V4L2 Media-Controller •(V4L2 API bypasses ISP)	•CSI0, CSI1, CSI2, CSI3, CSI4, CSI5 •Note: The media-controller driver model is adopted in the 24.1 release. the Soc_camera driver is provided, but deprecated. doublecheck with Sean or Shantanu
•Peripheral devices	•INA support	•Current monitoring for: CPU/GPU/VDD_IN
•Platform support	•P3310-B00 C03	•
•WiFi	•Multi-Region support	•Region Support: • US - 0x010000 • Tiawan - 0x020000 • Europe - 0x040000 • Japan - 0x080000 • Korea - 0x100000

	<ul style="list-style-type: none"> • Canada - 0x200000 • Israel - 0x400000 • default (lowest-common-denominator)
•Dual-band 2.4 GHz/5 GHz	•BCM 4354
•STA Mode	•BCM 4354
•HostAP Mode	•BCM 4354
•P2P Mode	•BCM 4354 (but WEP-only security)
•WPA2 Security	•BCM 4354

I/O

I/O Type	Feature	Notes
•SPI	•Physical Port: SPI1	•Maximum bus speed: 65 MHz
	•Physical Port: SPI2	•Maximum bus speed: 65 MHz
	•Physical Port: SPI 0/3	•Maximum bus speed: N/A
	•Packed/Unpacked	•SPI4
	•Full Duplex Mode	•SPI4
	•Both Enable Bit	•SPI4
	•Both Enable Byte	•SPI4
	•Bi-directional	•SPI4
	•Least Significant Bit	•SPI4
	•Least Significant Byte First	•SPI4
	•Software or Hardware Chip Select Polarity Section	•SPI4
	•Supported Modes 1/2/3/4	•SPI4
	•Dual SPI	•SPI MISO/MOSI can act as Rx and Tx
	•Multiple transfer request	•Multiple SPI transfer request from single call
•SDMMC	•I/O Speeds (Clock speed)	•SDMMC1 (SD card): 204 MHz
		•
		•

	•Hot Plug Support	•SDMMC1 (SD card)
	•SD High Speed Mode	•SDMMC1 (SD card)
	•SDR50	•SDMMC1 (SD card)
	•SDR104	•SDMMC1 (SD card)
	•HS533	•SDMMC4 (eMMC)
	•HS400	•-
	•HS200	•-
	•DDR Mode	•SDMMC1 (SD card)
	•Voltage Switching	•SDMMC1 (SD card)
	•Frequency Tuning	•SDMMC1 (SD card)
	•Packed Commands	•SDMMC4 (eMMC)
	•Cache Control	•SDMMC4 (eMMC)
	•Discard	•SDMMC4 (eMMC)
	•Sanitize	•SDMMC4 (eMMC)
	•RPMB	•SDMMC4 (eMMC)
	•HPI	•SDMMC4 (eMMC)
	•BKOPS	•SDMMC4 (eMMC)
	•Power Off Notification	•SDMMC4 (eMMC)
	•Sleep	•SDMMC4 (eMMC)
	•Field Firmware Upgrade	•SDMMC4 (eMMC)
	•	•
	•Device Life Estimation Type A	•SDMMC4 (eMMC)
	•Device Life Estimation Type B	•SDMMC4 (eMMC)
	•PRE EOL Information	•SDMMC4 (eMMC)
	•Power Management	•SDMMC4 (eMMC)
•SATA	•Speed	•GEN1 •GEN2
	•AHCI Mode	•1.3.1
	•SATA Specification	•3.1
	•HIPM	•Yes
	•	•
	•NCQ	•Yes
	•Port Multiplier Support	•CBS
	•Link Power Management States	•Partial •Slumber

	•Device Power Management States	•D0
		•D1
		•D2
	•Runtime Time Power Management	•Yes
	•S.M.A.R.T	•Self-Monitoring Analysis and Reporting Technology
	•Dev sleep support	•-
•USB 3.0	•Speeds	•USB0: HS/480 Mbps
	•Lanes	•USB 0: N/A
	•USB 3.0 Support	•USB0: Yes
	•Connector	•USB0: Micro AB
	•USB 2.0 Support	•USB 0: Yes
	•Remote Wakeup Support	•USB0: USB 2.0
	•Host - Auto Suspend Support	•USB 0: Yes
	•XOTG Support	•USB 0: Yes
	•XUSB SS/HS/FS/LS Host Mode	•USB 0: Yes
	•XUSB SS/HS/FS/LS Device Mode	•USB 0: Yes
	•XUSB Device Port U1/U2/U3 Transition	•USB 0: Yes
	•XUSB Host Port U1/U2/U3 Transition	•USB 0: Yes
	•XUSB Device ELPG	•USB 0: Yes
	•XUSB Host ELPG	•USB 0: Yes
	•Class Support	•Mass storage USB 0: Yes
		•USB video class USB 0: Yes
		•HID (USB 0: Yes
		•USB video class (USB 0: Yes
		•MTP USB 0: Yes
		•Ethernet USB 0: Yes
		•Thumb/Hard Drive USB 0: Yes
		•Mouse USB 0: Yes
		•CDC - NCM/ECM USB 0: Yes
•EQOS	•Ping	•
	•Speed	•

	•LP_IDDQ Mode Support	•
	•Suspend Resume over NFS Support	•
	•NFS Boot	•
•RTC	•Alarm	•
	•Wakeup from SC7	•
•Watchdog	•Tegra Watchdog	•Watchdog reboot from hang
	•Tegra Watchdog	•Watchdog kick
	•PMIC Watchdog	•Watchdog reboot from hang
	•PMIC Watchdog	•Watchdog kick
•GPIO	•System Programable GPIO Support	•-
	•System Programable Pinmux SupportWakeable GPIO	•-
	•Timestamping GPIO	•-
•UART	•Speed	•UART0 (Debug): 115200
	•Hardware Flow Control for Debug	•UART0: No
	•PIO Mode	•UART0: Yes
	•DMA Mode	•UART0: Yes
	•FIFO Mode	•UART0: Yes
•System	•UCM1 4/4/16	•
	•UCM2 24x7	•
	•Reboot Support	•
	•Shutdown Support	•
	•SC7	•
	•Wake from Idle	•
	•Wake from Sleep	•
	•cpuidle	•
	•cpufreq	•
	•DVFS	•
	•Dual USB 3.0 A Port	•
	•Dual PCIE Port	•
	•CPU Hotplug	•
	•EMC Scaling	•
	•initrd Support	•
	•CPU Load Behavior	•

	•System Boot with ATF as Secure Monitor	•
	•Secure boot with ATF	•
•GPU	•GPU @ POR Frequency	•

Note:

•PCIe: Tegra TX1 does not have any path from AHB-DMA or APB-DMA engines to PCIe IP as PCIe is connected directly to MSELECT and AHB-DMA and APB DMA engines only interact with IPs connected to respective AHB and APB buses. So it is not possible to use either AHB or APB engines for PCIe.

CUDA

Feature	Version
CUDA	•Version 8.0.64

Graphics

Graphics APIs	Notes
•OpenGL	•4.5
•OpenGL-ES	•3.2
•Vulkan	•1.0.1
•EGL	•1.4
•GLX	•
•GLVnd Version of EGL	•Vendor neutral functionality
•NVDC - Direct Rendering Manager (DRM)	•Compatibility with DRM 2.0
•EGL Stream	•
•X11 ABI-20	•Legacy from 24.2 using Ubuntu 16.04
API Support	Notes
•GL + EGL	•
•EGL without X11	•Content display without X11 usage

EGL and OpenGL ES Support

EGL is an interface between Khronos rendering APIs, such as OpenGL ES, and the underlying native platform window system. It handles graphics context management, surface/buffer binding, and rendering synchronization. EGL enables high-performance, accelerated, mixed-mode 2D and 3D rendering using other Khronos APIs.

L4T supports the EGL 1.4 specification, [Khronos Native Platform Graphics Interface \(EGL 1.4 Specification\)](#).

The OpenGL ES driver in this release supports the following OpenGL ES specifications:

- [OpenGL ES Common Profile Specification 2.0](#)
- OpenGL 4.5

For more information on OpenGL ES, see the [Khronos OpenGL ES API Registry](#).

Video Decoders

Video Decode	Output Formats	Sampling Frequency and Bit rate/Frame rate	Notes
•H.264	•NV12, NVMM:NV12	•3840 x 2160 at 60 fps Up to 120 Mbps	•Full-frame, Disable-DPB, Skip-Frames
•H.265	•NV12, NVMM:NV12	•3840 x 2160 at •60 fps •Up to 160 Mbps	•Decode Support in Gstreamer 1.4.5 and later
•JPEG	•I420, NVMM:I420	•600 MP/sec	•-
•VP8	NV12, NVMM:NV12	•3840 x 2160 at •60 fps •Up to 140 Mbps	•-
•VP9	NV12, NVMM:NV12	•3840 x 2160 at •60 fps •Up to 120 Mbps	•-

Video Encoders

Video Encode	Input Formats	Sampling Frequency and Bit rate/Frame rate	Notes
•H.264	•I420, NV12, NVMM:I420, NVMM:NV12	•3840 x 2160 at •30 fps •Up to 120 Mbps	•Supported features: control-rate, Bitrate, Iframeinterval, Quality-Level, Low-Latency, SliceIntrarefreshEnable, Sliceintrarefreshinterval, Bit-Packetization, VBV-Size, temporal-tradeoff, Insert-SPS-PPS, Slice-Header-Spacing, Profile, num-B-Frames, Force-IDR
•JPEG	•I420, NVMM:I420	•600 MP/sec	•-
•H.265	•I420, NVMM:I420, NVMM:NV12	•3840 x 2160 at •30 fps	•Supported features: control-rate, Bitrate,

Video Encode	Input Formats	Sampling Frequency and Bit rate/Frame rate	Notes
		•Up to 100 Mbps	Iframeinterval, Quality-Level, SliceIntrarefreshEnable, Sliceintrarefreshinterval, Bit-Packetization, VBV-Size, temporal-tradeoff, Insert-SPS-PPS, Force-IDR
•VP8	•I420, NV12, NVMM:I420, NVMM:NV12	•3840 x 2160 at •30 fps •Up to 120 Mbps	•Supported features: control-rate, Bitrate, Iframeinterval, Quality-Level, Force-IDR

Note: Use the `gst-inspect-1.0` utility to understand feature details. For example, the `gst-inspect-1.0 omxh264enc` command provides feature details of the H.264 encoder.

Display Outputs

nveglglessink	nvoverlaysink
•X11 Window	•Panel Overlay
•-	•Overlay
•-	•Overlay-Depth
•-	•Overlay-X
•-	•Overlay-Y
•-	•Overlay-W
•-	•Overlay-H

Conversion, Scaling, and Rotation Formats

•Input Formats	•Output Formats	•Notes
•I420	•I420	•Flip-Method
•UYVY	•UYVY	•Flip-Method
•NV12	•NV12	•Flip-Method
•GRAY8	•GRAY8	•Flip-Method
•NVMM:I420	•NVMM:I420	•Flip-Method
•NVMM:NV12	•NVMM:NV12	•Flip-Method

CSI and USB Camera Formats

Output Format	Options	Notes
•Gst-nvcamerasrc	•NVMM: I420, NV12	•
	•Whitebalance Mode	•
	•Color effect	•
	•Auto-exposure	•
	•Flicker	•
	•Contrast	•
	•Saturation	•
	•TNR strength	•
	•TNR Mode	•
	•Edge Enhancement	•
	•Intent	•Still/Video/Video snapshot/Preview
	•sensor-id	•
	•aeRegion	•
	•wbRegion	•
	•fpsRange	•
	•exposure-time	•
	•wbManualMode	•
	•wbGains	•
	•Sensor Orientation	•
	•Embedded Metadata	•
•EGL Producer	•	
•Simultaneous multi-camera	•	
•nveglstreamsrc	•NVMM: I420, NV12	•Input formats
•nvvideosink	•NVMM: I420, NV12	•Output formats
•nvegltransform	•NVMM: I420, NV12	•Input formats
		•Output formats
•videocuda	•EGL Image	•Input formats
		•Output formats
•nveglglessink	•Output: X11 window	•
•nvoverlaysink	•Output: Panel Overlay	•Overlay index
		•Overlay-depth

		<ul style="list-style-type: none"> •Overlay-x •Overlay-y •Overlay-w •Overlay-h
•nvhdmioverlaysink	•Output: HDMI Overlay	<ul style="list-style-type: none"> •Overlay Index •Overlay Depth •Overlay-x •Overlay-y •Overlay-w •Overlay-h
•nvvidconv	•Input Formats	<ul style="list-style-type: none"> •I420 •UYVY •NV12 •GRAY8 •NVMM: I420 •NVMM: NV12
	•Output Formats	<ul style="list-style-type: none"> •I420 •UYVY •NV12 •GRAY8 •BGRx, RGBA •NVMM: I420 •NVMM: NV12 •NVMM:RGBA
•nvvidconv	•Video Flip Methods	<ul style="list-style-type: none"> •Counterclockwise •Rotate •Counter-clockwise •90 degrees
		<ul style="list-style-type: none"> •Rotate-180 •Rotate 180 degrees
		<ul style="list-style-type: none"> •Clockwise •Rotate clockwise •90 degrees
		<ul style="list-style-type: none"> •Horizontal-flip •Flip horizontally
		<ul style="list-style-type: none"> •Upper-right-diagonal - Flip across
	•Interpolation Methods	<ul style="list-style-type: none"> •Nearest •Linear •Smart •Bilinear
•Video Crop	•Crop coordinates	
•Temporal Noise Reduction	•-	

U-Boot Customization

U-Boot is the default boot loader for NVIDIA[®] Tegra[®] Linux Driver Package (L4T).

Requirements

The software requirements and prerequisites include:

- Linux-based Host System

Functionality of the U-Boot build and flashing utilities are validated using an Ubuntu `<os_ver_host>` host system. Later versions of Ubuntu, or alternative Linux distributions, may work with host-specific modifications.

- Tegra Linux Driver Package (L4T)

Download and install the latest L4T package from the Tegra Developer Zone:

<http://developer.nvidia.com/linux-tegra>

- Device Tree Compiler (DTC)

The DTC is used to compile device tree files contained in the U-Boot source tree. Many of the DTC packages available from standard Linux distribution package management systems, like APT, are not updated with a version of DTC that supports the features required by the U-Boot makefile.

For a detailed example of building DTC from source see [Building Device Tree Compiler](#).

A pre-built DTC binary is included in the kernel directory of the release. This binary is built from the kernel sources in this release. The sources are located in the `scripts/dtc` directory. You build DTC by building the kernel `dtbs` target.

- Download the Linaro Aarch64 toolchain available at:

<https://www.linaro.org/downloads/>

- U-Boot source

For information, see [Downloading and Building U-Boot](#).

- Kernel source

For information, see:

- [Setting up the Root File System](#)
- [Downloading the Kernel Source](#)
- [Building the NVIDIA Kernel](#)
- [Adding a Compiled Kernel to the Root File System](#)

Downloading and Building U-Boot

Before flashing U-Boot to your reference platform, you must download and build it on your Linux host system. NVIDIA offers a Git repository containing the source code for a version of U-Boot suitable for L4T.

Prerequisite

Before over-writing the U-Boot binaries with the extracted L4T release package, backup all of the original U-Boot files in:

```
Linux_for_Tegra/bootloader/<platform>/<board_and_rev>/
```

Where:

- `<platform>` is the SOC name or platform.
- `<board_and_rev>` is:
 - Jetson TX1: p2371-2180
 - Jetson TX2: p2771-0000/500

To download and build U-Boot

1. Perform either of the following options:

- Execute the following command:

```
$ git clone -n git://nv-tegra.nvidia.com/3rdparty/u-boot.git
```

- Use the `source_sync.sh` script in the L4T release.

If you run `source_sync.sh -u` without parameters, the script prompts for the `<tag_name>`, which is the release tag name as provided in the *Release Notes*.

The `-k` option to `source_sync` syncs the kernel sources. A space between the `-u` and `-k` options is allowed. By default, if no option is provided, the script syncs the kernel and u-boot sources.

Also, you can run the script by passing the `<tag_name>` as follows:

```
$ cd <your_L4T_root>/Linux_for_Tegra
$ ./source_sync.sh -u <tag_name>
```

This syncs the source to:

```
<source_sync.sh_location>/sources/u-boot_source
```

The `<uboot_src_dir>` directory becomes:

```
<your_L4T_root>/Linux_for_Tegra/sources/u-boot_source
```

Note:

The rest of these instructions assume your current working directory is the U-Boot source tree.

2. Checkout the Git tag name:

```
$ git checkout -b <branch_name> <tag_name>
```

Where:

- <branch_name> is the name of your local branch.
- <tag_name> is the release tag name as provided in the *Release Notes*.

3. Set the build environment:

```
$ export CROSS_COMPILE=<your_toolchain_location>
```

For example, export CROSS_COMPILE=aarch64-linux-gnu-

4. Build U-Boot by executing:

```
$ make distclean
$ make <board_and_rev>_defconfig
$ make
```

Where: <board_and_rev> is:

- Jetson TX1: p2371-2180
- Jetson TX2: p2771-0000-500

Before Flashing

Before flashing, you must copy the updated U-Boot to the L4T tree.

To copy the updated U-Boot for flashing, to the reference platform

- Execute the following command on your Linux host system:

```
$ cp <uboot_src_dir>/u-boot{,.bin,.dtb,-dtb.bin} \
<your_L4T_root>/Linux_for_Tegra/bootloader/<platform>/<board_and_rev>
```

Where <board_and_rev> is:

- Jetson TX1: p2371-2180
- Jetson TX2: p2771-0000/500

Flashing U-Boot Full L4T Image

You must flash U-Boot to internal eMMC only. At boot time, U-Boot fetches the boot configuration file, kernel and device tree, which can reside on one of the following storage devices used for boot:

- Internal eMMC
- An SD card

When executing the script that flashes U-Boot, specify a command-line option to identify the storage device containing the root filesystem. This is required so that the appropriate boot configuration file is selected. The boot

configuration file contains kernel command line parameters that control where the Linux kernel looks for the root filesystem.

To flash U-Boot and select the script to mount the root filesystem from internal eMMC

- Select a boot configuration file that causes the kernel to mount the root filesystem from internal eMMC:

```
$ sudo ./flash.sh <platform> mmcblk0p1
```

To flash U-Boot and set the script to mount the root filesystem from an SD card

- Select a boot configuration file that causes the kernel to mount the root filesystem from an SD card:

```
$ sudo ./flash.sh <platform> mmcblk1p1
```

To flash U-Boot and select the script to mount the root filesystem from an NFS server

- Select a boot configuration file that causes the kernel to mount the root filesystem from an NFS server:

```
$ sudo ./flash.sh -N <IPA>:/<nfs directory> [-n <target IPA>:<host IPA>:<gateway IPA>:<
```

Where:

- <interface name> is eth0 for RJ45 connector and eth1 for a USB Ethernet dongle.
- <IPA> is the NFS server hosting the root filesystem.
- <nfs_directory> is the full pathname of exported root filesystem.
- <target IPA> is the static IP address for the target device.
- <host IPA> is the static IP address for the NFS server.
- <gateway IPA> is the static IP address for the gateway.
- <netmask> is the static netmask for the local network.

Note:

The -n option is recommended on point-to-point network connections where no DHCP server is configured.

Flashing U-Boot Only

To flash the full L4T image to the reference platform see [Flashing U-Boot Full L4T Image](#) above. To flash a new version of U-Boot, perform the following:

To flash the new U-Boot

- Execute the following command on Jetson TX1:

```
$ sudo ./flash.sh -k EBT <platform> mmcblk0p1
```

- Execute the following command on Jetson TX2:

```
$ sudo ./flash.sh -k kernel <platform> mmcblk0p1
```

Changing the eMMC Partition Layout

The following information is based on eMMC hardware and software layout information in the following files:

```
Linux_for_Tegra/<target_board>.conf
```

```
Linux_for_Tegra/bootloader/<platform>/cfg/gnu_linux_tegraboot_emmc_full.xml
```

```
Linux_for_Tegra/bootloader/<platform>/cfg/flash_l4t_t186.xml
```

Note:

The kernel is installed into the filesystem alongside the boot configuration file.

- On T210 systems, the kernel (LNX) partition is not used.
- On T186 systems, the kernel partition contains U-Boot.

Aside from this difference, U-Boot has the same internal eMMC partition layout as that used by cboot.

eMMC IC Parameter

The eMMC IC parameter is defined by two variables in the `Linux_for_Tegra/<target_board>.conf` file to:

- Limit the size of the total usable data area
- Determine the location of GPT partitions

The eMMC ID parameter includes:

- `BOOTPARTSIZE`: specifies the eMMC boot partition size which consists of boot0 partition size + boot1 partition size.
- `EMMCSIZE`: specifies the eMMC usable data size which consists of `BOOTPARTSIZE` + user partition size.

Note:

`boot0`, `boot1`, and user partition size can be obtained from the eMMC device datasheet.

Root Filesystem Size

The root filesystem partition is the largest of the partitions. The size of this partition is one of the key factors in partition layout determination.

To modify the root filesystem partition size

- Modify the value of the `ROOTFSSIZE` variable in the `Linux_for_Tegra/<target_board>.conf` file.

Note:

The total space used by all partitions cannot exceed `EMMCSIZE`.

GPT Partitions

The `flash.sh` script creates the primary and secondary GPT partitions automatically, based on the internal eMMC partition layout.

- The protective MBR contains device information to prevent traditional boot loaders from performing destructive actions. It is located at LBA 0.
- The primary GPT partition contains the GUID Partition Table. It is located at LBA 1.
- The secondary GPT partition contains the same information as the primary GPT and serves as the backup. It is located at the last LBA of the boot device.
- The last Logical Block Address (LBA) varies from device to device. Both U-Boot and the kernel are able to obtain the last LBA.

LNx Partition

The LNx partition nominally contains the kernel.

- On T210 systems, when using U-Boot, this partition is not used.
- On T186 systems, when using U-Boot, this partition is used to contain U-Boot.

APP Partition

If root filesystem storage is in eMMC, the root filesystem is flashed to this partition. U-Boot expects the boot configuration file and kernel files to exist in the `<rootfs>/boot` directory. The `flash.sh` flashes the following kernel files in the APP partition:

- kernel (Image)
- boot configuration file (`extlinux/extlinux.conf`)

Note:

The `flash.sh` script treats the root filesystem-on-IP-network configuration as a special case. It flashes these kernel files in the `<APP partition>:/boot` directory.

Full Internal eMMC Partition Layout

For an example of the configuration file, see the following:

- For TX1 systems: `gnu_linux_tegraboot_emmc_full.xml` file
- For TX2 systems: `flash_14t_t186.xml` file

Building the Device Tree Compiler

The Device Tree Compiler (DTC) must be built from source code.

Note:

If you do not wish to pass in the DTC filename as a parameter to the U-Boot environment, ensure a local command path that contains DTC, such as `/usr/local/bin` or another choice, is at the beginning of the shell command path.

To build DTC from source

1. Download the DTC source code by executing the following `git clone` command:

```
$ git clone git://git.kernel.org/pub/scm/utils/dtc/dtc.git
$ cd dtc
```

The rest of this procedure assumes your current working directory is the DTC source tree.

2. Build DTC by executing:

```
$ make
```

3. To install into the default directory `$HOME`, execute:

```
$ make install
```

4. To install to a specific directory, execute:

```
$ make install PREFIX=/usr/local
```

In either case, the following directories are created below the installation directory:

```
bin
lib
include
```

Replacing the Kernel

You can replace the kernel with your own custom kernel. Use the procedure that applies to the type of storage device from which your device boots.

To replace the kernel in systems that boot from internal eMMC

These steps can be used in SD card and USB storage device cases.

1. Boot the Jetson TX1/TX2 system.
2. Copy the new kernel files, using `scp`, into the `/boot` directory.
3. Reboot the Jetson TX1/TX2 system.

To replace the kernel in systems that boot from an SD Card or USB device

1. Connect the SD Card or USB device to your host system.
2. Copy the new kernel files to the `/boot` directory on the SD Card or USB device.
3. Disconnect the SD Card or USB device from the host system.
4. Connect the SD Card or USB device to the Jetson TX1/TX2 system.
5. Reboot the Jetson TX1/TX2 system.

To replace the kernel in systems that boot from an NFS server

1. Boot the Jetson TX1/TX2 system.

2. On the target system enter the following command:

```
$ sudo mount /dev/mmcblk0p1 /mnt
```

3. Copy the new kernel files, using `scp`, to the `mnt/boot` directory.
4. Reboot the Jetson TX1/TX2 system.

Copying the Root Filesystem

For external media, you must copy the root filesystem to the device **after** running the `flash.sh` command. Then you attach the device to the target system. This is required because the `flash.sh` script adds the required files to the root filesystem directory on the host when it runs.

Boot Sequence and Sysboot Configuration Files

The U-Boot functionality includes a default booting scan sequence. It scans bootable devices in the following order:

- External SD Card
- Internal eMMC
- USB Device
- NFS Device

It looks for an `extlinux.conf` configuration file in the following directory of the bootable partition in the device:

```
<rootfs>/boot/extlinux
```

Upon finding the `extlinux.conf` file, U-Boot performs the following.

- Uses the `sysboot` command to read out boot configuration from `extlinux.conf`
- Loads kernel Image file
- Boots the kernel

The Image is user-accessible in the `<rootfs>/boot` location after booting. The `extlinux.conf` file is user accessible in the `<rootfs>/boot/extlinux` location. Users can easily change these files to test their own kernel without flashing.

The file `extlinux.conf` is a standard text-format `sysboot` configuration file that contains all boot information. It indicates the kernel image filename, the device tree blob filename, and the kernel boot command line. There are four example `extlinux.conf` files provided in the L4T release for each supported board:

```
extlinux.conf.emmc
extlinux.conf.sdcard
extlinux.conf.usb
extlinux.conf.nfs
```

During flashing, `flash.sh` copies the appropriate variant to the following location on the target:

```
<rootfs>/boot/extlinux/extlinux.conf
```

The `extlinux.conf` files are very similar, except for different kernel boot command lines. You can find the `extlinux.conf` files on your host system at:

```
bootloader/<platform>/<board>/
```

Where:

- `<platform>` is the SOC name or platform.
- `<board>` is:
 - Jetson TX1: p2371-2180
 - Jetson TX2: p2771-0000

Optimizing U-Boot Boot Time

By default, U-Boot includes a default configuration that enables all supported hardware features. It searches the available devices for boot scripts. This enables out-of-the-box support for the widest possible variety of storage devices and boot configurations.

This flexibility delays execution of the final operating system because hardware support takes time to initialize and scanning all attached storage and network devices takes time. In constrained or pre-configured systems, this flexibility may not be necessary. You may know ahead of time, which storage device contains the required files, or that certain devices are not required to be initialized by the boot-loader. To optimize boot time, configure U-Boot to allow for these constraints and reduce system boot time.

Disabling PCIe Support

U-Boot delays enumerating and initializing PCIe devices until the user, or a boot script, explicitly attempts to access a PCIe device. Consequently, disabling PCIe support does not reduce boot time.

To disable PCIe support

1. Remove the following values from `include/configs/<board>.h`:

```
CONFIG_PCI_EHCI
CONFIG_PCI_TEGRA
CONFIG_CMD_PCI
```

2. Remove the following values from `configs/<board_and_rev>_defconfig`:

```
CONFIG_PCI_TEGRA
CONFIG_RTL8169
CONFIG_E1000
```

Where:

- `<board>` is:
 - Jetson TX1: p2371-2180
 - Jetson TX2: p2771-0000
- `<board_and_rev>` is:

- Jetson TX1: p2371-2180
- Jetson TX2: p2771-0000-500

Disabling USB Support

U-Boot delays enumerating and initializing USB devices until the user, or a boot script, explicitly attempts to access a USB device. Consequently, disabling USB support does not reduce boot time.

To disable USB support

1. Remove the following values from `configs/<board_and_rev>_defconfig`:

```
CONFIG_CMD_USB
CONFIG_CMD_USB_MASS_STORAGE
CONFIG_USB
CONFIG_USB_GADGET
CONFIG_USB_GADGET_DOWNLOAD
```

2. Edit `include/configs/<board>.h` to remove the `#include` of `tegra-common-usb-gadget.h`.
3. Edit `include/configs/tegra-common-post.h` to remove the USB entry from `BOOT_TARGET_DEVICES`.

Where:

- `<board>` is:
 - Jetson TX1: p2371-2180
 - Jetson TX2: p2771-0000
- `<board_and_rev>` is:
 - Jetson TX1: p2371-2180
 - Jetson TX2: p2771-0000-500

Environment Configuration

U-Boot runtime behavior is controlled by scripts contained in the U-Boot environment. When U-Boot begins execution, it waits for `<bootdelay>` seconds before executing the automatic boot sequence. During this time, the user may interrupt the boot process to access the U-Boot shell. If not interrupted, U-Boot executes `<bootcmd>` as a shell command. `<bootcmd>` contains a series of commands to search storage devices for boot scripts and execute them.

By modifying the values of these variables at compile or manufacturing time, U-Boot can be directed to boot from a specific device in a specific manner, thereby reducing boot time.

The following table identifies the variables that can be modified.

Variable	Description
• <code>bootdelay</code>	• Contains the number of seconds that U-Boot pauses to determine whether the user wishes to interrupt the boot sequence.

	<ul style="list-style-type: none"> •To avoid delay, set to 0. •<i>Note:</i> Although this value avoids delay, if the user has requested to interrupt the boot process before the U-Boot shell is reached, that request is honored. •To avoid a delay and user interruption of the boot process, set <code>bootdelay</code> to a negative value.
• <code>bootcmd</code>	<ul style="list-style-type: none"> •Contains a sequence of U-Boot shell commands to be executed automatically at boot. •For systems with custom requirements, this value can be completely replaced. However, simple customizations, such as selecting a specific storage device for booting, does not require editing this value.
• <code>boot_targets</code>	<ul style="list-style-type: none"> •Contains a space-separated list of storage devices or network protocols that U-Boot scans to find boot scripts. •Valid values include: <ul style="list-style-type: none"> • <code>mmc0</code> - the built-in eMMC. • <code>mmc1</code> - the SD card slot. • <code>usb0</code> - any attached USB Mass Storage device. • <code>pxe</code> - network, using DHCP to receive an IP address, then PXE to download a <code>syslinux</code> configuration file. • <code>dhcp</code> - network using DHCP to receive an IP address, then TFTP to download a U-Boot boot script. •This variable can be set to a single specific device, or a more restrictive list than the default.
• <code>scan_dev_for_boot_part</code>	<ul style="list-style-type: none"> •Contains a script to parse the device partition table and determines which partition U-Boot must scan for boot files. •If the partition number is known ahead of time, replace the script with a simpler script that hardcodes the value of <code><devplist></code> with a single partition number (represented in hexadecimal), and then runs either or both of the following variables: <ul style="list-style-type: none"> • <code>scan_dev_for_extlinux</code> or • <code>scan_dev_for_boot_scripts</code>
• <code>boot_prefixes</code>	<ul style="list-style-type: none"> •Contains a list of filesystem directories to scan for boot scripts or configuration files. •File system layouts vary between installations. •For example, <code>/boot</code> may be a separate partition containing boot scripts, or part of the root filesystem. •By default, U-Boot searches both locations for boot scripts. •In constrained cases, the user may set this variable to a single directory name so that U-Boot does not search unnecessary directories. •<i>Note:</i> All entries in this variable must contain both a leading and a trailing <code>/</code>.
• <code>scan_dev_for_extlinux</code>	<ul style="list-style-type: none"> •Contains a script to search for <code>extlinux</code> configuration files. If found, boots the system based on their content.

	<ul style="list-style-type: none"> •If the system is known not to use extlinux configuration files, replace this script with commands that do nothing. •For best results, set this variable to an innocuous value, such as “true”, rather than leaving it empty, so that U-Boot does not complain about attempts to execute an empty variable as a script.
•scan_dev_for_scripts	<ul style="list-style-type: none"> •Contains a script to search for U-Boot boot scripts. If found, loads and executes them. •If the system is known not to use U-Boot boot scripts, replace this script with commands that do nothing. •Note: The default boot scripts execute <scan_dev_for_extlinux> prior to executing <scan_dev_for_scripts>. Therefore, modifying this variable does not affect systems that boot using extlinux configuration files because this script will never be executed. •For best results, set this variable to an innocuous value, such as “true”, rather than leaving it empty, so that U-Boot does not complain about attempts to execute an empty variable as a script.
•boot_scripts	<ul style="list-style-type: none"> •Contains a space-separated list of U-Boot script names for <scan_dev_for_scripts> to search for. •If the script name is known ahead of time, set this variable to the desired value rather than the default list.

Setting Environment Variables

You can set environment variable at:

- Compile-Time
- Manufacturing Time
- Flashing Time

Compile-Time

When U-Boot starts executing, it attempts to initialize the environment variables from data stored in flash memory. The location of the data is determined by the U-Boot configuration file on the board. If the data is missing or corrupted, U-Boot uses the default set of values that are built into the U-Boot binary.

The default L4T flashing process does not write this data into flash memory, so the built-in copy is always used initially. The built-in default environment values are set in the following places in the U-Boot source code:

U-Boot Source Code	Variable Value
•include/config_distro_bootcmd.h	•This file defines the variables and scripts related to the automatic boot process.
•include/config_distro_defaults.h	•The value of CONFIG_BOOTDELAY determines the default value of <bootdelay>.

<ul style="list-style-type: none"> • include/configs/tegra-common-post.h 	<ul style="list-style-type: none"> • The value of BOOT_TARGET_DEVICES determines: <ul style="list-style-type: none"> • The default value of <boot_targets>. • The set of legal values that can appear in <boot_targets>. • Removing entries from this variable prevents the use of those values in <boot_targets> at all. • To modify this value at compile-time: <ul style="list-style-type: none"> • Modify include/config_distro_bootcmd.h to avoid setting <boot_targets> if the board-specific configuration file has already defined this value. • Modify include/config/<board>.h to set <boot_targets>.
---	--

Manufacturing and Flashing Time

You can modify your manufacturing flow to add an extra step that writes a saved copy of the environment variables to flash. The following provides a set of alternatives to use for modifying your manufacturing flow:

- After flashing the board, arrange for the board to execute U-Boot, and cause U-Boot to execute commands that modify and save the environment:
 - Reset the board so that U-Boot runs and sends commands to U-Boot using the serial console. You must write and execute some program on your host system to detect when U-Boot begins executing, interrupt the U-Boot automatic boot process, and then send the commands to U-Boot.
 - Execute a sequence of commands similar to the following:

```
env default -f -a; setenv boot_targets mmc0; saveenv
```

- Manually edit the U-Boot environment interactively, save the result, and extract the appropriate flash memory region to a file on your host system. Then, modify the L4T partition layout XML file to that the data is written to flash during any subsequent flashing process.

Alternatively, apply this technique if your manufacturing process programs flash chips directly, rather than using the L4T flashing tools on each board. You must merge the extracted saved environment data into your flash image file, rather than referencing it from the L4T partition layout XML file.

- Generate a set of saved environment data using the U-Boot `fw_env` tool located in the `tools/env` directory in the U-Boot source code. Then, modify the L4T partition layout XML file so that data is written to flash during any subsequent flashing process.

Alternatively, apply this technique if your manufacturing process programs flash chips directly, rather than using the L4T flashing tools on each board. You must merge the extracted saved environment data into your flash image file, rather than referencing it from the L4T partition layout XML file.

extlinux.conf Modifications

L4T includes an `extlinux.conf` file that tells U-Boot which kernel filenames to load. It also includes the command line to pass to the kernel. The file is configured to display a boot menu to the user for 3 seconds before automatically booting. To reduce boot processing time, you can remove this timeout.

To remove extlinux.conf timeout boot menu

- Edit `extlinux.conf` file to remove the lines containing the following keywords:

```
TIMEOUT
MENU
```

On an installed L4T target system, this file is located at:

```
/boot/extlinux/extlinux.conf
```

Prior to flashing, this file is located in the host L4T flashing directory at

```
bootloader/<platform>/<board_and_rev>/extlinux.conf.emmc
```

Where:

- `<platform>` is the SOC name or platform.
- `<board_and_rev>` is:
 - Jetson TX1: p2371-2180
 - Jetson TX2: p2771-0000/500

Debugging U-Boot Environment

Use these debugging tips to help you debug your U-Boot environment. These examples do not represent a comprehensive listing of U-Boot functionality. For a full list of supported commands and their usage by U-Boot, consult U-Boot documentation and source.

Interrupting U-Boot

You can interrupt U-Boot during boot.

To interrupt U-Boot

1. Connect a terminal application to the serial port of the board.
2. Power on or reset the board.
3. Wait for the U-Boot sign on message to appear. Key presses before this point in time may be ignored.
4. Press any key to interrupt the automatic boot process.

Getting Help

On the U-Boot terminal screen, type `help` at any time for the list of supported commands from the U-Boot terminal.

To see the U-Boot Help text

- To see the U-Boot help text enter the following command:

```
# help
```


Listing a Directory Structure

You can list the directory structure of a particular device.

To list the directory structure

- To list the directory structure of eMMC device 0 partition 1, enter the following command on the target:

```
# ls mmc 0:1
```

This command functions correctly on EXT2/3/4 and FAT file systems.

Example output follows:

```
<DIR>      4096 .
<DIR>      4096 ..
<DIR>      4096 bin
<DIR>      4096 boot
<DIR>      4096 dev
<DIR>      4096 etc
<DIR>      4096 home
<DIR>      4096 lib
<DIR>      4096 lost+found
<DIR>      4096 media
<DIR>      4096 mnt
<DIR>      4096 opt
<DIR>      4096 proc
<DIR>      4096 root
<DIR>      4096 sbin
<DIR>      4096 selinux
<DIR>      4096 srv
<DIR>      4096 sys
<DIR>      4096 tmp
<DIR>      4096 usr
<DIR>      4096 var
```

Listing the Contents of a Directory

You can list the contents of any directory.

To list the contents of a directory

- List directory contents with the following command:

```
# ls mmc 0:1 <directory>
```

Where `<directory>` is a pathname in the filesystem.

For example, to list contents of the `/boot` directory where the Image file should be, (as shown in the example output below), use the following command:

```
# ext2ls mmc 0:1 /boot
<DIR>      1024 .
<DIR>      1024 ..
           908 extlinux
          5910248 Image
```

Printing the U-Boot Environment

You can print the entire U-Boot environment.

To print the U-Boot environment

- Execute the following command:

```
# printenv
```

Printing/Setting Environment Variables

You can print and set environment variables.

To print an environment variable

- Execute the following command:

```
# printenv <environment_variable>
```

Where `<environment_variable>` refers to an environment variable in U-Boot.

For example, to print the list of devices U-Boot sends console output to, execute:

```
# printenv stdout
```

Output can be as follows:

```
stdout=serial
```

To set an environment variable

- Execute the following command:

```
# setenv <environment_variable> <new_value>
```

Where `<environment_variable>` refers to an environment variable in U-Boot and `<new_value>` is the new value for that variable.

For example, to modify the set of devices that U-Boot sends console output to, execute:

```
# setenv stdout serial
```

To save the modified environment

- Execute the following command:

```
# saveenv
```

The saved modified environment is preserved in case of resets and reboots.

Mass Storage Partitions Configuration

NVIDIA® Linux® for Tegra® releases support formatting mass storage media into multiple partitions for storing data, such as the device OS image, boot loader image, device firmware, and boot loader splash screens.

Partition Configuration File

The partition configuration file describes the entire mass storage devices and partitions in devices. For Linux® for Tegra® releases the default partition configuration file directory is available at:

```
<top>/Linux_for_Tegra/<board name>/cfg
```

During the flashing procedure, `flash.sh` reads in the partition configuration file and converts variables into real numbers or names as specified in `<device>.conf` files or in option parameters and save it in `bootloader/flash.xml` file. Then the `bootloader/tegraflash.py` reads in `bootloader/flash.xml` file and performs actual flashing as specified by `bootloader/flash.xml` file.

The partition configuration file is composed of the following XML format elements:

- Root element
- Device elements
- Partition elements

Root Element

There is only root element in a configuration file and it has following syntax:

```
<?xml version="1.0"?>

<!-- NVIDIA Tegra Partition Layout Version 1.0.0 -->

<partition_layout version="01.00.0000">
  <!-- Device elements go here. -->
</partition>
```

Device Element

There are multiple device elements in the root element and it has following syntax:

```
<device type="<device type>" instance="<instance>">
  <!-- Partition elements go here. -->
</device>
```

Device Attributes Table

The following table shows supported device attributes and their legal values.

Device Attributes	Values	Description
•device type	<ul style="list-style-type: none"> • • •sdmmc_boot •sdmmc_user 	<ul style="list-style-type: none"> •(Required) Specifies the type of mass storage device. •boot0 and boot1 linearly combined •EMMC user partition
•instance	•0-3	•(Required) 3 is the default primary

Partition Element

There are multiple partition elements in a device element and it has following syntax:

```
<partition name="<name>" type="<type>" oem_sign="true">
  <!-- partition properties go here -->
</partition>
```

Partition Attributes Table

The following table shows supported partition attributes and their legal values.

Partition Attributes	Values	Description
•name	<ul style="list-style-type: none"> •<name>such as: •BCT, APP, and mb1, or <keyword> •For a complete list, see default partition overview section below. 	<ul style="list-style-type: none"> •(Required) Specifies a name for the partition. This name is used when opening a partition for read/write access. The name can have a maximum of 36 characters. There are 3 kinds of partition names: mandatory, optional, and user defined. All mandatory and optional partition names should be associated with specific partition types for specific operations and all user specified partition names should be associated with partition type “data” for no specific action by bootloader.
•type	<ul style="list-style-type: none"> •<type> such as: •boot_config_table, wb0, and data, or <keyword> •For a complete list, see default partition overview section below. 	<ul style="list-style-type: none"> •(Required) Specifies the type of partition. Certain partition types are valid only when they are associated specific partition names. Detailed example is in default partition table section below.
•oem_sign	<ul style="list-style-type: none"> •True • 	Append generic signature header to the binary and perform OEM signing on that.

	<ul style="list-style-type: none"> • False/Undefined 	Append NULL signature header to the binary and perform NULL signing on that.
--	---	--

Partition Properties Table

The following table shows supported partition properties and their legal values.

Partition Elements	Values	Description
•allocation_policy	•sequential	•Specifies the type of allocation policy. Sequential partitions begin immediately after the preceding partition. (The legacy “absolute” allocation policy has been deprecated.)
•filesystem_type	•basic	•The basic filesystem_type specifies “raw” partition. (The legacy “ext2” filesystem has been deprecated.)
•size	<ul style="list-style-type: none"> •<byte size> •Or <keyword> 	<ul style="list-style-type: none"> •(Required) Specifies the size of the partition in bytes. Decimal and hexadecimal values are valid. Size can be the exact file size being programmed in bytes. However, to minimize changes to your CFG file size increases over time, you can set the partition size in the CFG file greater than actual file size • •If the partition size specified is less than an erase block size, it will align itself to the partition erase block boundary. • •If a partition filesystem_type is enhanced file system and size is on a block boundary, then size should be specified to be one additional sector. • •The size property of the secondary GPT partition, which is defined at the end of each device, should always be 0xffffffffffff and the actual size of the Secondary GPT partition is the size defined in the partition which immediately preceding the secondary GPT partition. The actual size of the partition immediately preceding secondary GPT is the size of remaining space left subtracted by the actual secondary GPT partition size.
•file_system_attribute	•-	•Locks the file system. Not currently implemented.

•allocation_attribute	•0x8 or 0x808	•The allocation attribute for all partitions must be 0x8. The partition immediately preceding secondary GPT should set this attribute to 0x808.
•percent_reserved	•-	•Reserved for future use.
•filename	•<filename> •or <Keyword>	•Specifies the name of the file to write into the partition.

Partition Property Keyword Table

The L4T releases adopted a <keyword> in the partition configuration file to abstract varying mandatory and optional partition names, types, and sizes. Upon detection of the <keyword> in the partition configuration file in `bootloader/<device>/cfg`, the `flash.sh` replaces <keyword> with sensible default value for the <device>. If a user replaces any of the <keyword> with a specific value, the value is honored by `flash.sh` and the partition is flashed as specified.

The following table shows supported partition property keywords and their legal values after translation into `bootloader/flash.xml`.

Partition Keywords	Default translated values in flash.xml
•MB1NAME	•mb1
•MB1TYPE	•mb1_bootloader
•SPENAME	•spe-fw
•SPETYPE	•spe_fw
•MB2NAME	•mb2
•MB2TYPE	•mb2_bootloader
•MB2FILE	•nvtboot.bin
•MPBNAME	•mts-preboot
•MPBTYPE	•mts_preboot
•MPBFILE	•preboot_d15_prod_cr.bin
•APPSIZE	•30064771072
•APPFILE	•system.img
•MBPNAME	•mts-bootpack
•MBPTYPE	•mts_bootpack
•MBPFILE	•mce_mts_d15_prod_cr.bin
•TBCNAME	•cpu-bootloader
•TBCTYPE	•bootloader

•TBCFILE	•cboot.bin
•TBCDTB-NAME	•bootloader-dtb
•TBCDTB-FILE	•Appropriate bootloader DTB file name
•TOSNAME	•secure-os
•TOSFILE	•tos.img
•EKSFILE	•eks.img
•BPFNAME	•bpmp-fw
•BPFSIGN	•true
•BPFFILE	•bpmp.bin
•BPFDTB-NAME	•Bpmp-fw-dtb
•BPFDTB-SIGN	•true
•BPFDTB-FILE	•Appropriate bpmp DTB file name
•SCENAME	•sce-fw
•SCESIGN	•true
•SCEFILE	•camera-rtcpu-sce.bin
•SC7NAME	•sc7
•WB0TYPE	•WB0
•WB0FILE	•warmboot.bin
•FBNAME	•FBNAME (Not used)
•FBTYPE	•data
•FBSIGN	•false
•FBFILE	•Removed/Unused
•SOSFILE	•Removed/Unused
•LNXNAME	•kernel
•LNXSIZE	•67108864
•LNXFILE	•boot.img
•KERNELDTB-NAME	•Kernel-dtb
•KERNELDTB-FILE	•Appropriate kernel DTB file name

Default Partition Overview

The following table summarizes the usage for each partition in `sdmmc_boot` device:

Name	Type	Description
•BCT	•boot_config_table	•(Mandatory) Contains the Boot Configuration Table.
•mb1	•mb1_bootloader	•(Mandatory) Contains NVIDIA signed MB1 binary.
•MB1_BCT	•mb1_boot_config_table	•(Mandatory) Contains the MB1 boot configuration table.
•spe-fw	•spe-fw	•Contains the BPMP spe-fw binary.
•mb2	•mb2_bootloader	•Contains the primary copy of the TegraBoot binary..
•mts-preboot	•mts_preboot	•Contains Denver preboot firmware.
•secondary_gpt	•secondary_gpt	•Contains the secondary GPT of <code>sdmmc_boot</code> device.

The following table summarizes the usage for each partition in `sdmmc_user` device:

Name	Type	Description
•master_boot_record	•protective_master_boot_record	•(Mandatory) Contains protective MBR.
•primary_gpt	•primary_gpt	•(Mandatory) Contains the primary GPT of <code>sdmmc_user</code> device. All partition defined after this entry will be configured in kernel and accessible by standard partition tools such as <code>gdisk</code> or <code>parted</code> .
•APP	•data	•(Mandatory) Contains the RootFS. This partition must be defined after primary GPT, so that it can be accessed as fixed known special device <code>/dev/mmcblk0p1</code> .
•mts-bootpack	•Mts_bootpack	•(Mandatory) Contains Denver boot pack firmware.
•cpu-bootloader	•bootloader	•(Mandatory) Contains the final boot stage CPU bootloader binary that loads the binary in kernel partition.
•bootloader-dtb	•data	•(Mandatory) Contains the bootloader device tree blob (DTB).
•secure-os	•data	•(Mandatory) Contains the trusted OS.

•eks	•data	•(Optional) Contains the encrypted keys.
•adsp-fw	•data	•(Mandatory) Contains ADSP software.
•bpmp-fw	•data	•(Mandatory) Contains BPMP firmware.
•bpmp-fw-dtb	•data	•(Mandatory) Contains the BPMP firmware device tree blob (DTB).
•sce-fw	•data	•(Optional) Contains camera-rtcpu-sce firmware.
•sc7	•WBO	•(Mandatory) Contains warm boot binary.
•FBNAME	•data	•(Optional) Reserved for future use. Removable.
•BMP	•data	•(Optional) Contains BMP images (if any) for display during boot.
•SOS	•data	•(Optional) Contains a partition used for recovery. Removable.
•kernel	•data	•(Mandatory) Contains the U-Boot secondary bootloader by default. The U-Boot searches the root device and loads the Linux kernel from rootdev's / boot directory. If the U-Boot is not selected during flashing, this partition contains the Linux kernel.
•kernel-dtb	•data	•(Mandatory) Contains the kernel device tree blob (DTB).
•CAC	•data	•(Optional) Reserved for future use. Removable.
•UDA	•data	•(Optional) This partition, with the allocation attribute set to 0x808, automatically take all of remaining space of the device except secondary GPT space. This partition can be mounted and used to store user data.
•secondary_gpt	•secondary_gpt	•(Mandatory) Contains the secondary GPT of sdmmc_user device.

Example CFGs

Following is example partition configuration file after `flash.sh` completes translations and substitutions:

```
<?xml version="1.0"?>
<!-- Nvidia Tegra Partition Layout Version 1.0.0 -->
```

```

<partition_layout version="01.00.0000">
  <device type="sdmmc_boot" instance="3">
    <partition name="BCT" type="boot_config_table">
      <allocation_policy> sequential </allocation_policy>
      <filesystem_type> basic </filesystem_type>
      <size> 32768 </size>
      <file_system_attribute> 0 </file_system_attribute>
      <allocation_attribute> 8 </allocation_attribute>
      <percent_reserved> 0 </percent_reserved>
    </partition>
    <partition name="mb1" type="mb1_bootloader">
      <allocation_policy> sequential </allocation_policy>
      <filesystem_type> basic </filesystem_type>
      <size> 262144 </size>
      <file_system_attribute> 0 </file_system_attribute>
      <allocation_attribute> 8 </allocation_attribute>
      <percent_reserved> 0 </percent_reserved>
      <filename> mb1_dev.bin </filename>
    </partition>
    <partition name="MB1_BCT" type="mb1_boot_config_table">
      <allocation_policy> sequential </allocation_policy>
      <filesystem_type> basic </filesystem_type>
      <size> 65536 </size>
      <file_system_attribute> 0 </file_system_attribute>
      <allocation_attribute> 8 </allocation_attribute>
      <percent_reserved> 0 </percent_reserved>
    </partition>
    <partition name="spe-fw" type="spe_fw" oem_sign="true">
      <allocation_policy> sequential </allocation_policy>
      <filesystem_type> basic </filesystem_type>
      <size> 131072 </size>
      <file_system_attribute> 0 </file_system_attribute>
      <allocation_attribute> 8 </allocation_attribute>
      <percent_reserved> 0 </percent_reserved>
      <filename> spe.bin </filename>
    </partition>
    <partition name="mb2" type="mb2_bootloader" oem_sign="true">

```

```

    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 262144 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> nvtboot.bin </filename>
</partition>
<partition name="mts-preboot" type="mts_preboot" oem_sign="true">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 262144 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> preboot_d15_dev_cr.bin </filename>
</partition>
<partition name="secondary_gpt" type="secondary_gpt">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 0xFFFFFFFFFFFFFFFF </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
</partition>
</device>

<device type="sdmmc_user" instance="3">
    <partition name="master_boot_record" type="protective_master_boot_record">
        <allocation_policy> sequential </allocation_policy>
        <filesystem_type> basic </filesystem_type>
        <size> 512 </size>
        <file_system_attribute> 0 </file_system_attribute>
        <allocation_attribute> 8 </allocation_attribute>
        <percent_reserved> 0 </percent_reserved>
    </partition>
    <partition name="primary_gpt" type="primary_gpt">

```

```

    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 2097152 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
</partition>
<partition name="APP" type="data">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 15032385536 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 0x8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> system.img </filename>
</partition>
<partition name="mts-bootpack" type="mts_bootpack" oem_sign="true">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 4194304 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> mce_mts_d15_dev_cr.bin </filename>
</partition>
<partition name="cpu-bootloader" type="bootloader" oem_sign="true">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 262144 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> cboot.bin </filename>
</partition>
<partition name="bootloader-dtb" type="data" oem_sign="true">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>

```

```

    <size> 262144 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> tegra186-quill-p3310-1000-a00-00-base.dtb </filename>
</partition>
<partition name="secure-os" type="data" oem_sign="true">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 3145728 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> tos.img </filename>
</partition>
<partition name="eks" type="data" oem_sign="true">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 2048 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> eks.img </filename>
</partition>

<partition name="adsp-fw" type="data" oem_sign="true">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 4194304 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
</partition>

<partition name="bpmp-fw" type="data" oem_sign="true">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>

```

```

    <size> 618144 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> bpmp.bin </filename>
</partition>
<partition name="bpmp-fw-dtb" type="data" oem_sign="true">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 512000 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> tegra186-a02-bpmp-quill-p3310-1000-a00-00-te770d-ucm2.dtb </file
</partition>
<partition name="sce-fw" type="data" oem_sign="true">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 2097152 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> camera-rtcpu-sce.bin </filename>
</partition>
<partition name="sc7" type="WB0">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 6291456 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> warmboot_dev.bin </filename>
</partition>
<partition name="FBNAME" type="data" oem_sign="false">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 2097152 </size>

```

```

    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 0x8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
</partition>
<partition name="BMP" type="data">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 134217728 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 0x8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
</partition>
<partition name="SOS" type="data">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 33554432 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 0x8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
</partition>
<partition name="kernel" type="data">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 67108864 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> boot.img </filename>
</partition>
<partition name="kernel-dtb" type="data">
    <allocation_policy> sequential </allocation_policy>
    <filesystem_type> basic </filesystem_type>
    <size> 262144 </size>
    <file_system_attribute> 0 </file_system_attribute>
    <allocation_attribute> 8 </allocation_attribute>
    <percent_reserved> 0 </percent_reserved>
    <filename> tegra186-quill-p3310-1000-a00-00-base.dtb </filename>

```



```

</partition>
<partition name="CAC" type="data">
  <allocation_policy> sequential </allocation_policy>
  <filesystem_type> basic </filesystem_type>
  <size> 268435456 </size>
  <file_system_attribute> 0 </file_system_attribute>
  <allocation_attribute> 0x8 </allocation_attribute>
  <percent_reserved> 0 </percent_reserved>
</partition>
<partition name="UDA" type="data">
  <allocation_policy> sequential </allocation_policy>
  <filesystem_type> basic </filesystem_type>
  <size> 18432 </size>
  <file_system_attribute> 0 </file_system_attribute>
  <allocation_attribute> 0x808 </allocation_attribute>
  <percent_reserved> 0 </percent_reserved>
</partition>
<partition name="secondary_gpt" type="secondary_gpt">
  <allocation_policy> sequential </allocation_policy>
  <filesystem_type> basic </filesystem_type>
  <size> 0xFFFFFFFFFFFFFFFF </size>
  <file_system_attribute> 0 </file_system_attribute>
  <allocation_attribute> 8 </allocation_attribute>
  <percent_reserved> 0 </percent_reserved>
</partition>
</device>
</partition_layout>

```

Kernel Boot Time Optimization

NVIDIA[®] Tegra[®] Linux Driver Package (L4T) provides a generic boot kernel for development of your product. To decrease kernel boot time, customize the provided kernel based on the requirements of your product.

The kernel includes a default configuration that enables all supported hardware features, and searches all available devices for boot scripts. This enables out-of-the box support for the widest possible variety of controllers, features, storage devices, and boot configurations.

This flexibility comes at a cost:

- Some hardware support takes time to initialize
- Enabling all software features, mostly over Advanced Peripheral Bus (APB), takes time
- Scanning all attached storage and network devices takes time thereby delaying execution of the final operating system

In constrained or pre-configured systems, this flexibility may not be necessary; the system designer may know ahead of time which storage device contains the required files, or that certain devices do not need to be initialized by the kernel. To reduce system boot time, can configure the kernel to respect these constraints.

For a Jetson Tegra X1 system running L4T with the default configuration, it takes 12 seconds from cold power-on to begin showing the login prompt. When the following optimization techniques are applied, that process can be reduced to approximately three seconds.

Device Tree Nodes

If you are not using any controller from Tegra SoC, disable the Device Tree nodes for those device tree entries. The device tree directory is available at:

```
<top>/hardware/nvidia/platform/t18x/  
<top>/hardware nvidia/soc/t18x
```

The DTB that is flashed is available at:

```
<top>/hardware/nvidia/platform/t18x/quill/kernel-dts/tegra186-quill-p3310-1000-a00-00-b
```

Real-time Clock

Two Real-time Clocks (RTC) are enabled by default:

- Tegra RTC
- PMIC RTC

Enable one, both, or none, based on your development needs. Be aware of the following differences:

- Tegra RTC is the fastest in response, but cannot work as a backup RTC when the system is turned off.
- PMIC RTC can work as a backup RTC, however it is slow in response because of the transfer average to I2C.

If you do not shutdown the system, or if you are not concerned about backup time, then you do not need to enable RTC. Disabling RTC can speed up boot time.

To disable the RTC

1. Remove the the following configuration from your `tegra21_defconfig` file.

```
CONFIG_RTC_DRV_MAX77620=y
```

2. Add the following configuration to the `tegra21_defconfig` file.

```
# CONFIG_RTC_HCTOSYS is not set
```

3. Set the Tegra RTC to disabled in the DTS file as follows:

```
rtc {
    compatible = "nvidia,tegra-rtc";
    reg = <0x00000000 0x00000003 0x00000002 0x0000009b>;
    interrupts = <0x00000000 0x00000005 0x00000002>;
    status = "disabled";
};
```

Environment Configuration

You can optimize boot time by modifying the environment configuration in the root file system.

Disable Console over UART

Prints over UART are a major bottleneck in kernel boot time. To reduce this time, you can remove `console=ttyS0` from the `extlinux.conf` configuration file.

Once the system is ready for deployment, you can remove the UART console logs or review the console logs over the framebuffer console. The console log is `console=tty1` in the `extlinux.conf` configuration file.

Secondary Bootloader

You can remove uboot as the secondary bootloader to remove boot time of u-boot. The following modifications are required before flashing:

1. Edit the common configuration file at:

```
<top>/tegra_for_linux/ p2771-000.conf
```

2. Set `USE_UBOOT` to 0.

```
# To configure whether to use U-Boot,
# do either of the following before running flash.sh:
# 1) Set environment variable USE_UBOOT to 0 or 1.
# 2) Edit the line below to set USE_UBOOT to 0 or 1.
if [ -z "${USE_UBOOT}" ]; then
```

```
USE_UBOOT=0;
fi;
```

Compile-Time Configuration

To reduce compile-time configuration, examine the generated configuration file to identify which configurations are required. Once the required configurations are defined, identify which ones to boot asynchronously. For those configurations, the drivers probe is executed asynchronously in a separate thread instead of the main initial thread.

Additionally, examine the required configurations and verify that they can be programmed as modules so that the drivers are loaded when it is called for use. When the drivers are not loaded, the kernel image is reduced and more RAM space is available.

The following topics provide examples of each of these conditions.

Asynchronous Probe

The asynchronous probe feature is available, by default, from Kernel Version 3.18:

<https://lkml.org/lkml/2015/1/16/576>

To move the driver to another thread

- Add the `probe_type` in your driver as follows:

```
static struct platform_driver sdhci_tegra_driver = {
    .driver          = {
        .name         = "sdhci-tegra",
        .of_match_table = sdhci_tegra_dt_match,
        .pm           = SDHCI_PLTFM_PMOPS,
        .probe_type  = PROBE_PREFER_ASYNCHRONOUS,
    },
    .probe           = sdhci_tegra_probe,
    .remove          = sdhci_tegra_remove,
    .shutdown        = sdhci_tegra_shutdown,
};
```

File System

To decrease boot time the filesystem, modify the following configurations to set them as modules:

```
CONFIG_FUSE_FS=m
CONFIG_VFAT_FS=m
CONFIG_NTFS_FS=m
```

Sound

Audio codec requires some time to initialize, To eliminate this initialization time, disable the audio configurations as follows:

```
# CONFIG_SND_SOC_TEGRA_ALT is not set
# CONFIG_SND_SOC_TEGRA_ALT_FORCE_CARD_REG is not set
# CONFIG_SND_SOC_TEGRA_T186REF_ALT is not set
# CONFIG_SND_SOC_TEGRA_T186REF_MOBILE_ALT is not set
```

Modularize Kernel Driver

To reduce boot time, modularize the drivers that are not required during boot, such as: HID, NET, QSPI, or USB.

Disable Debugging

To reduce boot time, disable debugging kernel hacks in the production configuration.

For example: FTRACE, TRACER, KMEMLEAK, etc.

Document Change History

Document number: DA_XXXX-001

Version	Date	Authors	Description of Change
1.0	21 June 2016	Nine Feng	Initial draft
2.0	4 Aug 2016	Chi-Jen Huang	Extensive updates

Note:

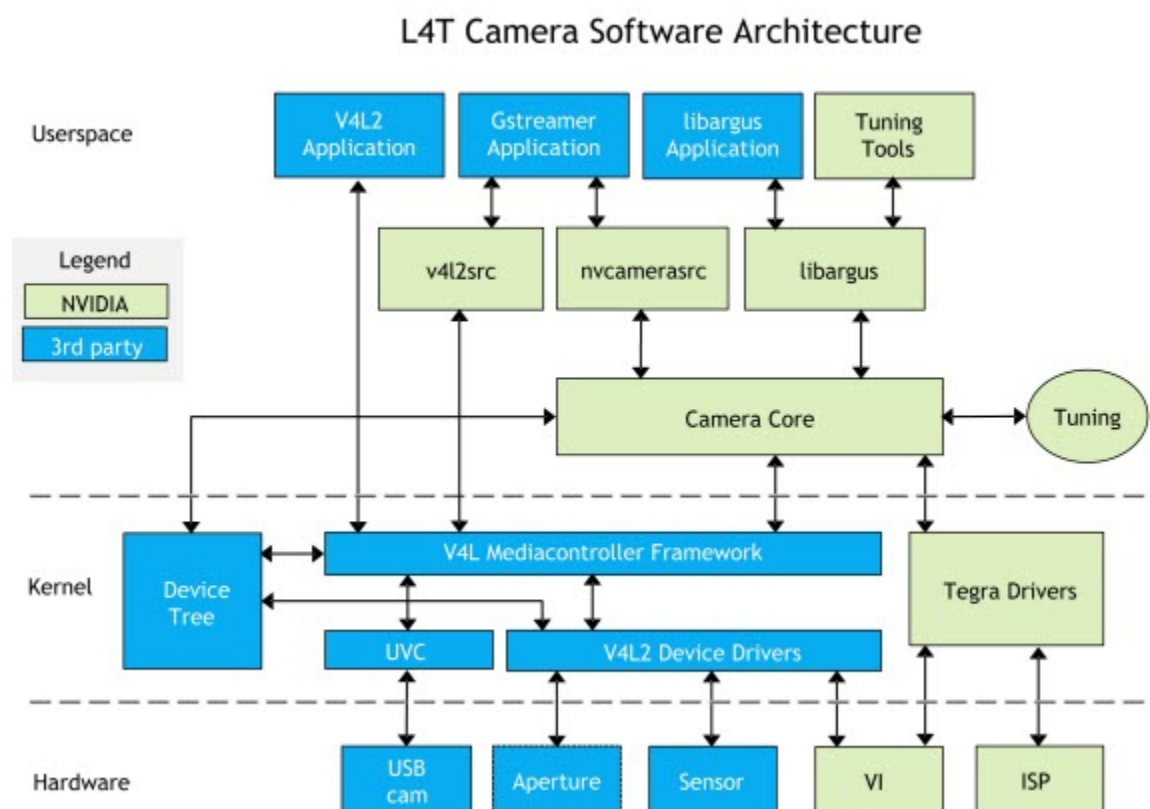
Hyperlinks in this document are a legacy of the HTML version and may not work in the PDF version.

Camera Software Development Solution

This document describes the NVIDIA® Jetson™ TX2 camera software solution and explains the NVIDIA supported and recommended camera software architecture for fast and optimal time to market. Development options are outlined and explained to customize the camera solution for USB, YUV, and Bayer camera support. Additionally, debugging approaches are explored to aid in the development process.

Camera Architecture Stack

The NVIDIA camera software architecture includes NVIDIA components that allow for ease of development and customization:



The camera architecture includes the following NVIDIA components:

- libargus—Provides a low-level API based on the camera core stack.
- nvcamerasrc—An NVIDIA camera GStreamer plugin that provides many options to control ISP properties.
- v4l2src—A standard Linux V4L2 application that uses direct kernel IOCTL calls to access V4L2 functionality.

NVIDIA provides OV5693 Bayer sensor as a sample. NVIDIA tunes this sensor for the Jetson platform. The drive code, based on the media controller framework, is available at:

```
./kernel/drivers/media/i2c/ov5693.c
```

NVIDIA provides additional sensor support for BSP software releases. Developers must work with NVIDIA certified camera partners for any Bayer sensor and tuning support. The work involved includes:

- Sensor driver development
- Custom tools for sensor characterization
- Image quality tuning

These tools and operating mechanisms are NOT part of the public Jetson Embedded Platform (JEP) Board Support Package release.

For more information on sensor driver development, see the *NVIDIA Tegra X1 V4L2 Sensor Driver Programming Guide*.

Camera API Matrix

The following table provides a matrix of the camera APIs available at each camera configuration.

	Uses Tegra ISP (CSI Interface)	Does not use Tegra ISP (CSI Interface)	USB (UVC) * (USB Interface)
Camera API	libargus GStreamer (GST-NVCamera)	V4L2	V4L2
<ul style="list-style-type: none"> •* Customer can support peripheral bus device such as: Ethernet Non-UVC USB 			

Note: The default OV5693 camera does not contain an integrated ISP. Use of the V4L2 API with the reference camera records “raw” Bayer data.

Approaches for Validating and Testing the V4L2 Driver

This topic provides various testing and debugging approaches to assist development and validation of the sensor functionality.

Once your driver development is complete, use the provided tools or application to validate and test the V4L2 driver interface.

For general GStreamer and multimedia operations, see the *Multimedia User Guide*, which is available from the NVIDIA Embedded Download Center at:

<https://developer.nvidia.com/embedded/downloads>

Applications Using libargus Low-level APIs

The NVIDIA Multimedia API provides samples that demonstrate how to use the libargus APIs to preview, capture, and record the sensor stream.

The Multimedia API can be installed only with NVIDIA JetPack. It cannot be installed as a stand-alone package. JetPack is available for download from the NVIDIA Embedded Download Center at:

<https://developer.nvidia.com/embedded/downloads>

To view the *Multimedia API Reference*, click the API tab at the top of this page.

Applications Using GStreamer with the nvcamerasrc Plugin

With GStreamer with the nvcamerasrc, you can:

- Enable ISP post-processing for Bayer sensors
- Perform format conversion
- Generate output directly for YUV sensor and USB camera

For example, for a Bayer sensor with the format 1080p/30/BGGR, you can:

- Save the preview into the file as follows:

```
gst-launch-1.0 nvcamerasrc num-buffers=200 sensor-id=0 ! 'video/x-raw(memory:NVMM),width=1920,height=1080'
```

- Render the preview to an HDMI screen as follows:

```
gst-launch-1.0 nvcamerasrc sensor-id=0 ! 'video/x-raw(memory:NVMM),width=1920,height=1080'
```

Applications Using GStreamer with V4L2 Source Plugin

Use this approach for YUV sensor or USB camera to output YUV images without ISP post-processing. This approach does not use any of the NVIDIA camera software stack.

For example, a USB camera with the format 480p/30/YUY2, you can:

- Save the preview into a file as follows (based on software converter):

```
gst-launch-1.0 v4l2src num-buffers=200 device=/dev/video0 ! 'video/x-raw, format=YUY2, width=640, height=480'
```

- Render the preview to a screen as follows:

```
//export DISPLAY=:0 if you are operating from remote console
gst-launch-1.0 v4l2src device=/dev/video0 ! 'video/x-raw, format=YUY2, width=640, height=480'
```

For a YUV sensor with the format 480p/30/YVYU, you can:

- Save the preview into a file as follows (based on hardware accelerated converter):

```
gst-launch-1.0 -v v4l2src device=/dev/video0 ! 'video/x-raw, format=(string)UYVY, width=640, height=480'
```

- Render the preview to a screen as follows:

```
//export DISPLAY=:0 if you are operating from remote console
gst-launch-1.0 v4l2src device=/dev/video0 ! 'video/x-raw, format=(string)UYVY, width=(i
```

Applications Using V4L2 IOCTL Directly

Use this approach to verify basic functionality during sensor bringup.

For example, capture from a Bayer sensor with the format 1080p/30/RG10:

```
v4l2-ctl --set-fmt-video=width=1920,height=1080,pixelformat=RG10 --stream-mmap --stream
```

ISP Support

To enable ISP support

- Built-in to the Camera Core where the release package includes initial ISP configuration files for reference sensors.
- Place the ISP configuration file into RootFS so that when the system boots, it scans the pre-defined folders to search for the ISP configuration files that match the sensor module. This can be performed at runtime to provide more flexibility.

Note:

CSI cameras, with integrated ISP and USB camera, can work in ISP bypass mode. Provided ISP support is available for Jetson Developer Kit (OV5693) RAW camera module.

Document Change History

Document number: DA_07725-001

Version	Date	Authors	Description of Change
•v1.0	•13 Jul 2015	•jbang/msum	•Initial release.
•v2.0	•05 Oct 2015	•jbang/hlang	•Change in document title, other content updates.
•v3.0	•20 Jan 2016	•jbang/gmead	•Adds device tree and other BSP-integration information.
•v4.0	•21 Apr 2016	•gigon bae/ hlang	•Added GNU licenses.
•v4.1	•03 Jun 2016	•fchen/alocsin	•Updated diagrams, added device tree verification code, and equations for property value calculations.

Note:

Hyperlinks in this document are a legacy of the HTML version and may not work in the PDF version.

V4L2 Sensor Driver Programming Guide

Sensor Driver Programming Guide

This topic provides advanced information for developing USB cameras and Bayer and YUV image sensor with NVIDIA[®] Tegra[®] Board Support Package (BSP) releases.

This topic describes the sensor driver architecture required by Tegra[®] platform software, and provides guidance on the implementation of drivers suitable for use with this software release. Implementation of a camera sensor driver will enable acquisition of camera data over the CSI bus, in the native format provided by the sensor, and is intended for use in enabling sensors that contain an ISP (e.g. YUV output).

The programming interfaces for the Tegra[®] ISP and associated image quality tuning are not covered in this document. NVIDIA provides a software implementation for supported camera modules. (For more information, refer to product documentation.)

NVIDIA Tegra BSP supports two types of camera programming paths, depending on the camera and the application.

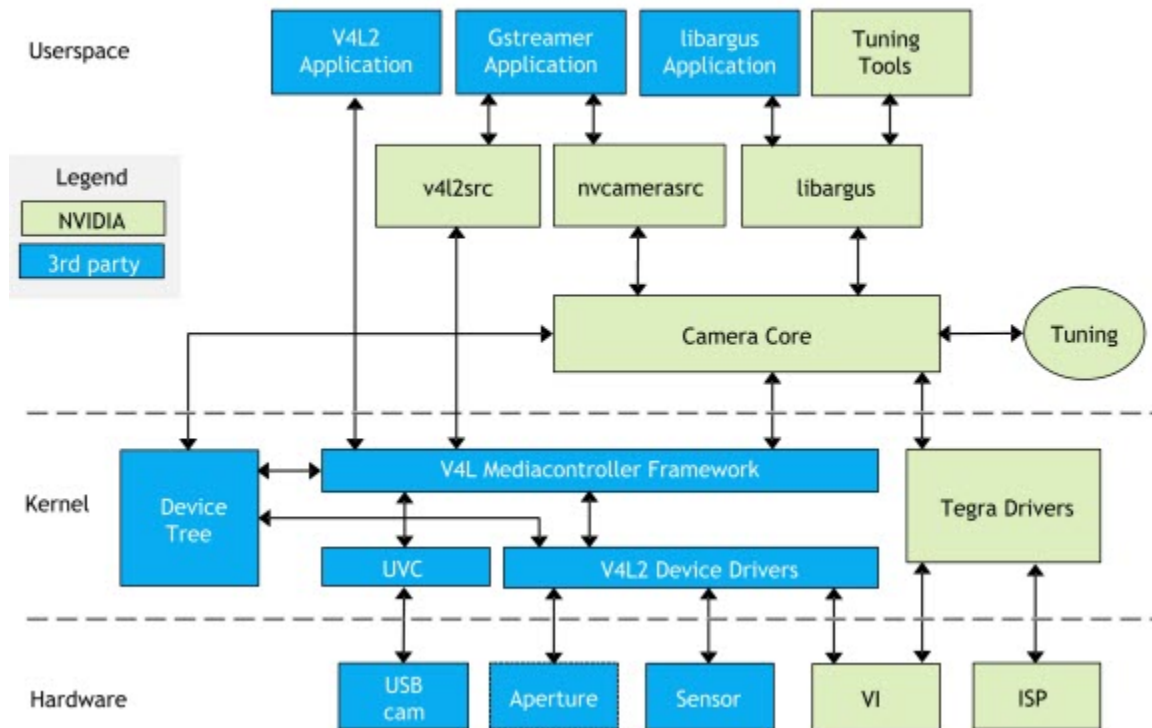
- Scalable Camera Framework
- Direct V4L2 Interface

Scalable Camera Framework

The Scalable Camera Framework (SCF) user mode library provides all the controls and data processing between the application and kernel-mode V4L2 drivers. In applications that use the Tegra ISP functionality, use the Scalable Camera Framework. Typically, for Bayer sensors, applications use the SCF library to convert RGB format to YUV format and to do various post processing tasks.

The following block diagram shows how this framework with the application and kernel mode V4L2 drivers.

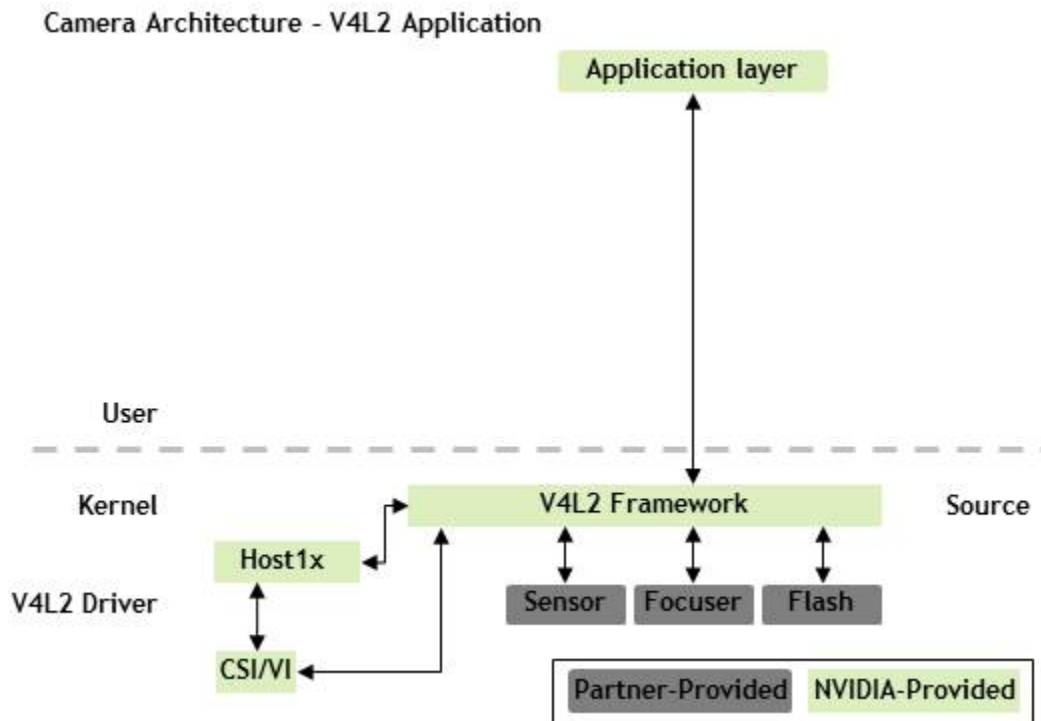
L4T Camera Software Architecture



Direct V4L2 Interface

In applications support a direct V4L2 interface, use this interface to communicate to the NVIDIA V4L2 driver without having to use the SCF library. Use this path for a YUV sensor since this sensor has a built-in ISP and frame does not need extra processing.

The following block diagram shows how an application uses the kernel mode V4L2 drivers:



This chapter describes how to bring up a Bayer sensor with the Tegra BSP. Bringing up the sensor requires customers to develop two things:

- Device Tree in the Linux kernel
- V4L2 sensor driver

Read the following sections to learn how to develop these; our examples use OmniVision OV5693 sensor, and the source code for OV5693 sensor is available to customers.

Camera Modules

A camera module installed on the target platform can consist of one or more devices. A typical rear camera module includes a complementary metal-oxide semiconductor (CMOS) sensor, combined with a Voice Coil Motor (VCM) focuser. A typical front camera module might include a single CMOS sensor only.

To add one or more camera modules to a device tree, find or create a `tegra-camera-platform` device node in the kernel source tree. Usually, that node is in the following directory:

```
arch/arm64/boot/dts
```

In a `tegra-camera-platform` device node, you must create a module table (modules) with one or more modules. Each module must contain its basic information and the definition of the devices that are inside that module.

Note:

All value fields in camera-related device nodes must use the string data type, except for the files that refer to other device nodes.

A typical device-tree node for a camera module looks like this:

```
tegra-camera-platform {
    compatible = "nvidia, tegra-camera-platform";
    modules {
        module0 {
            badge = "e3326_front_P5V27C";
            position = "rear";
            orientation = "1";
            drivernode0 {
                pcl_id = "v4l2_sensor";
                proc-device-tree = "/proc/device-tree/host1x/i2c@546c0000/ov5693_c@36";
            };
            drivernode1 {
                pcl_id = "v4l2_focuser_stub";
            };
        };
    };
};
```

Module Properties

The following table shows the information for `moduleX`: `module` (or `moduleX`: `modules`).

Property	Value
• <code>badge</code>	<ul style="list-style-type: none"> •A unique name that identifies this module. •Guidelines for naming the three parts of the <code>badge_info</code> property: <ul style="list-style-type: none"> • The first part is the camera board ID (<code>camera_board_id</code>) of the module. • The second part contains the position of the module, for example, <code>rear</code> or <code>front</code>. • The third part contains the last six characters of a part number, which you can find in the data sheet on the module from the vendor. •If your system has multiple modules that are identical, create a unique name for each module. •For example, if you have one module for a rear facing camera and an identical module for a front facing camera, you can call the rear camera module <code>e3326_rear_P5V27C</code> and the front camera module <code>e3326_front_P5V27C</code>.
• <code>position</code>	<ul style="list-style-type: none"> •The camera-facing information. The positions supported depends on the number of cameras in the system: <ul style="list-style-type: none"> • Two-camera system: <code>rear</code> and <code>front</code>.

	<ul style="list-style-type: none"> • Three-camera system: bottom, top, and center. • Six-camera system: bottomleft, bottomright, centerleft, centerright, topleft, and topright.
•orientation	<ul style="list-style-type: none"> •The orientation related to the display panel. If there is no display panel, just label each camera with 0-based index. <ul style="list-style-type: none"> • Rear facing: 0 • Front facing: 1
•drivernodeX	<ul style="list-style-type: none"> •The information on the driver node; X: 0-based index.

Driver Properties

The following table shows the information for `drivernodeX: device` (or `drivernodeX: devices`).

Property	Value
•pcl_id	•A unique name that identifies this device.

Individual Imaging Device

An imaging device (a component inside the camera module) can be a sensor, focuser, or flash. Be sure to add all the required information to the device-tree node to support the device operation.

For each device-tree node for a device, assign a device node that contains the name of the device, its slave address, and a compatible string that identifies the node.

Note: Except for the ones that refer to other device nodes, all value fields in camera-related device nodes must use the string data type.

An example device-tree node for the OV5693 V4L2 sensor driver:

```
ov5693_c@36 {
    compatible = "nvidia,ov5693";
    reg = <0x36>;
    devnode = "video2";

    physical_w = "3.674";
    physical_h = "2.738";
    sensor_model = "ov5693";

    avdd-reg = "vana";
    iovdd-reg = "vif";

    mode0 { // OV5693_MODE_2592X1944
        mclk_khz = "24000";
        num_lanes = "2";
```

```

tegra_sinterface = "serial_c";
discontinuous_clk = "no";
dpcm_enable = "false";
cil_settletime = "0";

active_w = "2592";
active_h = "1944";
pixel_t = "bayer_bggr";
readout_orientation = "90";
line_length = "2688";
inherent_gain = "1";
mclk_multiplier = "6.67";
pix_clk_hz = "160000000";

min_gain_val = "1.0";
max_gain_val = "16";
min_hdr_ratio = "1";
max_hdr_ratio = "64";
min_framerate = "1.816577";
max_framerate = "30";
min_exp_time = "34";
max_exp_time = "550385";
Embedded_metadata_height = "0";
};
...
};

```

Device Properties

For the device tree node for the V4L2 sensor-device, define the required hardware resource for the device, as shown in the following table.

Property	Value
•compatible	•Specifies the device identifier. The Linux kernel uses this keyword to bind the device driver to a specific device.
•reg	•Specifies the I2C slave address.
•mclk	•Specifies the name of the input clock for the device. • Note: CAM1_MCLK (GPIO_PS0) and CAM2_MCLK (GPIO_PS1)

	are both from <code>extperiph3_clk</code> , of which <code>clk</code> name must be <code>cam_mclk1</code> (or <code>mclk3</code>).
•XXX-gpio	<ul style="list-style-type: none"> •Specifies the general-purpose input/output (GPIO) pins for the device, where XXX is the GPIO pin for the camera. •There are four default gpio pins for camera: <ul style="list-style-type: none"> • S4 - Camera0 Reset • S7 - Camera0 Power Down • S5 - Camera1 Reset • T0 - Camera1 Power Down
•XXX-supply	<ul style="list-style-type: none"> •Specifies the regulator for the device, where XXX is the actual regulator name defined somewhere else in the device tree. The <code>-supply</code> suffix is mandatory. •The following are the defined regulators: <ul style="list-style-type: none"> • <code>vana-supply = <en_vdd_cam_hv_2v8>; // analog 2.8v</code> • <code>vdig-supply = <en-vdd-cam_1v2>; // digital 1.2v</code> • <code>vif-supply = <en-vdd-cam>; // interface 1.8v</code> • <code>vvcn-supply = <en_vdd_vcm_2v8>; // analog 2.8v for vcm</code>
•XXX-reg	<ul style="list-style-type: none"> •Specifies the name of the regulator, where XXX is the regulator name for the sensor driver. The value of this field is the regulator name with the suffix <code>-reg</code>. •The following are the defined regulators: <ul style="list-style-type: none"> • <code>avdd-reg = "vana";</code> • <code>dvdd-reg = "vdig"</code> • <code>iovdd-reg = "vif"</code> • <code>vcmvdd-reg = "vvcn"</code>
•physical_w	•Specifies the physical width (in millimeters) of the sensor.
•physical_h	•Specifies the physical height (in millimeters) of the sensor.
•sensor_model	Specifies which sensor is in this module
•devnode	Specifies a value used to derive the kernel device node. •/ <code>dev/<devnode></code>
•post_crop_frame_drop	Specifies number of frames to be dropped after applying sensor crop settings.
•use_decibel_gain	When set to true, analog gain value received by driver is expressed in decibels. $dB = 20 * \log(\text{Analog Gain})$

Property-Value Pairs

The following table describes the property-value pairs that apply to the sensor mode for the V4L2 implementation.

Property	Value
•modeX	•Specifies the sensor-mode information, that is, the X: 0-based index.

•ports	•Specifies the media controller graph binding information.
•mclk_khz	•Specifies the standard MIPI driving clock, which is typically 24MHz.
•num_lanes	•Specifies the number of lane channels the sensor is programmed to output.
•tegra_sinterface	•Specifies the base Tegra serial interface to which the lanes are connected.
•discontinuous_clk	•Specifies the indication that the sensor is programmed to use a discontinuous clock on MIPI lanes.
•cil_settletime	•Specifies the value of the settle time of the MIPI lane. A 0 value attempts to auto-calibrate according to the <code>mclk_multiplier</code> parameter.
•dpcm_enable	•Specifies whether to enable dpcm compression for this mode. •Set to <code>true</code> or <code>false</code> .
•active_h	•Specifies the height of the pixel-active region.
•pixel_t	•Specifies the readout pixel pattern of the sensor. •The following examples shows the values for Bayer sensors with the “bggr” pixel pattern: <ul style="list-style-type: none"> • 10 bit: <code>bayer_bggr</code> • 12 bit: <code>bayer_bggr12</code> • 14 bit: <code>bayer_bggr14</code>
•readout_orientation	•Specifies the readout orientation that is based on the orientation of the camera module. Change this property if you would like to program a different readout order for this mode. •Possible values: <ul style="list-style-type: none"> • 0 • 90 • 180 • 270
•line_length	•Specifies the pixel line length (width) for sensor mode for calibrating the features in our camera stack. •This value must be equal or larger than <code>active_w</code> .
•mclk_multiplier	•Specifies the multiplier to MCLK for timing the capture sequence of the hardware. Use the following equation to calculate this value: • <code>mclk_multiplier = desired ISP clock frequency / mclk</code> . •This value must be larger than <code>pixel_clk_hz / mclk</code> to prevent ISP underrun.
•pix_clk_hz	•Specifies the sensor pixel clock for calculating the exposure, frame rate, and so forth.

	<ul style="list-style-type: none"> This value is calculated based on input clock (<code>mclk</code>) and PLL settings from sensor mode table. Please refer to sensor data sheet for how to calculate this value.
• <code>inherent_gain</code>	<ul style="list-style-type: none"> Specifies the gain obtained inherently from the mode, that is, pixel binning. Set to 1 if you do not know this value.
• <code>min_gain_val</code>	<ul style="list-style-type: none"> Specifies the minimum gain limit for the mode.
• <code>max_gain_val</code>	<ul style="list-style-type: none"> Specifies the maximum gain limit for the mode. This can be increased to include digital gain, if that is supported by the sensor.
• <code>min_exp_time</code>	<ul style="list-style-type: none"> Specifies the minimum exposure time limit for the mode in microseconds.
• <code>max_exp_time</code>	<ul style="list-style-type: none"> Specifies the maximum exposure time limit for the mode in microseconds.
• <code>min_hdr_ratio</code>	<ul style="list-style-type: none"> Specifies the minimum high-dynamic-range (HDR) ratio limit for the mode (for interleave HDR sensors). For non-HDR sensors, set <code>min_hdr_ratio</code> to an empty string.
• <code>max_hdr_ratio</code>	<ul style="list-style-type: none"> Specifies the maximum HDR ratio limit for the mode (for HDR sensors).
• <code>min_framerate</code>	<ul style="list-style-type: none"> Specifies the minimum frame-rate limit for the mode in frames per second (fps). Use the following equation to calculate this value: <ul style="list-style-type: none"> $\text{min_framerate} = \text{pix_clk_hz} / (\text{line_length} * \text{maximum frame length})$
• <code>max_framerate</code>	<ul style="list-style-type: none"> Specifies the maximum frame-rate limit for the mode in fps. Use the following equation to calculate this value: <ul style="list-style-type: none"> $\text{max_framerate} = \text{pix_clk_hz} / (\text{line_length} * \text{minimum frame length})$
• <code>embedded_metadata_height</code>	<ul style="list-style-type: none"> Specifies how many extra embedded metadata rows for each frame. Set to 0 to disable embedded metadata support.

Example Focuser Driver Properties

The following table shows the required information for the example of the LC898212 focuser driver.

```
lc898212@72 {
    compatible = "nvidia,lc898212";
    reg = <0x72>;

    devnode = "video6";
```

```

type = "default";
ports {
    #address-cells = <1>;
    #size-cells = <0>;
    port@0 {
        reg = <0>;
        lc898212_out0: endpoint {
            remote-endpoint = <&vi_in1>;
        };
    };
};
};
};

```

The following table describes the focuser driver properties.

Property	Value
•compatible	•Specifies the device identifier.
•reg	•Specifies the I2C slave address.
•type	•Specifies the focuser type: <ul style="list-style-type: none"> •• default: VCM focuser. •• steppermotor: Stepper motor focuser.
•ports	•Media controller graph binding info.

V4L2 Kernel Driver

The content of this chapter is based on the Video for Linux 2 (V4L2) driver for the OmniVision OV5693 sensor (`ov5693.c`) at:

```
kernel/drivers/media/i2c/ov5693.c
```

NVIDIA suggests that you look at that source before reading this the rest of this chapter.

Macro Definitions

Following are the sensor-specific macro values:

- The minimum and maximum values for each control.
- The default value for each control.
- The macro values that is required for sensor timing or general functionality.

Sensor-Private Data

The following structure contains the private data that are specific to the sensor:

```

struct ov5693 {
    struct camera_common_power_rail power;
    int numctrls;
    struct v4l2_ctrl_handler ctrl_handler;
    struct camera_common_eeprom_data eeprom[OV5693_EEPROM_NUM_BLOCKS];
    u8 eeprom_buf[OV5693_EEPROM_SIZE];
    struct i2c_client *i2c_client;
    struct v4l2_subdev *subdev;
    struct media_pad pad;
    s32 group_hold_prev;
    bool group_hold_en;
    struct regmap *regmap;
    struct camera_common_data *s_data;
    struct camera_common_pdata *pdata;
    struct v4l2_ctrl *ctrls[];
};

```

The following table describes the sensor properties.

Property	Value
•power	•Holds generic power controls, include regulators, clks, and GPIOs.
•numctrls	•Holds the number of v4l2 controls for the sensor.
•ctrl_handler	•Holds the required v4l2 handler to controls, needed for v4l2 control init.
•i2c_client	•Holds a handle to i2c_client, used to access to sensor i2c client instance.
•subdev	•Holds a handle for v4l2 sub-device, needed to run subdev operations (ops).
•eeprom	•Holds common EEPROM device data.
•eeprom_buf	•Holds EEPROM buffer storage.
•pad	•Holds media pad used for media controller initialization for a device to work as SINK or SOURCE.

• <code>group_hold_prev</code>	•Holds previous state use by group hold control handler to check for change of state.
• <code>group_hold_en</code>	•Holds group hold enable flag directly related to group hold control.
• <code>reg_map</code>	•Holds a register map setup for I2C read and write.
• <code>s_data</code>	•Holds a handle to common data, see documentation for <code>camera_common_data</code> .
• <code>pdata</code>	•Holds a handle to common platform data, populated by read Device Tree.
• <code>ctrls</code>	•Holds handles to initialized v4l2 controls, dynamic array, MUST BE LAST FIELD.

Configuring Regmap

Depending on the I2C interface of the sensor, you should update the values of `reg_bits` and `val_bits`.

```
regmap_config {
    reg_bits = 16;
    val_bits = 8;
};
```

The following table describes the `regmap_config` properties.

Property	Value
• <code>reg_bits</code>	•Specifies the number of bits needed to represent the I2C register offset.
• <code>val_bits</code>	•Specifies the number of bits in the buffer to store data to be transferred over I2C.

Check the vendor register programming table of your sensor to determine the size of `reg_bits` and `val_bits`.

Configuring Controls

To link the controls to their control handlers, set up the function pointers:

- Point `g_volatile_ctrl` to the internal `get volatile` control handler of the sensor.
- Point `s_ctrl` to the internal `set` control handler of the sensor.

```
static const struct v4l2_ctrl_ops ov5693_ctrl_ops = {
    .g_volatile_ctrl = ov5693_g_volatile_ctrl,
    .s_ctrl = ov5693_s_ctrl,
};
```

The following code example lists the controls and their initialized values. This list is looped through during the

`ctrls_init` call to initialize each of the controls. Each control is then accessible through the `ctrls` handler in the private data. The set of controls defined for OV5693 are the standard ones used by the user-mode PCL V4L2 driver.

Note: Additional controls require a change in the PCL user mode driver.

Three types of controls are defined for the OV5693 sensor.

```
static struct v4l2_ctrl_config ctrl_config_list[] = {
/* Integer Control: setting integer values such as gain, coarse
 * time, *and frame length.
 */
{
    .ops = &ov5693_ctrl_ops, //pointer to control ops
                          //function

    .id = V4L2_CID_GAIN,    //id, defined in
                          //camera_common.h

    .name = "Gain",        //string name of control
    .type = V4L2_CTRL_TYPE_INTEGER, //type of control
    .flags = V4L2_CTRL_FLAG_SLIDER, //control flags

    // The following three are the value most likely
    // needs changing
    .min = OV5693_MIN_GAIN, //control value lower bound
    .max = OV5693_MAX_GAIN, //control value upper bound
    .def = OV5693_DEFAULT_GAIN, //default control value
    .step = 1,                //increment step size for
                          //value
},
...
/* String Data Control: converts data to string format then
 * send to PCL user mode
 * drivers, used for EEPROM, OTP, and fuse id.
 */
{
    .ops = &ov5693_ctrl_ops,
```

```

        .id = V4L2_CID_EEPROM_DATA,
        .name = "EEPROM Data",
        .type = V4L2_CTRL_TYPE_STRING,
        .flags = V4L2_CTRL_FLAG_VOLATILE,
        .min = 0,

        /* the following one is the value that likely needs
         * changing, the string size is 2 times actual
         * buffer size
         */
        .max = OV5693_EEPROM_STR_SIZE,
        .step = 2,
    },
...

/* Menu Control: used as on/off switch for group hold and HDR.
 * switch_ctrl_qmenu is used to define the states on/off there
 * shouldn't be a need to change these controls, unless a
 * completely new one is being added.
 */
{
    .ops = &ov5693_ctrl_ops,
    .id = V4L2_CID_GROUP_HOLD,
    .name = "Group Hold",
    .type = V4L2_CTRL_TYPE_INTEGER_MENU,
    .min = 0,
    .max = ARRAY_SIZE(switch_ctrl_qmenu) - 1,
    .menu_skip_mask = 0,
    .def = 0,
    .qmenu_int = switch_ctrl_qmenu,
},
};

```

Setting Up Control Registers

Set up register writes for integer controls with the following functions. `addr` depends on the sensor being used; `val` is the source from the control handler. These functions are called by each control handler to set up register writes for each of the controls.

- `ov5693_get_frame_length_regs`

- `ov5693_get_coarse_time_regs`
- `ov5693_get_coarse_time_short_regs`
- `ov5693_get_gain_reg`
- `ov5693_get_gain_short_reg`

Read-Write Wrapper in the Register

The following functions are wrappers for the read-write interface of the I2C register. For OV5693, use the `regmap` interface. However, you can modify these functions for other interfaces.

- `ov5693_read_reg`
- `ov5693_write_reg`
- `ov5693_write_table`

Power Functions

The following table describes the functions for power-related controls.

Function	Description
• <code>ov5693_power_on</code>	<ul style="list-style-type: none"> •Contains the power-on sequence. You must modify this function according to the specification sheets. •Calls <code>regulator_enable()</code> to turn on regulators and use <code>gpio_set_value()</code> to toggle the GPIO pins.
• <code>ov5693_power_off</code>	<ul style="list-style-type: none"> •Contains the power-off sequence. You must modify this function according to the specification sheets. •Calls <code>regulator_enable()</code> to turn on regulators and use <code>gpio_set_value()</code> to toggle the GPIO pins.
• <code>ov5693_power_put</code>	•Calls <code>regulator_put()</code> on all regulators.
• <code>ov5693_power_get</code>	•Calls <code>camera_common_regulator_get()</code> helper function or <code>regulator_get()</code> to acquire all regulators.

Setting Up V4L2 Subdevice and Camera Common

The `ov5693_s_stream` function is mainly for writing mode tables by making calls to register the `write_table` function. You set up mode tables in the `ov5693_mode_tbls.h` file. For details, see [Mode Tables](#) in this chapter.

In addition to writing mode tables and enabling the stream through stream-enable register writes, `ov5693_s_stream` also writes the initial integer-control values to the register through direct calls to the integer-control handlers. For details, see the section [Control Handlers](#).

```
control.id = V4L2_CID_GAIN;
err = v4l2_g_ctrl(&priv->ctrl_handler, &control);
err |= ov5693_set_gain(priv, control.value);
```

.....

If a test pattern is supported by the sensor and you can create a register table for that pattern, you can add a `test_mode` flag to write the test-mode table here.

Camera common is a set of functions that are common to camera drivers of the NVIDIA kernel, to which this driver refers. Camera common sets up most of the V4L2 framework, requiring linkage from the driver in the form of the following:

```
struct camera_common_sensor_ops
struct camera_common_data
```

For details on modifying and adding new modes, see [Mode Tables](#) in this chapter.

The following code snippets set up the V4L2 subdevices and camera common for registering the sensor driver with the V4L2 framework. The `s_stream` pointer must point to the internal `s_stream` function of the sensor; you can leave the other pointers as is.

```
static struct v4l2_subdev_video_ops ov5693_subdev_video_ops = {
    .s_stream = ov5693_s_stream,
    ...
};
```

You need not modify this structure:

```
static struct v4l2_subdev_core_ops ov5693_subdev_core_ops = {
    ...
};
```

Match the name for the pointer for the `core` and `video` operations function with the two from above, as follows:

```
static struct v4l2_subdev_ops ov5693_subdev_ops = {
    .core = &ov5693_subdev_core_ops,
    .video = &ov5693_subdev_video_ops,
};
```

During the parsing of the device tree, `of_device_id` matches the `compatible` field with the one in the Device Tree.

```
static struct of_device_id ov5693_of_match[] = {
    { .compatible = "nvidia,ov5693", },
    { },
};
```

This structure is required for camera common and you must set up the function pointers appropriately. Link the `power_on`, `power_off`, `write_reg`, and `read_reg` functions here:

```
static struct camera_common_sensor_ops ov5693_common_ops = {
    .power_on = ov5693_power_on,
```

```

.power_off = ov5693_power_off,
.write_reg = ov5693_write_reg,
.read_reg = ov5693_read_reg,
};

```

Control Handlers

This section describes the control handlers.

Set Control

The two subsections describe the handlers for set control.

V4L2 Set-Control Operation

The V4L2 set-control function contains a `switch` statement to redirect set-control calls to their appropriate control handlers.

Note: Read-only controls, such as `fuse_id` and One-Time Programmable Read-Only Memory (OTP ROM), do not have a case statement in the control ID `switch` statement.

```

ov5693_s_ctrl {
    ...
    switch (ctrl->id) {
        case V4L2_CID_GAIN:
            ...
        case V4L2_CID_EEPROM_DATA:
            ...
        case V4L2_CID_HDR_EN:
            ...
    }
    ...
}

```

Note: For `EEPROM_DATA`, the string control must have a pre-allocated string passed in. Hence, a null check is necessary before passing the string to the control handler.

HDR_EN is a pure software control. No control handler writes to the hardware so simply break out of the switch statement.

Setter-Control Handlers (Writes)

Setter-control handlers are the control handlers called by `s_ctrl`. They perform additional control-handling operations, such as writes to registers. The majority of these controls make calls to the control register setup functions (see the section [Setting Up Control Registers](#)). The exception is `write_eeprom`, which acts as a separate I2C device with its own I2C write interface.

- `ov5693_set_group_hold`
- `ov5693_set_gain`
- `ov5693_set_frame_length`
- `ov5693_set_coarse_time`
- `ov5693_set_coarse_time_short`
- `ov5693_write_eeprom`

Note a special case for the gain-control handler, which contains a function call to:

```
ov5693_calculate_gain(val, OV5693_GAIN_SHIFT);
```

The function takes a binary-coded decimal from user space as input and computes the gain-register value according to the vendor-provided gain-calculation formula. Because that formula can vary from sensor to sensor, you must rewrite this function in the V4L2 sensor driver for each of the sensors.

The input of this function is the gain value passed in from user space. The value is a binary-coded decimal ranging from 1 to 16. The binary-coded decimal is divided into a six-byte integer representation and a two-byte decimal representation. The most significant six bytes of `val` is the integer representation. The least significant two bytes is the decimal representation, which is actually the numerator of a fraction over the maximum decimal representation of 0xFF.

The output of this function is the actual gain-register value programmed over I2C. That value depends on the gain-calculation formula provided by the sensor vendor (usually found in the data sheets for the sensor). The goal of this function is to convert the decimal `val` input into the gain-register value with as little truncation as possible.

For the OV5693 formula on which the `ov5693_calculate_gain` (or `to_gain`) function is based, see the *OV5693 Software Reference Manual*.

Exposure/Framerate Controls

Sensor drivers controlling sensor exposure and frame rate must implement `V4L2_CID_FRAME_LENGTH` and `V4L2_CID_COARSE_TIME` controls. These controls are used to update sensor frame length and coarse integration time registers with the values calculated from user-mode camera stack.

If the sensor driver has special requirements or direct calculation of the exposure and frame rate is required, `V4L2_CID_FRAME_RATE` and `V4L2_CID_EXPOSURE` controls can be implemented instead. When these two controls are implemented, they receive unprocessed frame rate and exposure values in fixed point format. The driver must perform any necessary determination of final register settings.

Note:

The `V4L2_CID_FRAME_LENGTH` and `V4L2_CID_COARSE_TIME` controls take 32-bit input parameters and `V4L2_CID_FRAME_RATE` and `V4L2_CID_EXPOSURE` take 64-bit input

parameters. Use either set of controls but do not mix or use both sets at once.

Fixed Point Format

When the driver implements `V4L2_CID_FRAME_RATE` and `V4L2_CID_EXPOSURE` controls, the received control values are in Q10.22 fixed point format. In Q10.22 fixed point format the upper 10 bits are for storing integer values and the lower 22 bits are for fraction values. This is a workaround for lacking floating point in the kernel.

This formatting is achieved by multiplying the original value with $1 \ll 22$. For your convenience, a `FIXED_POINT_SCALING_FACTOR` value is defined.

```
#define FIXED_POINT_SCALING_FACTOR (1ULL << 22)
```

Caution should be taken before using these fixed point format values. The value must be divided back to original precision for the calculation. For best results, multiply first, and then divide the value with `FIXED_POINT_SCALING_FACTOR` last, to avoid unexpected truncation/rounding issues.

Get-Volatile Control

Following are the get-volatile controls:

- **V4L2 get-volatile control operation:** The `ov5693_g_volatile_ctrl` function contains a switch statement for redirecting get-control calls to their appropriate control handlers.

Note: For non-volatile controls, get-control simply returns the previously written value stored in the control handler.

- **Get-control handler (reads):** The `ov5693_read_eeprom` control handler is an example of a get-volatile control handler. This handler reads the volatile value directly from the EEPROM register and updates the value that is read back every time get is called on this control.

Other Control-Related Functions

Following are two other control-related functions:

- **EEPROM device-related controls:** These two functions are for setting up EEPROM as a separate I2C device with its own `regmap` interface:

```
ov5693_eeprom_device_init
ov5693_eeprom_device_release
```

- **Handlers called only on `ctrls_init`:** These control handlers are called only once from the `ctrls_init` function (see the section [Boot-Time Initialization](#)). That is because OTP and `fuse_id` controls are read-only and their values only need to be read once during boot time.

Boot-Time Initialization

This section describes the functions for initializing boot time.

Control Initialization

The `ov5693_ctrls_init` function iterates through `ctrl_config_list` (see [Configuration Controls](#)) and

registers each control as a new custom control with the V4L2 framework. `s_ctrl` is also called for each control to set it to its default value defined in `ctrl_config_list`.

You need not modify the function except for the calls for initializing the value for the read-only controls. See the calls to `otp_setup` and `fuse_id_setup` in the section [Control Handlers](#).

```
ov5693_ctrls_init {
    ...
    err = ov5693_otp_setup(priv);
    ...
    err = ov5693_fuse_id_setup(priv);
    ...
}
```

Device Tree Parser

The `ov5693_parse_dt` function, which parses device trees, takes the Device Tree node and looks for the parameters (according to the sensor-related private data) required by the sensor driver.

The values include but are not limited to the following:

```
mclk
pwn-gpios
reset-gpios
af-gpios
avdd-reg
dvdd-reg
iovdd-reg
```

For details on how to set up device trees with the appropriate values, see the documentation. You must match the name list with the names in the Device Tree for their respective name-value pairs.

Port binding

```
vi {
    ports {
        port@0 {
            reg = <0>;
            vi_in0: endpoint {
                bus-width = <2>;
                remote-endpoint = <&ov5693_out0>;
            };
        };
    };
};
```



```

ov5693_c@36 {
    ports {
        port@0 {
            reg = <0>;
            ov5693_out0: endpoint {
                csi-port = <2>;
                bus-width = <2>;
                remote-endpoint = <&vi_in0>;
            };
        };
    };
};

```

The following table describes the port binding properties.

Function	Description
•port	•Specifies the mediapad that the port is connected. All imager devices will have only one media pad for binding the connection with VI.
•csi-port	•This value defines the CSI port sensor is connected. For imager devices like focuser or flash this field is not required.
•bus-width	•Bus width defines the number of CSI lanes connected to sensor.
•remote-endpoint	•Remote end point is the label used for binding two ports. The binding expects one port is the sink and the other one is the source.

Media Controller Setup

A few additional components are needed to setup the media controller for V4L2 use. The open call is a placeholder operation for satisfying the V4L2 sub device internal operation requirements. The setup likely will not need to be change besides a name changed to match the devices.

```
ov5693_open
```

Sub device function ops need to link to the open operation:

```

static const struct v4l2_subdev_internal_ops ov5693_subdev_internal_ops = {
    .open = ov5693_open,
};

```

Media entity function ops need to link to the V4L2 sub device link validation method:

```
static const struct media_entity_operations ov5693_media_ops = {
    .link_validate = v4l2_subdev_link_validate,
};
```

Sensor-Driver Probing

The following subsections explain the probe function of the sensor driver.

Entry Point for Initialization During Boot

The `ov5693_probe` function is the entry point of the driver. The function starts off by allocating memory for common data and sensor-private data (see the section [Sensor-Private Data](#)).

```
common_data = devm_kzalloc(&client->dev,
    sizeof(struct camera_common_data), GFP_KERNEL);

priv = devm_kzalloc(&client->dev,
    sizeof(struct ov5693) + sizeof(struct v4l2_ctrl *) *
    ARRAY_SIZE(ctrl_config_list),
    GFP_KERNEL);
```

The function then initializes `regmap`:

```
priv->regmap = devm_regmap_init_i2c(client, &sensor_regmap_config);
```

Next, the function calls the other initialization functions, including the following:

```
// See the section Parser for Device Trees: stored in private
// data pdata field.
ov5693_parse_dt(client);

// See the section Power Functions.
ov5693_power_get(priv);

// Link the appropriate subdev ops (see the section
// Setting Up V4L2 Subdevice and Camera Common.
v4l2_i2c_subdev_init(&common_data->subdev, client, &subdev_ops);

// See the section Initialization of Controls.
ov5693_ctrls_init(priv);

// See the section Other Control-Related Functions.
```

```
ov5693_eeprom_device_init(priv);
```

Next, the function links the common data and sensor-private values:

```
// Link to ops. See Setting Up V4L2 Subdevice and Camera Common.
common_data->ops = &ov5693_common_ops;

// Control handler linking
common_data->ctrl_handler = &priv->ctrl_handler;

// I2C client passed in to probe
common_data->i2c_client = client;

// Default to frame format 0. See Mode Tables in this chapter.
common_data->frmfmt = &ov5693_frmfmt[0];

// Based on default data format definition, generally defaults
// to the same format. See the section Macro Definitions.
common_data->colorfmt = camera_common_find_datafmt(
    OV5693_DEFAULT_DATAFMT);

// Power-handler linking
common_data->power = &priv->power;

// Sensor-private data linking
common_data->priv = (void *)priv;

// Number of format is frame format. See the section Macro Definitions.
common_data->numfmts = ARRAY_SIZE(ov5693_frmfmt);

// Set up of port information for device
camera_common_parse_ports(...)
// Port information is also used to create the name for debugfs node setup
sprintf(debugfs_name, "ov5693_%c", common_data->csi_port + 'a');
camera_common_create_debugfs(common_data, debugfs_name);
```

Setup of Default Values for Common Data

See the section [Macro Definitions](#).

```

common_data->def_mode = OV5693_DEFAULT_MODE;
common_data->def_width = OV5693_DEFAULT_WIDTH;
common_data->def_height = OV5693_DEFAULT_HEIGHT;
common_data->def_clk_freq = OV5693_DEFAULT_CLK_FREQ;
debugfs_name

// I2C client passed in to probe
priv->i2c_client = client;

// Link to common data above
priv->s_data = common_data;

// Link to V4L2 subdevice handler
priv->subdev = &common_data->subdev;
// Link to subdevice dev to i2c_client dev (for media controller usage)
priv->subdev->dev = &client->dev;

// Initialize previous group hold state to 0
priv->group_hold_prev = 0;

```

Setup of Media Controller

All imager devices must register themselves as sub devices to media controller framework. VI acts as the master device which controls the binding, parsing the DT and establishes the media links once all the sub devices are registered.

Each sub device can register to media controller framework by defining the entity and pads information.

Entity: Media entity is the unit device represented by media controller framework for establishing connections. Each media entity can have multiple pads. Framework provides a list of known entity types and the corresponding media operations.

Pad: Pad represents the device behaves as SINK or SOURCE port. Imager device must have a SOURCE port which is bound to VI. If imager device has a SINK port then the SOURCE port which binds the device SINK port must be represented in DT to complete the binding.

```

// Link subdevice internal and media entity operations
priv->subdev->internal_ops = &ov5693_subdev_internal_ops;
priv->subdev->entity.ops = &ov5693_media_ops;

```

```
// Setup subdevice flags and media entity type
priv->subdev->flags |= V4L2_SUBDEV_FL_HAS_DEVNODE | V4L2_SUBDEV_FL_HAS_EVENTS;
priv->subdev->entity.type = MEDIA_ENT_T_V4L2_SUBDEV_SENSOR;

// Setup media controller pad flags
priv->pad.flags = MEDIA_PAD_FL_SOURCE;

// Initialize and register subdevice and media entity with media controller framework.
media_entity_init(&priv->subdev->entity, 1, &priv->pad, 0);
v4l2_async_register_subdev(priv->subdev);
```

Removing Sensor Drivers

The `ov5693_remove` function removes the sensor-device instance and calls on a device shutdown.

This function makes calls to various `free`, `put`, and `destroy` functions that match up with several calls in `probe`. It must remain largely the same.

```
// See the section Initialization of Controls for details on
// freeing the control handler.
v4l2_ctrl_handler_free(&priv->ctrl_handler);

// See the section Power Functions.
ov5693_power_put(priv);
```

Device Registration

After driver development is complete, you must add your device information to the system kernel device tree so it can be instantiated when the kernel boots. There are two ways of registering your device.

Using Plugin Manager

If your camera module has onboard EEPROM and has a valid camera ID programmed, Plugin Manager can be used. If your device module does not meet this requirement, use the main platform device tree instead by following the directions in [Using Main Platform Device Tree File](#).

Plugin Manager automatically links devices when the system kernel boots. Plugin Manager uses the camera module ID returned by the boot loader to update the device tree entries with proper information at runtime. Plugin Manager allows a single system image to support multiple camera devices. To change camera modules, power down the device, replace the camera module, and then reboot. The new module works automatically.

For Plugin Manager support, add your DTSI file to the camera configuration DTSI file, and then update the camera plugin manager DTSI with proper override information.

Prerequisites

- You have obtained the kernel source files. For more information, see [Synchronizing the Kernel in Getting Started](#).

To enable Plugin Manager support

- Add your device DTSI to the camera lists in:

```
kernel/arch/arm64/boot/dts/tegra210-platforms/tegra210-jetson-cv-camera-modules.dtsi
```

- Add the override information to:

```
kernel/arch/arm64/boot/dts/tegra210-plugin-manager/tegra210-jetson-cv-camera-plugin-man
```

Using Main Platform Device Tree File

Register your new device by updating main platform DTSI file to include your new device DTSI file. Because Tegra uses Plugin Manager by default, you must first unregister Plugin Manager support first and then add your device information to the main device tree DTSI file.

Prerequisites

- You have obtained the kernel source files. For more information, see [Synchronizing the Kernel in Getting Started](#).

To register a device using main-platform device tree files

- Locate and edit the following file:

```
kernel/arch/arm64/boot/dts/tegra210-plugin-manager/tegra210-jetson-cv-plugin-manager.dt
```

- In the DTSI, remove the following line:

```
#include "tegra210-jetson-cv-camera-plugin-manager.dtsi"
```

- Locate and edit the following file:

```
kernel/arch/arm64/boot/dts/tegra210-jetson-cv-base-p2597-2180-a00.dts
```

- In the DTS file, replace the following line:

```
#include "tegra210-platforms/tegra210-jetson-cv-camera-modules.dtsi"
```

With an `#include` statement specifying the DTSI file for your new device.

Note:

Use `tegra210-jetson-cv-camera-modules.dtsi` as a model for generating your DTSI. In your file, change `status = disable` to `status = okay`.

How to Verify the V4L2 Sensor Driver

When the driver is complete, run `v4l2-compliance` and `v4l2-ctl` to ensure that your driver can capture raw data through the V4L2 interface.

To run a v4l2-compliance test

- Enter the command:

```
v4l2-compliance -d /dev/video0
```

Specify the device node path with option `-d`.

- If you have a single camera, the device path is `/dev/video0`.
- If you have multiple cameras, the number after `video` indicates the index of one of the camera available in the system. The program will use this camera for the test.

To run a v4l2-ctl test

- Enter the command:

```
v4l2-ctl --set-fmt-video=width=2592,height=1944,pixelformat=RG10 --stream-mmap --stream
```

You can use `v4l2-ctl` to capture RAW data. You must provide appropriate parameter values for the camera driver you are using.

Both `v4l2-compliance` and `v4l2-ctl` are available as open source projects. Documentation for these commands is available from LinuxTV at:

<https://www.linuxtv.org/wiki/index.php/V4l-utils>

Debugging Tips

When the driver is complete, before running any camera applications, you should first check to see if the camera device node(s) got populated correctly in the system.

If the driver is loaded properly, you should see a `/dev/<video#>` node. The device node is the file I/O interface with which the user-space driver accesses the driver.

Problems might occur in the probing process and require debugging. Typically, problems occur in the clock, GPIO, and regulator setup. See below for a few tips.

Verify that driver name matches the name in the Device Tree

Your device name should be the same in both the Device Tree and your driver. This is how kernel binds the driver to your device.

In Device Tree:

```
compatible = "nvidia,ov5693";
```

In driver:

```
static struct of_device_id ov5693_of_match[] = {
    { .compatible = "nvidia,ov5693", },
    { },
};
...
.driver = {
    .name = "ov5693",
    .owner = THIS_MODULE,
    .of_match_table = of_match_ptr(ov5693_of_match),
},
```

Verify that all names match the Device Tree

To debug those problems, first verify that the names that are being read through the `parse_dt` function matches those in the Device Tree.

Verify that the Device-Tree values match the hardware

Ensure that the values assigned to the Device Tree fields match the hardware they describe. Oftentimes, regulator names, GPIO numbers, and clock names are out-of-date in the Device Tree, causing probing to fail.

Verify that functions run to completion

Ensure that the `power_get` calls runs to completion. For details, see the sections [Power Functions](#) and [Boot-Time Initialization](#).

Another common problem that occurs during probe is in `control init`.

Verify that default values are correctly linked

Verify that default values are all linked correctly in the control configuration (see the section [Sensor-Private Data](#)) and that the default macros (see the section [Macro Definitions](#)) are updated appropriately with the values in the mode table of the sensor.

If new controls have been added to the control configuration list (on rare occasions), ensure that they contain the appropriate control handlers and default values.

After probing succeeds, problems could still occur with the control setting. A common problem is in the values of the register writes.

Verify that control-register values are correct

Ensure that the control-register setup (see the section [Setting Up Control Registers](#)) contains the correct register address and formats according to the mode table and data sheets.

For gain control, be sure to create a new `calculate_gain` (`to_gain`) function for calculating the gain value according to the gain formula in the data sheets.

Verify that mode-specific settings are correct

Double-check the mode-specific settings in the Device Tree. Update them if necessary. The most common issues that cause sensor timeout are due to the wrong values for the following settings:

- The value for `num_lanes` must be 1, 2, or 4, depending on your sensor configuration.
- The `tegra_sinterface` setting must be set to the proper CSI port where the sensor is connected.
- The `discontinuous_clk` setting specifies for Tegra that the sensor is running in continuous clock mode (free running), or discontinuous clock mode (gated). If unsure, first set this value to “no”.
- Set the `mclk_multiplier` to equal or larger than `pix_clk_hz / mclk` to make sure the ISP is running fast enough to process the data from sensor.

Verify that i2c accesses are working properly

Look closely in kernel logs for any i2c errors for the camera sensor. If there is an i2c error, check the power-on function to make sure the sensor power-on sequence is correct. The order of enabling power, clock, GPIO, and timing must comply with sensor specifications.

Mode Tables

This topic explains mode tables and describes the procedure for adding them.

The modes tables of the register reside in a separate header file called `sensor_mode_tbls.h` and are included by the main driver. Those tables can be a list of `reg_8` or `reg_16` address-value pairs. Mode tables are separated by resolution. For the `ov5693` driver, a common mode table also exists, which contains the common register values among all the resolutions.

```
static ov5693_reg mode_table_common[] = {
    ...
};
```

The `start` and `stop` stream-register values must be in a separate mode table. When you separate them, delete the `start` stream-register values at the end of the resolution mode tables.

```
static ov5693_reg ov5693_start[] = {
    { 0x0100, 0x01 }, /* mode select streaming on */
    { OV5693_TABLE_END, 0x00 }
};

static ov5693_reg ov5693_stop[] = {
    { 0x0100, 0x00 }, /* mode select streaming off */
    { OV5693_TABLE_END, 0x00 }
};
```

Also, a table for the color bars of the test pattern is used if the `test_mode` flag is enabled, activating the test-pattern mode of the sensor. That table is required only if test pattern is supported by the sensor.

```
static ov5693_reg tp_colorbars[] = {
```

```

    ...
};

```

At the end of the header of the mode tables is an enumeration of all the mode tables, as well as a list that maps the enumeration to table pointers.

```

enum {
    OV5693_MODE_4096X3072,
    ...
};
static ov5693_reg *mode_table[] = {
    [OV5693_MODE_2593X1944] = mode_2593X1944,
    ...
};

```

The `camera_common_frmfmt` array list is a required table that sets up the format for the V4L2 framework. Each of the elements on that list contains the resolutions, the `is_hdr` flag, and the enumeration for the mode.

```

static const struct camera_common_frmfmt ov5693_frmfmt[] = {
    {{2592, 1944}, 0, OV5693_MODE_2593X1944},
    ...
};

```

Adding a register mode table is a relatively simple and straightforward process.

To add a register mode table

1. Obtain the address-value pairs for the mode table you would like to add.

Note: If separate stream-on and stream-off mode tables exist, you can omit them from the mode table.

2. Format the pairs according to the register structure and initialize a static array with the mode resolution as part of the name. For example:

```

static ov5693_reg mode_####x####[] = {
    { 0x0100, 0x01 }, /* mode select streaming on */
    ...
    /* your addr and val pairs */
};

```

The mode table is now in place. Be sure to end the table array with the following:

```

{ OV5693_TABLE_END, 0x00 }

```

3. Create a new enumeration in the list of enumerations and add it to the array of mode tables:

```

enum {
    ...

```

```

    OV5693_MODE_####X####,
}

static ov5693_reg *mode_table[] = {
    ...
    [OV5693_MODE_####X####] = mode_####x####,
};

```

4. Add the new mode to the `camera_common_frmfmt` array:

```

static const struct camera_common_frmfmt ov5693_frmfmt[] = {
    ...
    {{####, ####}, 0, OV5693_MODE_####X####},
};

```

Open Source Software Licenses

This topic provides license information about the open source software libraries included in this NVIDIA product.

GNU General Public License 2.0

mtd-utils

mtd-utils is a collection of utilities licensed under GPLv2 and copyrighted by various individuals. Vibrante Android uses the following utilities from this collection.

mkfs

Copyright (C) 2008 Nokia Corporation.

Copyright (C) 2008 University of Szeged, Hungary

ubinize

Copyright (C) 2008 Nokia Corporation

Copyright (c) International Business Machines Corp., 2006

FreeType 6 version 2.4.12

Copyright 1996-2002, 2006 by David Turner, Robert Wilhelm, and Werner Lemberg.

ddccontrol version 0.2

Copyright(C) 2004 Oleg I. Vdovikin (oleg@cs.msu.su)

Copyright(C) 2004-2005 Nicolas Boichat (nicolas@boichat.ch)

Copyright(C) 2014 NVIDIA CORPORATION.

FreeImage 3 version 3.15.1

Copyright (C) 2003-2007 FreeImage Group

FreeImage is currently maintained by HarvÃ© Drolon.

FreeImage is licensed under the GNU General Public License (GPL) and the FreeImage Public License (FIPL). You can choose the license that has the most advantages for you.

License Information

This product includes copyrighted third-party software licensed under the terms of the GNU General Public License. All third-party software packages are copyright by their respective authors. Third Party Terms are hereby incorporated into the Agreement by this reference

GNU General Public License

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients

to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
- b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
- c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the

whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that

contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED

INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

<one line to give the program's name and a brief idea of what it does.> Copyright (C) <year><name of author>

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) year name of author Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'. This is free software, and you are welcome to redistribute it under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program `Gnomovision' (which makes passes at compilers) written by James Hacker.

<signature of Ty Coon>, 1 April 1989 Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License.

GNU Lesser General Public License 2.1

This product includes copyrighted third-party software licensed under the terms of the GNU Lesser General Public License. All third-party software packages are copyright by their respective authors.

GNU Lesser General Public License

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc. 51 Franklin Street, Fifth Floor, Boston, MA02110-1301, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent

holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) The modified work must itself be a software library.

b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.

c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.

d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of

the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

- b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

- a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
- b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot

distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY

TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

<one line to give the library's name and a brief idea of what it does.> Copyright (C) <year><name of author>

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA02110-1301USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the library 'Frob' (a library for tweaking knobs) written by James Random Hacker.

<signature of Ty Coon>, 1 April 1990 Ty Coon, President of Vice

That's all there is to it!

Tegra ASoC Driver

The NVIDIA® Tegra® ASoC driver is implemented for the Android and Linux operating systems and is intended to work seamlessly with different Tegra devices, using an existing framework called Advanced Linux Sound Architecture (ALSA), which is maintained by the upstream Linux community.

This document describes hardware functionality present on Tegra X2 and Jetson TX2. Hardware functionality may not be supported by the software. Review your software documentation to determine availability of software for audio features.

On Jetson TX2, audio output is supported using HDMI audio. Audio input hardware in the development kit is NOT supported.

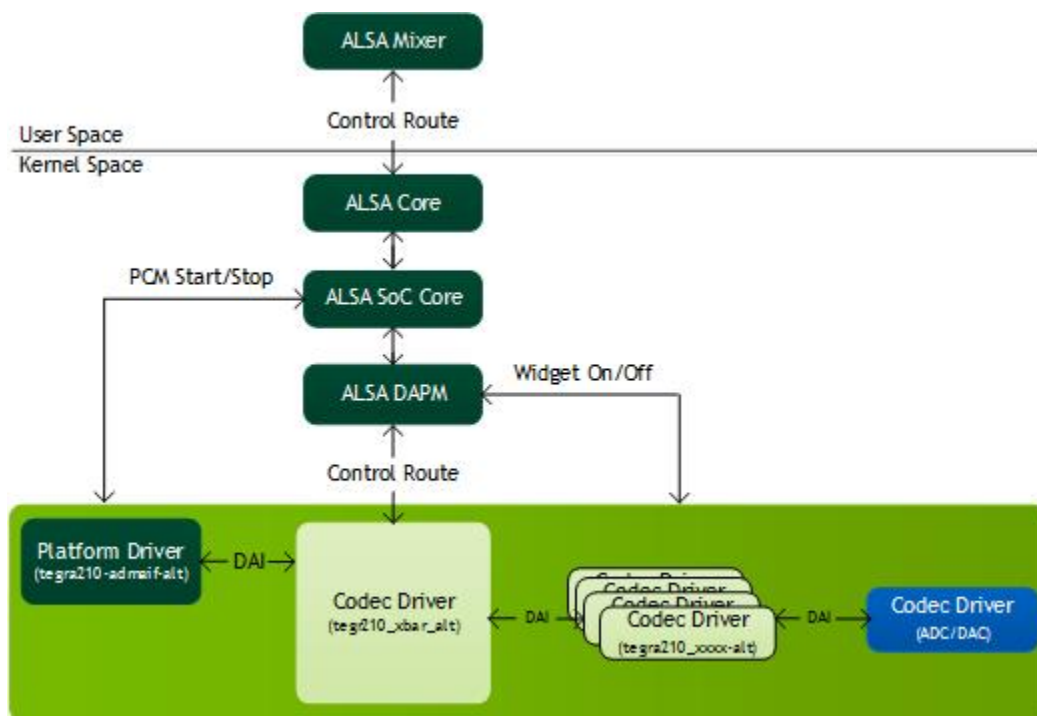
ALSA

The ALSA framework is a part of the Linux kernel that is supported and maintained by the larger Linux community. This makes it feasible to adapt to the framework, designing a driver that leverages NVIDIA audio routing support. ALSA has a very good collection of sound card drivers, including actual codec drivers for platform ADC and DAC support, and can support adding new codec drivers as well.

ALSA also includes libraries and utilities that enable more refined audio control in Linux and Android user space. These include alsamixer, aplay, arecord, tinycap, and others. This enables you to control audio applications without having to interact with kernel space drivers.

The following diagram shows the ALSA software hierarchy.

Tegra ASoC Driver Overview



User space ALSA applications interact with ALSA core (kernel space) through APIs provided by userspace libraries that initialize the actual hardware codecs at the back end of the audio pipeline.

More information on the ALSA framework is available at the following link:

http://www.alsa-project.org/main/index.php/Main_Page

DAPM

ALSA is designed to support various functionalities including but not limited to dynamic audio routing to the available PCM devices. The component of ALSA core that provides this support is Dynamic Audio Power Management (DAPM). DAPM controls the power flow into and out of various codec blocks in the audio subsystem, thereby minimizing power consumption. DAPM introduces switches or kernel controls in the form of widgets to turn ON/OFF the power of a module and help manipulate the required bit of the specific register dynamically using user space applications such as aplay, arecord, or alsamixer. The widgets are classified into various groups.

More information on the widgets and their applications is available at:

<http://www.alsa-project.org/main/index.php/DAPM>

<https://www.kernel.org/doc/Documentation/sound/alsa/soc/DPCM.txt>

In terms of software hierarchy, DAPM is part of the ALSA core as shown in the [Tegra ASoC Driver Overview](#) diagram, helping to manage the codec module power efficiently.

For details on the clocking and power management in Tegra ASoC driver, see [Clocking and Power Management](#) topic.

For more information see the [Dynamic Audio Routing](#) topic.

Device Tree

The Device Tree is a data structure which describes devices on the platform. It is passed to the operating system at boot time and allows you to avoid hard coding component details in the operating system. In this way it makes it easier to change hardware configurations without rebuilding the kernel.

The data structure contains the name of nodes and properties. Each node can have properties or child nodes, and each property is comprised of a name and more than one value. Device Tree structures must be written in the correct format and according to format rules so that the data structure can be parsed by the operating system.

More details on usage of DTS and its script format can be found at:

http://www.devicetree.org/Device_Tree_Usage

For a simple device tree example see the [Codec Driver Instantiation via Device Tree](#) section of this chapter.

Audio Driver

The Tegra Audio Driver leverages Tegra Audio Hub (AHUB) hardware acceleration in the form of platform and codec drivers. The ARM peripheral bus interface (ADMAIF) is implemented as a platform driver with PCM interfaces for playback/record and the rest of the AHUB modules such as the Audio Cross Bar (XBAR), Audio

Multiplexer (AMX), Audio Demultiplexer (ADX) and Inter-IC sound (I2S) implemented as codec drivers. Each of the drivers is connected to XBAR through Digital Audio Interfaces (DAIs), inside a machine driver, forming an audio hub.

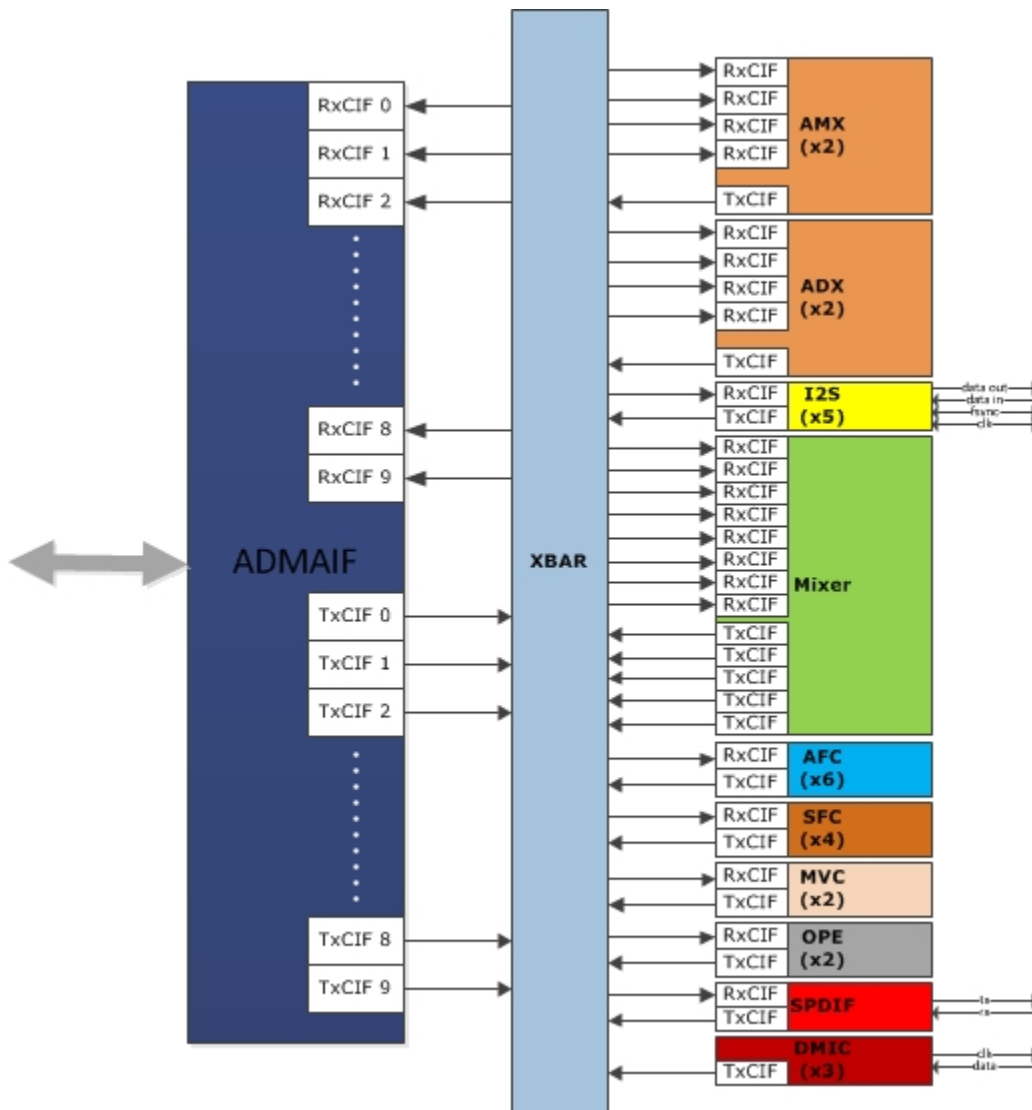
The machine driver probe instantiates the sound card device and registers all the PCM interfaces as exposed by ADMAIF. After booting, before we can use these interfaces to playback or record audio, you must set up the audio paths inside XBAR. By default, XBAR has no routing connections at boot, and no DAPM path is complete to power on the corresponding widgets. The XBAR driver introduces MUX widgets for all the audio components and enables you to create any custom routing through kcontrol from user space using the alsamixer utility. If the audio path is not complete, the DAPM path is not closed. Hardware settings are not applied and you might not hear any audio output.

For more details on how to set up the route and how to play or record on the PCM interfaces, see the [Audio Playback/Record Examples](#) section of this document.

Tegra Audio Hub

In this section, we provide an overview of the audio hub hardware architecture inside the Tegra SoC and describe the software architecture of the driver.

Tegra Audio Hub Architecture



The audio hub contains the modules I2S, SPDIF, and the Digital MIC Controller (DMIC) that interface with the external world. The audio hub also contains Mixer, Sampling Frequency Converter (SFC), Master Volume Control (MVC), Output Processing Engine (OPE), Audio Multiplexers (AMX), Audio Demultiplexers (ADX) and Audio Flow Controllers (AFC). An Audio Direct Memory Access (ADMA) component is included (ADMAIF-DMA) to communicate with memory. The Crossbar (XBAR) facilitates the routing of audio samples through these modules using a proprietary protocol called Audio Client Interface (ACIF).

The modules in the audio hub support various kinds of audio devices that are expected to interface with the application processor, such as cellular baseband devices, different types of audio CODECs, Bluetooth modules, and A/V receivers. The audio hub is capable of supporting the different interface and signal quality requirements of these devices.

As shown in the [Tegra Audio Hub Architecture](#) diagram, each of the AHUB modules has at least one RX port and one TX port, and some have more than one, depending on whether they are for duplex flow such as I2S. The RX ports feed off from XBAR and the TX ports are fed back into XBAR, except in the case of I2S where one end feeds

of the actual codec. This configuration makes XBAR a switch where an audio input can be fed to multiple outputs depending on use case.

For dynamic audio routing examples see the [Audio Playback/Record Examples](#) topic.

Each ADMAIF has both TX and RX FIFOs that support simultaneous recording and playback. ADMA transfers the data to the ADMAIF FIFO for all audio routing scenarios, as shown in the [Tegra Audio Hub Architecture](#) diagram.

For details on hardware configuration of each module, refer to the *Tegra Reference Manual*.

Software Architecture

The software architecture of the Tegra ASoC driver is very similar to the hardware architecture of the AHUB itself, as shown in the [Tegra Audio Hub Architecture](#) diagram. This enables you to leverage all the features supported by the hardware and still conform to the ALSA framework. This ALSA System on a Chip (ASoC) driver is comprised of platform, codec and machine drivers.

- Platform driver—This driver is responsible for PCM registration and interfaces with the PCM driver. The ADMAIF is the platform driver.
- Codec driver—Any driver that registers `snd_soc_codec_driver` structure with the ASoC core can be viewed as the codec driver. The module must have at least one input and one output. In addition, the structure provides a way to define your own DAPM widgets for power management and also kcontrols for register setting from user space. All other modules except ADMAIF are implemented as codec drivers.
- Machine driver—This driver connects one or more codec drivers and a PCM driver together for a given platform.

For details on how to write a machine driver and identify a sound card, see the [Tegra Machine Driver](#) topic and the kernel documentation for ASoC framework at:

<https://www.kernel.org/doc/Documentation/sound/alsa/soc/>

Tegra Platform Driver

The Tegra platform driver realizes the number of ports of playback and capture possible inside the AHUB. Some or all of these ports can be connected to form a full audio routing path. You must complete these paths as described in [Audio Playback/Record Examples](#).

ADMAIF

In the Tegra ASoC driver, ADMAIF is implemented as a platform driver and interfaces with the PCM driver. The PCM driver helps perform DMA operations by overriding the function pointers exposed by the `snd_pcm_ops` structure. The PCM driver is platform agnostic and interacts only with the SOC DMA engine upstream APIs. The DMA engine then interacts with the platform specific DMA driver to get the correct DMA settings. The ADMAIF platform driver defines DAIs and registers the same with ASoC core.

In Tegra X2, there are 10 ADMAIFs. Each ADMAIF is bi-directional, facilitating 10 streams of playback and 10 streams of capture.

Playback Hardware Devices in the Tegra ASoC Driver

**** List of PLAYBACK Hardware Devices ****

```

card 0: tegrasndt210ref [tegra-snd-t210ref], device 0: ADMAIF1 CIF ADMAIF1-0 []
  Subdevices: 1/1
  Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 1: ADMAIF2 CIF ADMAIF2-1 []
  Subdevices: 1/1
  Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 2: ADMAIF3 CIF ADMAIF3-2 []
  Subdevices: 1/1
  Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 3: ADMAIF4 CIF ADMAIF4-3 []
  Subdevices: 1/1
  Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 4: ADMAIF5 CIF ADMAIF5-4 []
  Subdevices: 1/1
  Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 5: ADMAIF6 CIF ADMAIF6-5 []
  Subdevices: 1/1
  Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 6: ADMAIF7 CIF ADMAIF7-6 []
  Subdevices: 1/1
  Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 7: ADMAIF8 CIF ADMAIF8-7 []
  Subdevices: 1/1
  Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 8: ADMAIF9 CIF ADMAIF9-8 []
  Subdevices: 1/1
  Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 9: ADMAIF10 CIF ADMAIF10-9 []
  Subdevices: 1/1
  Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 93: ADSP PCM ADSP-FE1-93 []
  Subdevices: 1/1
  Subdevice #0: subdevice #0

```

Capture Hardware Devices in the Tegra ASoC Driver

```
**** List of CAPTURE Hardware Devices ****
```

```
card 0: tegrasndt210ref [tegra-snd-t210ref], device 0: ADMAIF1 CIF ADMAIF1-0 []
```

```

Subdevices: 1/1
Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 1: ADMAIF2 CIF ADMAIF2-1 []
Subdevices: 1/1
Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 2: ADMAIF3 CIF ADMAIF3-2 []
Subdevices: 1/1
Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 3: ADMAIF4 CIF ADMAIF4-3 []
Subdevices: 1/1
Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 4: ADMAIF5 CIF ADMAIF5-4 []
Subdevices: 1/1
Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 5: ADMAIF6 CIF ADMAIF6-5 []
Subdevices: 1/1
Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 6: ADMAIF7 CIF ADMAIF7-6 []
Subdevices: 1/1
Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 7: ADMAIF8 CIF ADMAIF8-7 []
Subdevices: 1/1
Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 8: ADMAIF9 CIF ADMAIF9-8 []
Subdevices: 1/1
Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 9: ADMAIF10 CIF ADMAIF10-9[]
Subdevices: 1/1
Subdevice #0: subdevice #0
card 0: tegrasndt210ref [tegra-snd-t210ref], device 93: ADSP PCM ADSP-FE1-93 []
Subdevices: 1/1
Subdevice #0: subdevice #0

```

Tegra Codec Driver

A small overview of codec drivers is presented in the [Tegra Audio Hub](#) section of this document. In the Tegra ASoC driver implementation, the rest of the AHUB modules, except ADMAIF, are implemented as codec drivers.

Their responsibilities include:

- Interface to other modules by defining DAIs.
- Define DAPM widgets and establish DAPM routes for dynamic power switching.
- Expose additional kcontrols (kernel controls) as needed for user space utilities to dynamically control module behavior.

XBAR

The XBAR codec driver defines RX, TX and MUX widgets for all the interfacing modules such as ADMAIF, AMX, ADX, I2S, SPDIF, DMIC, Mixer, SFC, MVC, OPE and AFC. MUX widgets are permanently routed to the corresponding TX widgets inside the `snd_soc_dapm_route` structure.

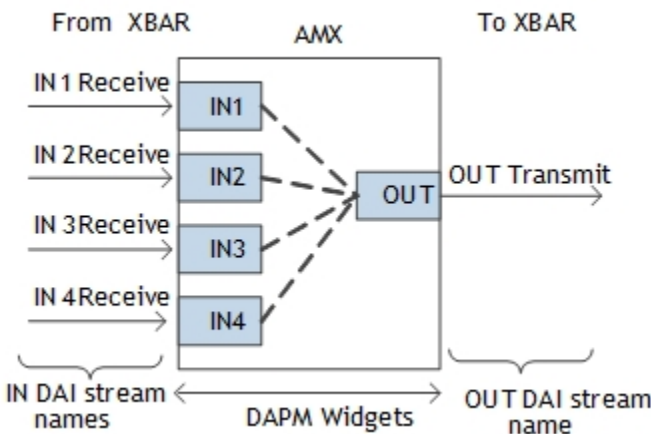
However, the XBAR interconnections are made by connecting any RX widget block to any MUX widget block as needed using the alsamixer utility. The `get/put` handlers for these widgets are implemented so that audio connections are stored by setting the appropriate bit in the hardware MUX register.

For more information see the [Audio Playback/Record Examples](#) topic and the [Dynamic Audio Routing](#) topic.

AMX

An Audio Multiplexer (AMX) module can multiplex up to 4 streams of 16 channels, 32 bits each, into one time-division multiplexed (TDM) stream of 16 channels and 32 bits. The 4 RX ports of AMX originate from XBAR and 1 TX port feeds into XBAR. The DAPM widgets exposed are as shown in the following diagram. The DAPM routes established using these widgets are shown in the dotted lines.

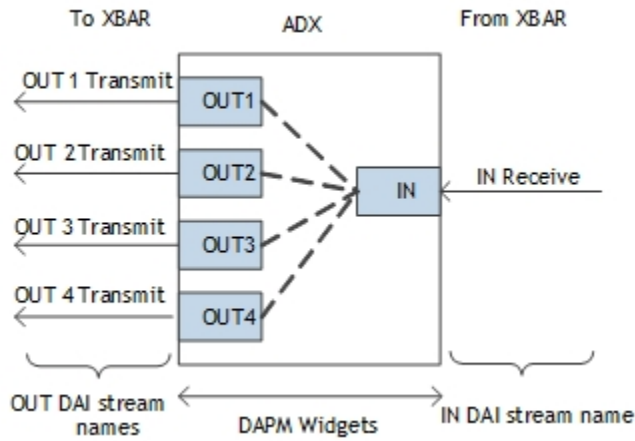
AMX Codec Driver Internals



ADX

An Audio Demultiplexer (ADX) module can de-multiplex a single TDM stream of 16 channels and 32 bits into 4 streams of up to 16 channels, 32 bits each. The 1 RX port of ADX originates from XBAR and 4 TX ports feeds into XBAR. The DAPM widgets exposed are as shown in the following diagram. The DAPM routes established using these widgets are shown in the dotted lines.


ADX Codec Driver Internals



I2S

An I2S module supports 6 different modes, the most useful of which are I2S and TDM. When I2S is configured in TDM mode, we can use AMX and ADX for multiplexing and demultiplexing audio streams. I2S module can also be in I2S mode if the physical codec interfaced does not support TDM mode of operation. An I2S codec driver supports bidirectional data flow and thus defines CIF and DAP RX/TX widgets as shown in the following diagram. The CIF side of I2S, interfaces with XBAR and DAP side is meant to interface with the physical codec on the given platform. The DAPM routes established using these widgets are shown in the dotted lines. I2S modules also expose kernel control to enable internal I2S loopback.

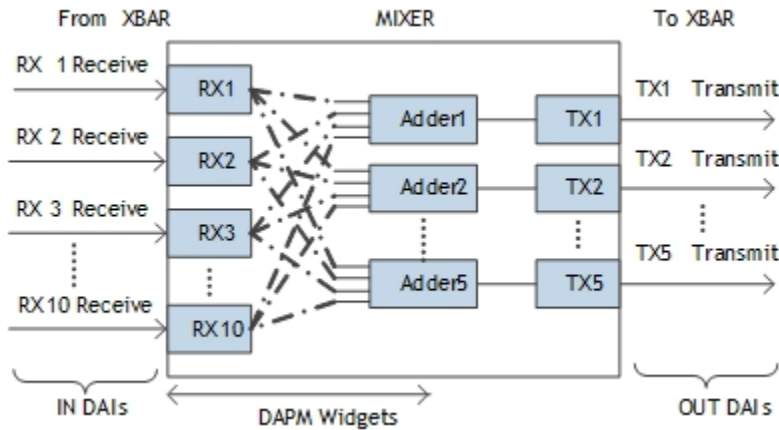
I2S Codec Driver Internals



Mixer

The Mixer can mix audio streams from any of the 10 XBAR-originating input ports to any of the 5 output ports. The DAPM widgets and routes for Mixer are shown in the following diagram. The Mixer driver also exposes Rx Gain and Mixer Enable as additional kcontrols to set the volume of each input stream and to globally enable or disable the Mixer itself.


Mixer Codec Driver Internals



SFC

The Sampling Frequency Converter(SFC) can convert the input sampling frequency to the required sampling rate. SFC has one input port and one output port which are connected to XBAR. The DAPM widgets and routes for SFC are shown in the following diagram.


SFC Codec Driver Internals



SPDIF

The SPDIF driver is very similar to I2S in terms of the DAPM widgets that are exposed and the routing setup. SPDIF is also bidirectional and adapts to the sampling rate of the incoming data. It is implemented as described in the following diagram.


SPDIF Codec Driver Internals



DMIC

It is beneficial to interface directly to digital microphones and speakers via pulse-density modulation (PDM), thus avoiding the need for a PDM-capable external codec. The DMIC controller implements a converter to convert PDM (Pulse density modulation) signals to PCM (Pulse code modulation) signals. The DAPM widgets and routes are as shown below.


DMIC Codec Driver Internals



MVC

MVC provides gain or attenuation to a digital signal path. The digital volume control block is a generic block. It can be used in input or output digital signal path. It can also be used for per-stream volume control and master volume control. The DAPM widgets and routes as shown below.


MVC Codec Driver Internals



OPE

The Output Processing Engine (OPE) is a client of AHUB. OPE contains PEQ and MBDRRC which have a scalable number of BiQuad stages, support stereo, 5p1 and 7p1 channels, and meet ultra-low power(ULP) audio requirements. The DAPM widgets and routes are shown in the following diagram.

OPE Codec Driver Internals



Tegra Machine Driver

The Tegra machine driver connects the codec drivers by linking the DAIs exposed by each module described in previous chapter. It defines the `snd_soc_dai_link` structure and instantiates the sound card.

In general, machine driver responsibilities include:

- Populate the `snd_soc_dai_link` structure with appropriate CPU and CODEC DAIs
- Physical codec clock setting (if any) and codec initializations
- Master/slave configurations (if any)
- Define DAPM widgets to route through the physical codec internals and complete DAPM path as needed
- Propagate the runtime sampling frequency to the individual codec drivers as needed

This section describes how to use the generic machine driver for a given platform and instantiate a sound card.

To adapt the generic machine driver template to a given platform, you must initialize the platform data structure in the device tree structure file shown below. It is important to gather some information about the hardware platform. You must identify the physical codecs present on board and the I2S instance they interface with in the hardware. Refer to hardware schematics for the platform to obtain this information.

```
sound_ref {
    compatible = " ";
    nvidia,model = " ";
    nvidia,num-codec-link = < >;
    nvidia,num-amx = < >;
    nvidia,num-adx = < >;
    nvidia,amx-slot-size = < >;
    nvidia,adx-slot-size = < >;
    nvidia,amx-slot-map = < >;
    nvidia,adx-slot-map = < >;

    nvidia,audio-routing = ;
    nvidia,xbar = <&tegra_axbar>;

    nvidia,dai-link-1 {
        link-name = " ";
        cpu-dai = < >;
        codec-dai = < >;
        cpu-dai-name = " ";
        codec-dai-name = " ";
        tx-mask = < >;
    }
}
```

```

        rx-mask = < >;
        format = " ";
        bitclock-slave;
        frame-slave;
        bitclock-noninversion;
        frame-noninversion;
        bit-format = " ";
        bclk_ratio = < >;
        srate = < >;
        num-channel = < >;
        name-prefix = " ";
    };
};

```

The sound node is added to the DT file for sound card registration and passing platform related data. To bind a device driver to a specific device in a DT file, each node must define a `compatible` property, which is a list of strings. The optional `nvidia,model` property is used for the sound card name. The `nvidia,num-codec-link` property shows the number of dai links exposed outside of the Tegra devices such as I2S, SPDIF or DMIC. This is the same as the number of `nvidia,dai-link-<number>` child nodes.

In Tegra X2 devices, each AMX and ADX has two modules, `nvidia,num-amx` and `nvidia,num-adx`. These indicate the total number of AMX and ADX to be registered by the machine driver. The `nvidia,amx-slot-size` and `nvidia,adx-slot-size` properties indicate the slot map size of each module, which is related to the size of `nvidia,amx-slot-map` and `nvidia,adx-slot-map`. The machine driver uses the default value for each property to support stereo 16 bits while initiating AMX and ADX modules unless the above properties are overridden.

The `nvidia,dai-link-<number>` node is initialized based on I2S and physical codec links as shown above. You can replace the codec with a dummy `spdif-dit.0` codec. In that case, you must program the physical codecs at the required sampling rate and operational modes either through a separate utility or by adding necessary `hw_params` API calls inside the machine driver. The `link-name` property represents each DAI link, and differentiates the codecs connected to the different I2S instances.

The `bitclock-slave`, `frame-slave`, `format`, `bitclock-noninverstion` and `frame-noninversion` properties set I2S in I2S/TDM formats and/or master/slave modes. The `bclk-ratio` property is available to configure the I2S bit clock sample rate. The `srate` and `num-channel` properties indicate I2S LRCK and the number of channels in one frame. DAPM routes are initialized as established in `nvidia,audio-routing` property.

To populate the `snd_soc_dai_link` structure in the machine driver

1. Initialize all XBAR related DAI links common to any machine driver. Use the APIs in the utility file `tegra_asoc_machine_alt.h` in the machine driver probe function as shown in the following example:

```

/* get the xbar dai link structure */
tegra_machine_dai_links=
    tegra_machine_get_dai_link();

/* set AMX/ADX dai_init */
tegra_machine_set_dai_init(TEGRA210_DAI_LINK_AMX1,

```

```

    &tegra_t210ref_amx1_dai_init);
tegra_machine_set_dai_init(TEGRA210_DAI_LINK_ADX1,
    &tegra_t210ref_adx1_dai_init);
tegra_machine_set_dai_init(TEGRA210_DAI_LINK_AMX2,
    &tegra_t210ref_amx2_dai_init);
tegra_machine_set_dai_init(TEGRA210_DAI_LINK_ADX2,
    &tegra_t210ref_adx2_dai_init);

/* set sfc dai_init */
tegra_machine_set_dai_init(TEGRA210_DAI_LINK_SFC1_RX,
    &tegra_t210ref_sfc1_init);
tegra_machine_set_dai_init(TEGRA210_DAI_LINK_SFC2_RX,
    &tegra_t210ref_sfc2_init);
tegra_machine_set_dai_init(TEGRA210_DAI_LINK_SFC3_RX,
    &tegra_t210ref_sfc3_init);
tegra_machine_set_dai_init(TEGRA210_DAI_LINK_SFC4_RX,
    &tegra_t210ref_sfc4_init);

```

2. The generic machine driver handles DAI link management based on the initialized platform specific data in the DT file.

Because this step is platform/machine dependent, it is discussed with an example in subsequent sections of this chapter.

Tegra X2

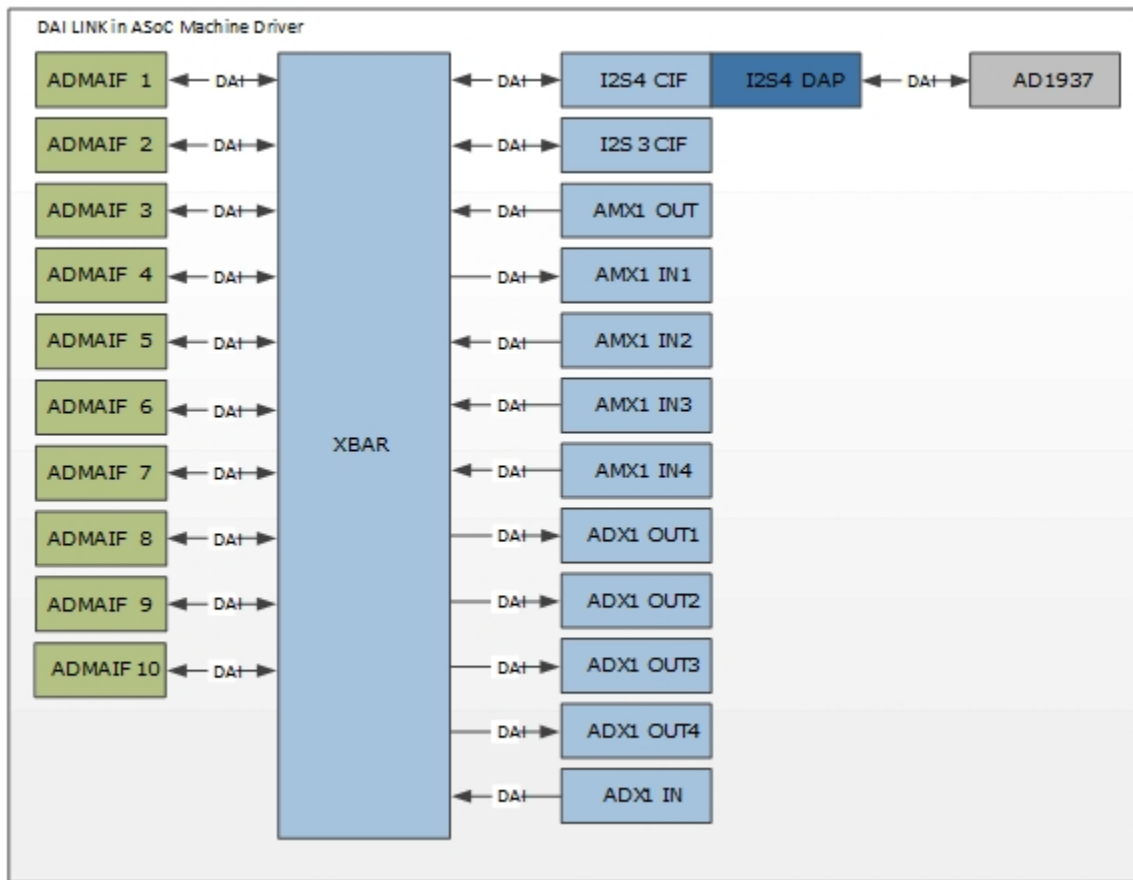
In the Tegra X2 combination of boards, there is one available AD1937s codec. AD1937 supports eight channels in TDM and I2S modes, respectively. The AD1937 codec is connected to I2S4 on P1892.

Machine Specific DAI links

The platform data structure must be initialized in the DT file so that the machine driver can retrieve the device specific information during the platform probe. The illustration below shows the DAI link diagram for PEDP+ with Tegra X2 where the left side of each DAI is CPU DAI and right side of each DAI is codec DAI in `snd_soc_dai_link`.

Note: The following diagram shows the DAI LINK for PEDP+ platform; check the Release Notes for platforms supported in your release.

DAI LINK In the Machine Driver of PEDP with Tegra X2



```

sound_ref {
    compatible = "nvidia,tegra-audio-t210ref";
    nvidia,model = "tegra-snd-t210ref";
    nvidia,num-codec-link = <1>;
    nvidia,num-amx = <1>;
    nvidia,num-adx = <1>;
    nvidia,amx-slot-size = <32 32>;
    nvidia,adx-slot-size = <32 32>;
    nvidia,addr-max9485 = <112>;
    nvidia,amx-slot-map = <
        /* jack 0 */
        TDM_SLOT_MAP(0, 0, 0)
        TDM_SLOT_MAP(0, 0, 0)
        TDM_SLOT_MAP(0, 1, 0)
        TDM_SLOT_MAP(0, 1, 1)
        TDM_SLOT_MAP(0, 0, 0)
  >
}

```

```

TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 2, 0)
TDM_SLOT_MAP(0, 2, 1)
/* jack 1 */
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(1, 1, 0)
TDM_SLOT_MAP(1, 1, 1)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(1, 2, 0)
TDM_SLOT_MAP(1, 2, 1)
/* jack 2 */
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(2, 1, 0)
TDM_SLOT_MAP(2, 1, 1)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(2, 2, 0)
TDM_SLOT_MAP(2, 2, 1)
/* jack 3 */
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(3, 1, 0)
TDM_SLOT_MAP(3, 1, 1)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(3, 2, 0)
TDM_SLOT_MAP(3, 2, 1)>;
nvidia,adx-slot-map = <
/* jack 0 */
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 1, 0)
TDM_SLOT_MAP(0, 1, 1)
TDM_SLOT_MAP(0, 0, 0)

```

```

TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 2, 0)
TDM_SLOT_MAP(0, 2, 1)
/* jack 1 */
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(1, 1, 0)
TDM_SLOT_MAP(1, 1, 1)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(1, 2, 0)
TDM_SLOT_MAP(1, 2, 1)
/* jack 2 */
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(2, 1, 0)
TDM_SLOT_MAP(2, 1, 1)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(2, 2, 0)
TDM_SLOT_MAP(2, 2, 1)
/* jack 3 */
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(3, 1, 0)
TDM_SLOT_MAP(3, 1, 1)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(0, 0, 0)
TDM_SLOT_MAP(3, 2, 0)
TDM_SLOT_MAP(3, 2, 1)>;

nvidia, audio-routing =
    "Headphone-z", "z DAC1OUT",
    "Headphone-z", "z DAC2OUT",
    "Headphone-z", "z DAC3OUT",
    "Headphone-z", "z DAC4OUT",
    "z ADC1IN", "LineIn-z";

```



```

nvidia,xbar = <&tegra_axbar>;

nvidia,dai-link-1 {
    link-name = "ad-playback-z";
    cpu-dai = <&tegra_i2s4>;
    codec-dai = <&ad1937z>;
    cpu-dai-name = "I2S4";
    codec-dai-name = "ad193x-hifi";
    tx-mask = <0xFF>;
    rx-mask = <0xFF>;
    format = "dsp_a";
    bitclock-slave;
    frame-slave;
    bitclock-noninversion;
    frame-noninversion;
    bit-format = "s32_le";
    bclk_ratio = <1>;
    srate = <48000>;
    num-channel = <8>;
    name-prefix = "z";
};
};

```

The `dai_fmt` sets I2S4 in TDM master mode. It also connects codec AD1937 with I2S4 DAPM routes and AMX/ADX slot maps initialized. The sound card is named `tegra-snd-t210ref`.

The DAPM widgets supported in the generic machine driver include Headphone-x, Headphone-y, Headphone-z, Headphone-s, LineIn-x, LineIn-y, LineIn-z, and LineIn-s. Based on how `name_prefix` and `dai_fmt` are initialized, the DAPM route must be set by parsing the IN/OUT DAPM widgets of the physical codec. Without this, the DAPM path is not complete and audio does not function correctly. The Tegra X2 audio map is shown below.

```

nvidia, audio-routing =
    "Headphone-z", "z DAC1OUT",
    "Headphone-z", "z DAC2OUT",
    "Headphone-z", "z DAC3OUT",
    "Headphone-z", "z DAC4OUT",
    "z ADC1IN", "LineIn-z",

```

Though we can assign any ADMAIF to any AMX via XBAR, we can fix ADMAIF1 to ADMAIF4 as the inputs of AMX1 are connected to I2S4.

The machine driver creates three `snd_soc_dai_links` for each `dai-link-<number>` node initialized in the DT file and appends them to XBAR DAI links using the utility API. It also initializes the `.dai_ops` for the ADMAIF interfaces accordingly as shown below.

```

/* set ADMAIF dai_ops */
for (i = TEGRA210_DAI_LINK_ADMAIF1;
     i <= TEGRA210_DAI_LINK_ADMAIF10; i++)
    tegra_machine_set_dai_ops(i, &tegra_t210ref_spdif_ops);

for (i = 0; i < machine->num_codec_links; i++) {
    if (tegra_t210ref_codec_links[i].name) {
        if (strstr(tegra_t210ref_codec_links[i].name,
                  "ad-playback-z")) {
            for (j = TEGRA210_DAI_LINK_ADMAIF1;
                 j <= TEGRA210_DAI_LINK_ADMAIF4; j++)
                tegra_machine_set_dai_ops(j,
                                           &tegra_t210ref_ad1937_z_ops);
        } else if (strstr(tegra_t210ref_codec_links[i].name,
                          "ad-playback-x")) {
            for (j = TEGRA210_DAI_LINK_ADMAIF5;
                 j <= TEGRA210_DAI_LINK_ADMAIF8; j++)
                tegra_machine_set_dai_ops(j,
                                           &tegra_t210ref_ad1937_x_ops);
        } else if (strstr(tegra_t210ref_codec_links[i].name,
                          "spdif-playback")) {
            tegra_machine_set_dai_ops(
                TEGRA210_DAI_LINK_ADMAIF9,
                &tegra_t210ref_spdif_ops);
        }
    }
}

/* append t210ref specific dai_links */
card->num_links =
    tegra_machine_append_dai_link(tegra_t210ref_codec_links,
    2 * machine->num_codec_links);

```

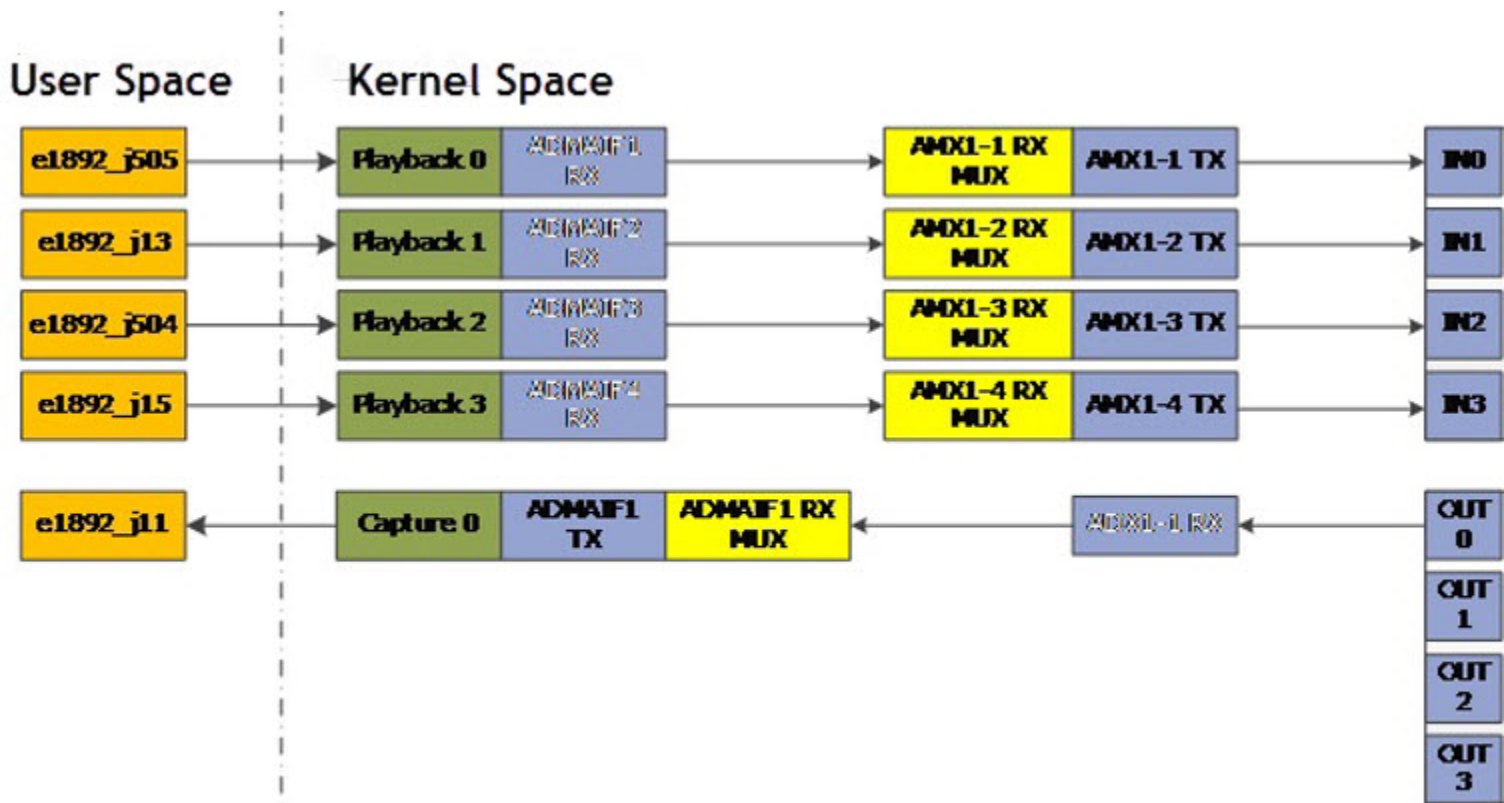
```
tegra_machine_dai_links = tegra_machine_get_dai_link();
card->dai_link = tegra_machine_dai_links;
```

Audio Path

The dai_link connection in the machine driver only connects the platform and codec drivers to XBAR driver. However, to realize any audio output from the physical codecs, you must complete the DAPM path by routing through the internals of XBAR, i.e., connect the RX and MUX widgets from user space using alsamixer utility as follows.

Note: The following commands are specific to Tegra X2 and represent the audio path shown in the following diagram. This diagram shows the DAI LINK for PEDP+ platform; check the *Release Notes* for platforms supported in your release.

Tegra X2 Audio Path



Block with white letters is input of MUX

XBAR Route Setting for Tegra X2

```
amixer -c 0 sset 'I2S4 Mux' 'AMX1'
amixer -c 0 sset 'AMX1-1 Mux' 'ADMAIF1'
amixer -c 0 sset 'AMX1-2 Mux' 'ADMAIF2'
```

```

amixer -c 0 sset 'AMX1-3 Mux' 'ADMAIF3'
amixer -c 0 sset 'AMX1-4 Mux' 'ADMAIF4'
amixer -c 0 sset 'ADX1 Mux' 'I2S4'
amixer -c 0 sset 'ADMAIF1 Mux' 'ADX1-1'

```

Dynamic Audio Routing

Instantiation of the sound card after boot indicates that all codec drivers and the platform driver are interconnected using the machine driver. The remaining step before obtaining the audio output on the physical codecs involves routing through the XBAR internals via MUX widgets to complete the DAPM path using the alsamixer utility. To support audio routing dynamically between AHUB modules, this route setting step is performed from user space. This provides flexibility for complicated use cases.

For example, `[amixer -c 0 sset 'I2S1 Mux' 'I2S1']` realizes the internal AHUB path I2S1 RX -> XBAR -> I2S1 TX. Similarly, the following are use cases that emphasize dynamic audio route control from user space.

Case 1: Internal AHUB TDM Path

Path: I2S -> ADX -> AMX -> I2S

Commands:

```

amixer -c 0 sset 'ADX1 Mux' 'I2S4'
amixer -c 0 sset 'AMX1-1 Mux' 'ADX1-1'
amixer -c 0 sset 'AMX1-2 Mux' 'ADX2-2'
amixer -c 0 sset 'AMX1-3 Mux' 'ADX3-3'
amixer -c 0 sset 'AMX1-4 Mux' 'ADX4-4'
amixer -c 0 sset 'I2S4 Mux' 'AMX1'

```

Modify Case 1 to Record on I2S3 (I2S Mode) And Output On I2S4 (TDM Mode)

Path: [I2S -> ADX -> AMX -> I2S] + [I2S3 -> AMX -> I2S4]

Initially audio records from I2S4 and outputs on I2S4.

Commands:

```

amixer -c 0 sset 'AMX1-1 Mux' 'None'
amixer -c 0 sset 'AMX1-1 Mux' 'I2S3'

```

After this step, audio records from I2S3 and outputs on I2S4

Codec Driver Instantiation via Device Tree

Based on architecture, the Makefile in the following directory conditionally compiles the required DTS files into DTB files:

```
$KERNEL_TOP/arch/arm64/boot/dts/
```

When the kernel is flashed, the flash script chooses the required DTS file for parsing during boot, and the ASoC codecs listed in DTS are instantiated. To add any new module instantiation as a requirement, identify and edit the Device Tree script as reported in the `dmesg` log on the target as shown in the following example:

```
<6>[ 0.000000] Tegra reserved memory:
<6>[ 0.000000] LP0:                f25ff000 - f25fffff
<6>[ 0.000000] Bootloader framebuffer: 00000000 - 00000000
<6>[ 0.000000] Bootloader framebuffer2: 00000000 - 00000000
<6>[ 0.000000] Framebuffer:          ef800000 - f09fffff
<6>[ 0.000000] 2nd Framebuffer:      f0a00000 - f19fffff
<6>[ 0.000000] Carveout:             00000000 - 00000000
<6>[ 0.000000] Vpr:                 f4600000 - ffffffff
<6>[ 0.000000] Tsec:                 00000000 - 00000000
<7>[ 0.000000] On node 0 totalpages: 980992
<7>[ 0.000000]   DMA32 zone: 6244 pages used for memmap
<7>[ 0.000000]   DMA32 zone: 0 pages reserved
<7>[ 0.000000]   DMA32 zone: 456704 pages, LIFO batch:31
<7>[ 0.000000]   Normal zone: 7168 pages used for memmap
<7>[ 0.000000]   Normal zone: 524288 pages, LIFO batch:31
<6>[ 0.000000] psci: probing function IDs from device-tree
<6>[ 0.000000] DTS File Name: /home/juskim/git/android/t210_main/kernel/arch/arm64/b
<6>[ 0.000000] Tegra21: Speedo/IDDQ fuse revision 0
```

To add new devices for instantiation

- Add the device name with the base address and status as “okay”, as shown in the following example.

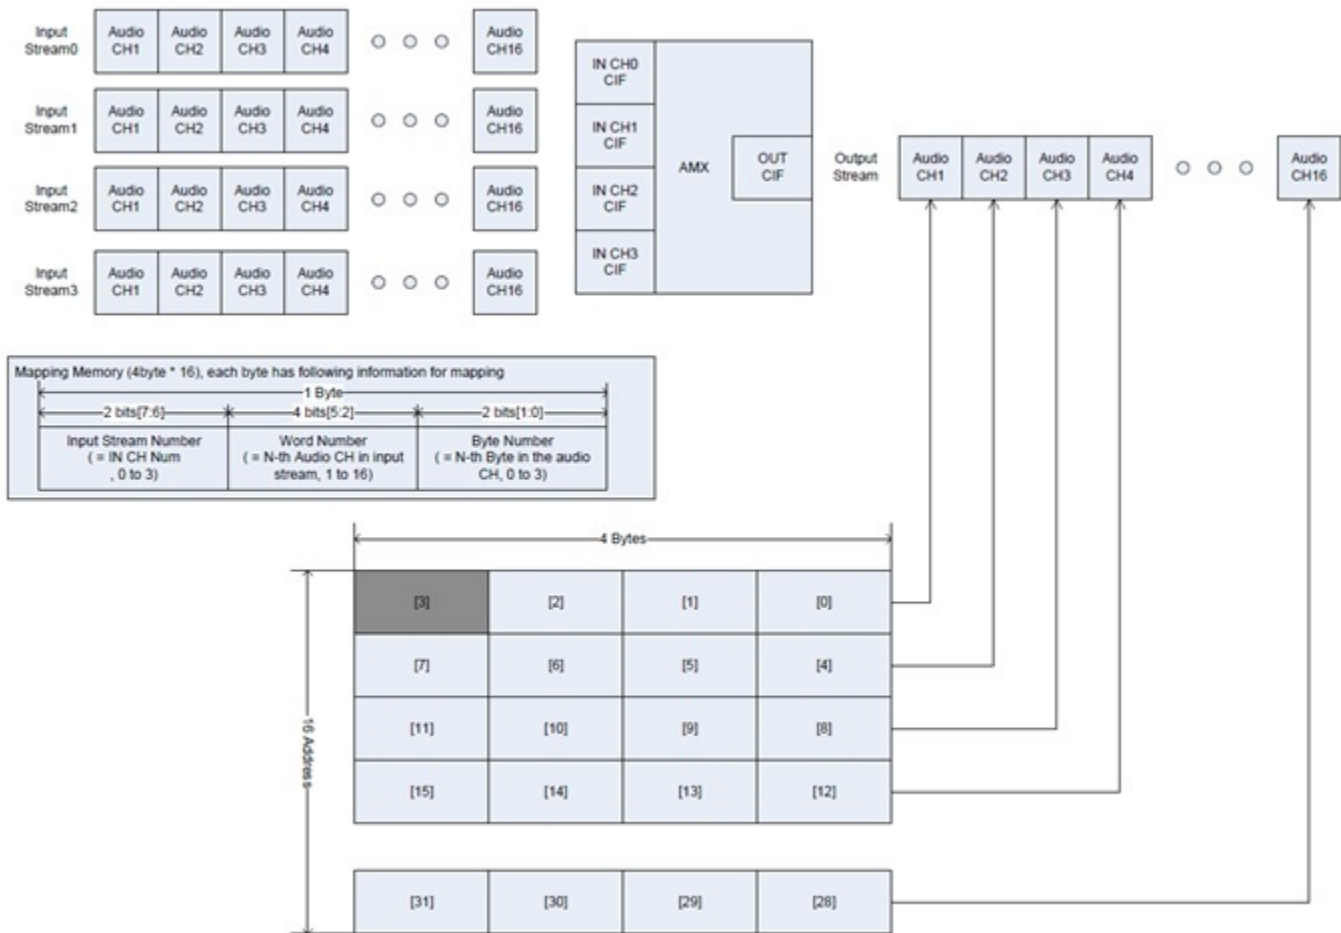
```
ahub {
    status = "okay";
    i2s@702d1000 {
        pinctrl-names = "dap_active", "dap_inactive";
        pinctrl-0 = <>;
        pinctrl-1 = <>;
        status = "okay";
    };
};
```

TDM Slot Mapping

The Tegra X2 AHUB consists of 2 hardware audio multiplexers (AMX) and demultiplexers (ADX) as described in

[Tegra Audio Hub](#). Initialize slot mapping in the platform data as shown in [Machine Specific DAI links](#). Without any initialization, the machine driver defaults the slot map for AMX inputs and ADX outputs at 16-bit stereo and 8-channel, 32 bit TDM output. To change the slot mapping of hardware TDMs, adjust the slot map structures in the DT file.

The following diagram shows the AMX can multiplex up to 4 input streams, each of which can be up to 16 channels, 32 bits, into 1 single output stream of 16 channels, 32 bits to form a TDM signal. The capability of ADX is the reverse of AMX.



The slot map is an integer array of size 32 that controls the mapping of these input streams to the output stream. The following macro is available:

```
#define TDM_SLOT_MAP(stream_id, nth_channel, nth_byte) \
    ((stream_id << 16) | (nth_channel << 8) | (nth_byte))
```

where `stream_id` can vary from 0 to 3, `nth_channel` can vary from 1 to 16, and `nth_byte` can vary from 0 to 3.

Clocking and Power Management

The clock tree of the Tegra ASoC driver in the idle state, when no audio record or playback is in progress, is as shown below.

clock		state	ref	div	rate
i2s4_sync		on	1		24000000
i2s3_sync		on	1		24000000
*audio3		off	0		24000000
i2s2_sync		on	1		24000000
i2s1_sync		on	1		24000000
*audio1		off	0		24000000
i2s0_sync		on	1		24000000
spdif_in_sync		on	1		24000000
*audio2_dmic		off	0		24000000
*audio1_dmic		off	0		24000000
*audio0_dmic		off	0		24000000
*audio		off	0		24000000
*audio_2x		off	0	x2	48000000
*audio4		off	0		24000000
*audio2		off	0		24000000
*audio0		off	0		24000000
osc		on	3		38400000
pll_ref		on	7	1.0	38400000
pll_a		on	1	x9.5	368639844
pll_a_out0	on	1	15.0		24575990
extern1	on	3	1.0		24575990
clk_out_1	on	2	1.0		24575990
d_audio		off	0	2.0	12287995
*dmic3		off	0	11.0	2234181
*dmic2		off	0	11.0	2234181
*dmic1		off	0	11.0	2234181
spdif_out		off	0	21.0	1170286
i2s4		off	0	5.5	4468362
i2s3		off	0	5.50	4468362
i2s1		off	0	5.50	4468362
i2s0		off	0	10.50	2340571
pll_p		on	15	x10.6	408000000
*spdif_in		off	0	8.50	48000000
ape		off	0	2.0	204000000
xbar.ape		off	0		204000000

adsp.ape	off	0	204000000
adma.ape	off	0	204000000

The clocks of the individual modules, AMX, ADX, AFC, SFC, MIXER, and others, are internally handled by the APE clock. The clock for the codec drivers I2S and XBAR are switched OFF in idle. They are turned ON when audio playback or recording is in progress. The `idle_bias_off` option provided by ASoC core (`snd_soc_codec_driver`) is set to 1 in the individual codec drivers for dynamic audio power management. With this option, the ASoC core calls the appropriate `resume()` and `suspend()` functions in the codec driver that was registered using `SET_RUNTIME_PM_OPS` during platform probe.

The suspend and resume APIs are called when playback or record stops or starts. The suspend API caches the register map and disables the clock. The resume API enables the clock and then syncs the register map from the cache.

Note: Volatile registers are not synced in this process. Those registers must be reprogrammed after the clock is re-enabled. The internal mapping RAM FIFOs, if any, are cleared during every suspend/resume cycle. The RAM configuration must be restored as well during this process.

Audio Playback/Record Examples

To test audio playback and record scenarios, make sure the XBAR route is set up using `alsamixer` commands, based on the platform. The following provides the sample testing commands for audio playback and record scenarios.

Command	Result
<code>• aplay -D AUDIO_OUT0 sample.wav</code>	• Audio output from AUDIO_OUT0 on e1892.
<code>• aplay -D AUDIO_OUT1 sample.wav</code>	• Audio output from AUDIO_OUT1 on e1892.
<code>• aplay -D AUDIO_OUT2 sample.wav</code>	• Audio output from AUDIO_OUT2 on e1892.
<code>• aplay -D AUDIO_OUT3 sample.wav</code>	• Audio output from AUDIO_OUT3 on e1892.
<code>• arecord -D AUDIO_IN -f dat -d 30 record_result.wav</code>	• Audio recorded as input from AUDIO_IN on E1892 and saved to file record_result.wav.

Troubleshooting

The following issues can be overcome using the described solutions.

Issue 1: No Sound Cards Found

- Identify the ASoC error that lead to no sound card detection with the `dmesg [dmesg | grep ASoC]` command. Output is similar to the following:

```
[4.874720] tegra-audio-t210ref tegra-audio-t210ref.0: ASoC: no source widget found for
[4.874724] tegra-audio-t210ref tegra-audio-t210ref.0: ASoC: Failed to add route x OUT->
[4.874736] tegra-audio-t210ref tegra-audio-t210ref.0: ASoC: no sink widget found for x
```



```
[4.874739] tegra-audio-t210ref tegra-audio-t210ref.0: ASoC: Failed to add route LineIn-
```

In this case, `x OUT` and `x IN` are the widgets of the `spdif-dit` dummy codec. It might have not been instantiated in ASoC. Confirm that with the following command:

```
cat /sys/kernel/debug/asoc/codecs
```

if `spdif-dit` is instantiated. The `spdif` device must be instantiated using `platform_register` in the board-specific file.

- If output from the `dmesg [dmesg | grep ASoC]` command is similar to the following:

```
[4.874720] tegra-audio-t210ref tegra-audio-t210ref.0: ASoC: CPU DAI DAP not registered
```

In this case, “DAP” is the CPU DAI for the I2S to codec dai link. The I2S codec may not be instantiated in ASoC. Confirm that with the following command:

```
cat /sys/kernel/debug/asoc/codecs
```

if `tegra30-i2s` is instantiated.

Identification of the DAI link at the point of failure would give a clue on the I2S instance number that failed to instantiate. Accordingly, the I2S codec driver can be instantiated by providing a suitable entry point in DTS file as described in [Codec Driver Instantiation via Device Tree](#).

Issue 2: Sound Not Audible

1. Confirm if the DAPM path is completed. Check the path using `alsamixer` utility as described in [Audio Path](#).
2. Confirm the pinmux setting for I2S master/slave mode. Check the following files.

```
arch/arm64/boot/dts/tegra210-foster-e-p2530-common.dtsi
```

```
arch/arm64/boot/dts/tegra210-platforms/tegra210-foster-e-pinmux-p2530-0930-e00.dtsi
```

3. Confirm the `status` property of `i2s` is set to `okay` in the following file:

```
arch/arm64/boot/dts/tegra210-common.dtsi
```

4. Confirm the clock settings for I2S master/slave mode. Probe the Frame sync (FS) and bit clock (BCLK) of I2S using a scope.
5. If I2S is configured in TDM mode, please check whether the AMX slot map is configured correctly based on the description in [TDM Slot Mapping](#).
6. You can also form an internal DAPM path connecting I2S with the same I2S instance using the following command:

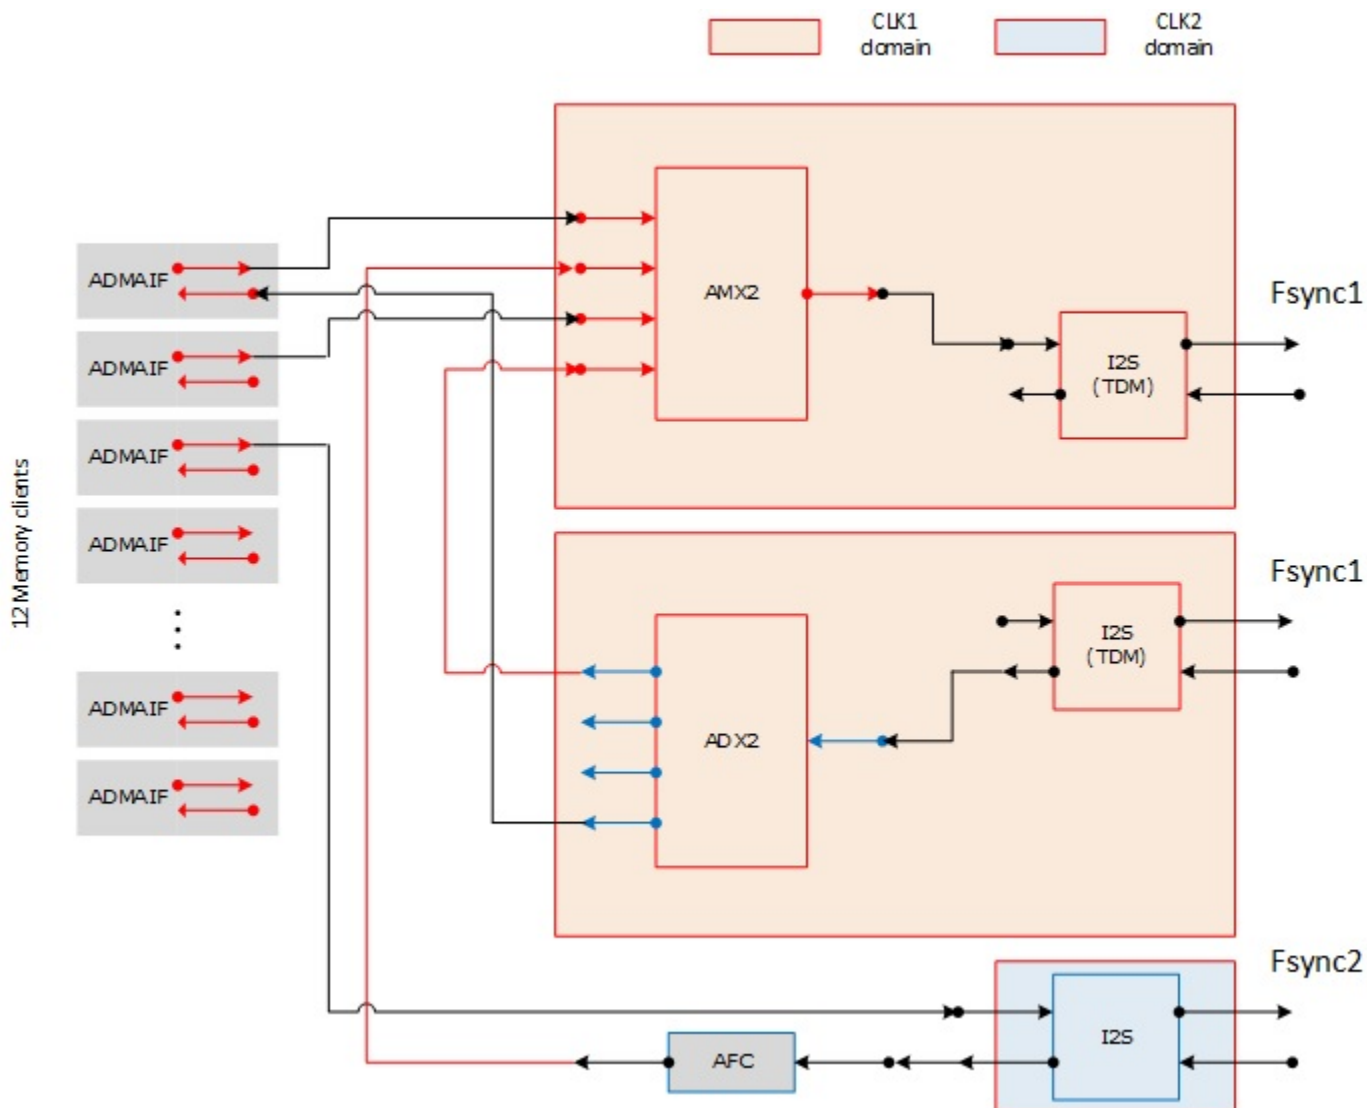
```
amixer -c 0 sset 'I2S Mux' 'I2S'
```

If the codec record/playback functions correctly, that confirms I2S and onboard codecs are configured correctly. It also isolates the problem to rest of the AHUB path.

Miscellaneous Examples

This section describes setting up the XBAR internal audio path from user space for complicated Tegra use cases.

Simple Internal Audio Path



Routing Commands

- AMX connection:

```
amixer -c 0 sset 'AMX2-1 Mux' 'ADMAIF5'
amixer -c 0 sset 'AMX2-2 Mux' 'AFC1'
amixer -c 0 sset 'AMX2-3 Mux' 'ADMAIF6'
amixer -c 0 sset 'AMX2-4 Mux' 'ADX2-1'
amixer -c 0 sset 'I2S3 Mux' 'AMX2'
```

- ADX connection:

```
amixer -c 0 sset 'ADX2 Mux' 'I2S3'
amixer -c 0 sset 'ADMAIF5 Mux' 'ADX2-4'
```

- AFC connection:

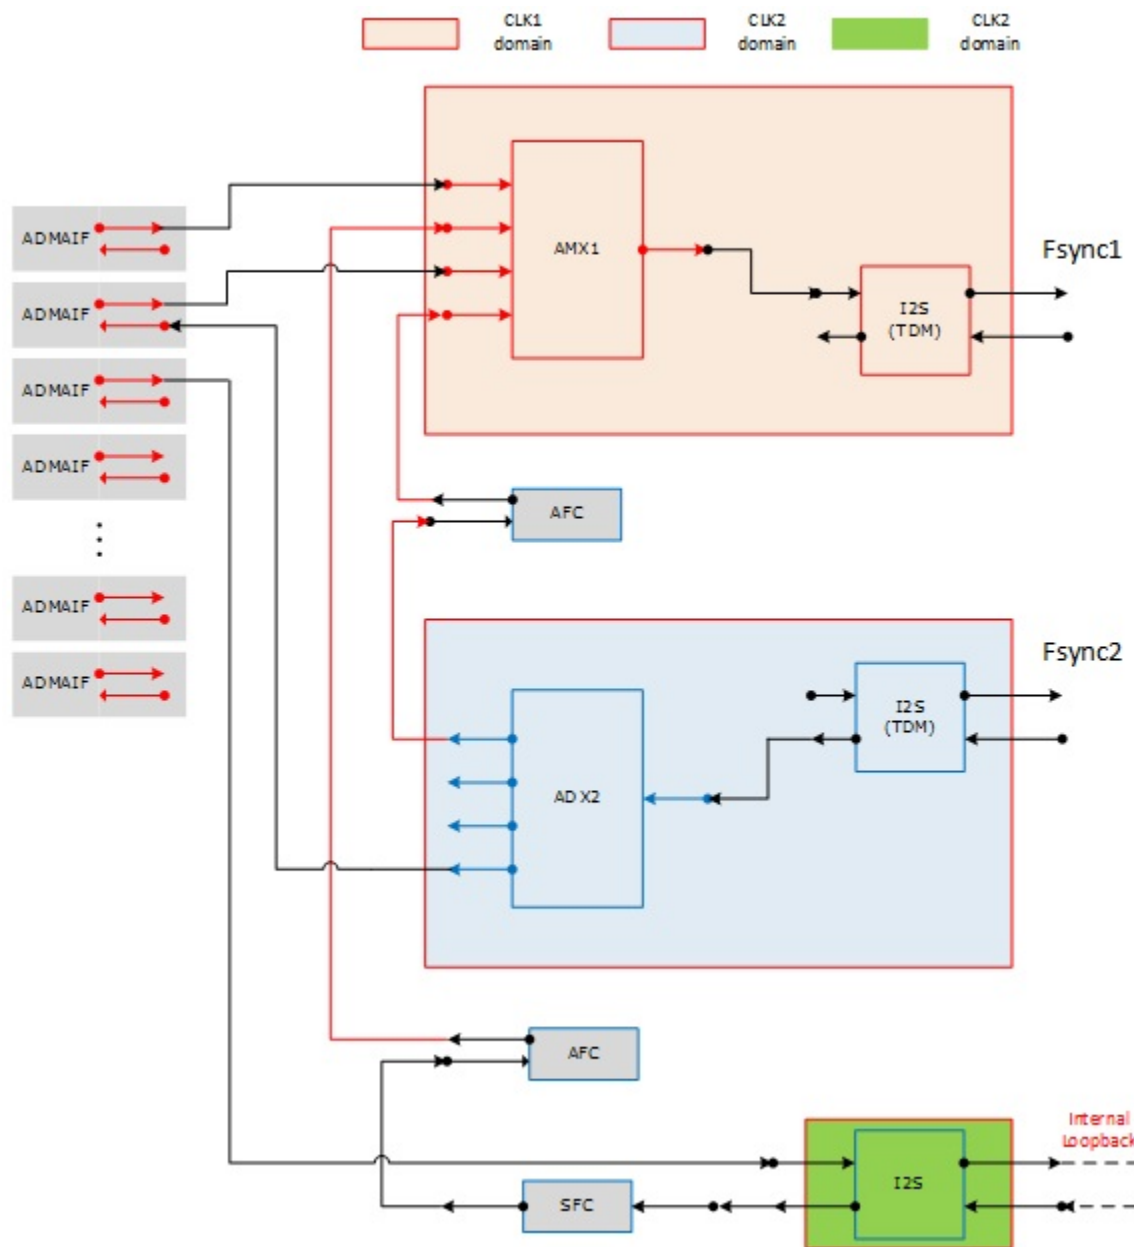
```
amixer -c 0 sset 'AFC1 Mux' 'I2S4'
amixer -c 0 sset 'I2S4 Mux' 'ADMAIF7'
```

Testing Commands

```
aplay -c 0 -D hw:0,4 -f S16_LE -r 48000 sample1.wav&
aplay -c 0 -D hw:0,5 -f S16_LE -r 48000 sample2.wav&
aplay -c 0 -D hw:0,6 -f S16_LE -r 48000 sample3.wav&
arecord -c 0 -D hw:0,4 -f wav -c 2 -r 48000 -b 16 sample4.wav &
```

I2S-x and I2S-y Under Same Clock Domain

If I2S-x and I2S-y are under the same clock domain for the sake of case study, AFC1 insertion between ADX1-1 and AMX1-4 does not help. However, the audio routing path can still be constructed as illustrated below and described in [Routing Commands](#) and [Testing commands](#).



Routing Commands

- AMX connection:

```
amixer -c 0 sset 'AMX1-1 Mux' 'ADMAIF1'
amixer -c 0 sset 'AMX1-2 Mux' 'AFC1'
amixer -c 0 sset 'AMX1-3 Mux' 'ADMAIF2'
amixer -c 0 sset 'AMX1-4 Mux' 'AFC2'
amixer -c 0 sset 'I2S4 Mux' 'AMX1'
```

- ADX connection:

```
amixer -c 0 sset 'ADX2 Mux' 'I2S3'
```

```
amixer -c 0 sset 'ADMAIF2 Mux' 'ADX2-4'  
amixer -c 0 sset 'AFC1 Mux' 'ADX2-1'
```

- AFC/SFC connection:

```
amixer -c 0 sset 'SFC1 Mux' 'I2S5'  
amixer -c 0 sset 'I2S5 Mux' 'ADMAIF3'  
amixer -c 0 sset 'AFC2 Mux' 'SFC1'  
amixer -c 0 cset iface=MIXER,name='output rate' '48kHz'  
amixer -c 0 sset 'I2S5 Loopback' 1
```

Testing Commands

```
aplay -c 0 -D hw:0,0 -f S16_LE -r 48000 sample1.wav &  
aplay -c 0 -D hw:0,1 -f S16_LE -r 48000 sample2.wav &  
aplay -c 0 -D hw:0,2 -f S16_LE -r 48000 sample3.wav &  
arecord -c 0 -D hw:0,1 -f wav -c 2 -r 48000 -b 16 sample4.wav &
```

Building the AARCH 64 Toolchain and glibc

The NVIDIA® Tegra® Linux Driver Package contains the script and patches used to create the set of toolchains that NVIDIA uses to produce the AARCH64 L4T binaries. This topic describes how to obtain and build these toolchains.

Toolchain Information

The AARCH64 toolchain is built as aarch64-unknown-linux-gnu. The toolchain binaries are statically built.

The toolchain contains the following components:

- GCC version: 4.8.5
- Glibc Version: 2.17

Downloading the Pre-Built Toolchain

Download the pre-built toolchain binaries from the Embedded Download Center on the NVIDIA Developer page at:

```
https://developer.nvidia.com/embedded/downloads
```

The AARCH64 toolchain is available at:

```
http://developer.nvidia.com/embedded/dlc/l4t-gcc-toolchain-64-bit=27-1
```

Copy the contents of the tarball to the location of your toolchain and start using it.

Building the Toolchain

The toolchain build script and patches are provided in the Jetson Toolchain Build package.

To build the toolchain

1. Obtain the Toolchain Scripts package as described in [Downloading the Pre-built Toolchain](#).
2. Uncompress the TAR file with the following command:

```
tar xpf jetson-tx1-toolchain-build.tbz2
```

The toolchain source tarball contains the following:

- The required sources to build the toolchain.
 - The `make-aarch64-toolchain.sh` script that builds the toolchain.
 - A `patches` folder containing patches to the toolchain sources.
3. Run the toolchain building script.

Build products are located in the `install` directory.

The `get_src` function in the `make-aarch64-toolchain.sh` script downloads all the relevant sources required for building this toolchain. These sources are also provided in the tarball. Users can comment the `get_src` function in the `make-aarch64-toolchain.sh` script.

Troubleshooting

Each step in the build process logs its progress and errors to a file named for that stage:

Build Process Step	Log File
•Downloading sources	• <code>get_src.log</code>
•Applying patches	• <code>apply_patches.log</code>
•Build binutils	• <code>build_binutils.log</code>
•Installing Linux headers	• <code>install_linux_headers.log</code>
•Building GCC stage 1	• <code>build_gcc_stage1.log</code>
•Building Glibc stage 1	• <code>build_glibc_stage1.log</code>
•Building final Glibc	• <code>build_final_glibc.log</code>

If the build succeeds, the script prints “Success!” at the end. Otherwise check log files for the last step shown to find the error.

The following are two common build errors:

- If using Ubuntu 14.04, install the necessary packages, in addition to the ones mentioned in the README, with the following commands:

```
$ sudo apt-get install gawk
$ sudo apt-get install texinfo
$ sudo apt-get install automake
$ sudo apt-get install libtool
$ sudo apt-get install g++
```

- The build script downloads sources for the toolchain components from well-known locations. These URLs may become outdated over time. Verify that the URLs in the `get_src()` function are up-to-date. The toolchain components are also provided as separate tarball that users can use directly.

MB1 Platform Configuration

This topic provides guidance on customizing static platform-specific settings in the MB1 boot configuration tables (BCT).

In the boot sequence, MB1 uses the MB1 BCT to configure platform-specific static settings. MB1 executes before all other CPUs are enabled. The MB1 stage is owned by NVIDIA and signed by NVIDIA and the OEM.

For more information, see [Boot Flow \(Tegra code-name Parker\)](#).

About MB1 BCT

MB1 BCT specifies platform-specific data. When Tegraflash is invoked to flash a platform, it calls the `tegrabct_v2` tool to create the MB1 BCT. It uses the following data:

- Platform configuration files
- `tegrab1_mb1_bct.h` header file

The BCT required by this stage is signed by the OEM. The MB1 stage offers to perform platform specific initialization. It also sets up the secure control register (SCR).

Platform-specific configuration files specify:

- Pinmux and GPIOs configuration
- Prod setting
- Pad voltage setting
- PMIC setting
- Secure register configuration

Location of Configuration Files

These configuration files located at:

```
<top>/bootloader/<platform|ver>/BCT/
```

The following sections describe these configuration files.

Pinmux, and GPIO Configuration

The pinmux configuration file provides pinmux and GPIO configuration information. The typical format for this data is register address and data, as a pair. MB1 only allows writes to the pinmux and GPIO address range, from this table.

The pinmux configuration file is located at:

```
<top>/bootloader/<platform|ver>/BCT/tegra<t-arch>-mb1-bct-pinmux-gpio-<board>-<board_re
```


Usage

```
pinmux.<address> = <value>;
```

Where:

- `pinmux` is the domain name for GPIO and pinmux configuration data.
- `<address>` is the absolute register address.
- `<value>` is the 32-bit data value.

Device-side implementation

```
write(value, address);
```

Example

```
#### Pinmux for used pins ####
pinmux.0x02434060 = <value1>; # gen1_i2c_scl_pc5.PADCTL_CONN_GEN1_I2C_SCL_0
pinmux.0x02434064 = <value2>; # gen1_i2c_scl_pc5.PADCTL_CONN_CFG2TMC_GEN1_I2C_SCL_0
pinmux.0x02434068 = <value1>; # gen1_i2c_sda_pc6.PADCTL_CONN_GEN1_I2C_SDA_0
pinmux.0x0243406C = <value2>; # gen1_i2c_sda_pc6.PADCTL_CONN_CFG2TMC_GEN1_I2C_DA_0
    ::::
#### Pinmux for unused pins for low-power configuration ####
pinmux.0x02434040 = <value1>; # gpio_wan4_ph0.PADCTL_CONN_GPIO_WAN4_0
pinmux.0x02434044 = <value2>; # gpio_wan4_ph0.PADCTL_CONN_CFG2TMC_GPIO_WAN4_0
pinmux.0x02434048 = <value1>; # gpio_wan3_ph1.PADCTL_CONN_GPIO_WAN3_0
pinmux.0x0243404C = <value2>; # gpio_wan3_ph1.PADCTL_CONN_CFG2TMC_GPIO_WAN3_0
```

Prod Configuration

The `prod` setting is the configuration of system characterization and interface and controller settings. This is needed for the interface to work in given platform reliably. The `prod` setting is set at the controller level and separately at the pinmux level. The examples below describe pinmux configuration only.

The format of this configuration is a tuple of register address, mask, and data value. MB1 reads data from address, modifies it based on mask and value, and then writes it back to address.

The prod configuration file is located at:

```
<top>/bootloader/<platform|ver>/BCT/tegra<t-arch>-mb1-bct-prod-<board>-<board_revision>
```

Usage

```
prod.<address>.<mask> = <value>;
```

Where

- `prod` is the domain name prefix for the setting.
- `<address>` is the pad control register address.
- `<mask>` is the mask value (4 bytes, unsigned).
- `<value>` is the data value (4 bytes, unsigned).

Device-side implementation

```
val = read(address)
val = (val & ~mask) | (value & mask);
write(val, address);
```

Example

```
prod.0x02436010.0x00006000 = 0x00002000; # SDMMC4_DAT7, DRV_TYPE: DRIVE_2X
prod.0x02436014.0x00006000 = 0x00002000; # SDMMC4_DAT6, DRV_TYPE: DRIVE_2X
prod.0x02436018.0x00006000 = 0x00002000; # SDMMC4_DAT5, DRV_TYPE: DRIVE_2X
prod.0x0243601c.0x00006000 = 0x00002000; # SDMMC4_DAT4, DRV_TYPE: DRIVE_2X
prod.0x02436020.0x00006000 = 0x00002000; # SDMMC4_DAT3, DRV_TYPE: DRIVE_2X
```

Pad Voltage Configuration

Tegra pins and pads are designed to support multiple voltage levels at a given interface. They can operate at 1.2 volts (V), 1.8 V or 3.3 V. Based on the interface and power tree of a given platform, SW needs to write the correct voltage of these pads to enable interface. If pad voltage is higher than IO power rail then pin will not work on that level. If pad voltage is lower than IO power rail then it can damage the SoC pads. Hence it is require configuring the correct pad voltage based on power tree.

The pad configuration file is located at:

```
<top>/bootloader/<platform|ver>/BCT/tegra<t-arch>-mb1-bct-pad-quill-<board>-<board_rev|
```

Usage

```
pmc.<address> = <value>;
```

Where:

- `pmc` is the domain name prefix for the setting.
- `<address>` is the absolute register address.
- `<value>` is the 32-bit data value.

Device-side implementation

```
write(value, address);
```

Example

```
pmc.0x0c36003c = 0x0000003e; # PMC_IMPL_E_18V_PWR_0
```

```
pmc.0x0c360040 = 0x00000079; # PMC_IMPL_E_33V_PWR_0
```

PMIC Configuration

During system boot, MB1 enables system power rails such as CPU, SRAM, and CORE as well as some system PMIC configurations. The typical configurations are:

- Enabling rails
- Setting rail voltages
- FPS configurations

Enabling and setting of voltages of rails may require:

- I2C command to devices
- MMIO access to Tegra registers, either read-modify-write or write-only
- Delay after the commands

Rail-specific configurations, such as I2C commands, MMIO access, and delays, are platform-specific. The MB1 BCT configuration file must provide configuration information.

The MB1-CFG format supports:

- I2C, PWM commands and MMIO commands on any sequence.
- Any I2C/PWM controller instance.
- Any 7-bit slave address of the device.
- MMIO commands on read-modify-write format to support read only and Read-modify-write format.
- I2C commands are read-modify-write format to support read only and Read-modify-write format.
- PWM commands are for enabling and configuring the PWM.
- Any amount of delay between commands.
- Write only commands for PWM/I2C/MMIO.
- Any size of device registers address and data size for i2c commands.
- I2c command on the 400KHz.
- The sequence may be
 - 1 MMIO, 1 I2C
 - 1 I2C, 1 MMIO
 - 2 MMIO, 1 I2C
 - 1 MMIO, 2 I2C

The typical rail/configurations are divided into following PMIC command domains:

- Generic: General PMIC configurations.
- CPU: Command related to CPU rails only.
- GPU: Commands related to GPU only.

- SRAM: Commands related to SRAM only.
- CORE: Commands related to CORE only.
- MEM: Commands related to Memory only.
- THERMAL: For thermal configurations.
- SHUTDOWN: For shutdown related configurations.

If there is not any configuration for given rail then there is no need to provide the command sequence of that rails. MB1 device side code ignores the configuration of that rail.

Each rail is defined unique ID to make the parsing and BCT binary to easier. The unique IDs are described in the following table.

Rail Name	ID
•GENERIC	•1
•CPU	•2
•CORE	•3
•SRAM	•4
•GPU	•5
•MEM	•6
•THERMAL	•7
•SHUTDOWN	•8

The PMIC configuration file is located at:

```
<top>/bootloader/<platform|ver>/BCT/tegra<t-arch>-mb1-bct-pmic-vcm31-<board>-<board_rev>
```

Usage

Common parameters:

```
pmic.<parameter> = <value>;
```

Rail-specific parameters:

```
pmic.<rail-names>.<rail-id>.<parameters> = <value>;
```

Where common <parameters> is one of the following:

Parameter	Description
•command-retries-count	•The number of allowed command attempts.
•wait-before-start-bus-clear-us	•Wait timeout, in microseconds before issuing the bus clear command. The wait time is calculated

	as $1 \ll n$ microseconds where n is provided by this parameter.
•rail-count	•Number of rails in this configuration file that need to be configured.

Example

```
pmic.command-retries-count = <value>;
pmic.wait-before-start-bus-clear-us = <value>;
pmic.rail-count = <value>;
```

Rail-Specific Parameters

Rail-specific parameters take the following format:

- The rail specific commands are divided into blocks.
- Each rail can have one or more blocks. Each block of given rails are indexed starting from 0.
- Each block contains either MMIO or I2C commands. If both MMIO and I2C commands are required then commands are broken into multiple blocks.
- If block contains I2C type of commands then all commands are sent to same device. If it is require having i2c commands for multiple devices then it needs to split into multiple blocks.
- If commands on given blocks are I2C type the device address, register address size, register data size are parameters which is not needed for MMIO commands.
- Given block can contain more than one commands but all commands are same type.
- Delay is provided after each commands of a given blocks. The delay is the same for all commands. If different delay are required then it need to split into multiple blocks.

Rail specific parameters are prefixed by the following:

```
pmic.<rail-name>.<rail-id>.
```

block-count

```
pmic.<rail-name>.<rail-id>.block-count = <value>;
```

Where <value>, for block-count, is the number of command blocks for a given rail.

block

The block identification parameter. All blocks are indexed, starting from 0.

```
pmic.<rail-name>.<rail-id>.block[index]
```

type

The type parameter specifies the command type, either MMIO (0) or I2C (1).

delay

The delay parameter specifies the delay in microseconds after each command in a given block.

count

The count parameter is the number of commands in a block.

I2C type-specific parameters

The following are I2C type specific parameters.

Parameter	Description
•i2c-controller-id	•Controller ID of I2C.
•slave-add	•7-bit slave address.
•reg-data-size	•Register size in bits: 0 or 8:1 byte •16: 2 byte
•reg-add-size	•Register address size in bits: 0 or 8:1 byte •16: 2 byte

Commands

Commands can be either MMIO or I2C. The information is in the format `<address>.<mask> = <data>`, to support the read-modify-write sequence. All commands are indexed, to facilitate multiple commands in a given block. Commands are sent to device in sequence, starting from index 0, in the following format:

```
commands[command-index].<addr>.<mask> = <data>;
```

PWM specific commands

The following are the PWM specific commands.

Command	Value
•type	•2 (for PWM =2)
•controller-id	•<0 to 7> /* Based on the platform */
•source-frq-hz	•<PWM clock source freq> /* In Hz */
•period-ns	•<Period in Nano Seconds> /* PWM periods */
•min-microvolts	•<Vout from PWM regulator if duty cycle is 0>
•max-microvolts	•<Vout from PWM regulator if duty cycle is 100>
•init-microvolts	•<Vout from PWM regulator after initialization.

•enable	•<1/0> /* Enable the PWM or just configure */
---------	---

For example:

```
# 1. Set 950mV voltage.
pmic.core.3.block[0].type = 2; # PWM Type
pmic.core.3.block[0].controller-id = 5; #GP_PWM6
pmic.core.3.block[0].source-frq-hz = 102000000; #102MHz
pmic.core.3.block[0].period-ns = 2600; # 384K is period.
pmic.core.3.block[0].min-microvolts = 600000;
pmic.core.3.block[0].max-microvolts = 1200000;
pmic.core.3.block[0].init-microvolts = 950000;
pmic.core.3.block[0].enable = 1;
```

Generic Format

The code excerpts in this section show the common and rail specific parameters in a generic format.

The common parameters shown are:

```
pmic.command-retries-count = <u32>;
pmic.wait-before-start-bus-clear-us = <u32>;
pmic.rail-count = <u32>;
```

The rail-specific parameters shown are:

```
pmic.<rail-name>.block-count
```

The generic format is as follows:

```
##### BLOCK 0 #####
pmic.<rail-name>.<rail-id>.block[0].type = <0 for MMIO, 1 for I2C>
pmic.<rail-name>.<rail-id>.block[0].delay = <u32>
pmic.<rail-name>.<rail-id>.block[0].count = <calculated>;

#For I2C specific
pmic.<rail-name>.<rail-id>.block[0].I2c-controller-id = <u32>;
pmic.<rail-name>.<rail-id>.block[0].slave-add = <u32>;
pmic.<rail-name>.<rail-id>.block[0].reg-data-size = <u32>;
pmic.<rail-name>.<rail-id>.block[0].reg-add-size = <u32>;

#I2C and MMIOs
pmic.<rail-name>.<rail-id>.block[0].commands[0].<addr>.<mask> = <data>;
```

```

pmic.<rail-name>.<rail-id>.block[0].commands[1].<addr>.<mask> = <data>;
pmic.<rail-name>.<rail-id>.block[0].commands[2].<addr>.<mask> = <data>;
pmic.<rail-name>.<rail-id>.block[0].commands[3].<addr>.<mask> = <data0>;
::::

##### BLOCK 1 #####
pmic.<rail-name>.<rail-id>.block[1].type = <0 for MMIO, 1 for I2C>
pmic.<rail-name>.<rail-id>.block[1].delay = <u32>
pmic.<rail-name>.<rail-id>.block[1].count = <Calculated>
#For I2C
pmic.<rail-name>.<rail-id>.block[1].I2c-controller-id
pmic.<rail-name>.<rail-id>.block[1].slave-add
pmic.<rail-name>.<rail-id>.block[1].reg-data-size
pmic.<rail-name>.<rail-id>.block[1].reg-add-size

#I2C and MMIOs
pmic.<rail-name>.<rail-id>.block[1].commands[0].<addr>.<mask> = <data>;
pmic.<rail-name>.<rail-id>.block[1].commands[1].<addr>.<mask> = <data>;
pmic.<rail-name>.<rail-id>.block[1].commands[2].<addr>.<mask> = <data>;
pmic.<rail-name>.<rail-id>.block[1].commands[3].<addr>.<mask> = <data0>;
::::

```

For example, in usage:

```

pmic.command-retries-count = 1;
pmic.wait-before-start-bus-clear-us = 0;
pmic.rail-count = 6;

#####GENERIC RAIL (ID = 1) DATA #####
pmic.generic.1.block-count = 1;

# 1. Set PMIC MBLDP = 1, CNFGGLBL1 bit 6 = 1
pmic.generic.1.block[0].type = 1; # I2C Type
pmic.generic.1.block[0].i2c-controller-id = 4;
pmic.generic.1.block[0].slave-add = 0x78; # 7Bit:0x3c
pmic.generic.1.block[0].reg-data-size = 8;
pmic.generic.1.block[0].reg-add-size = 8;
pmic.generic.1.block[0].delay = 10;

```



```

pmic.generic.1.block[0].count = 1;
pmic.generic.1.block[0].commands[0].0x00.0x40 = 0x40;

##### #CORE RAIL (ID = 3) DATA #####
pmic.core.3.block-count = 2;

# 1. Set 950mV voltage.
pmic.core.3.block[0].type = 1; # I2C Type
pmic.core.3.block[0].i2c-controller-id = 4;
pmic.core.3.block[0].slave-add = 0x70; # 7Bit:0x38
pmic.core.3.block[0].reg-data-size = 8;
pmic.core.3.block[0].reg-add-size = 8;
pmic.core.3.block[0].delay = 1000;
pmic.core.3.block[0].count = 1;
pmic.core.3.block[0].commands[0].0x07.0xFF = 0x2E;

# 2. Set GPIO3 Power down slot to 6.
pmic.core.3.block[1].type = 1; # I2C Type
pmic.core.3.block[1].i2c-controller-id = 4;
pmic.core.3.block[1].slave-add = 0x78; # 7Bit:0x3c
pmic.core.3.block[1].reg-data-size = 8;
pmic.core.3.block[1].reg-add-size = 8;
pmic.core.3.block[1].delay = 10;

```

PWM specific example:

```

pmic.core.3.block-count = 1;

# 1. Set 950mV voltage.
pmic.core.3.block[0].type = 2; # PWM Type
pmic.core.3.block[0].controller-id = 5; #GP_PWM6
pmic.core.3.block[0].source-frq-hz = 102000000; #102MHz
pmic.core.3.block[0].period-ns = 2600; # 384K is period.
pmic.core.3.block[0].min-microvolts = 600000;
pmic.core.3.block[0].max-microvolts = 1200000;
pmic.core.3.block[0].init-microvolts = 950000;
pmic.core.3.block[0].enable = 1;

```

: : : :

Secure PMC Scratch Register Configurations for BootROM

Note: In the Tegra Parker cold boot path for L1 and L2, it may be required to program I2C and MMIO commands to bring PMIC rails to OTP values. Contact NVIDIA for more information.

The following table lists typical L1 and L2 reset levels.

Reset Cause	Reset Level	Context Preservation Required	Trigger New Reset Event?
•Cold Power On	•L0	•No	•No
•Fifth WDT expiration	•L1	•No	•Yes, Trigger L0 BR
•Sensor/AO tag	•L1	•No	•Yes, trigger shut down from boot ROM
•VF Sensor	•L1	•No	•No
•SW_MAIN	•L1	•No	•Depends on AO scratch setting
•HSM	•L1	•AO scratch	•Yes, Trigger L0/L1 from MB1
•Fourth WDT expiration	•L2	•DRAM, internal RAM, L1 reset logic	•No
•SC7	•Warm	•DRAM, AO scratch	•No
•SC8	•Warm	•AO scratch, AO context	•No

The I2C commands are provided to the boot ROM via AO scratch registers. MB1 device-side code configures these AO scratch registers based on the platform specific information for boot ROM reset paths in the MB1 BCT.

1. The boot ROM requires the data on AO secure scratch registers.
2. There are 8 reset cases on which bootrom issue the commands. These are:
 - Watchdog 5 expiry
 - Watchdog 4 expiry
 - SC7 exit
 - SC8 exit
 - SW-Reset
 - AO-Tag/sensor reset
 - VF Sensor
 - HSM

3. Each reset case can have three sets of AO blocks of commands.
4. Each AO block has multiple blocks and each block can have multiple commands

In configuration format, AO blocks come first, then all reset conditions are initialized with the ID of the AO blocks.

Boot ROM commands are prefixed by `bootrom`. For each reset condition there are three sets of AO blocks, called `aocommand`.

```
bootrom.<reset-name>.aocommand[0] = <ao block ID>
bootrom.<reset-name>.aocommand[1] = <ao block ID>
bootrom.<reset-name>.aocommand[2] = <ao block ID>
```

Absence of `aocommand` or reset name in configuration data is treated as if there is no command for that reset condition.

AO Block Parameters

Parameters used in AO block configuration data are described in the following table.

Parameter	Description
•aoblock-count	•Number of AO sets described in the file.
•command-retries-count	•The number of allowed command attempts.
•delay-between-commands-us	•Wait timeout, in microseconds, before issuing the bus clear command. The wait time is calculated as $1 \ll n$ microseconds where n is provided by this parameter.
•wait-before-start-bus-clear-us	•Wait time in microseconds before issuing the bus clear command.
•block-count	•Number of blocks in the AO block.
•type	•Command type whether MMIO or I2C. Use value as 0 for MMIO and 1 for I2C.
•count	•Number of commands in this block.

I2C type-specific parameters

The following are I2C type specific parameters.

Parameter	Description
•i2c-controller-id	•Controller ID of I2C.
•slave-add	•7-bit slave address.
•reg-data-size	•Register size in bits: 0 or 8:1 byte

	<ul style="list-style-type: none"> •16: 2 byte
•reg-add-size	<ul style="list-style-type: none"> •Register address size in bits: 0 or 8:1 byte •16: 2 byte

Commands

Commands can be either MMIO or I2C. The information is in the format `<address> = <data>`, to support the write-only sequence. All commands are indexed, to facilitate multiple commands in a given block. Commands are sent to the device in sequence, starting from index 0, in the following format

```
commands[command-index].<addr> = <data>;
```

All reset conditions support three AO blocks, initialized as follows:

```
bootrom.<reset-name>.aocommand[0] = <ao block ID>
bootrom.<reset-name>.aocommand[1] = <ao block ID>
bootrom.<reset-name>.aocommand[2] = <ao block ID>
```

Where `<reset_name>` is one of the following: `watchdog5`, `watchdog4`, `sc7`, `sc8`, `soft-reset`, `sensor-aotag`, `vfsensor`, or `hsm`.

Example

```
bootrom.aoblock-count = 2;

# Automatic power cycling: Set MAX77620
# Register ONOFFCNFG2, bit SFT_RST_WK = 1 (default is "0" after cold boot),
# Register ONOFFCNFG1, bit SFT_RST = 1
bootrom.aoblock[0].command-retries-count = 1;
bootrom.aoblock[0].delay-between-commands-us = 1;
bootrom.aoblock[0].wait-before-start-bus-clear-us = 1;
bootrom.aoblock[0].block-count = 1;

bootrom.aoblock[0].block[0].type = 0; # I2C Type
bootrom.aoblock[0].block[0].slave-add = 0x3c; # 7Bit:0x3c
bootrom.aoblock[0].block[0].reg-data-size = 8;
bootrom.aoblock[0].block[0].reg-add-size = 8;
bootrom.aoblock[0].block[0].count = 2;
bootrom.aoblock[0].block[0].commands[0].0x42 = 0xda;
bootrom.aoblock[0].block[0].commands[1].0x41 = 0xf8;

# Shutdown: Set MAX77620
```

```

# Register ONOFFCNFG2, bit SFT_RST_WK = 0
# Register ONOFFCNFG1, bit SFT_RST = 1
bootrom.aoblock[1].command-retries-count = 1;
bootrom.aoblock[1].delay-between-commands-us = 1;
bootrom.aoblock[1].wait-before-start-bus-clear-us = 1;
bootrom.aoblock[1].block-count = 1;

bootrom.aoblock[1].block[0].type = 0; # I2C Type
bootrom.aoblock[1].block[0].slave-add = 0x3c; # 7Bit:0x3c
bootrom.aoblock[1].block[0].reg-data-size = 8;
bootrom.aoblock[1].block[0].reg-add-size = 8;
bootrom.aoblock[1].block[0].count = 2;
bootrom.aoblock[1].block[0].commands[0].0x42 = 0x5a;
bootrom.aoblock[1].block[0].commands[1].0x41 = 0xf8;

# Shutdown in sensor/ao-tag
#reset in soft reset.
# no commands for other case
bootrom.sensor-aotag.aocommand[0] = 1;
bootrom.soft-reset.aocommand[0] = 0;

```

Security Configuration Register Configurations

Tegra code-name Parker has separate registers for configuring bridge client security, bridge firewalls, known as security configuration registers (SCRs). SCRs are either configured by NVIDIA for Tegra code-name Parker platforms or re-configured for custom platforms. Custom configuration is via MB1 BCT at the MB1 stage.

The list of SCRs for re-configuration for custom platforms is the same in MB1 device-side code and SCR platform data. SCR register addresses hard coded in MB1 can only contain data from the platform. The data cannot be masked, and can only be written to the registers as is.

The SCR configuration file is located at:

```
<top>/bootloader/<platform|ver>/BCT/auto_scr.cfg
```

Usage

```
scr.<reg_index>.<exclusion-info> = <32 bit value>; # <reg_name>
```

Where `scr` is the domain name prefix for the setting `<reg_index>` is the matching MB1 and CFG file sequence, beginning at 0; and `exclusion-info` is one of the following values:

Value	Description
•0	• Include: regular SCRs loaded from BCT in cold boot, from stored context in warm boot.
•1	• Exclude: Present data in the CFG file but do not load data from the BCT. Allows SCR programming in MB2 or later.
•2	• SC7 resume: Program from BCT in cold boot, but exclude for warm boot.

MB1 code lists SCR register absolute addresses in an indexed list.

Example

```
# SCR register configurations
scr.134.1 = 0x20000000; # GPIO_M_SCR_00_0
scr.135.1 = 0x20000000; # GPIO_M_SCR_01_0
```

Miscellaneous Configurations

This section documents several settings that do not fit into any of the other categories.

The miscellaneous configuration file is located at:

```
<top>/bootloader/<platform|ver>/BCT/tegra186-mb1-bct-misc-vcm31-p2382.cfg
```

The following table describes the fields contained in `misc.cfg`.

Field	Description	Configuration Example
•enable_can_boot	•Controls early CAN initialization. If set, spe-can firmware loading spe-r5 processor boot is done.	•enable_can_boot = 1;
•enable_blacklisting	•Controls DRAM ECC blacklisting: •0: Disable ECC blacklisting 1: Enable ECC blacklisting	•enable_blacklisting = 0;
•disable_sc7	•Controls SC7 state entry: •0: Enable sc7 1: Disable sc7	•disable_sc7 = 0;
•fuse_visibility	•Certain fuses cannot be read or written by default because they are not visible. If this field is set, MB1 enables fuse visibility for such fuses.	•fuse_visibility = 1;
•enable_vpr_resize	•Controls enablement of VPR resize functionality.	•enable_vpr_resize=0
•Disable_el3_bl	•Used for eliminate execution of EL3	•Disable_el3_bl = 1

	Bl after secure os and can start EL2 bootloader	
--	---	--

Debug

There are debug functionalities which can be enabled disabled via BCT flag.

Debug Control Fields	Description	Configuration Example
•uart_instance	•Configures the UART instance for console prints.	•debug.uart_instance = 1;
•enable_log	•Enables/disables console logging.	•debug.enable_log = 1;
•enable_secure_settings	•Unused.	•-

AOTAG

The AO-TAG register is programmed in the MB1, which controls the maximum temperature at which Tegra Parker is allowed to operate. If the temperature exceeds that limit, auto-shutdown is triggered.

AOTag Control Fields	Description	Configuration Example
•boot_temp_threshold	•Boot temperature threshold in millicentigrade. If temperature is higher than the temperature specified in this field, MB1 waits or shuts down the device.	•aotag.boot_temp_threshold = 105000;
•cooldown_temp_threshold	•Cool down temperature threshold in millicentigrade. MB1 resumes booting when the device has cooled to this threshold temperature.	•aotag.cooldown_temp_threshold = 85000;
•cooldown_temp_timeout	•Contains max time MB1 should wait for system temperature to go down below “cooldown_temp_threshold”.	•Cooldown_temp_timeout = 30000
•enable_shutdown	•If set to 1, enables shutdown using aotag if temperature is above boot temperature threshold.	•aotag.enable_shutdown = 1;

Clock

The clock control fields in the following table hold the clock divider values for the various modules that MB1 programs.

Clock Control Fields	Description	Configuration Example
•bpmp_cpu_nic_divider	<ul style="list-style-type: none"> •Program the cpu nic divider to control the BPMP CPU frequency. •A value 1 less than the value in the field is directly written to the register. 	•clock.bpmp_cpu_nic_divider = 1;
•bpmp_apb_divider	<ul style="list-style-type: none"> •Program the apb divider to control the APB bus frequency. •A value 1 less than the value in the field is directly written to the register. 	•clock.bpmp_apb_divider = 1;
•axi_cbb_divider	<ul style="list-style-type: none"> •Program the axi_cbb divider to control the AXI-CBB bus frequency. •A value 1 less than the value in the field is directly written to the register. 	•clock.axi_cbb_divider = 1;
•se_divider	<ul style="list-style-type: none"> •Program the se divider to control the SE Controller frequency. •A value 1 less than the value in the field is directly written to the register. 	•clock.se_divider = 1;
•aon_cpu_nic_divider	<ul style="list-style-type: none"> •Program the cpu_nic divider to control the AON(SPE) CPU frequency. •A value 1 less than the value in the field is directly written to the register. 	•clock.aon_cpu_nic_divider = 1;
•aon_apb_divider	<ul style="list-style-type: none"> •Program the apb divider to control the AON(SPE) APB frequency. •A value 1 less than the value in the field is directly written to the register. 	•clock.aon_apb_divider = 1;
•aon_can0_divider	<ul style="list-style-type: none"> •Program the can0 divider to control the CAN0 controller frequency. •A value 1 less than the value in the field is directly written to the register. 	•clock.aon_can0_divider = 1;
•aon_can1_divider	<ul style="list-style-type: none"> •Program the can1 divider to control the CAN1 controller frequency. •A value 1 less than the 	•clock.aon_can1_divider = 1;

	value in the field is directly written to the register.	
•osc_drive_strength	•Unused	•-
•pllaon_divp	•Program the P value of PLL-AON. •A value 1 less than the value in the field is directly written to the register.	•clock.pllaon_divp = 2;
•pllaon_divn	•Program the N value of PLL-AON. •A value 1 less than the value in the field is directly written to the register.	•clock.pllaon_divn = 25;
•pllaon_divm	•Program the M value of PLL-AON. •A value 1 less than the value in the field is directly written to the register.	•clock.pllaon_divm = 1;

CPU Parameters

These settings contain the initial settings passed to CPU-Init FW. Contact NVIDIA before changing any of these settings.

Field	Description	Configuration Example
•Bootcpu	•Specify Boot CPU. 4 means A57 cpu0 and 0 mean Denver0. For automotive applications use A57-cpu0.	•cpu.bootcpu = 4
•ccplex_platform_features	•Platform feature passed to the CPU-Init FW.	•cpu.ccplex_platform_features = 0x581;
•lsr_dvcomp_params_b_cluster	•Cluster setting for initializing ADC and DVC, which need to be functional before CPU rails are brought up	•cpu.lsr_dvcomp_params_b_cluster = 0xC0780F05C;
•lsr_dvcomp_params_m_cluster	•Cluster setting for initializing ADC and DVC, which need to be functional before CPU rails are brought up	•cpu.lsr_dvcomp_params_m_cluster = 0xC0780F05C;
•nafll_m_cluster_data	•Initial NAFLL settings for cluster for Denver	•cpu.nafll_m_cluster_data = 0x11F04461;
•nafll_b_cluster_data	•Initial NAFLL settings for cluster for A57	•cpu.nafll_b_cluster_data = 0x11F04461;

AST settings

This contains the AST settings for various firmware loaded by MB1/MB2. These are the virtual addresses of the firmware. MB1/MB2 programs corresponding physical addresses based on the location where it loaded the firmware in memory (DRAM). Normally there is no need to change these settings.

Fields	Description	Configuration Example
•bpmp_fw_va	•Virtual address for BPMP-FW	•ast.bpmp_fw_va = 0x50000000;
•mb2_va	•Virtual address for MB2-FW	•ast.mb2_va = 0x52000000;
•sce_fw_va	•Virtual address for SCE-FW	•ast.sce_fw_va = 0x70000000;
•apr_va	•Virtual address for Audio-protected region used by APE-FW	•ast.apr_va = 0xC0000000;
•ape_fw_va	•Virtual address for APE-FW	•ast.ape_fw_va = 0x80000000;

SW Carveout

These settings specify the address and size for BL carveout.

Field	Description	Configuration Example
•cpubl_carveout_addr	•Start location of the CPU-BL Carveout	•sw_carveout.cpubl_carveout_addr = 0x96000000;
•cpubl_carveout_size	•Size of the CPU-BL Carveout	•sw_carveout.cpubl_carveout_size = 0x02000000;
•mb2_carveout_size	•Size of the MB2 Carveout	•sw_carveout.mb2_carveout_size = 0x00400000;

I2C Settings

These settings specify the operating frequency of the I2C bus in MB1/MB2 (default is 100 KHz).

Field	Description	Configuration Example
•0	•Specify the clock for I2C controller instance 0	•i2c.0 = 400;
•4	•Specify the clock for I2C controller instance 4	•i2c.4 = 1000;

Dev Parameters

These are the device settings used by MB1/MB2.

Field	Description	Configuration Example
•qspi.clk_src	<ul style="list-style-type: none"> Specify the clock source. The value corresponds to what is mentioned in the QSPI CLK SRC register. 0: pll_out0 4: pll4_muxed 	<ul style="list-style-type: none"> devinfo.qspi.clk_src = 0; # •
•qspi.clk_div	<ul style="list-style-type: none"> clk_div = N+1; Hence N = 3 & clk_rate = 163.2 MHz = (408 MHz / ((N / 2) + 1)) 	<ul style="list-style-type: none"> devinfo.qspi.clk_div = 4; •
•qspi.width	<ul style="list-style-type: none"> Specify the width of the QSPI BUS during transfer 0 : 1 bit (x1 mode) 1 : 2 bit (x2 mode) 2 : 4 bit (x4 mode) 	<ul style="list-style-type: none"> devinfo.qspi.width = 2
•qspi.dma_type	<ul style="list-style-type: none"> Specify which DMA to use for transfer if mode of transfer is DMA. For QSPI, in MB1/MB2, BPMP-DMA should be used. 0 : GPC-DMA 1 : BPMP-DMA 	<ul style="list-style-type: none"> devinfo.qspi.dma_type = 1
•qspi.xfer_mode	<ul style="list-style-type: none"> Specify mode of transfer 0: PIO 1: DMA 	<ul style="list-style-type: none"> devinfo.qspi.xfer_mode = 1;
•qspi.read_dummy_cycles	<ul style="list-style-type: none"> The dummy cycles allow the device internal circuits additional time for accessing the initial address location. During the dummy cycles the data value on IOs are “don’t care” and may be high impedance. 	<ul style="list-style-type: none"> devinfo.qspi.read_dummy_cycles = 9
•qspi.trimmer_val1	<ul style="list-style-type: none"> tx_clk_tap_delay for QSPI 	<ul style="list-style-type: none"> devinfo.qspi.trimmer_val1 = 0
•qspi.trimmer_val2	<ul style="list-style-type: none"> rx_clk_tap_delay for QSPI 	<ul style="list-style-type: none"> devinfo.qspi.trimmer_val2 = 0

Watchdog timer controller settings

These settings specify watchdog timer controller register values. These values will be configured by MB1.

Field	Description	Configuration Example
-------	-------------	-----------------------

•bpmp_wdtcr	•Contains the bpmp processor watchdog timer register value	•wdt.bpmp_wdtcr = 0x710640; configures for 100sec
•Sce_wdtcr	•Contains the SCE processor watchdog timer register value	•wdt.sce_wdtcr = 0x707103;
•aon_wdtcr	•Contains aon's watchdog timer register value	•wdt.aon_wdtcr = 0x700000;
•rtc2_ao_wdtcr	•Contains rtc2_ao watchdog timer register value	•wdt.rtc2_ao_wdtcr = 0x700000;
•top_wdt0_wdtcr	•Contains top_wdt0 watchdog timer register value	•wdt.top_wdt0_wdtcr = 0x715016;
•top_wdt1_wdtcr	•Contains top_wdt1 watchdog timer register value	•wdt.top_wdt1_wdtcr = 0x710640;
•top_wdt2_wdtcr	•Contains top_wdt2 watchdog timer register value	•wdt.top_wdt2_wdtcr = 0x707103;

Generating MB1 Configuration Files

This topic provides details on how to generate different MB1 configuration files.

Pinmux, GPIO, and Pad

For information on configuring these components, refer to [Configuring Pinmux, GPIO and PAD](#).

PMIC

The PMIC configuration file is created manually. Here are the steps:

1. Get information about the set of commands to enable and setting voltage of each rail. If OTP values are in desired voltage, do not reprogram voltage register. Also do not enable all rails, only do what is recommended by boot sequence. The voltage for rail need to be set per the boot sequence recommendation.
2. Get information about generic setting needed from the MB1.
3. Once all information is there, split these commands per rail.
4. Make the list of commands sequence, delay between commands and then make blocks. Blocks can contain multiple commands if they are same type (i2C or MMIO), having same delay and communicating to same device. If any think different then it will be in the different blocks.
5. Create configuration file based on above details.

BootROM

The BOOTROM configuration file is created with manually. Here are the steps:

1. Get information about the set of commands to send to device in different reset path. There may be possibility that there is same type for commands for different reset paths. Collect all such information from system team.
2. Make the sets of commands needed for each reset path independently.
3. Pickup commands for one reset path, make the list of commands sequence, delay between commands and then make blocks out of these. Blocks can contain multiple commands if they are same type (i2C or MMIO), having same delay and communicating to same device. If any think different then it will be in the different blocks.
4. Put all blocks in one `aoblocks`.
5. Similarly make all `aoblocks` for all reset paths.
6. Initialize the different reset path's `aocommand` with these `aoblock` indexes.
7. Create configuration file based on above details.

Flashing

For information on flashing, refer to [Flashing the Boot Loader and Kernel](#).

Configuring Pinmux, GPIO and PAD

The pinmux and PAD configuration files provide pinmux and GPIO, and PAD configuration information.

For each board and chip configuration, a pinmux spreadsheet is provided to:

- Show the locations and default pinmux settings
- Define the pinmux settings in the source code or device tree

The spreadsheet is available at:

```
https://developer.nvidia.com/embedded/downloads
```

For information on the pinmux and GPIO, and PAD configuration files, refer to [MB1 Platform Configuration](#).

Pinmux Configuration Process

To configure pinmux settings, perform the following tasks:

- Modify the pinmux spreadsheet for changes in your configurable fields.
- Generate a DTSI file from spreadsheet.
- Generate a CFG file from DTSI file.

Generating DT File

For an initial pinmux configuration, generate a DTSI file.

To generate DT file

1. Open the pinmux spreadsheet in Microsoft Excel.
2. Click the Generate DT File button.

In response to the "Would you like to generate device tree file for pinmux table" prompt, click 'Yes'.

In response to the board name prompt, specify the board name. If board name is not specified, the spreadsheet uses the default value.

If the fields are error free, the script displays a message indicating success and creates the DTSI files in the current working directory.

For example:

```
tegra18x-p3310-1000-a03-cvm-config-pinmux.dtsi
tegra18x-p3310-1000-a03-cvm-config-gpio-default.dtsi
tegra18x-p3310-1000-a03-cvm-config-padvoltage-default.dtsi
```

3. Copy the files from Windows to the Linux system.

Generating Configuration files using dts2cfg

Note: Use the `dos2unix` command to convert the `.dts` file to be readable by the script. Make sure there are no blank lines at the end of the address info file.

The pinmux `dts2cfg` tool converts pinmux, GPIO, and the PAD DTS files to `.cfg` format. The tool is located in the following directory:

```
Linux_for_Tegra/kernel/pinmux/
```

To convert pinmux and GPIO files to CFG format

Usage

```
python pinmux-dts2cfg.py [--pinmux] <ADDRESS_FILE> <GPIO_ADDRESS_FILE> <POR_VAL_FILE> <
```

The following table describes the supported options.

Options	Description
<code>--pinmux <arguments></code>	•(Optional) Specifies pinmux arguments. By default, the script assumes the data is for pinmux unless the command line includes the <code>--pad</code> option.
<code>--mask</code>	•(Optional) Specifies the output (useful to build GR sheets).
<code>--help</code>	•Prints usage information.

The following table describes the `--pinmux` arguments.

<code>--pinmux</code> Arguments	Description
•ADDRESS_FILE	•Specifies the address info file.
•GPIO_ADDRESS_FILE is	•Specifies the GPIO address info file.
•POR_VAL_FILE	•Specifies the por val pair info file.
•MANDATORY_PINMUX_FILE	•Specifies the mandatory pinmux values info file.
•PINMUX_DTS_FILE	•Specifies the device tree source file for pinmux.

•GPIO_DTS_FILE	•Specifies the device tree source file for GPIO.
•VERSION	•Specifies the numeric version, like 1.0.

The resulting `.cfg` file is printed to standard output. You can save this file by redirecting the output to a file. Here is an example:

```
python pinmux-dts2cfg.py \
  --pinmux \
  addr_info.txt gpio_addr_info.txt por_val.txt \
  mandatory_pinmux.txt \
  tegra186-quill-p3310-1000-a00-pinmux.dtsi \
  tegra186-quill-p3310-1000-a00-gpio-default.dtsi \
  1.0 \
> tegra186-quill-p3310-1000-a00-pinmux-gpio.cfg
```

Copy the configuration file to the default location:

```
<top>/bootloader/t186ref/BCT/
```

To convert PAD file to CFG format

The `pinmux-dts2cfg.py` tool converts pad file to CFG format.

Usage

```
python pinmux-dts2cfg.py --pad <PAD_FILE> <PAD_DTS_FILE> <VERSION> [--mask] [--help]
```

The following table describes the supported options and arguments.

Options and Arguments	Description
--help	Prints usage information.
--pad <arguments>	Specifies pad attributes.
•--mask	Controls the output (useful to build GR sheets).

The following table describes the `--pad` option arguments.

--pad Arguments	Description
•PAD_FILE	Specifies the pad info file.

•PAD_DTS_FILE	Specifies the device tree source file for pads.
•VERSION	Specifies the numeric version, like 1.0.

Because the resulting CFG file is printed to standard output, it is recommended to redirect the output to a file. Here is an example:

```
python pinmux-dts2cfg.py \
--pad pad_info.txt \
tegra186-quill-char-e3301-1080-a00-padvoltage-default.dtsi \
1.0
> tegra186-mb1-bct-pad-quill-char-e3301-1080-a00.cfg
```

Copy the configuration file to the default location:

```
<top>/bootloader/t186ref/BCT/
```

Boot Flow

Note: This chapter is currently under review. We recommend that before you use this information, check with your NVIDIA® customer engineer (CE) for verification of its accuracy.

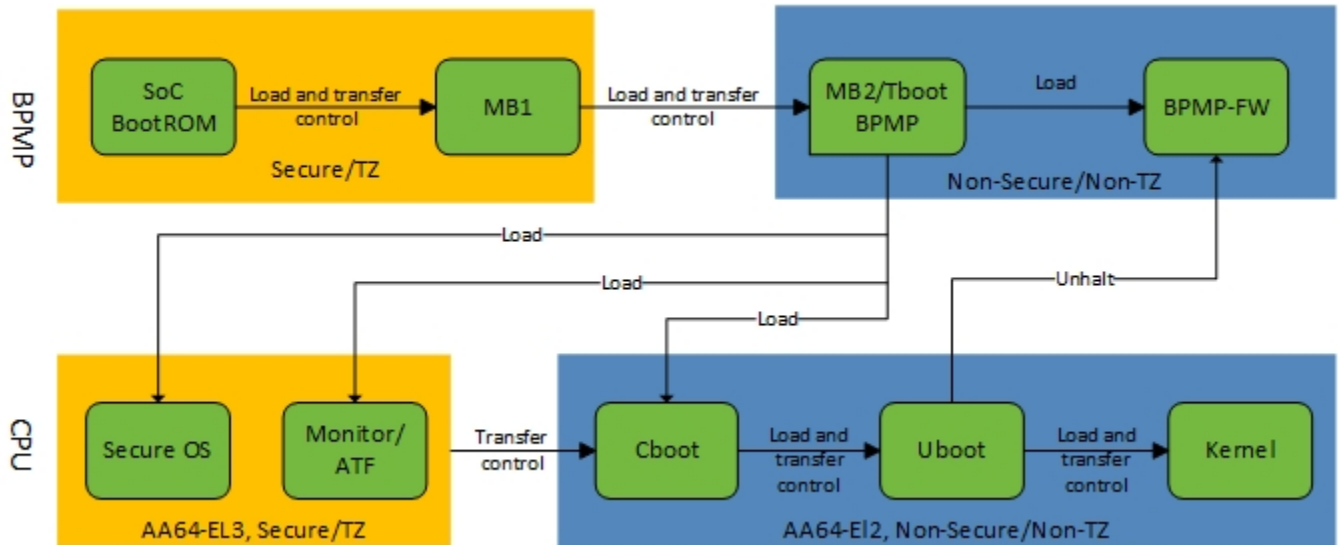
This topic describes the boot flow for NVIDIA® Tegra® code-name Parker for native environments.

Responsibilities of the Boot Software

The primary functionality of the different components of boot software is initialization of SOC (System On a Chip), including MC/EMC initialization and CPU initialization; setting up the security parameters; loading different firmware; maintaining Chain of Trust; setting memory carveouts for different firmware; flashing the device; and booting to the OS. Boot software also performs other operations defined by product requirements, including but not limited to display (HDMI/DSI), displaying the boot logo, boot OS specific update mechanisms and OS specific boot methods.

High Level Tegra Software Boot

The following diagram illustrates the flow of control in the boot software.



BootROM

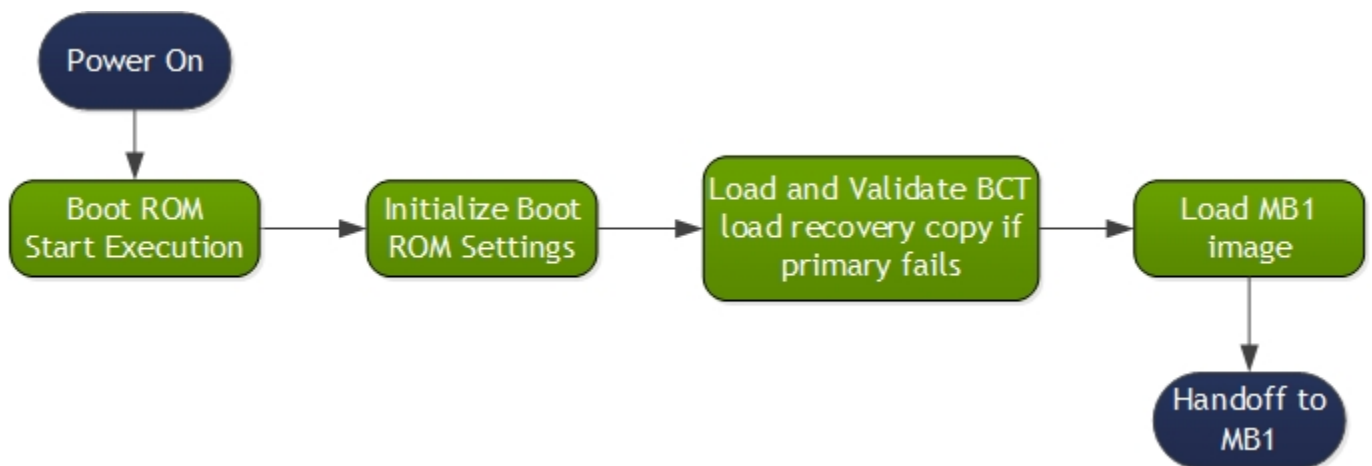
The BootROM (BR) is hard-wired in the Tegra chip. It initializes the Boot Media and loads Microboot1 (MB1) from the Boot Media.

Multiple copies of the BR-BCT (Boot Configuration Table) are stored at the start of the Boot-Media. The BR-BCT contains configuration parameters used by the BootROM for hardware initialization.

The BCT also contains information about the boot loader (BL), including:

- Size
- Entry point
- Load address
- Hash

The BootROM uses this information to verify and load the boot loader. The following diagram shows the boot flow.



Boot Loader Components

The boot loader and flash components are:

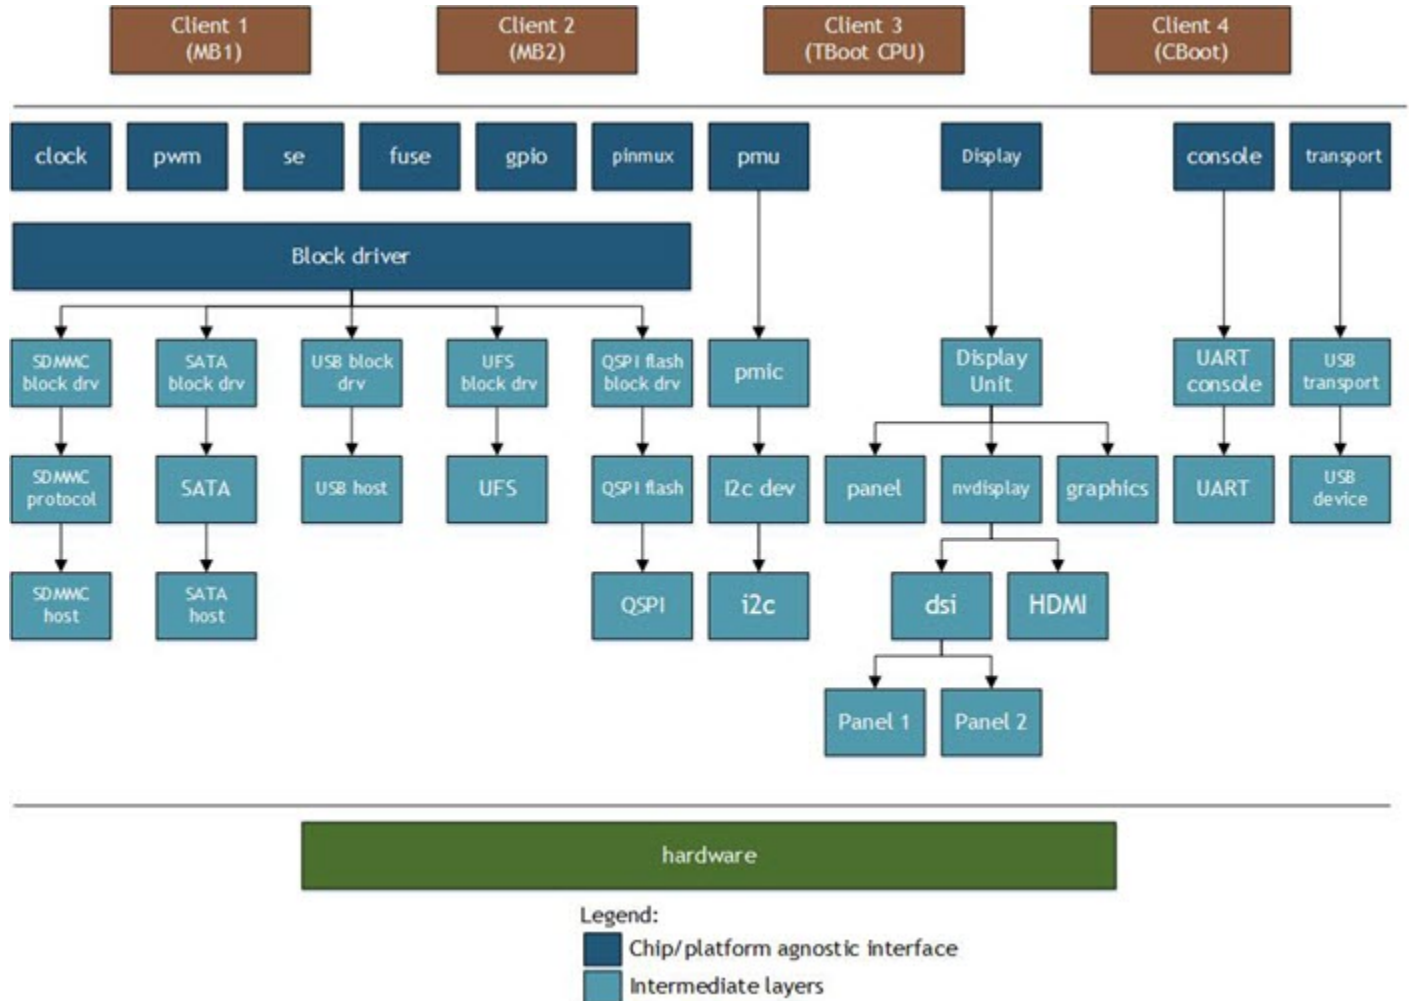
- Common Driver Framework (CDF)
- Microboot1 (MB1)
- TBoot-BPMP (MB2)
- CPU-BL
 - TBoot-CPU
 - CBoot
 - U-Boot

Common Driver Framework

MB1, TegraBoot-BPMP, and TegraBoot-CPU/CBoot each needs to use one of several drivers. Moreover, some libraries are common across the boot loader components. Instead of using separate sets of drivers and libraries, the binaries share a common pool of drivers and libraries called the Common Driver Framework (CDF).

Note: U-Boot does not support or use CDF as it's built separately from the NVIDIA-proprietary bootloader software.

The following diagram shows the software design architecture using CDF:



CDF consists of the following:

- Storage drivers like eMMC, QSPI, SATA
- Host interface drivers like UART, USB
- Display drivers
- Debug and Console library
- Libraries of other modules like clock, SE, PMIC, timer, fuse, GPIO, PWM, EEPROM, keyboard, and PSCI
- Software libraries like the cryptographic, transport, partition manager, sparse, and transport libraries
- A decompression library that supports LZF, Zlib, and LZ4
- `clib`, a dynamic memory allocation library, and a cache library

The CDF is in the following location:

[<top>/vendor/nvidia/tegra/bootloader/partner/common/](#)

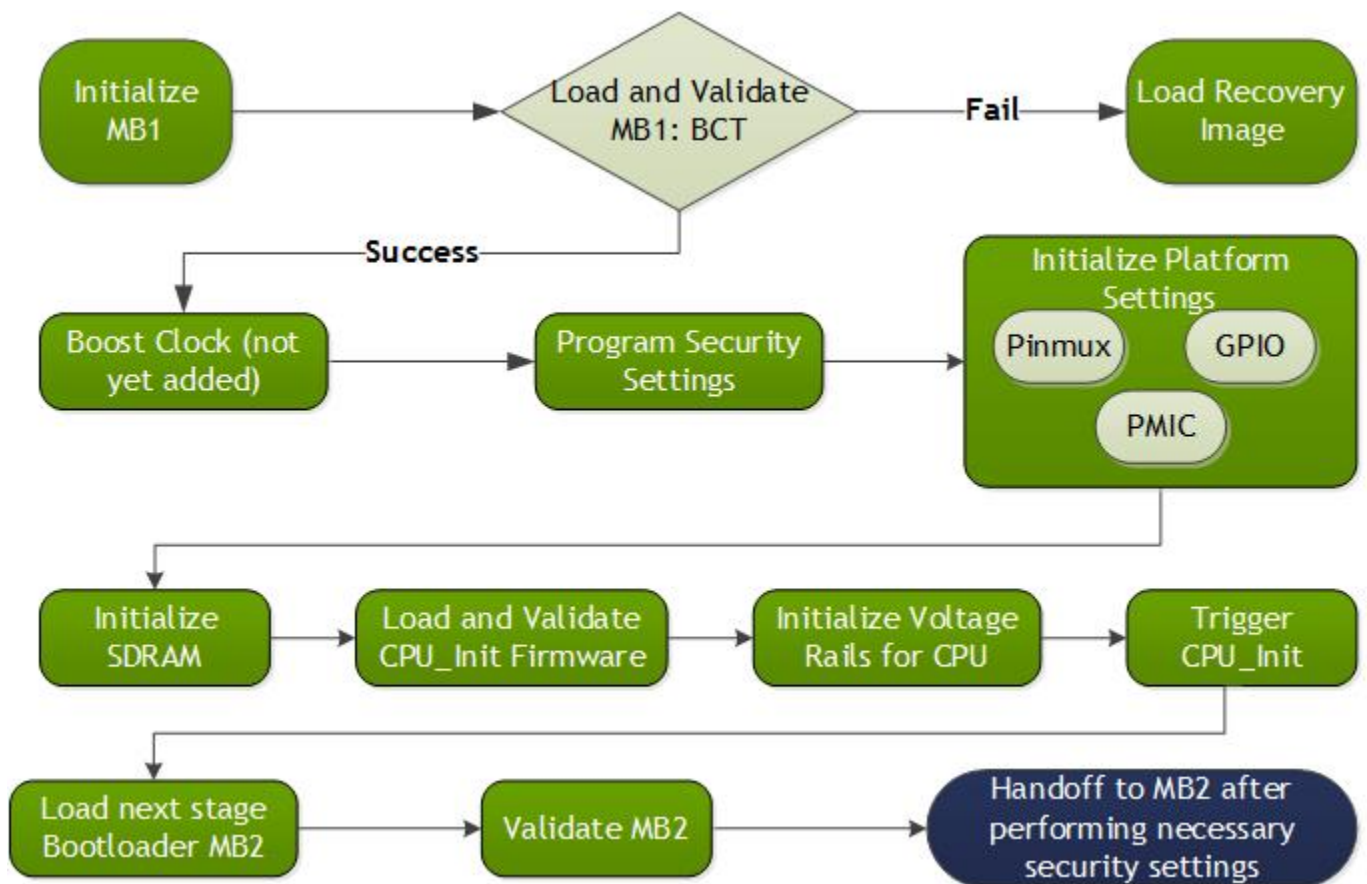
and

```
<top>/vendor/nvidia/tegra/bootloader/partner/t18x/common/
```

Microboot1

Microboot1 (MB1) is the first boot software component loaded by BR in SysRAM, and runs on BPMP. MB1 is considered as an extension to BR. MB1 provides extensive flexibility to alter or fix the boot sequence. This component also implements some delicate functionality like Denver initialization and security fabric changes.

Because MB1 is an extension to BootROM, it is signed and encrypted by an NVIDIA owned key (separate keys residing in BootROM). The following diagram shows the flow of control in MB1.



MB1 is responsible for:

- Firmware initialization, include MTS preboot, SPE, and CAN
- Platform configuration, which includes configuring pinmux/GPIO configuration, pad settings, prod settings, initializing certain rails including rails required by CPU, memory, GPU etc., PMIC commands required by BR, and SCR configuration
- Initialize the SDRAM based on the MB1 boot configuration table (MB1-BCT)
- Load the firmware that initializes the CPU complex (CCplex)
- Set the security settings

- Configure pinmux/GPIO for the SoC
- Program the PMIC for the VDD_CPU and VDD_DDR rails
- Creating carveouts
- Support flashing (available as a different `mb1_recovery.bin` binary)
- Support “RCM boot” by downloading a blob that contains the binaries, and loading binaries from that blob.
- Load the next stage boot loader, TBoot-BPMP (MB2)
- Support SC7

MB1 is owned by NVIDIA, and so is provided as a binary in the BSP package. Although it is provided as a binary, you can configure its behavior for a specific platform using its Boot Configuration Table, called MB1-BCT. MB1 reads MB1-BCT to configure the platform.

TegraBoot

TegraBoot is the boot loader component that executes after MB1. This component is divided into two portions: TBoot-BPMP (MB2) and TBoot-CPU, depending on the processor on which they are executing. Both are released to OEMs as source code for reference implementation.

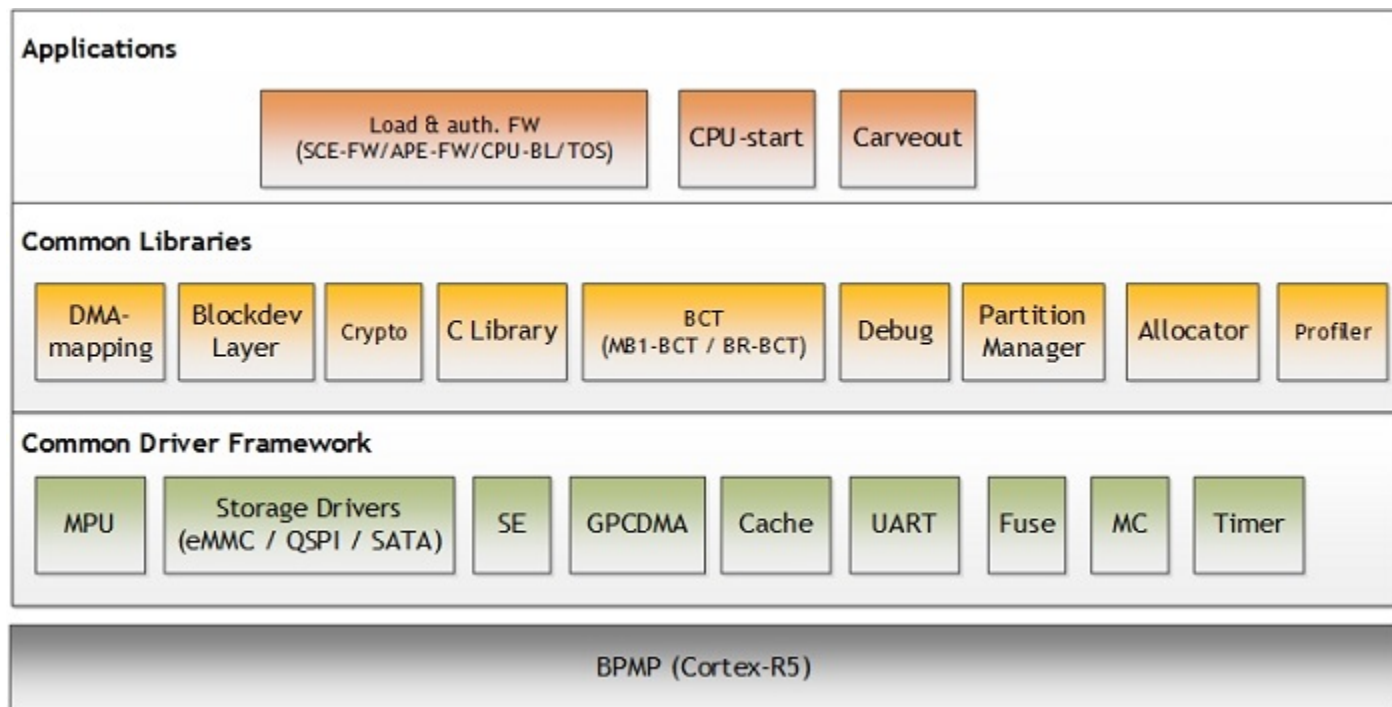
TegraBoot-BPMP(tboot-BPMP)

This portion of TegraBoot runs on BPMP. There are two variants of TegraBoot-BPMP (MB2), one used for cold boot and the other for recovery boot. Both are generated using different configuration files and make files.

TBoot-BPMP is responsible for:

- Loads and initializes firmware (FW) components such as SCE-FW, APE-FW, BPMP-FW, TOS
- Creates carveouts
- Completes CPU initialization
- Loads the next stage boot loader
- Supports flashing
- Supports “RCM boot”
- Reads PMIC reset reason
- Loads the bootloader device tree and passes the device tree load address to CPU-BL
- Gives control to BPMP-FW

The following block diagram shows the components of TBoot-BPMP:



TegraBoot-CPU

TegraBoot-CPU is a component of TegraBoot that runs on the CPU. Use TegraBoot-CPU to flash the device. You can also use TegraBoot-CPU for “RCM boot,” whose flow is very similar to a cold boot except that the binaries are transferred from the host over USB and are loaded directly to SDRAM. TOS and BPMP-FW are not loaded in this path.

TegraBoot-CPU is responsible for:

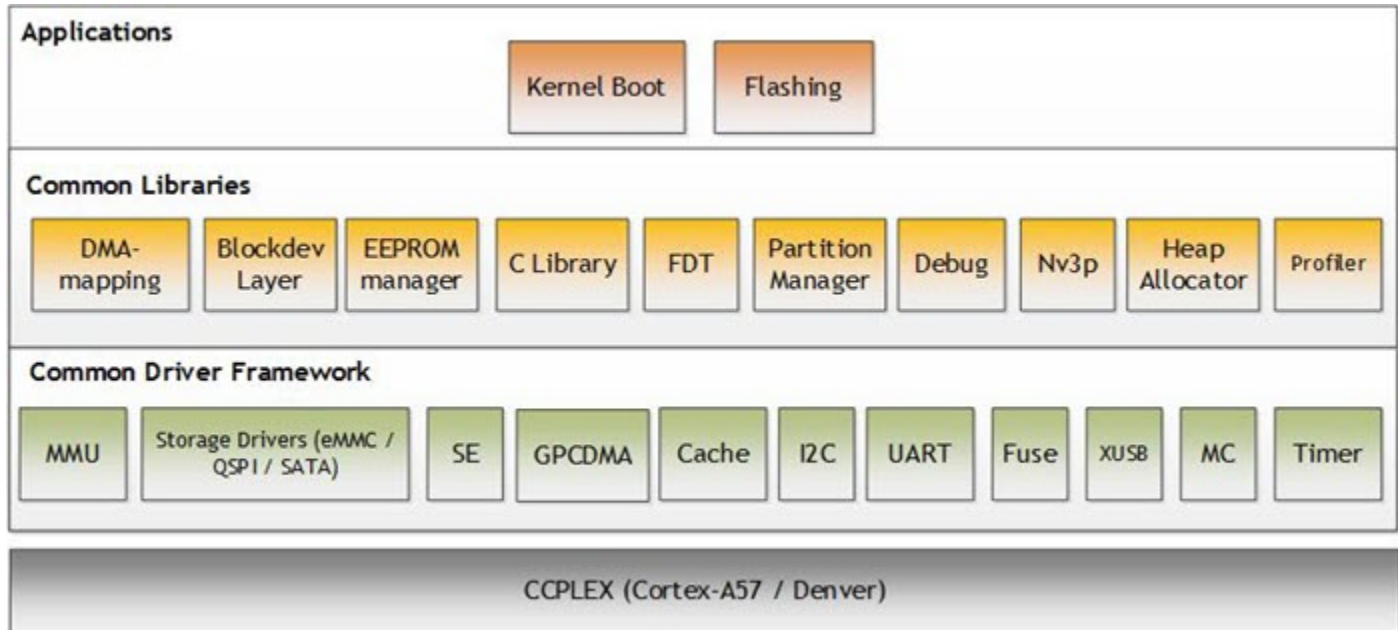
- Flashing
- RCM boot

As soon as the CPU starts running, the CPU starts in EL3 mode. TegraBoot-CPU executes a trampoline code which changes the processor mode from EL3 to EL2. It then initializes USB and starts the 3P protocol to flash the device.

For normal flashing TegraBoot-CPU takes the binaries from the host one by one and flashes them to the device. In case of RCM boot, it takes the binaries from the blob downloaded by MB1. The binaries are not flashed on to the device.

TBoot-CPU can also support cold boot by integrating the relevant libraries (`linuxboot`, etc.). However this version of TBoot-CPU is not tested or maintained.

The following block diagram shows the components of TBoot-CPU:



CBoot

CBoot is the primary CPU boot loader used on mobile platforms in the cold boot path. CBoot is feature-rich boot loader used for cold boot. In addition to booting the kernel, it also supports display, boot logo, and verified boot.

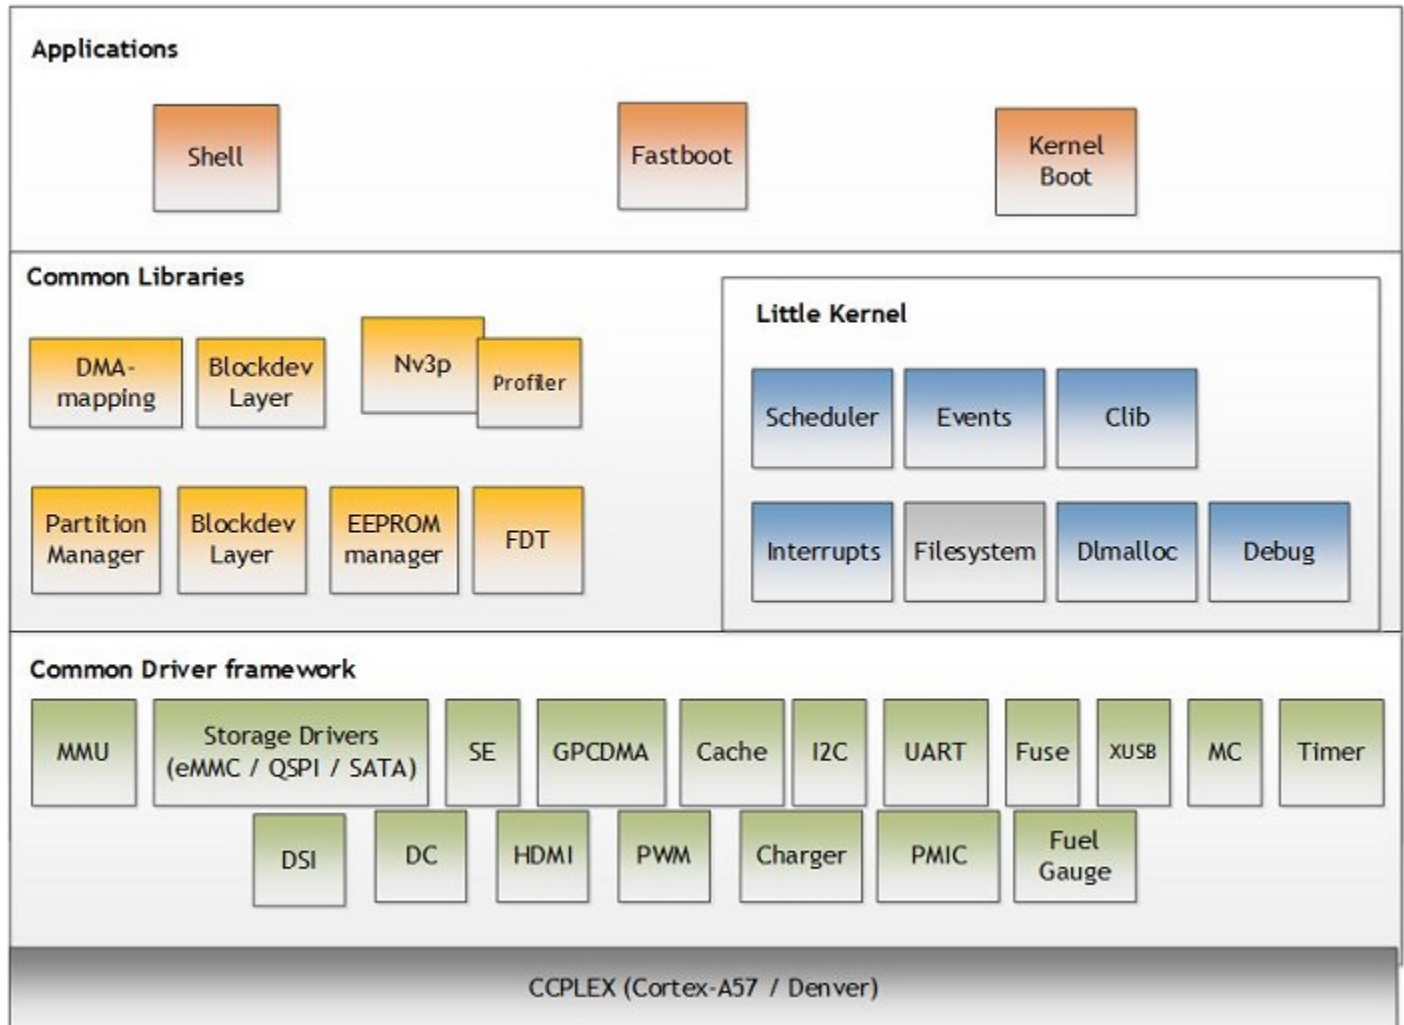
CBoot is based on the Little Kernel (LK) open source boot loader. CBoot primarily uses the interrupt and scheduling frameworks of LK. Cboot uses CDF for frameworks, drivers, and libraries.

As soon as the CPU starts running, CBoot starts in EL3 mode and begins executing TOS. TOS completes its initialization and passes control to CBoot in EL2 mode.

BL and Kernel will use separate device trees stored in separate partitions. CBoot is responsible for:

- Parsing the CPU-BL parameters and initializing the boot loader device tree
- Booting to the kernel or chaining to U-Boot to boot the kernel
- Supporting the update mechanism
 - Loading the recovery kernel
- Verified boot

The following block diagram shows the components of CBoot:



U-Boot

U-Boot is the default boot loader for NVIDIA® Tegra® Linux Driver Package (L4T). For information on U-Boot, refer to [U-Boot Customization](#).

TegraFlash

TegraFlash is the host software component primarily responsible for flashing binaries onto the device.

For more information about TegraFlash, see the following:

- [Boot Configuration Tables](#)

Security

This section explains the security in the BDK.

Validation of Binaries

All of the binaries involved in the boot process have a generic header (except MB1). The binary is appended to the generic header. The header stores the signature of the binary, which is used for validation in secure boot.

Boot Binaries and Associated Signing

The following table summarizes the boot binaries available in Tegra Code-name Parker and their respective signing authorities.

Boot Binary	Signing Authority
•BRBCT	•OEM generated binary, signed by OEM
•MB1BCT	•OEM generated binary, signed by OEM
•MB1	•NVIDIA provided binary, signed/encrypted by NVIDIA, OEM will sign/encrypt again
•MTS-Preboot	•NVIDIA provided binary, signed/encrypted by NVIDIA, OEM will sign/encrypt again
•MTS-Package (MTS-DMCE + MTS-Proper)	•Signed/encrypted separately by NVIDIA, OEM will sign/encrypt again.
•SPE-CAN-FW	•OEM owned binary, signed/encrypted by OEM.
•TBoot-BPMP	•NVIDIA provided reference code, ultimately OEM owned, signed/encrypted by OEM.
•TBoot-CPU	•NVIDIA provided reference code, ultimately OEM owned, signed/encrypted by OEM.
•APE-FW	•NVIDIA offers a choice between an NVIDIA signed/encrypted binary and an unencrypted binary for the OEM to sign/encrypt with their board-specific DTB choice.
•SPE-FW	•NVIDIA offers a choice between an NVIDIA signed/encrypted binary and an unencrypted binary for the OEM to sign/encrypt with their board-specific DTB choice.
•SCE-FW	•NVIDIA offers a choice between an NVIDIA signed/encrypted binary and an unencrypted binary for the OEM to sign/encrypt with their board specific DTB choice.
•CPE-FW	•NVIDIA offers a choice between an NVIDIA signed/encrypted binary and an unencrypted binary for the OEM to sign/encrypt with their board-specific DTB choice.
•BPMP-FW	•NVIDIA offers a choice between an NVIDIA signed/encrypted binary and an unencrypted binary for the OEM to sign/encrypt with their board-specific DTB.

•CBoot	•NVIDIA provided reference code, ultimately OEM owned, signed/encrypted by OEM.
•TOS	•NVIDIA provided reference code, ultimately OEM owned, signed/encrypted by OEM.
•SC7-FW	•NVIDIA provided binary, signed by NVIDIA, OEM will sign/encrypt again.
•BL-DTB	•NVIDIA provided reference code, ultimately OEM owned, signed/encrypted by OEM.
•U-Boot	•NVIDIA-adapted OSS code. Unsigned/Unencrypted?

Secure Boot

Secure boot provides a system of cryptographic checks at each stage of the boot process. These checks ensure the integrity of all of the system software that is executed, preventing any unauthorized or maliciously modified software from being run.

The secure BootROM is responsible for initializing the hardware root of trust (RoT) in the Tegra BSP.

Cryptography is enabled as follows:

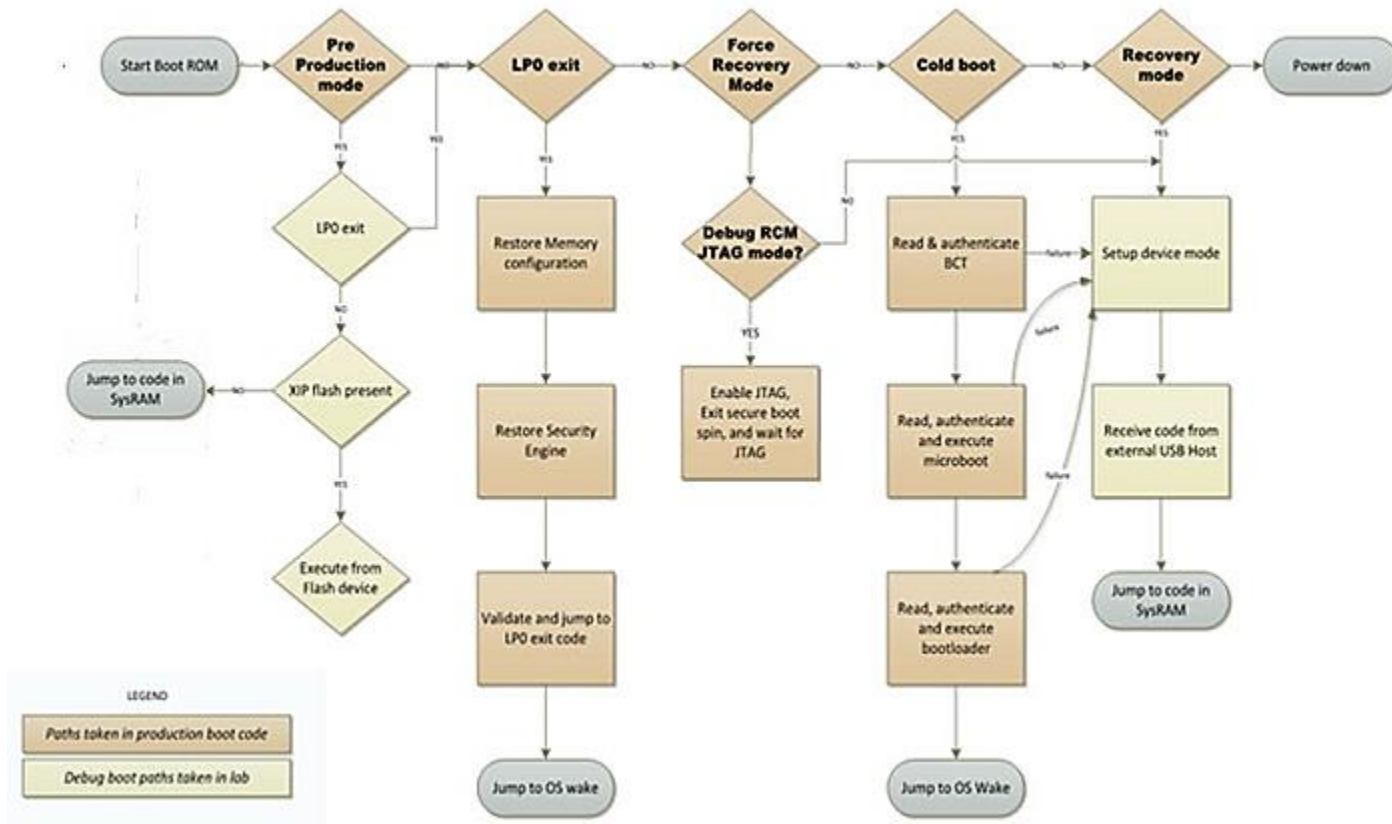
- Authentication: RSA2048
- Encryption: AES-CBC128

Supported boot devices include:

- eMMC
- QSPI

Secure Boot Overview

The following diagram shows the flow of secure boot functionality.

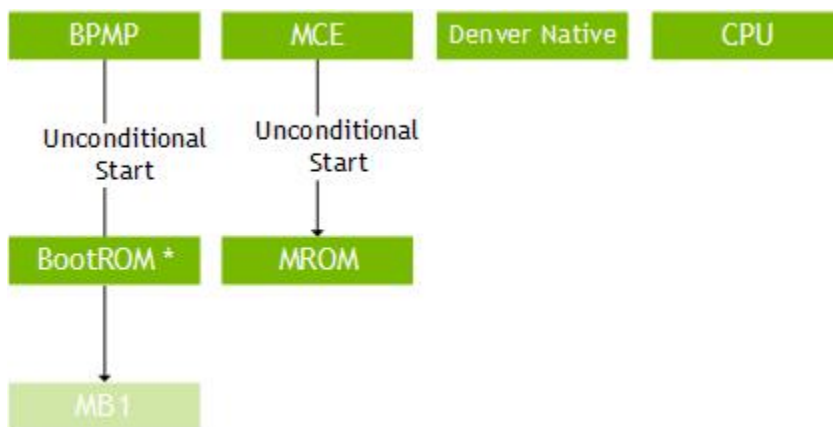


If authentication fails, the Boot ROM automatically goes into USB Recovery mode and waits for input via the USB. The Boot ROM never yields control from Recovery Mode except to authenticated code or error messages.

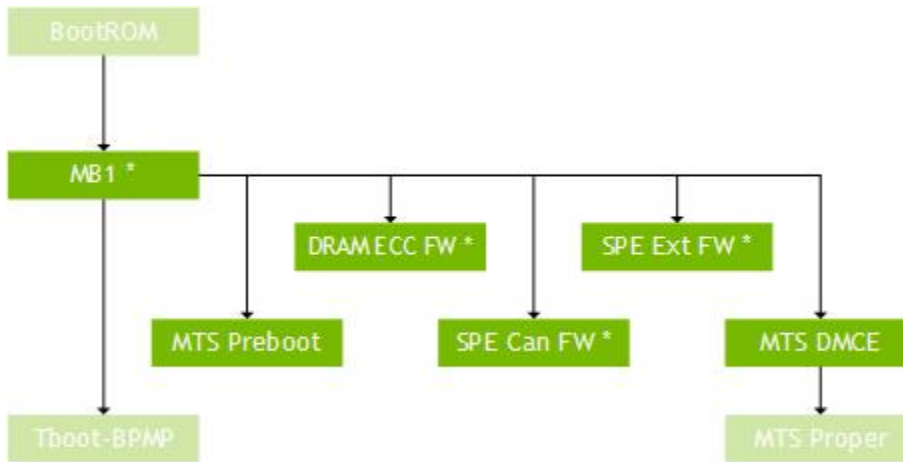
Secure Boot Flow

The following three diagrams show the flow of secure boot stage by stage. Light colored blocks indicate transitions from one diagram to the next.

Each element marked with an asterisk (“*”) has a note below the diagram.



- **BootROM:** Loads BRBCT and establishes a trusted OEM public key by computing the hash and comparing it to the fuse settings. It then verifies the BCT using the OEM authority (i.e. the public key).



- **MB1:** Loads MTS Preboot and verifies it using the double verification mechanism. It first verifies the OEM credential, then the NVIDIA credential (using MAK/MDK).

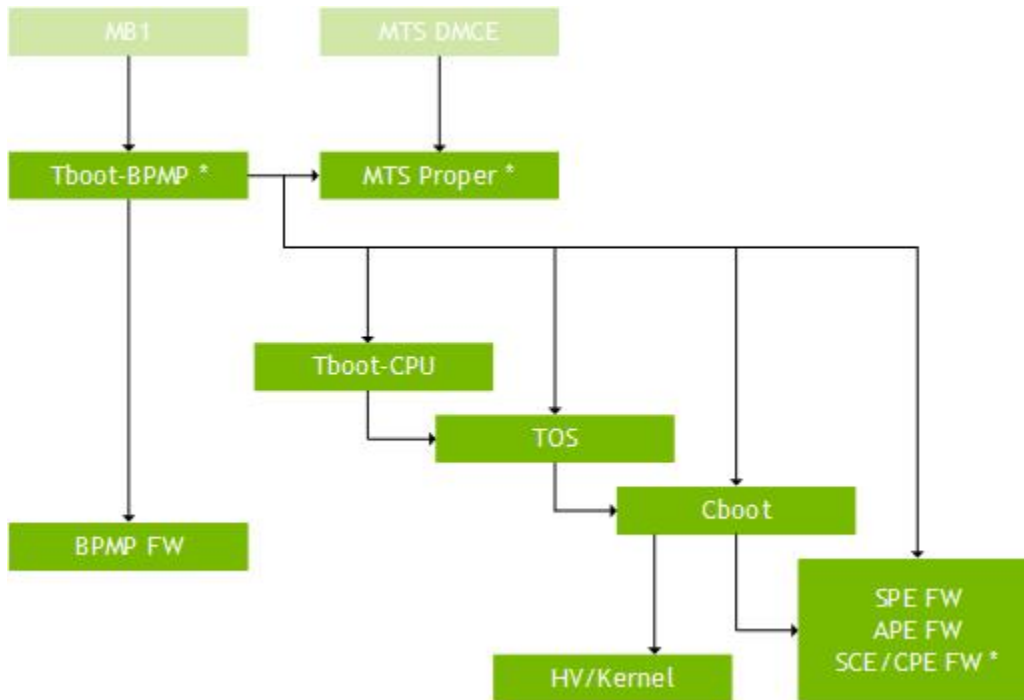
Once the MTS Preboot starts to execute by DB0, MTS initializes SDRAM and then loads MTS Package (MTS DMCE and MTS Proper). It verifies and decrypts the entire package, except that it does NVIDIA verification and decryption only on the MTS DMCE portion of the package.

When MTS-related processing is complete, MB1 loads TBoot-BPMP and verifies it using the OEM authority only, and then passes control to it.

- **DRAM ECC FW:** If required by the platform, MB1 checks whether the boot was started by a reset caused by a two-bit ECC error. If it was, MB1 runs DRAM ECC FW to update the blacklist page and reset the system again. DRAM ECC FW is verified by the OEM authority.
- **SPE FWs (SPE Can and SPE Ext):** MB1 optionally loads these components after MTS Preboot starts, DRAM is operating, and MTS Package has been loaded and authenticated. In most cases only the smaller SPE Can FW is loaded here; SPE Ext FW is loaded at a later stage.

The SPE FW is typically signed by the OEM authority, but may also be signed by the NVIDIA authority. MB1 identifies the signing requirement from the `SignType` field of the generic signature header.

SPE Authentication/Decryption keys are used for the NVIDIA authority.



TBoot-BPMP: This is the main binary loader in the boot flow. It loads the binaries that are required to start CPU execution. These binaries are owned by the OEM, and are verified and decrypted using the OEM authority.

After the CPU has started executing the main boot path, TBoot-BPMP loads the coprocessor FWs. Each of these FWs may be verified and decrypted by the OEM authority or the NVIDIA authority, but not by both.

After all components have been loaded, the boot path BPMP code exits and the CPU software resets it into execution of the BPMP FW.

In the distributed BSP, TBoot-BPMP loads the following binaries:

- TBoot-CPU/CBoot
- TOS
- Coprocessor firmware:
 - SPE FW
 - APE FW
 - SCE/CPE FW
 - BPMP FW

The OEM controls loading of these binaries, and may include or exclude components as appropriate.

SPE, APE, and SCE/CPE FWs: This firmware is loaded by TBoot-BPMP, but is later started by CBoot. Each firmware may be verified and decrypted by the OEM authority or the NVIDIA authority, but not by both.

Build System

TBD - Need build instructions in **tmp** mentioned by Tom Warren for this section.

Display Configuration and Bringup

This topic provides information to help you configure display resolution. Tegra BSP supports a variety of modes on HDMI and DP monitors, including the CEA modes and detailed timing modes from the display EDID.

Setting HDMI or DP Screen Resolution

The screen resolution can be modified using `xrandr` utility or RandR protocol available at:

<https://www.x.org/releases/X11R7.5/doc/man/man1/xrandr.1.html>

<https://cgит.freedesktop.org/xorg/proto/randrproto/tree/randrproto.txt>

To change the default HDMI/DP screen resolution

1. Export the `DISPLAY` variable.

```
export DISPLAY=:0
```

2. Obtain the applicable resolutions list.

```
xrandr
```

The resulting output shows a list of the HDMI, DP, or DSI display settings, if connected.

```
Screen 0: minimum 8 x 8, current 6960 x 2160, maximum 32767 x 32767
DSI-0 connected primary 1200x1920+0+0 (normal left inverted right x axis y axis) 0mm x 0mm
1200x1920 60.00*+
HDMI-0 connected 3840x2160+1200+0 (normal left inverted right x axis y axis) 1600mm x 900mm
3840x2160 60.02*+ 59.98 50.01 30.00 29.97 25.00 24.00 23.98
4096x2160 60.02 59.98 50.01 30.00 29.97 25.00 24.00 23.98
1920x1080 60.00 59.95 50.00 30.00 29.97 25.00 24.00 23.98 61.25 51.04
1440x576 52.08
1360x768 60.02
1280x1024 60.00
1280x720 60.00 59.94 50.00
1152x864 59.97
1024x768 60.01
800x600 60.32
720x576 50.00
720x480 59.94
720x400 70.04
640x480 59.94 59.94
DP-0 connected 1920x1200+5040+0 (normal left inverted right x axis y axis) 520mm x 320mm
1920x1200 59.95*+ 59.89
1920x1080 60.00 59.95 50.00 24.00 23.98 61.25 60.05 51.04
1600x1200 60.01
1440x576 52.08
1440x480 62.69
1280x1024 75.03 60.00
1280x720 60.00 59.94 50.00
1152x864 75.00
1024x768 75.03 60.01
800x600 75.00 60.32
720x576 50.00
720x480 59.94
720x400 70.04
640x480 75.00 59.94 59.94
```

3. Switch the resolution to the desired display resolution.

```
xrandr --output HDMI-0 --mode <res>
```

```
xrandr --output DP-0 --mode <res>
```

Where `<res>` is the desired resolution, for example, `640x480`.

The highest rate supported, for the specified mode, is automatically chosen.

4. Select the desired refresh rate.

```
xrandr --output HDKI-0 --moe <res> --rate <refresh_rate>
```

Where `<refresh_rate>` is the desired refresh rate, for example, 60.

Use `xrandr` to display all the supported refresh rates for a mode.

Mirroring or Extending Displays

When multiple displays are connected, you can choose the relative position of each monitor (HDMI or DP) or panel (DSI).

To choose the relative position of each monitor

1. Execute the following command to mirror the HDMI monitor with the DSI panel.

```
xrandr --output HDMI-0 --same-as DSI-0
```

2. Execute the following command to extend the display environment by placing the HDMI monitor to the left of the DSI panel.

```
xrandr --output HDMI-0 --left-of DSI-0
```

3. Execute the following command to place the DP monitor to the right of the DSI panel.

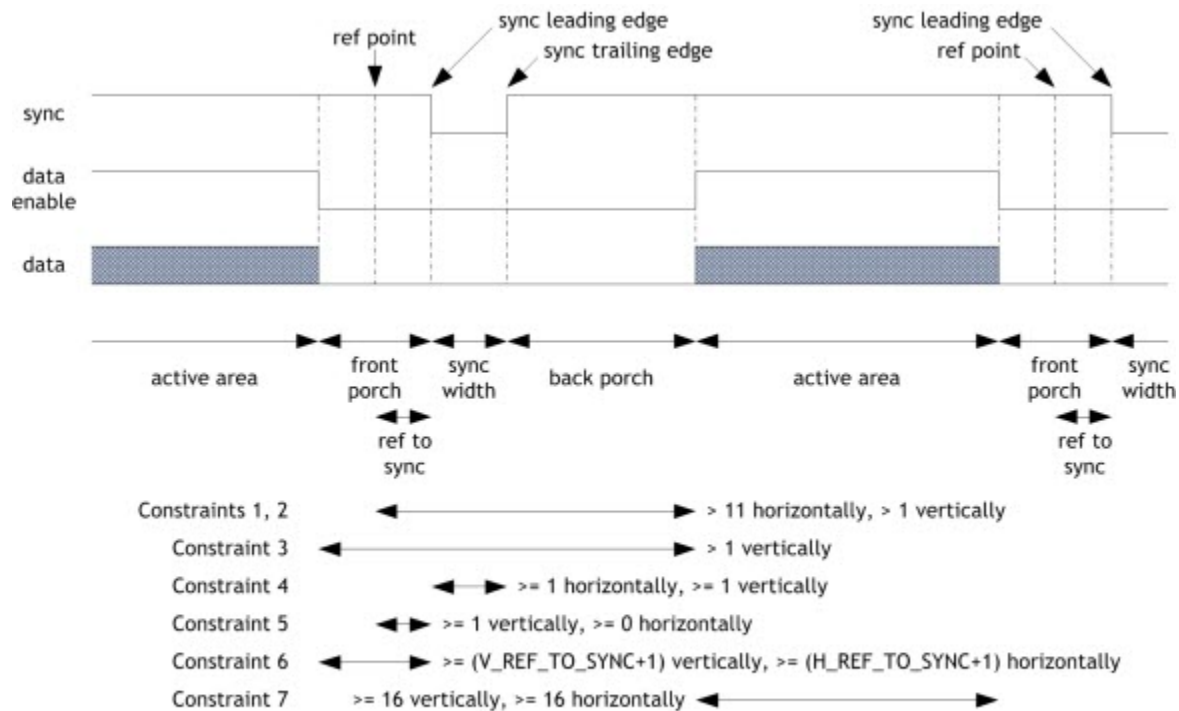
```
xrandr --output DP-0 --right-of DSI-0
```

Determining Display Timing Values

This topic describes how to determine the timing values when referring to the LCD specification and the signal polarities. It also describes the constraints for picking the `ref_to_sync` values.

Display Controller Timing Diagram

The following diagram shows display controller timing. The diagram relates to both vertical and horizontal timing. Only the back porch can be a negative value, though it is typically positive.



Where the constraints are described in the “DC_DISP_DISP_TIMING_OPTIONS_0” section of the *Tegra Technical Reference Manual (TRM)* for your chip.

- **Back porch**—the distance between the trailing edge of the sync and the beginning of the active area. Some LCD specifications define the back porch from the leading edge of the sync rather than the trailing edge.
- **Front porch**—the distance between the end of the active area and the leading edge of the sync.
- **Vertical timing**—is in terms of lines.
- **Horizontal timing**—is in terms of pixels.
- **Sync polarity**—is in terms of the sync width; active low means the sync width pulse will be low, as in the diagram.
- **Pixel clock polarity**—is in terms of when data will transition; an active low pixel clock means that the data changes on the falling edge of the clock, to be latched on the rising edge.
- **Data enable polarity** (also called **display enable** or **DE**)—selects the level of the signal during the active area of display. The diagram has DE active high.

Configuring the Display Serial Interface

This topic describes how to configure the Display Serial Interface (DSI). Both the DSI driver and the DSI panel must be configured. Configuring the DSI driver is done through the device tree. Details about display device tree bindings is available at:

```
<top>/kernel/kernel-4.4/Documentation/devicetree/bindings/video/nvidia,tegra<t-arch>-dc
<top>/kernel/kernel-4.4/Documentation/devicetree/bindings/video/nvidia,tegra<t-arch>-ds
```

Note: Configuring the DSI panel requires the panel specification from the

vendor. You must use the correct display configuration file for your particular platform.

To configure the DSI driver

- Locate and edit the appropriate display configuration file depending on your release at:

```
<top>/hardware/nvidia/platform/tegra/common/kernel-dts/ /panel-<panel_type>.dtsi
```

Where <panel_type> specifies the characteristics encoded for the panel, using the format:

```
brand-scrRes-scrSize
```

For example, the file `panel-p-wuxga-10-1.dtsi` configures the following panel:

```
Filename: panel-p-wuxga-10-1.dtsi
```

```
brand    = p        = Panasonic
scrRes   = wuxga    = WUXGA 1920 x 1200
scrSize  = 10-1     = 10.1 inches
```

The `dsi_bindings` structure encapsulates all the controller configuration options, specified in files at:

```
<top>/kernel/kernel-4.4/Documentation/devicetree/bindings/video
```

The naming format is:

```
nvidia,tegra<t-arch>-<node_type>.txt
```

Where <node_type> is `dc`, `hdmi`, `dp`, or `dsi`.

To configure the DSI panel

1. Configure the regulators.

The details depend on the board schematics.

2. Configure the General-Purpose Input/Outputs (GPIOs) required for resetting and backlighting the panel.

Some panels are more demanding and may require other GPIOs as well.

3. Set up the initialize/suspend sequence.

These are passed as device tree bindings in the power tree for the platform and read in the following file:

```
<top>/kernel/display/drivers/video/tegra/dc/panel/board-panel.c
```

Configuring the Display Port

This topic describes how to enable the Display Port (DP), to map DP to a particular display controller (dc), and assign SOR bindings or enable dpaux appropriately.

The base device tree file is available at:

```
<top>/hardware/nvidia/platfom/t18x/quill/kernel-dts/tegra186-quill-p3310-1000-c03-00-ba
```

Enable the following nodes in this file:

```
host1x {
    nvdisplay@15220000 {
        status = "okay";
        nvidia,dc-or-node = "/host1x/sor";
    };

    sor {
        status = "okay";
        dp-display {
            status = "okay";
        };
    };

    dpaux@155c0000 {
        status = "okay";
    };
};
```

Configuring Pulse Width Modulation (PWM)

Backlight control to the panel is done through Pulse Width Modulation (PWM) or Parametric Four-Wave Mixing (PWFM) or through the backlight device embedded in the panel.

The PFW and PFWM driver located at:

```
<top>/kernel/kernel-4.4/drivers/video/backlight/pwm_bl.c
```

The driver file for the backlight device embedded in the panel is:

```
<top>/kernel/kernel-4.4/drivers/video/backlight/xxx.c
```

For example, the backlight for the LG 5 inch 720p panel on `pluto` is driven by the embedded max8831 integrated circuit (IC). That driver file is:

```
<top>/kernel/kernel-4.4/drivers/video/backlight/max8831_bl.c
```

To make backlight increment in a linear manner to the brightness increment, backlight is calibrated and the corresponding values are stored in the backlight response curve.

Panel Driver and Device Tree

This section points to examples of source files and configuration files for a DSI panel driver.

Internal DSI panel driver and device tree examples

- Driver:

```
<top>kernel/display/drivers/video/tegra/dc/panel/panel-a-1200-1920-8-0.c
```

- Device tree:

Appropriate nodes must be enabled in the device tree. The base device tree file is available at:

```
<top>/hardware/nvidia/platform/t18x/quill/kernel-dts/tegra186-quill-p3310-1000-c03-00-b
```

- Enable the DC, DSI, and appropriate panel node.
- Update the `dc-or-node` property to map DSI to a particular controller as follows:

```
host1x {
    nvdisplay@15200000 {
        status = "okay";
        nvidia,dc-or-node = "/host1x/dsi";
    };
    dsi {
        status = "okay";

        panel-s-wuxga-8-0 {
            status = "okay";
        };
    };
};
```

PDF Documents

The following documentation is included in the NVIDIA[®] Tegra[®] Linux Driver Package (L4T).

- [Tegra Linux Driver Package Development Guide](#) (PDF) — this document
- [Jetson TX2 Accelerated GStreamer User Guide](#) (PDF)
- [Jetson TX2 Developer Kit User Guide](#) (PDF)
- [Runtime Boot Loader Update Process for Jetson TX2](#) (PDF)
- [Platform Adaptation and Bring-Up Guide: Tegra Linux Driver Package for Jetson TX2](#) (PDF)
- [NVIDIA Jetson TX1/TX2 Module EEPROM Layout](#) (PDF)

Appendix

This appendix contains supplementary materials.

Tegra Stats Utility

L4T provides the Tegra Stats utility (tegrastats), which reports memory usage and processor usage for NVIDIA[®] Tegra[®] devices.

JetPack installer or the flashing script places the script in the home directory on the target at:

```
$ (HOME) /tegrastats
```

On the host system, the script is delivered in the TAR file at:

```
Linux_for_Tegra/nv_tegra/nv_tools.tbz2
```

Reported Statistics

The following table shows the statistics the Tegra Stats utility reports.

Statistic	X	Y	Z
<ul style="list-style-type: none"> •RAM X/Y (lfb NxZ) •Largest Free Block (lfb) is a statistic about the memory allocator, and refers to the largest contiguous block of physical memory that can currently be allocated. It is at most 4 MB, and can become smaller with memory fragmentation. • •Note: The word physical in the above; allocations made in virtual memory can be bigger. 	<ul style="list-style-type: none"> •Amount of RAM in use in MB 	<ul style="list-style-type: none"> •Total amount of RAM available for applications 	<ul style="list-style-type: none"> • Z is the size of the largest free block, N the number of free blocks of this size. •
<ul style="list-style-type: none"> •CPU [X%,Y%, ,]@Z • •Note: T30 has up to four CPU cores in the G-Cluster and one CPU core in the LP-Cluster. 	<ul style="list-style-type: none"> •Load statistics for each of the CPU cores relative to the current running frequency Z, or 'off' in case a core is currently powered down. • •This is a rough approximation based on time spent in system idle process as 	<ul style="list-style-type: none"> •Load statistics for each of the CPU cores relative to the current running frequency Z, or 'off' in case a core is currently powered down. • •This is a rough approximation based on time spent in system idle process as reported by the 	<ul style="list-style-type: none"> •CPU frequency in MHz, which dynamically goes up or down depending on the CPU workload

	reported by the Linux kernel in /proc/stat file.	Linux kernel in /proc/stat file.	
<ul style="list-style-type: none"> •avp X%@Y • •AVP is the audio/video processor, which is not visible to the OS or applications. However, some forms of video decode/encode make heavy use of it. •Note: In Tegra X1 and later devices, AVP hardware is replaced by ARM7 	<ul style="list-style-type: none"> •Percent of the AVP or ARM7 that is being used (the AVP load) relative to the current running frequency 	<ul style="list-style-type: none"> •AVP or ARM7 frequency in MHz, which dynamically goes up or down depending on the AVP workload 	<ul style="list-style-type: none"> •N/A
<ul style="list-style-type: none"> •vde Y • •VDE is the video hardware engine. •Note: VDE is deprecated beginning with R24.1. 	<ul style="list-style-type: none"> •N/A 	<ul style="list-style-type: none"> •VDE frequency in MHz 	<ul style="list-style-type: none"> •N/A
<ul style="list-style-type: none"> •edp limit X 	<ul style="list-style-type: none"> •Limit of the CPU maximum frequency 	<ul style="list-style-type: none"> •N/A 	<ul style="list-style-type: none"> •N/A
<ul style="list-style-type: none"> •emc X%@Y • •EMC is the external memory controller, through which all system/carve-out/GART memory accesses go. 	<ul style="list-style-type: none"> •Percent of the EMC memory bandwidth that is being used, relative to the current running frequency 	<ul style="list-style-type: none"> •EMC frequency in MHz 	<ul style="list-style-type: none"> •N/A
<ul style="list-style-type: none"> •GR3D X%@Y • •GR3D is the GPU engine. 	<ul style="list-style-type: none"> •Percent of the GR3D that is being used, relative to the current running frequency 	<ul style="list-style-type: none"> •GR3D frequency in MHz 	<ul style="list-style-type: none"> •N/A
<ul style="list-style-type: none"> •MSENC Y • •MSENC is the video hardware encoding engine. 	<ul style="list-style-type: none"> •N/A 	<ul style="list-style-type: none"> •MSENC frequency in MHz 	<ul style="list-style-type: none"> •N/A
<ul style="list-style-type: none"> •NVDEC Y • •NVDEC is the video hardware decoding engine. 	<ul style="list-style-type: none"> •N/A 	<ul style="list-style-type: none"> •NVDEC frequency in MHz 	<ul style="list-style-type: none"> •N/A

Example Output

The following example output is extracted from running a 1080p video playback use case. CPU frequency is running between 300 and 750 MHz and only one CPU is active.

Notes: The examples provided are for previous Tegra devices.

```
E/TegraStats( 7484): RAM 92/184MB (1fb 1x4MB) | DFS(%@MHz): cpu [52%,35%]@527 avp 18%@1
E/TegraStats( 7484): RAM 93/184MB (1fb 1x4MB) | DFS(%@MHz): cpu [51%,51%]@750 avp 8%@10
E/TegraStats( 7484): RAM 93/184MB (1fb 1x4MB) | DFS(%@MHz): cpu [58%,45%]@600 avp 8%@10
E/TegraStats( 7484): RAM 93/184MB (1fb 1x4MB) | DFS(%@MHz): cpu [48%,49%]@463 avp 8%@10
E/TegraStats( 7484): RAM 93/184MB (1fb 1x4MB) | DFS(%@MHz): cpu [60%,41%]@411 avp 8%@10
E/TegraStats( 7484): RAM 93/184MB (1fb 1x4MB) | DFS(%@MHz): cpu [60%,48%]@384 avp 9%@10
E/TegraStats( 7484): RAM 93/184MB (1fb 1x4MB) | DFS(%@MHz): cpu [60%,44%]@438 avp 10%@1
E/TegraStats( 7484): RAM 93/184MB (1fb 1x4MB) | DFS(%@MHz): cpu [44%,53%]@411 avp 8%@10
E/TegraStats( 7484): RAM 93/184MB (1fb 1x4MB) | DFS(%@MHz): cpu [57%,47%]@359 avp 8%@10
E/TegraStats( 7484): RAM 93/184MB (1fb 1x4MB) | DFS(%@MHz): cpu [55%,44%]@750 avp 8%@10
E/TegraStats(7484): RAM 98/184MB (1fb 1x1MB) | DFS(%@MHz): cpu [77%,68%]@300 avp 2%@100
E/TegraStats(7484): RAM 94/184MB (1fb 1x2MB) | DFS(%@MHz): cpu [35%,29%]@300 avp 2%@100
E/TegraStats(7484): RAM 95/184MB (1fb 1x2MB) | DFS(%@MHz): cpu [35%,29%]@300 avp 2%@100
...
```

Options

TegraStats options and usage are described in the following table.

Option	Description
•<delay>	•The delay between outputs in milliseconds (ms). For example, a value of 500 causes output to display twice per second. If no delay value is specified, the default delay is

	1000 ms. A <delay> value, if specified, must be the first value after the command, before any other options.
•--start	•Run TegraStats as a daemon, in the background.
•--stop	•Stop running any running instances of tegrastats.
•--log-file <out_file>	•Write the output of TegraStats to <out_file>.
•-once	•Print data only once and then exit.
•-throttle	•Track the throttle count.
•<thermal_zone>	•Print temperature in specified <thermal_zone>. The value for <thermal_zone> must match one of the strings in "/sys/devices/virtual/thermal/thermal_zone*/type.

Running Tegra Stats

When you run `tegrastats` on Linux devices, it prints statistics to stdout.

Example Log Print

```
E/TegraStats( 3627): RAM 250/490MB | DFS(%@MHz): cpu [61%,45%]@835 avp 17%@100 vde 34%
```

To run tegrastats

- On the target, execute the following commands:

```
sudo $(HOME)/tegrastats <delay> &
```

Where <delay> is the frequency of log prints, expressed in milliseconds (1000 means print every second).

To stop tegrastats

- On the target, execute the following commands:

```
sudo $(HOME)/tegrastats --stop
```

Re-Deploying Tegra Stats

The Tegra Stats Utility is preinstalled in the default root file system in the home directory. If you have modified the root file system configuration and removed `tegrastats` from the build, you can re-deploy it on the target at runtime.

To re-deploy tegrastats

- Execute the following commands from the debug host PC:

```
scp tegrastats nvidia@<TARGET_DEVICE_IP>:/home/nvidia/
```

FAQ

This section provides answers to frequently asked questions about your release. Use it as the first step in troubleshooting problems. You can also try searching the Index in this document, contacting your support engineer, or filing a bug.

Linux FAQs

Are ARMv7 binaries compatible with aarch64 binaries?

No, while the kernel supports both ARMv7 (32bi) and aarch64 binaries, distros currently are exclusively aarch64 or ARMv7. ARMv7 binaries are not compatible in an aarch64 distro. The NVIDIA PDK supports aarch64 distro where ARMv7 binaries are not compatible.

How do I use display mode and resolution configuration with the X RandR application?

You can use the X Resize, Rotate and Reflect Extension (RandR) extension to manipulate and configure the attached displays (both the internal panel and any externally connected HDMI panel). The `xrandr(1)` utility is the most common way to do this.

You can find a tutorial on `xrandr` on the following website:

```
http://www.thinkwiki.org/wiki/Xorg\_RandR\_1.2
```

Are there generated ssh host keys for the sample file system?

There are no keys in the `/etc/ssh` directory of the provided sample file system. For information about creating the ssh host keys, see the `ssh-keygen` man page.

How do I determine the X driver ABI of the X server used in the root file system?

All `tegra_drv.abi*.so` files are in the driver package. By default the `apply_binaries.sh` script creates a sym-link from `tegra_drv.so` to the X ABI driver compatible with the provided sample file system.

How do I prevent the system display from blanking out?

Linux kernel 3.1 added a power saving feature that may blank the display of an idle system even when applications are running. The feature is called console blank (screen saver). It is defined as:

```
consoleblank= [KNL]
```

Where `[KNL]` is the console blank (screen saver) timeout in seconds. This defaults to $10 * 60 = 10$ mins. A value of 0 disables the blank timer.

By passing arguments to the kernel command line, you can:

- Disable this feature, or
- Set the timeout to a longer interval.

With the `flash.sh` script, you can override the kernel command line options passed from fastboot to the kernel.

To disable the console blank (screen saver) from the kernel command line

1. In the grub configuration add the following line to the kernel parameters:

```
consoleblank=0
```

2. View the current `consoleblank` value with the following command:

```
$ cat /sys/module/kernel/parameters/consoleblank
```

To disable the console blank feature with an escape sequence

- Enter the following escape sequence:

```
$ echo -ne "\033[9;0]"
```

To change the console blank timeout value with an escape sequence

- Enter the following escape sequence:

```
$ echo -ne "\033[9;<timeout>]"
```

Where `<timeout>` is the timeout in seconds.


For more information on this escape sequence, see the `console_codes(4)` man page documents. For information on the input/output controls that provide some of the same functionality, see the `console_ioctl(4)` man page.

Glossary

This section provides definitions to frequently used terms in this release.


To browse the glossary by keyword

- Click the Glossary button on the tool bar (shown below).



To browse the glossary page-by-page

- Click the NEXT button (right arrow) on the toolbar.



<ABI_directory>: aarch64-linux-gnu [More...]

See Section 1.0 of the *Release Notes* to confirm your ABI directory info.

BCT: NVIDIA boot configuration table, a binary file used for flashing and updating.
[More...]

Stored on a secondary boot device, such as eMMC, which primarily stores information required to:

- initialize and configure the secondary boot device
- initialize and configure SDRAM
- locate and verify the boot loader

<board>: t186ref [More...]

Valid values depend on the supported board.

- Jetson TX1: p2371-2180
- Jetson TX2: p2771-000-500

See Section 1.0 of the *Release Notes* for your release to confirm your board information.

<board_and_rev>: p2371-2180 or p2771-0000 [More...]

Valid values depend on the supported board.

- Jetson TX1: p2371-2180
- Jetson TX2: p2771-000-500

See Section 1.0 of the *Release Notes* for your release to confirm your board and revision info.

BPMP: Boot and power management processor. [\[More...\]](#)

Boot and power management processor (BPMP) is firmware that runs on Cortex R5. During boot, BPMP executes the boot ROM code and controls the SOC boot sequence. After boot, BPMP runs power management functions.

CUDA: An NVIDIA parallel computing platform and programming model. [\[More...\]](#)

NVIDIA[®] CUDA[®] technology enables dramatic increases in computing performance by harnessing the power of the graphics processing unit (GPU).

<EGL|GL_ver>: 4.1 [More...]

See Section 1.0 of the *Release Notes* to confirm your NVIDIA EGL/GL version info.

`<lnx_ver>`: **Kernel 4.4.15.** [\[More...\]](#)

See Section 1.0 of the *Release Notes* for your release to confirm your Linux version info.

MB1: Microboot 1.[\[More...\]](#)

MB1 is the first boot loader to run after the boot ROM. Its major tasks are to:

- Initialize the SDRAM based on MB1 boot configuration table (MB1-BCT)
- Load the firmware that initializes the CPU complex
- Set the security settings
- Configure pinmux/GPIO for the SoC
- Program the PMIC for the VDD_CPU and VDD_DDR rails
- Load the next stage boot loader ([MB2](#))

MB1 runs on the [BPMP](#). It is considered an extended boot ROM because the security level of MB1 is the same as boot ROM.

MB2: Microboot 2 also called nvtboot-bpmp.[\[More...\]](#)

MB2 runs after [MB1](#). Its major tasks are to:

- Load the CPU_BL
- Load the BPMP-FW and its device tree binary (DTB)
- Load the Trusted Operating System (TOS), which is also called SecureOS

MB2 runs on the [BPMP](#).

`<os_ver_host>`: **14.04 (amd64 distribution)** [\[More...\]](#)

See Section 1.0 of the *Release Notes* for your release to confirm your host Ubuntu operating system version info.

`<os_ver_target>`: **14.04 (amd64 distribution)** . [\[More...\]](#)

See Section 1.0 of the *Release Notes* for your release to confirm your target Ubuntu operating system version information.

<platform>: jetson-tx1, jetson-tx2 [More...]

The target board supported by the release. See Section 1.0 of the *Release Notes* to confirm your platform.

<platform | ver>: t186ref [More...]

Valid values depend on the supported [platform](#) for your product.

See Section 1.0 of the *Release Notes* to determine your platform info.

<platform | ver>: t210ref [More...]

Valid values depend on the supported [platform](#) for your product.

See Section 1.0 of the *Release Notes* to determine your platform info.

<platform | ver>: t186ref [More...]

Valid values depend on the supported [platform](#) for your product.

See Section 1.0 of the *Release Notes* to confirm your platform info.

<release>: R27.1[More...]

See Section 1.0 of the *Release Notes* to confirm your release info.

`<release_num>`: **27.1.0** [\[More...\]](#)

See Section 1.0 of the *Release Notes* to confirm your release number.

release_tag: **tegra-l4t-r27.1** [\[More...\]](#)

See Section 1.0 of the *Release Notes* for your release to confirm your release tag information.

`<release_type>`: **aarch64** [\[More...\]](#)

See Section 1.0 of the *Release Notes* to confirm your release type.

<t-arch|ver>: 186 [More...]

Valid values for the Tegra architecture version depend on the supported product. Use:

- 186 for Tegra X1 devices.

See Section 1.0 of the *Release Notes* for your release to confirm your Tegra architecture version info.

<t-arch|ver>: 186 [More...]

Valid values for the Tegra architecture version depend on the supported product. Use:

- 186 for Tegra Code-name Parker devices.

See Section 1.0 of the *Release Notes* for your release to confirm your Tegra architecture version info.

<top>: Host directory where you installed the PDK/SDK and Foundation release (or other software). [More...]

Depending on the context, <top> may refer to the Host directory where you installed the SDK/PDK release, or it may refer to where you installed other software, such as a utility or add-on package.

<top>: Host directory where you installed the Linux Driver Package release (or other software). [More...]

Depending on the context, <top> may refer to the Host directory where you installed the L4T release, or it may refer to where you installed other software, such as a utility or add-on package.

Licenses

This section provides license information for the NVIDIA[®] Tegra[®] Linux Driver Package. Different components of NVIDIA[®] Tegra[®] Linux Driver Package involve different licenses, granted variously by NVIDIA, other vendors, and GPL copyright holders. Please read the following carefully so you understand your rights and obligations under these licenses.

NVIDIA Software

Note: This software license applies to software in the "nv_tegra" directory (NVIDIA binary drivers and supporting software), the files in the "nv_tegra/nv_sample_apps" directory: "nvgstcapture", "nvgstplayer", and "libgstnvidconv.so" (those files included in the nvgstapps.tbz2 file), and the files "bootloader/mkgpt", "bootloader/mkbootimg", and "bootloader/mkubootscript".

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, 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.

Sample File System

The sample root file system is derived from Ubuntu Linux, version 14.04 for the hardware floating point (hardfp) release. Information on re-creating the root file system is provided in the *Tegra Linux Driver Package Developers' Guide*. The license agreement for each software component is located in the software component's source code, made available from the same location from which this software was downloaded, or by request to oss-requests@nvidia.com.

GST OpenMAX

The software listed below is licensed under the terms of the LGPLv2.1 (see below). To obtain source code, contact oss-requests@nvidia.com.

gst-openmax (libgstomx.so, libgstegl-1.0.so.0, and libnvgstjpeg.so)

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc. 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights. We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library. To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or

table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License.

Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will

still fall under Section 6.) Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6.

Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications. You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)
- b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

- a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License¹¹. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances. It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our

decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

<one line to give the library's name and a brief idea of what it does.>

Copyright (C) <year> <name of author>

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the library `Frob' (a library for tweaking knobs) written by James Random Hacker.

<signature of Ty Coon>, 1 April 1990

Ty Coon, President of Vice

That's all there is to it!

GST EGL

GStreamer EGL/GLES Sink

Copyright (C) 2012 Collabora Ltd.

@author: Reynaldo H. Verdejo Pinochet <reynaldo@collabora.com>

@author: Sebastian Dröge <sebastian.droege@collabora.co.uk>

Copyright (c) 2014, NVIDIA CORPORATION. All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Linux Kernel

The Linux kernel in this release is licensed under the terms of the GPLv2 (see below). The revision of Linux kernel source code used to build this binary can be retrieved by running the 'source_sync.sh' script or by request to oss-requests@nvidia.com.

The device-tree-compiler (dtc) binary located in the "kernel" directory was built from the Linux kernel source code provided by this release. It also is licensed under the terms of the GPLv2 (see below). The revision of the Linux kernel source code which was used to build this dtc binary can be retrieved by request to oss-requests@nvidia.com.

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with

modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
- b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
- c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no

more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

<one line to give the program's name and a brief idea of what it does.>

Copyright (C) <year> <name of author>

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) year name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'. This is free software,
and you are welcome to redistribute it under certain conditions; type `show c' for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the program `Gnomovision' (which makes passes at
compilers) written by James Hacker.
```

```
<signature of Ty Coon>, 1 April 1989
Ty Coon, President of Vice
```

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License.

mkbooting and mkubootscript

mkbooting and mkubootscript are provided under the following terms:

Copyright 2007, The Android Open Source Project

Apache License

Apache License
Version 2.0, January 2004
<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a

lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

(a) You must give any other recipients of the Work or Derivative Works a copy of this License; and

(b) You must cause any modified files to carry prominent notices stating that You changed the files; and

(c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

(d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof,

You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

```
Copyright [yyyy] [name of copyright owner]
Licensed under the Apache License, Version 2.0 (the "License"); you may not use this fi
http://www.apache.org/licenses/LICENSE-2.0
Unless required by applicable law or agreed to in writing, software distributed under t
```

Copyright: WIDE Project

Copyright (C) 1995, 1996, 1997, and 1998 WIDE Project.
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, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the project nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE PROJECT 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 PROJECT 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.

FIPS pub 180-1: Secure Hash Algorithm (SHA-1)
based on: <http://csrc.nist.gov/fips/fip180-1.txt>
implemented by Jun-ichiro itojun Itoh <itojun@itojun.org>

GNU General Public License

For the complete text of this license, see [GNU GENERAL PUBLIC LICENSE](#) in this section.

Copyright: Regents of the University of California

Copyright (c) 1992, 1993

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, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

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.

Copyright 2007, The Android Open Source Project

U-Boot and mkimage

U-Boot is Free Software. It is copyrighted by Wolfgang Denk and many others who contributed code (see the actual source code for details). You can redistribute U-Boot and/or modify it under the terms of version 2 of the GNU General Public License as published by the Free Software Foundation. Most of it can also be distributed, at your option, under any later version of the GNU General Public License -- see individual files for exceptions.

NOTE! This license does *not* cover the so-called "standalone" applications that use U-Boot services by means of the jump table provided by U-Boot exactly for this purpose - this is merely considered normal use of U-Boot, and does *not* fall under the heading of "derived work" -- see file Licenses/Exceptions for details.

Also note that the GPL and the other licenses are copyrighted by the Free Software Foundation and other organizations, but the instance of code that they refer to (the U-Boot source code) is copyrighted by me and others who actually wrote it.

-- Wolfgang Denk

Like many other projects, U-Boot has a tradition of including big blocks of License headers in all files. This not only blows up the source code with mostly redundant information, but also makes it very difficult to generate License Clearing Reports. An additional problem is that even the same licenses are referred to by a number of

slightly varying text blocks (full, abbreviated, different indentation, line wrapping and/or white space, with obsolete address information, ...) which makes automatic processing a nightmare.

To make this easier, such license headers in the source files will be replaced with a single line reference to Unique License Identifiers as defined by the Linux Foundation's SPDX project [1]. For example, in a source file the full "GPL v2.0 or later" header text will be replaced by a single line:

```
SPDX-License-Identifier: GPL-2.0+
```

Ideally, the license terms of all files in the source tree should be defined by such License Identifiers; in no case a file can contain more than one such License Identifier list.

If a "SPDX-License-Identifier:" line references more than one Unique License Identifier, then this means that the respective file can be used under the terms of either of these licenses, i. e. with

```
SPDX-License-Identifier: GPL-2.0+ BSD-3-Clause
```

you can chose between GPL-2.0+ and BSD-3-Clause licensing.

We use the SPDX Unique License Identifiers here; these are available at [2].

[1] <http://spdx.org/>

[2] <http://spdx.org/licenses/>

Full Name	SPDX ID	OSI approved	File Name	URI
GNU General Public License v2.0 only	GPL-2.0	Y	gpl-2.0.txt	http://www.gnu.org/licenses/gpl-2.0.txt
GNU General Public License v2.0 or later	GPL-2.0+	Y	gpl-2.0.txt	http://www.gnu.org/licenses/gpl-2.0.txt
GNU Library General Public License v2 or later	LGPL-2.0+	Y	lgpl-2.0.txt	http://www.gnu.org/licenses/old-licenses/lgpl-2.0.txt
GNU Lesser General Public License v2.1 or later	LGPL-2.1+	Y	lgpl-2.1.txt	http://www.gnu.org/licenses/old-licenses/lgpl-2.1.txt
eCos license version 2.0	eCos-2.0		eCos-2.0.txt	http://www.gnu.org/licenses/ecos-license.html

BSD 2-Clause License	BSD-2-Clause	Y	bsd-2-clause.txt	http://spdx.org/licenses/BSD-2-Clause
BSD 3-clause "New" or "Revised" License	BSD-3-Clause	Y	bsd-3-clause.txt	http://spdx.org/licenses/BSD-3-Clause#licenseText
IBM PIBS (PowerPC Initialization and Boot Software) license	IBM-pibs	-	ibm-pibs.txt	
ISC License	ISC	Y	isc.txt	https://spdx.org/licenses/ISC
X11 License	X11	-	x11.txt	https://spdx.org/licenses/X11.html

GNU GENERAL PUBLIC LICENSE, Version 2, June 1991

For the complete text of this license, see [GNU GENERAL PUBLIC LICENSE](#) in this section.

mkbctpart

The mkbctpart library is covered by the License For Customer Use of NVIDIA Software. For the complete text of that license, see [License For Customer Use of NVIDIA Software](#) in this section.

brcm_patchram_plus

The brcm_patchram_plus is installed at the following location:

```
/usr/sbin/brcm_patchram_plus
```

Copyright (C) 2009-2011 Broadcom Corporation

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed

on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

For the complete text of the Apache license, see [Apache License](#) in this section.

libnvcam_imageencoder.so

Copyright (c) 2015, NVIDIA CORPORATION. All rights reserved.

NVIDIA CORPORATION and its licensors retain all intellectual property and proprietary rights in and to this software, related documentation and any modifications thereto. Any use, reproduction, disclosure or distribution of this software and related documentation without an express license agreement from NVIDIA CORPORATION is strictly prohibited.

This product incorporates software provided under the following terms:

NOTICE file in this case for the JPEGcode.

The authors make NO WARRANTY or representation, either express or implied, with respect to this software, its quality, accuracy, merchantability, or fitness for a particular purpose. This software is provided "AS IS", and you, its user, assume the entire risk as to its quality and accuracy.

This software is copyright (C) 1991-1998, Thomas G. Lane.

All Rights Reserved except as specified below.

Permission is hereby granted to use, copy, modify, and distribute this software (or portions thereof) for any purpose, without fee, subject to these conditions:

- (1) If any part of the source code for this software is distributed, then this README file must be included, with this copyright and no-warranty notice unaltered; and any additions, deletions, or changes to the original files must be clearly indicated in accompanying documentation.
- (2) If only executable code is distributed, then the accompanying documentation must state that "this software is based in part on the work of the Independent JPEG Group".
- (3) Permission for use of this software is granted only if the user accepts full responsibility for any undesirable consequences; the authors accept NO LIABILITY for damages of any kind.

These conditions apply to any software derived from or based on the IJG code, not just to the unmodified library. If you use our work, you ought to acknowledge us.

Permission is NOT granted for the use of any IJG author's name or company name in advertising or publicity relating to this software or products derived from it. This software may be referred to only as "the Independent JPEG Group's software".

We specifically permit and encourage the use of this software as the basis of commercial products, provided that all warranty or liability claims are assumed by the product vendor.

See also [RECITALS](#) in [NVIDIA Software](#).

libscf

libscf is covered by the NVIDIA license. For the complete text of this license, see [License For Customer Use of NVIDIA Software](#) in this section.

Portions of this software are provided under the following terms:

This folder contains libraries and headers of a few very popular still image codecs used by highgui module. The libraries and headers are preferably to build Win32 and Win64 versions of OpenCV. On UNIX systems all the libraries are automatically detected by configure script. In order to use these versions of libraries instead of system ones on UNIX systems you should use BUILD_<library_name> CMake flags (for example, BUILD_PNG for the libpng library).

libjpeg 8d (8.4) - The Independent JPEG Group's JPEG software.

Copyright (C) 1991-2012, Thomas G. Lane, Guido Vollbeding.

See IGJ home page <http://www.ijg.org> for details and links to the source code

HAVE_JPEG preprocessor flag must be set to make highgui use libjpeg. On UNIX systems configure script takes care of it.

libpng 1.5.12 - Portable Network Graphics library.

Copyright (c) 2004, 2006-2012 Glenn Randers-Pehrson. See libpng home page <http://www.libpng.org> for details and links to the source code

HAVE_PNG preprocessor flag must be set to make highgui use libpng. On UNIX systems configure script takes care of it.

libtiff 4.0.2 - Tag Image File Format (TIFF) Software

Copyright (c) 1988-1997 Sam Leffler

Copyright (c) 1991-1997 Silicon Graphics, Inc.

See libtiff home page <http://www.remotesensing.org/libtiff/> for details and links to the source code

HAVE_TIFF preprocessor flag must be set to make highgui use libtiff. On UNIX systems configure script takes care of it. In this build support for ZIP (LZ77 compression) is turned on.

zlib 1.2.7 - General purpose LZ77 compression library

Copyright (C) 1995-2012 Jean-loup Gailly and Mark Adler.

See zlib home page <http://www.zlib.net> for details and links to the source code

No preprocessor definition is needed to make highgui use this library - it is included automatically if either libpng or libtiff are used.

jasper-1.900.1 - JasPer is a collection of software (i.e., a library and application programs) for the coding and manipulation of images. This software can handle image data in a variety of formats. One such format supported by JasPer is the JPEG-2000 format defined in ISO/IEC 15444-1.

Copyright (c) 1999-2000 Image Power, Inc.

Copyright (c) 1999-2000 The University of British Columbia

Copyright (c) 2001-2003 Michael David Adams

The JasPer license can be found in src/libjasper.

OpenCV on Windows uses pre-built libjasper library (lib/libjasper*). To get the latest source code, please, visit the project homepage: <http://www.ece.uvic.ca/~mdadams/jasper/>

openexr-1.7.1 - OpenEXR is a high dynamic-range (HDR) image file format developed by Industrial Light & Magic for use in computer imaging applications.

Copyright (c) 2006, Industrial Light & Magic, a division of Lucasfilm Entertainment Company Ltd. Portions contributed and copyright held by others as indicated. All rights reserved.

The project homepage: <http://www.openexr.com>

ffmpeg-0.8.0 - FFmpeg is a complete, cross-platform solution to record, convert and stream audio and video. It includes libavcodec - the leading audio/video codec library, and also libavformat, libavutils and other helper libraries that are used by OpenCV (in highgui module) to read and write video files.

The project homepage: <http://ffmpeg.org/>

=====

== LICENSE file ==

== ==

== in this case for the Open Computer Vision code. ==

=====

IMPORTANT: READ BEFORE DOWNLOADING, COPYING, INSTALLING OR USING.

By downloading, copying, installing or using the software you agree to this license.

If you do not agree to this license, do not download, install, copy or use the software.

License Agreement for Protocol Buffers

This license applies to all parts of Protocol Buffers except the following:

- Atomicops support for generic gcc, located in:

```
src/google/protobuf/stubs/atomicops_internals_generic_gcc.h.
```

This file is copyrighted by Red Hat Inc.

- Atomicops support for AIX/POWER, located in

```
src/google/protobuf/stubs/atomicops_internals_power.h.
```

This file is copyrighted by Bloomberg Finance LP.

Copyright 2014, Google Inc. All rights reserved.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS 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 COPYRIGHT OWNER 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.

Code generated by the Protocol Buffer compiler is owned by the owner of the input file used when generating it. This code is not standalone and requires a support library to be linked with it. This support library is itself covered by the above license.

License Agreement for Open Source Computer Vision Library

Copyright (C) 2000-2008, Intel Corporation, all rights reserved.

Copyright (C) 2008-2011, Willow Garage Inc., all rights reserved.

Third party copyrights are property of their respective owners.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The name of the copyright holders may not be used to endorse or promote products derived from this software without specific prior written permission.

This software is provided by the copyright holders 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 Intel Corporation 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.

License Agreement for OpenCV Tutorial Library

```

/*          License Agreement
 *
 *          For OpenCV Tutorial Library
 *          http://computer-vision-talks.com
 *
 * Third party copyrights are property of their respective owners.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 *
 * Redistribution's of source code must retain the above copyright
 * notice, this list of conditions and the following disclaimer.
 *
 * Redistribution's in binary form must reproduce the above copyright
 * notice, this list of conditions and the following disclaimer in the
 * documentation and/or other materials provided with the distribution.
 *
 * The name of the copyright holders may not be used to endorse or
 * promote products derived from this software without specific prior
 * written permission.
 *

```

```
* This software is provided by the copyright holders 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 Intel Corporation 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
*/
```

NOTICE file in this case for the Open Computer Vision code.

IMPORTANT: READ BEFORE DOWNLOADING, COPYING, INSTALLING OR USING.

By downloading, copying, installing or using the software you agree to this license. If you do not agree to this license, do not download, install, copy or use the software.

License Agreement for Open Source Computer Vision Library

Copyright (C) 2000-2008, Intel Corporation, all rights reserved.
Copyright (C) 2008-2011, Willow Garage Inc., all rights reserved.
Third party copyrights are property of their respective owners.

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

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The name of the copyright holders may not be used to endorse or promote products derived from this software without specific prior written permission.

This software is provided by the copyright holders 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 Intel Corporation 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.

NOTICE file in this case for the OpenCV Tutorial code.

License Agreement for OpenCV Tutorial Library

```

/*          License Agreement
*          For OpenCV Tutorial Library
*          http://computer-vision-talks.com
*
* Third party copyrights are property of their respective owners.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
*
* Redistribution's of source code must retain the above copyright
* notice, this list of conditions and the following disclaimer.
*
* Redistribution's in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in the
* documentation and/or other materials provided with the
* distribution.
*
* The name of the copyright holders may not be used to endorse or
* promote products derived from this software without specific prior
* written permission.
*
* This software is provided by the copyright holders 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 Intel Corporation or contributors be liable X
* limited to, the implied 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.
*/
```

This package was debianized by Christopher L Cheney <ccheney@debian.org> on Fri, 22 Aug 2003 01:33:34 -0500.

The current maintainer is Roland Stigge <stigge@antcom.de>

It was downloaded from <http://www.ece.uvic.ca/~mdadams/jasper/>

Upstream Author: Michael Adams <mdadams@ece.uvic.ca>

License:

JasPer License Version 2.0

Copyright (c) 1999-2000 Image Power, Inc.

Copyright (c) 1999-2000 The University of British Columbia

Copyright (c) 2001-2003 Michael David Adams

All rights reserved.

Permission is hereby granted, free of charge, to any person (the "User") obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

1. The above copyright notices and this permission notice (which includes the disclaimer below) shall be included in all copies or substantial portions of the Software.
2. The name of a copyright holder shall not be used to endorse or promote products derived from the Software without specific prior written permission.

THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF THE SOFTWARE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER. THE SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE. NO ASSURANCES ARE PROVIDED BY THE COPYRIGHT HOLDERS THAT THE SOFTWARE DOES NOT INFRINGE THE PATENT OR OTHER INTELLECTUAL PROPERTY RIGHTS OF ANY OTHER ENTITY. EACH COPYRIGHT HOLDER DISCLAIMS ANY LIABILITY TO THE USER FOR CLAIMS BROUGHT BY ANY OTHER ENTITY BASED ON INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS OR OTHERWISE. AS A CONDITION TO EXERCISING THE RIGHTS GRANTED HEREUNDER, EACH USER HEREBY ASSUMES SOLE RESPONSIBILITY TO SECURE ANY OTHER INTELLECTUAL PROPERTY RIGHTS NEEDED, IF ANY. THE SOFTWARE IS NOT FAULT-TOLERANT AND IS NOT INTENDED FOR USE IN MISSION-CRITICAL SYSTEMS, SUCH AS THOSE USED IN THE OPERATION OF NUCLEAR FACILITIES, AIRCRAFT NAVIGATION OR COMMUNICATION SYSTEMS, AIR TRAFFIC CONTROL SYSTEMS, DIRECT LIFE SUPPORT MACHINES, OR WEAPONS SYSTEMS, IN WHICH THE FAILURE OF THE SOFTWARE OR SYSTEM

COULD LEAD DIRECTLY TO DEATH, PERSONAL INJURY, OR SEVERE PHYSICAL OR ENVIRONMENTAL DAMAGE ("HIGH RISK ACTIVITIES"). THE COPYRIGHT HOLDERS SPECIFICALLY DISCLAIM ANY EXPRESS OR IMPLIED WARRANTY OF FITNESS FOR HIGH RISK ACTIVITIES.

JasPer License Version 2.0

Copyright (c) 2001-2006 Michael David Adams

Copyright (c) 1999-2000 Image Power, Inc.

Copyright (c) 1999-2000 The University of British Columbia

All rights reserved.

Permission is hereby granted, free of charge, to any person (the "User") obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

1. The above copyright notices and this permission notice (which includes the disclaimer below) shall be included in all copies or substantial portions of the Software.
2. The name of a copyright holder shall not be used to endorse or promote products derived from the Software without specific prior written permission.

THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF THE SOFTWARE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER. THE SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE. NO ASSURANCES ARE PROVIDED BY THE COPYRIGHT HOLDERS THAT THE SOFTWARE DOES NOT INFRINGE THE PATENT OR OTHER INTELLECTUAL PROPERTY RIGHTS OF ANY OTHER ENTITY. EACH COPYRIGHT HOLDER DISCLAIMS ANY LIABILITY TO THE USER FOR CLAIMS BROUGHT BY ANY OTHER ENTITY BASED ON INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS OR OTHERWISE. AS A CONDITION TO EXERCISING THE RIGHTS GRANTED HEREUNDER, EACH USER HEREBY ASSUMES SOLE RESPONSIBILITY TO SECURE ANY OTHER INTELLECTUAL PROPERTY RIGHTS NEEDED, IF ANY. THE SOFTWARE IS NOT FAULT-TOLERANT AND IS NOT INTENDED FOR USE IN MISSION-CRITICAL SYSTEMS, SUCH AS THOSE USED IN THE OPERATION OF NUCLEAR FACILITIES, AIRCRAFT NAVIGATION OR COMMUNICATION SYSTEMS, AIR TRAFFIC CONTROL SYSTEMS, DIRECT LIFE SUPPORT MACHINES, OR WEAPONS SYSTEMS, IN WHICH THE FAILURE OF THE SOFTWARE OR SYSTEM COULD LEAD DIRECTLY TO DEATH, PERSONAL INJURY, OR SEVERE PHYSICAL OR ENVIRONMENTAL DAMAGE ("HIGH RISK ACTIVITIES"). THE COPYRIGHT HOLDERS SPECIFICALLY DISCLAIM ANY EXPRESS OR IMPLIED WARRANTY OF FITNESS FOR HIGH RISK ACTIVITIES.

This package was debianized by Christopher L Cheney <ccheney@debian.org> on

Fri, 22 Aug 2003 01:33:34 -0500.

The current maintainer is Roland Stigge <stigge@antcom.de>

It was downloaded from <http://www.ece.uvic.ca/~mdadams/jasper/>

Upstream Author: Michael Adams <mdadams@ece.uvic.ca>

License:

JasPer License Version 2.0

Copyright (c) 1999-2000 Image Power, Inc.

Copyright (c) 1999-2000 The University of British Columbia

Copyright (c) 2001-2003 Michael David Adams

All rights reserved.

Permission is hereby granted, free of charge, to any person (the "User") obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

1. The above copyright notices and this permission notice (which includes the disclaimer below) shall be included in all copies or substantial portions of the Software.
2. The name of a copyright holder shall not be used to endorse or promote products derived from the Software without specific prior written permission.

THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF THE SOFTWARE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER. THE SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE. NO ASSURANCES ARE PROVIDED BY THE COPYRIGHT HOLDERS THAT THE SOFTWARE DOES NOT INFRINGE THE PATENT OR OTHER INTELLECTUAL PROPERTY RIGHTS OF ANY OTHER ENTITY. EACH COPYRIGHT HOLDER DISCLAIMS ANY LIABILITY TO THE USER FOR CLAIMS BROUGHT BY ANY OTHER ENTITY BASED ON INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS OR OTHERWISE. AS A CONDITION TO EXERCISING THE RIGHTS GRANTED HEREUNDER, EACH USER HEREBY ASSUMES SOLE RESPONSIBILITY TO SECURE ANY OTHER INTELLECTUAL PROPERTY RIGHTS NEEDED, IF ANY. THE SOFTWARE IS NOT FAULT-TOLERANT AND IS NOT INTENDED FOR USE IN MISSION-CRITICAL SYSTEMS, SUCH AS THOSE USED IN THE OPERATION OF NUCLEAR FACILITIES, AIRCRAFT NAVIGATION OR COMMUNICATION SYSTEMS, AIR TRAFFIC CONTROL SYSTEMS, DIRECT LIFE SUPPORT MACHINES, OR WEAPONS SYSTEMS, IN WHICH THE FAILURE OF THE SOFTWARE OR SYSTEM COULD LEAD DIRECTLY TO DEATH, PERSONAL INJURY, OR SEVERE PHYSICAL OR ENVIRONMENTAL DAMAGE ("HIGH RISK ACTIVITIES"). THE COPYRIGHT HOLDERS SPECIFICALLY DISCLAIM ANY EXPRESS OR IMPLIED WARRANTY OF FITNESS FOR HIGH RISK ACTIVITIES.

The Independent JPEG Group's JPEG software

README for release 6b of 27-Mar-1998

This distribution contains the sixth public release of the Independent JPEG Group's free JPEG software. You are welcome to redistribute this software and to use it for any purpose, subject to the conditions under LEGAL ISSUES, below.

Serious users of this software (particularly those incorporating it into larger programs) should contact IJG at jpeg-info@uunet.uu.net to be added to our electronic mailing list. Mailing list members are notified of updates and have a chance to participate in technical discussions, etc.

This software is the work of Tom Lane, Philip Gladstone, Jim Boucher, Lee Crocker, Julian Minguillon, Luis Ortiz, George Phillips, Davide Rossi, Guido Vollbeding, Ge' Weijers, and other members of the Independent JPEG Group.

IJG is not affiliated with the official ISO JPEG standards committee.

DOCUMENTATION ROADMAP

This file contains the following sections:

OVERVIEW General description of JPEG and the IJG software.

LEGAL ISSUES Copyright, lack of warranty, terms of distribution.

REFERENCES Where to learn more about JPEG.

ARCHIVE LOCATIONS Where to find newer versions of this software.

RELATED SOFTWARE Other stuff you should get.

FILE FORMAT WARS Software **not** to get.

TO DO Plans for future IJG releases.

Other documentation files in the distribution are:

User documentation:

install.doc How to configure and install the IJG software.

usage.doc Usage instructions for cjpeg, djpeg, jpegtran, rdjpgcom, and wrjpgcom.

*.1 Unix-style man pages for programs (same info as usage.doc).

wizard.doc Advanced usage instructions for JPEG wizards only.

change.log Version-to-version change highlights.

Programmer and internal documentation:

libjpeg.doc How to use the JPEG library in your own programs.

example.c Sample code for calling the JPEG library.

structure.doc Overview of the JPEG library's internal structure.

filelist.doc Road map of IJG files.

coderrules.doc Coding style rules --- please read if you contribute code.

Please read at least the files install.doc and usage.doc. Useful information can also be found in the JPEG FAQ (Frequently Asked Questions) article. See ARCHIVE LOCATIONS below to find out where to obtain the FAQ article.

If you want to understand how the JPEG code works, we suggest reading one or more of the REFERENCES, then looking at the documentation files (in roughly the order listed) before diving into the code.

OVERVIEW

This package contains C software to implement JPEG image compression and decompression. JPEG (pronounced "jay-peg") is a standardized compression method for full-color and gray-scale images. JPEG is intended for compressing "real-world" scenes; line drawings, cartoons and other non-realistic images are not its strong suit. JPEG is lossy, meaning that the output image is not exactly identical to the input image. Hence you must not use JPEG if you have to have identical output bits. However, on typical photographic images, very good compression levels can be obtained with no visible change, and remarkably high compression levels are possible if you can tolerate a low-quality image. For more details, see the references, or just experiment with various compression settings.

This software implements JPEG baseline, extended-sequential, and progressive compression processes. Provision is made for supporting all variants of these processes, although some uncommon parameter settings aren't implemented yet. For legal reasons, we are not distributing code for the arithmetic-coding variants of JPEG; see LEGAL ISSUES. We have made no provision for supporting the hierarchical or lossless processes defined in the standard.

We provide a set of library routines for reading and writing JPEG image files, plus two sample applications "cjpeg" and "djpeg", which use the library to perform conversion between JPEG and some other popular image file formats. The library is intended to be reused in other applications.

In order to support file conversion and viewing software, we have included considerable functionality beyond the bare JPEG coding/decoding capability; for example, the color quantization modules are not strictly part of JPEG decoding, but they are essential for output to colormapped file formats or colormapped displays. These extra functions can be compiled out of the library if not required for a particular application. We have also included "jpegtran", a utility for lossless transcoding between different JPEG processes, and "rdjpgcom" and "wrjpgcom", two simple applications for inserting and extracting textual comments in JFIF files.

The emphasis in designing this software has been on achieving portability and flexibility, while also making it fast enough to be useful. In particular, the software is not intended to be read as a tutorial on JPEG. (See the REFERENCES section for introductory material.) Rather, it is intended to be reliable, portable, industrial-strength code. We do not claim to have achieved that goal in every aspect of the software, but we strive for it.

We welcome the use of this software as a component of commercial products. No royalty is required, but we do ask for an acknowledgement in product documentation, as described under LEGAL ISSUES.

LEGAL ISSUES

In plain English:

1. We don't promise that this software works. (But if you find any bugs, please let us know!)
2. You can use this software for whatever you want. You don't have to pay us.
3. You may not pretend that you wrote this software. If you use it in a program, you must acknowledge somewhere in your documentation that you've used the IJG code.

In legalese:

The authors make NO WARRANTY or representation, either express or implied, with respect to this software, its quality, accuracy, merchantability, or fitness for a particular purpose. This software is provided "AS IS", and you, its user, assume the entire risk as to its quality and accuracy.

This software is copyright (C) 1991-1998, Thomas G. Lane.

All Rights Reserved except as specified below.

Permission is hereby granted to use, copy, modify, and distribute this software (or portions thereof) for any purpose, without fee, subject to these conditions:

- (1) If any part of the source code for this software is distributed, then this README file must be included, with this copyright and no-warranty notice unaltered; and any additions, deletions, or changes to the original files must be clearly indicated in accompanying documentation.
- (2) If only executable code is distributed, then the accompanying documentation must state that "this software is based in part on the work of the Independent JPEG Group".
- (3) Permission for use of this software is granted only if the user accepts full responsibility for any undesirable consequences; the authors accept NO LIABILITY for damages of any kind.

These conditions apply to any software derived from or based on the IJG code, not just to the unmodified library. If you use our work, you ought to acknowledge us.

Permission is NOT granted for the use of any IJG author's name or company name in advertising or publicity relating to this software or products derived from it. This software may be referred to only as "the Independent JPEG Group's software".

We specifically permit and encourage the use of this software as the basis of commercial products, provided that all warranty or liability claims are assumed by the product vendor.

ansi2knr.c is included in this distribution by permission of L. Peter Deutsch, sole proprietor of its copyright holder, Aladdin Enterprises of Menlo Park, CA. ansi2knr.c is NOT covered by the above copyright and conditions, but instead by the usual distribution terms of the Free Software Foundation; principally, that you must include source code if you redistribute it. (See the file ansi2knr.c for full details.) However, since ansi2knr.c is not needed as part of any program generated from the IJG code, this does not limit you more than the foregoing paragraphs do.

The Unix configuration script "configure" was produced with GNU Autoconf. It is copyright by the Free Software

Foundation but is freely distributable. The same holds for its supporting scripts (`config.guess`, `config.sub`, `ltconfig`, `ltmain.sh`). Another support script, `install-sh`, is copyright by M.I.T. but is also freely distributable.

It appears that the arithmetic coding option of the JPEG spec is covered by patents owned by IBM, AT&T, and Mitsubishi. Hence arithmetic coding cannot legally be used without obtaining one or more licenses. For this reason, support for arithmetic coding has been removed from the free JPEG software. (Since arithmetic coding provides only a marginal gain over the unpatented Huffman mode, it is unlikely that very many implementations will support it.) So far as we are aware, there are no patent restrictions on the remaining code.

The IJG distribution formerly included code to read and write GIF files. To avoid entanglement with the Unisys LZW patent, GIF reading support has been removed altogether, and the GIF writer has been simplified to produce "uncompressed GIFs". This technique does not use the LZW algorithm; the resulting GIF files are larger than usual, but are readable by all standard GIF decoders.

We are required to state that

"The Graphics Interchange Format(c) is the Copyright property of CompuServe Incorporated. GIF(sm) is a Service Mark property of CompuServe Incorporated."

REFERENCES

=====

We highly recommend reading one or more of these references before trying to understand the innards of the JPEG software.

The best short technical introduction to the JPEG compression algorithm is

Wallace, Gregory K. "The JPEG Still Picture Compression Standard",

Communications of the ACM, April 1991 (vol. 34 no. 4), pp. 30-44.

(Adjacent articles in that issue discuss MPEG motion picture compression, applications of JPEG, and related topics.) If you don't have the CACM issue handy, a PostScript file containing a revised version of Wallace's article is available at <ftp://ftp.uu.net/graphics/jpeg/wallace.ps.gz>. The file (actually a preprint for an article that appeared in IEEE Trans. Consumer Electronics) omits the sample images that appeared in CACM, but it includes corrections and some added material. Note: the Wallace article is copyright ACM and IEEE, and it may not be used for commercial purposes.

A somewhat less technical, more leisurely introduction to JPEG can be found in "The Data Compression Book" by Mark Nelson and Jean-loup Gailly, published by M&T Books (New York), 2nd ed. 1996, ISBN 1-55851-434-1. This book provides good explanations and example C code for a multitude of compression methods including JPEG. It is an excellent source if you are comfortable reading C code but don't know much about data compression in general. The book's JPEG sample code is far from industrial-strength, but when you are ready to look at a full implementation, you've got one here...

The best full description of JPEG is the textbook "JPEG Still Image Data Compression Standard" by William B. Pennebaker and Joan L. Mitchell, published by Van Nostrand Reinhold, 1993, ISBN 0-442-01272-1. Price US\$59.95, 638 pp. The book includes the complete text of the ISO JPEG standards (DIS 10918-1 and draft DIS 10918-2). This is by far the most complete exposition of JPEG in existence, and we highly recommend it.

The JPEG standard itself is not available electronically; you must order a paper copy through ISO or ITU. (Unless you feel a need to own a certified official copy, we recommend buying the Pennebaker and Mitchell book instead; it's much cheaper and includes a great deal of useful explanatory material.) In the USA, copies of the standard

may be ordered from ANSI Sales at (212) 642-4900, or from Global Engineering Documents at (800) 854-7179. (ANSI doesn't take credit card orders, but Global does.) It's not cheap: as of 1992, ANSI was charging \$95 for Part 1 and \$47 for Part 2, plus 7% shipping/handling. The standard is divided into two parts, Part 1 being the actual specification, while Part 2 covers compliance testing methods. Part 1 is titled "Digital Compression and Coding of Continuous-tone Still Images, Part 1: Requirements and guidelines" and has document numbers ISO/IEC IS 10918-1, ITU-T T.81. Part 2 is titled "Digital Compression and Coding of Continuous-tone Still Images, Part 2: Compliance testing" and has document numbers ISO/IEC IS 10918-2, ITU-T T.83.

Some extensions to the original JPEG standard are defined in JPEG Part 3, a newer ISO standard numbered ISO/IEC IS 10918-3 and ITU-T T.84. IJG currently does not support any Part 3 extensions.

The JPEG standard does not specify all details of an interchangeable file format. For the omitted details we follow the "JFIF" conventions, revision 1.02. A copy of the JFIF spec is available from:

Literature Department

C-Cube Microsystems, Inc.

1778 McCarthy Blvd.

Milpitas, CA 95035

phone (408) 944-6300, fax (408) 944-6314

A PostScript version of this document is available by FTP at <ftp://ftp.uu.net/graphics/jpeg/jfif.ps.gz>. There is also a plain text version at <ftp://ftp.uu.net/graphics/jpeg/jfif.txt.gz>, but it is missing the figures.

The TIFF 6.0 file format specification can be obtained by FTP from <ftp://ftp.sgi.com/graphics/tiff/TIFF6.ps.gz>. The JPEG incorporation scheme found in the TIFF 6.0 spec of 3-June-92 has a number of serious problems. IJG does not recommend use of the TIFF 6.0 design (TIFF Compression tag 6). Instead, we recommend the JPEG design proposed by TIFF Technical Note #2 (Compression tag 7). Copies of this Note can be obtained from [ftp.sgi.com](ftp://ftp.sgi.com) or from <ftp://ftp.uu.net/graphics/jpeg/>. It is expected that the next revision of the TIFF spec will replace the 6.0 JPEG design with the Note's design. Although IJG's own code does not support TIFF/JPEG, the free libtiff library uses our library to implement TIFF/JPEG per the Note. libtiff is available from <ftp://ftp.sgi.com/graphics/tiff/>.

ARCHIVE LOCATIONS

=====

The "official" archive site for this software is <ftp.uu.net> (Internet address 192.48.96.9). The most recent released version can always be found there in directory [graphics/jpeg](ftp://ftp.uu.net/graphics/jpeg/). This particular version will be archived as <ftp://ftp.uu.net/graphics/jpeg/jpegsrvc.v6b.tar.gz>. If you don't have direct Internet access, UUNET's archives are also available via UUCP; contact help@uunet.uu.net for information on retrieving files that way.

Numerous Internet sites maintain copies of the UUNET files. However, only <ftp.uu.net> is guaranteed to have the latest official version.

You can also obtain this software in DOS-compatible "zip" archive format from the SimTel archives (<ftp://ftp.simtel.net/pub/simtelnet/msdos/graphics/>), or on CompuServe in the Graphics Support forum (GO CIS:GRAPHSUP), library 12 "JPEG Tools". Again, these versions may sometimes lag behind the <ftp.uu.net> release.

The JPEG FAQ (Frequently Asked Questions) article is a useful source of general information about JPEG. It is updated constantly and therefore is not included in this distribution. The FAQ is posted every two weeks to Usenet newsgroups comp.graphics.misc, news.answers, and other groups. It is available on the World Wide Web at <http://www.faqs.org/faqs/jpeg-faq/> and other news.answers archive sites, including the official news.answers

archive at rtfm.mit.edu: <ftp://rtfm.mit.edu/pub/usenet/news.answers/jpeg-faq/>. If you don't have Web or FTP access, send e-mail to mail-server@rtfm.mit.edu with body

send usenet/news.answers/jpeg-faq/part1

send usenet/news.answers/jpeg-faq/part2

RELATED SOFTWARE

=====

Numerous viewing and image manipulation programs now support JPEG. (Quite a few of them use this library to do so.) The JPEG FAQ described above lists some of the more popular free and shareware viewers, and tells where to obtain them on Internet.

If you are on a Unix machine, we highly recommend Jef Poskanzer's free PBMPLUS software, which provides many useful operations on PPM-format image files. In particular, it can convert PPM images to and from a wide range of other formats, thus making cjpeg/djpeg considerably more useful. The latest version is distributed by the NetPBM group, and is available from numerous sites, notably <ftp://wuarchive.wustl.edu/graphics/graphics/packages/NetPBM/>. Unfortunately PBMPLUS/NETPBM is not nearly as portable as the IJG software is; you are likely to have difficulty making it work on any non-Unix machine.

A different free JPEG implementation, written by the PVRG group at Stanford, is available from <ftp://havefun.stanford.edu/pub/jpeg/>. This program is designed for research and experimentation rather than production use; it is slower, harder to use, and less portable than the IJG code, but it is easier to read and modify. Also, the PVRG code supports lossless JPEG, which we do not. (On the other hand, it doesn't do progressive JPEG.)

FILE FORMAT WARS

=====

Some JPEG programs produce files that are not compatible with our library. The root of the problem is that the ISO JPEG committee failed to specify a concrete file format. Some vendors "filled in the blanks" on their own, creating proprietary formats that no one else could read. (For example, none of the early commercial JPEG implementations for the Macintosh were able to exchange compressed files.)

The file format we have adopted is called JFIF (see REFERENCES). This format has been agreed to by a number of major commercial JPEG vendors, and it has become the de facto standard. JFIF is a minimal or "low end" representation. We recommend the use of TIFF/JPEG (TIFF revision 6.0 as modified by TIFF Technical Note #2) for "high end" applications that need to record a lot of additional data about an image. TIFF/JPEG is fairly new and not yet widely supported, unfortunately.

The upcoming JPEG Part 3 standard defines a file format called SPIFF. SPIFF is interoperable with JFIF, in the sense that most JFIF decoders should be able to read the most common variant of SPIFF. SPIFF has some technical advantages over JFIF, but its major claim to fame is simply that it is an official standard rather than an informal one. At this point it is unclear whether SPIFF will supersede JFIF or whether JFIF will remain the de facto standard. IJG intends to support SPIFF once the standard is frozen, but we have not decided whether it should become our default output format or not. (In any case, our decoder will remain capable of reading JFIF indefinitely.)

Various proprietary file formats incorporating JPEG compression also exist. We have little or no sympathy for the existence of these formats. Indeed, one of the original reasons for developing this free software was to help force convergence on common, open format standards for JPEG files. Don't use a proprietary file format!

TO DO

=====

The major thrust for v7 will probably be improvement of visual quality. The current method for scaling the quantization tables is known not to be very good at low Q values. We also intend to investigate block boundary smoothing, "poor man's variable quantization", and other means of improving quality-vs-file-size performance without sacrificing compatibility.

In future versions, we are considering supporting some of the upcoming JPEG Part 3 extensions --- principally, variable quantization and the SPIFF file format.

As always, speeding things up is of great interest.

Please send bug reports, offers of help, etc. to jpeg-info@uunet.uu.net.

This copy of the libpng notices is provided for your convenience. In case of any discrepancy between this copy and the notices in the file `png.h` that is included in the libpng distribution, the latter shall prevail.

COPYRIGHT NOTICE, DISCLAIMER, and LICENSE:

If you modify libpng you may insert additional notices immediately following this sentence.

This code is released under the libpng license.

libpng versions 1.2.6, August 15, 2004, through 1.5.12, July 11, 2012, are Copyright (c) 2004, 2006-2012 Glenn Randers-Pehrson, and are distributed according to the same disclaimer and license as libpng-1.2.5 with the following individual added to the list of Contributing Authors

Cosmin Truta

libpng versions 1.0.7, July 1, 2000, through 1.2.5 - October 3, 2002, are Copyright (c) 2000-2002 Glenn Randers-Pehrson, and are distributed according to the same disclaimer and license as libpng-1.0.6 with the following individuals added to the list of Contributing Authors

Simon-Pierre Cadieux

Eric S. Raymond

Gilles Vollant

and with the following additions to the disclaimer:

There is no warranty against interference with your enjoyment of the library or against infringement. There is no warranty that our efforts or the library will fulfill any of your particular purposes or needs. This library is provided with all faults, and the entire risk of satisfactory quality, performance, accuracy, and effort is with the user.

libpng versions 0.97, January 1998, through 1.0.6, March 20, 2000, are Copyright (c) 1998, 1999 Glenn Randers-Pehrson, and are distributed according to the same disclaimer and license as libpng-0.96, with the following individuals added to the list of Contributing Authors:

Tom Lane

Glenn Randers-Pehrson

Willem van Schaik

libpng versions 0.89, June 1996, through 0.96, May 1997, are Copyright (c) 1996, 1997 Andreas Dilger Distributed according to the same disclaimer and license as libpng-0.88, with the following individuals added to the list of Contributing Authors:

John Bowler

Kevin Bracey

Sam Bushell

Magnus Holmgren

Greg Roelofs

Tom Tanner

libpng versions 0.5, May 1995, through 0.88, January 1996, are Copyright (c) 1995, 1996 Guy Eric Schalnat, Group 42, Inc.

For the purposes of this copyright and license, "Contributing Authors" is defined as the following set of individuals:

Andreas Dilger

Dave Martindale

Guy Eric Schalnat

Paul Schmidt

Tim Wegner

The PNG Reference Library is supplied "AS IS". The Contributing Authors and Group 42, Inc. disclaim all warranties, expressed or implied, including, without limitation, the warranties of merchantability and of fitness for any purpose. The Contributing Authors and Group 42, Inc. assume no liability for direct, indirect, incidental, special, exemplary, or consequential damages, which may result from the use of the PNG Reference Library, even if advised of the possibility of such damage.

Permission is hereby granted to use, copy, modify, and distribute this source code, or portions hereof, for any purpose, without fee, subject to the following restrictions:

1. The origin of this source code must not be misrepresented.
2. Altered versions must be plainly marked as such and must not be misrepresented as being the original source.
3. This Copyright notice may not be removed or altered from any source or altered source distribution.

The Contributing Authors and Group 42, Inc. specifically permit, without fee, and encourage the use of this source code as a component to supporting the PNG file format in commercial products. If you use this source code in a product, acknowledgment is not required but would be appreciated.

A "png_get_copyright" function is available, for convenient use in "about" boxes and the like:

```
printf("%s",png_get_copyright(NULL));
```

Also, the PNG logo (in PNG format, of course) is supplied in the files "pngbar.png" and "pngbar.jpg (88x31) and "pngnow.png" (98x31).

Libpng is OSI Certified Open Source Software. OSI Certified Open Source is a certification mark of the Open Source Initiative.

Glenn Randers-Pehrson

glennrp at users.sourceforge.net

July 11, 2012

Copyright (c) 1988-1997 Sam Leffler

Copyright (c) 1991-1997 Silicon Graphics, Inc.

Permission to use, copy, modify, distribute, and sell this software and its documentation for any purpose is hereby granted without fee, provided that (i) the above copyright notices and this permission notice appear in all copies of the software and related documentation, and (ii) the names of Sam Leffler and Silicon Graphics may not be used in any advertising or publicity relating to the software without the specific, prior written permission of Sam Leffler and Silicon Graphics.

THE SOFTWARE IS PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EXPRESS, IMPLIED OR OTHERWISE, INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

IN NO EVENT SHALL SAM LEFFLER OR SILICON GRAPHICS BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT OR CONSEQUENTIAL DAMAGES OF ANY KIND, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER OR NOT ADVISED OF THE POSSIBILITY OF DAMAGE, AND ON ANY THEORY OF LIABILITY, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Copyright (c) 2006, Industrial Light & Magic, a division of Lucasfilm Entertainment Company Ltd. Portions contributed and copyright held by others as indicated. All rights reserved.

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

* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

* Neither the name of Industrial Light & Magic nor the names of any other contributors to this software may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS 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 COPYRIGHT OWNER 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.

== LICENSE file ==

== ==

== in this case for the Threading Building Blocks code ==

GNU GENERAL PUBLIC LICENSE

For the complete text of this license, see [GNU GENERAL PUBLIC LICENSE](#) in this section.

Threading Building Blocks

The source code of Threading Building Blocks is distributed under version 2 of the GNU General Public License, with the so-called "runtime exception," as follows (or see any header or implementation file):

As a special exception, you may use this file as part of a free software library without restriction. Specifically, if other files instantiate templates or use macros or inline functions from this file, or you compile this file and link it with other files to produce an executable, this file does not by itself cause the resulting executable to be covered by the GNU General Public License. This exception does not however invalidate any other reasons why the executable file might be covered by the GNU General Public License.

ZLIB DATA COMPRESSION LIBRARY

zlib 1.2.7 is a general purpose data compression library. All the code is thread safe. The data format used by the zlib library is described by RFCs (Request for Comments) 1950 to 1952 in the files <http://tools.ietf.org/html/rfc1950> (zlib format), [rfc1951](http://tools.ietf.org/html/rfc1951) (deflate format) and [rfc1952](http://tools.ietf.org/html/rfc1952) (gzip format).

All functions of the compression library are documented in the file `zlib.h` (volunteer to write man pages welcome, contact zlib@gzip.org). A usage example of the library is given in the file `test/example.c` which also tests that the library is working correctly. Another example is given in the file `test/minigzip.c`. The compression library itself is composed of all source files in the root directory.

To compile all files and run the test program, follow the instructions given at the top of `Makefile.in`. In short `./configure; make test`, and if that goes well, `make install` should work for most flavors of Unix. For Windows, use one of the special makefiles in `win32/` or `contrib/vstudio/`. For VMS, use `make_vms.com`.

Questions about zlib should be sent to [<zlib@gzip.org>](mailto:zlib@gzip.org), or to Gilles Vollant [<info@winimage.com>](mailto:info@winimage.com) for the Windows DLL version. The zlib home page is <http://zlib.net/>. Before reporting a problem, please check this site to verify that you have the latest version of zlib; otherwise get the latest version and check whether the problem still exists or not.

PLEASE read the zlib FAQ http://zlib.net/zlib_faq.html before asking for help.

Mark Nelson <markn@ieee.org> wrote an article about zlib for the Jan. 1997 issue of Dr. Dobb's Journal; a copy of the article is available at <http://marknelson.us/1997/01/01/zlib-engine/>.

The changes made in version 1.2.7 are documented in the file ChangeLog.

Unsupported third party contributions are provided in directory contrib/.

zlib is available in Java using the `java.util.zip` package, documented at <http://java.sun.com/developer/technicalArticles/Programming/compression/>.

A Perl interface to zlib written by Paul Marquess <pmqs@cpan.org> is available at CPAN (Comprehensive Perl Archive Network) sites, including <http://search.cpan.org/~pmqs/IO-Compress-Zlib/>.

A Python interface to zlib written by A.M. Kuchling <amk@amk.ca> is available in Python 1.5 and later versions, see <http://docs.python.org/library/zlib.html>.

zlib is built into tcl: <http://wiki.tcl.tk/4610>.

An experimental package to read and write files in .zip format, written on top of zlib by Gilles Vollant <info@winimage.com>, is available in the contrib/minizip directory of zlib.

Notes for some targets:

- For Windows DLL versions, please see `win32/DLL_FAQ.txt`
- For 64-bit Irix, `deflate.c` must be compiled without any optimization. With `-O`, one `libpng` test fails. The test works in 32 bit mode (with the `-n32` compiler flag). The compiler bug has been reported to SGI.
- zlib doesn't work with `gcc 2.6.3` on a DEC 3000/300LX under OSF/1 2.1 it works when compiled with `cc`.
- On Digital Unix 4.0D (formerly OSF/1) on AlphaServer, the `cc` option `-std1` is necessary to get `gzprintf` working correctly. This is done by `configure`.
- zlib doesn't work on HP-UX 9.05 with some versions of `/bin/cc`. It works with other compilers. Use "make test" to check your compiler.
- `gzdopen` is not supported on RISCOS or BEOS.
- For PalmOs, see <http://palmzlib.sourceforge.net/>

Acknowledgments:

The deflate format used by zlib was defined by Phil Katz. The deflate and zlib specifications were written by L. Peter Deutsch. Thanks to all the people who reported problems and suggested various improvements in zlib; they are too numerous to cite here.

Copyright notice:

(C) 1995-2012 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly Mark Adler

jloup@gzip.org madler@alumni.caltech.edu

If you use the zlib library in a product, we would appreciate *not* receiving lengthy legal documents to sign. The sources are provided for free but without warranty of any kind. The library has been entirely written by Jean-loup Gailly and Mark Adler; it does not include third-party code.

If you redistribute modified sources, we would appreciate that you include in the file ChangeLog history information documenting your changes. Please read the FAQ for more information on the distribution of modified source versions.

gstvideocuda

Copyright (c) 2015, NVIDIA CORPORATION. All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

bpmp and tos-img

The bpmp and tos-img libraries are covered by the following copyrights and license. This information is also provided in the `LICENSE.bpmp_and_tos-img` file.

Copyright (c) 2013-2015, NVIDIA CORPORATION. All rights reserved.

NVIDIA CORPORATION and its licensors retain all intellectual property and proprietary rights in and to this software, related documentation and any modifications thereto. Any use, reproduction, disclosure or

distribution of this software and related documentation without an express license agreement from NVIDIA CORPORATION is strictly prohibited.

This product incorporates software provided under the following terms:

- Copyright (c) 2008-2010 Travis Geiselbrecht

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Legal Information

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 NON-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, CUDA, Jetson, Tegra, TensorRT, and Vibrante 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.

The Android robot is reproduced or modified from work created and shared by Google and is used according to terms described in the Creative Commons 3.0 Attribution License.

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

ARM, AMBA, and ARM Powered are registered trademarks of ARM Limited. Cortex, MPCore and Mali are trademarks of ARM Limited. All other brands or product names are the property of their respective holders. "ARM" is used to represent ARM Holdings plc; its operating company ARM Limited; and the regional subsidiaries ARM Inc.; ARM KK; ARM Korea Limited.; ARM Taiwan Limited; ARM France SAS; ARM Consulting (Shanghai) Co. Ltd.; ARM Germany GmbH; ARM Embedded Technologies Pvt. Ltd.; ARM Norway, AS and ARM Sweden AB.

Copyright

© 2017 by NVIDIA Corporation. All rights reserved.