

# ***TMS320C6000 Network Developer's Kit (NDK) Support Package Ethernet Driver Design Guide***

Literature Number: SPRUFP2  
January 2009



## IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, modifications, enhancements, improvements, and other changes to its products and services at any time and to discontinue any product or service without notice. Customers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All products are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its hardware products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by government requirements, testing of all parameters of each product is not necessarily performed.

TI assumes no liability for applications assistance or customer product design. Customers are responsible for their products and applications using TI components. To minimize the risks associated with customer products and applications, customers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any TI patent right, copyright, mask work right, or other TI intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information published by TI regarding third-party products or services does not constitute a license from TI to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

Reproduction of information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. Reproduction of this information with alteration is an unfair and deceptive business practice. TI is not responsible or liable for such altered documentation.

Resale of TI products or services with statements different from or beyond the parameters stated by TI for that product or service voids all express and any implied warranties for the associated TI product or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

Following are URLs where you can obtain information on other Texas Instruments products and application solutions:

### Products

Amplifiers	<a href="http://amplifier.ti.com">amplifier.ti.com</a>
Data Converters	<a href="http://dataconverter.ti.com">dataconverter.ti.com</a>
DSP	<a href="http://dsp.ti.com">dsp.ti.com</a>
Interface	<a href="http://interface.ti.com">interface.ti.com</a>
Logic	<a href="http://logic.ti.com">logic.ti.com</a>
Power Mgmt	<a href="http://power.ti.com">power.ti.com</a>
Microcontrollers	<a href="http://microcontroller.ti.com">microcontroller.ti.com</a>
Low Power Wireless	<a href="http://www.ti.com/lpw">www.ti.com/lpw</a>

### Applications

Audio	<a href="http://www.ti.com/audio">www.ti.com/audio</a>
Automotive	<a href="http://www.ti.com/automotive">www.ti.com/automotive</a>
Broadband	<a href="http://www.ti.com/broadband">www.ti.com/broadband</a>
Digital Control	<a href="http://www.ti.com/digitalcontrol">www.ti.com/digitalcontrol</a>
Military	<a href="http://www.ti.com/military">www.ti.com/military</a>
Optical Networking	<a href="http://www.ti.com/opticalnetwork">www.ti.com/opticalnetwork</a>
Security	<a href="http://www.ti.com/security">www.ti.com/security</a>
Telephony	<a href="http://www.ti.com/telephony">www.ti.com/telephony</a>
Video & Imaging	<a href="http://www.ti.com/video">www.ti.com/video</a>
Wireless	<a href="http://www.ti.com/wireless">www.ti.com/wireless</a>

Mailing Address: Texas Instruments  
Post Office Box 655303 Dallas, Texas 75265

## ***About This Guide***

This document describes the design of the Ethernet driver architecture introduced in NDK 2.0. This design differs from that in previous versions of NDK. All Ethernet drivers packaged as a part of the Network Developer's Kit (NDK) Support Package in NDK 2.0 follow the generic architecture described in this document.

### **Important Note:**

- ❑ This document covers only the Ethernet driver architecture. The serial driver is no longer supported as of NDK 2.0. For serial driver information, see documentation from earlier releases of the NDK.
- ❑ The setup and installation steps for each NDK Support Package (NSP) are provided in the Release Notes provided with that NSP.

## ***Intended Audience***

This document is intended for writers of Ethernet mini-drivers. This document assumes you have knowledge of Ethernet concepts.

## ***Related Documents***

The following books describe the TMS320C6x devices and related support tools. To obtain a copy of any of these TI documents, call the Texas Instruments Literature Response Center at (800) 477-8924. When ordering, please identify the book by its title and literature number. Many of these documents can be found on the Internet at <http://www.ti.com>.

- ❑ SPRU189 - *TMS320C6000 CPU and Instruction Set Reference Guide*. Describes the CPU architecture, pipeline, instruction set, and interrupts for the TMS320C6000 DSPs.
- ❑ SPRU190 - *TMS320C6000 DSP Peripherals Overview Reference Guide*. Provides an overview and briefly describes peripherals available on the TMS320C6000 family of DSPs.
- ❑ SPRU197 - *TMS320C6000 Technical Brief*. Provides an introduction to the TMS320C62x and TMS320C67x digital signal processors (DSPs) of the TMS320C6000 DSP family. Describes the CPU architecture, peripherals, development tools, and third-party support for the C62x and C67x DSPs.

- ❑ SPRU198 - *TMS320C6000 Programmer's Guide*. Reference for programming the TMS320C6000 digital signal processors (DSPs). Before you use this manual, you should install your code generation and debugging tools. Includes a brief description of the C6000 DSP architecture and code development flow, includes C code examples and discusses optimization methods for the C code, describes the structure of assembly code and includes examples and discusses optimizations for the assembly code, and describes programming considerations for the C64x DSP.
- ❑ SPRU509 - *Code Composer Studio Development Tools v3.3 Getting Started Guide* introduces some of the basic features and functionalities in Code Composer Studio to enable you to create and build simple projects.
- ❑ SPRU523 - *TMS320C6000 Network Developer's Kit (NDK) Software User's Guide*. Describes how to use the NDK libraries, how to develop networking applications on TMS320C6000 platforms, and ways to tune the NDK to fit a particular software environment.
- ❑ SPRU524 - *TMS320C6000 Network Developer's Kit (NDK) Software Programmer's Reference Guide*. Describes the various API functions provided by the stack libraries, including the low level hardware APIs.

## Notational Conventions

This document uses the following conventions:

- ❑ Program listings, program examples, and interactive displays are shown in a mono-spaced font. Examples use **bold** for emphasis, and interactive displays use **bold** to distinguish commands that you enter from items that the system displays (such as prompts, command output, error messages, etc.).
- ❑ Square brackets ( [ and ] ) identify an optional parameter. If you use an optional parameter, you specify the information within the brackets. Unless the square brackets are in a **bold** typeface, do not enter the brackets themselves.

## Trademarks

The Texas Instruments logo and Texas Instruments are registered trademarks of Texas Instruments. Trademarks of Texas Instruments include: TI, Code Composer, Code Composer Studio, DSP/BIOS, TMS320, TMS320C6000, TMS320C64x, TMS320DM644x, and TMS320C64x+.

All other brand, product names, and service names are trademarks or registered trademarks of their respective companies or organizations.

January 5, 2009

# Contents

---

---

---

---

<b>1</b>	<b>Architecture Overview</b> .....	<b>1-1</b>
	<i>This chapter provides an overview of the terminology and components involved in the Network Developer's Kit Support Package (NSP) Ethernet driver. It also describes the architecture of such drivers.</i>	
1.1	Acronyms .....	1-2
1.2	Ethernet Driver Architecture .....	1-3
1.2.1	NIMU-Specific Layer .....	1-4
1.2.2	Ethernet Mini-Driver .....	1-4
1.2.3	Generic EMAC/MDIO Chip Support Library .....	1-5
1.3	Flow Charts .....	1-6
1.4	Background .....	1-8
1.4.1	Network Control (NETCTRL) Module .....	1-8
1.4.2	Stack Event (STKEVENT) Object .....	1-8
1.4.3	Packet Buffer (PBM) Object .....	1-8
1.4.4	NDK Interrupt Manager .....	1-9
1.4.5	Data Alignment .....	1-9
1.5	API Overview .....	1-10
<b>2</b>	<b>NIMU Layer</b> .....	<b>2-1</b>
	<i>This chapter describes Network Interface Management Unit (NIMU) layer API.</i>	
2.1	Overview of the NIMU Layer .....	2-2
2.2	NIMU APIs .....	2-2
<b>3</b>	<b>Ethernet Mini-Driver Layer</b> .....	<b>3-1</b>
	<i>This chapter describes Ethernet mini-driver layer interface.</i>	
3.1	Overview .....	3-2
3.2	Data Structures .....	3-2
3.3	Ethernet Mini-Driver APIs .....	3-4
3.3.1	HwPktInit — Initialize Packet Driver Environment .....	3-5
3.3.2	HwPktOpen — Open Ethernet Device Instance .....	3-5
3.3.3	HwPktClose — Close Ethernet Device and Disable Interrupts .....	3-5
3.3.4	HwPktSetRx — Configure the Ethernet Receive Filter Settings .....	3-6
3.3.5	HwPktIoctl — Execute Driver-Specific IOCTL Commands .....	3-6
3.3.6	HwPktTxNext — Transmit Next Buffer in the Transmit Queue .....	3-6
3.3.7	_HwPktPoll — Mini-Driver Polling Function .....	3-7
3.4	Configuration Variables .....	3-7

<b>4</b>	<b>Generic EMAC/MDIO CSL Layer</b> . . . . .	<b>4-1</b>
	<i>This chapter describes the EMAC/MDIO CSL layer interface.</i>	
4.1	Overview . . . . .	4-2
4.2	CSL Data Structures . . . . .	4-2
4.3	EMAC APIs . . . . .	4-2
4.4	Callback Functions . . . . .	4-3
4.4.1	pfcBGetPacket . . . . .	4-3
4.4.2	pfcBFreePacket . . . . .	4-4
4.4.3	pfcBRxPacket . . . . .	4-4
4.4.4	pfcBStatus . . . . .	4-4
4.4.5	pfcBStatistics . . . . .	4-4

# Architecture Overview

---

---

---

This chapter provides an overview of the terminology and components involved in the Network Developer's Kit Support Package (NSP) Ethernet driver. It also describes the architecture of such drivers.

<b>Topic</b>	<b>Page</b>
<b>1.1 Acronyms</b> .....	<b>1-2</b>
<b>1.2 Ethernet Driver Architecture</b> .....	<b>1-3</b>
<b>1.3 Flow Charts</b> .....	<b>1-6</b>
<b>1.4 Background</b> .....	<b>1-8</b>
<b>1.5 API Overview</b> .....	<b>1-10</b>

## 1.1 Acronyms

The following acronyms are used in this document:

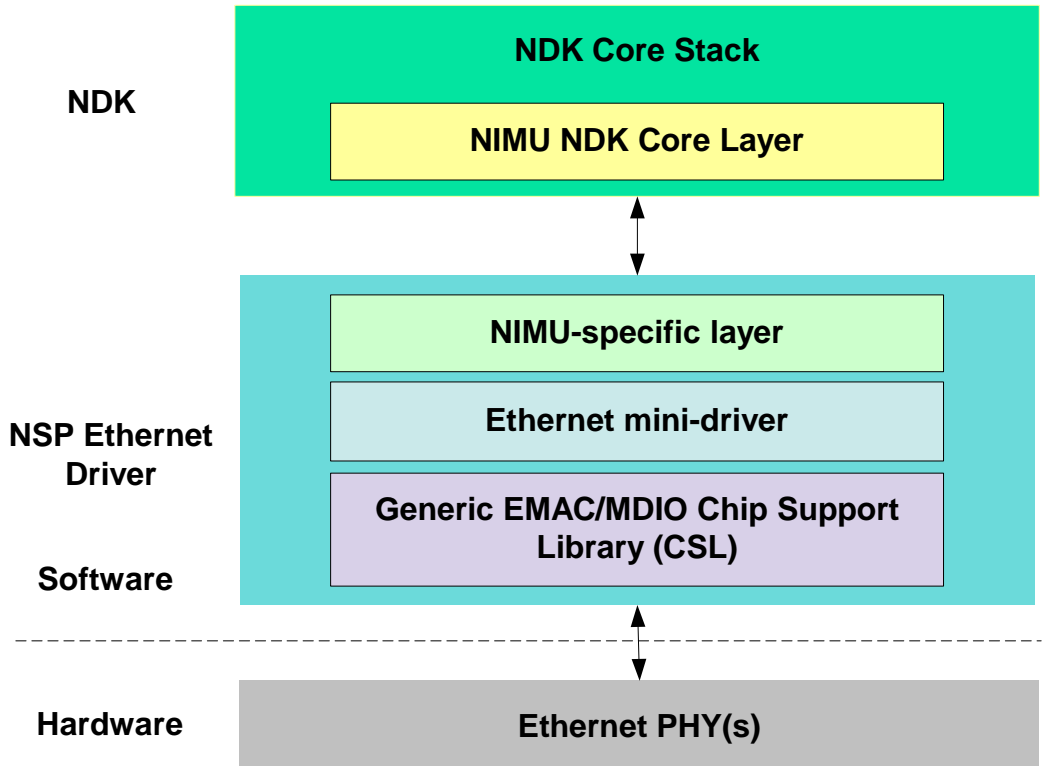
*Table 1–1 Acronyms*

<b>Acronym</b>	<b>Description</b>
API	Application Programming Interface
BD	Buffer Descriptor
CSL	Chip Support Library
DSP	Digital Signal Processor
EMAC	Ethernet Medium Access Protocol
LL	Low Level Packet Driver
MDIO	Management Data Input/Output Interface
NDK	Network Developer's Kit
NIMU	Network Interface Management Unit
NSP	NDK Support Package
OS AL	Operating Systems Abstraction Layer
Rx	Receive Operation
SGMII	Serial Gigabit Media Independent Interface
Tx	Transmit Operation



## 1.2 Ethernet Driver Architecture

The following diagram shows the architecture of the Ethernet driver design in the NDK 2.0 Support Package (NSP).



This new NSP Ethernet driver architecture consists of the following components:

- ❑ **NIMU-specific layer**, which acts as the interface between the NDK stack and the Ethernet driver. See Section 1.2.1.
- ❑ **Ethernet mini-driver**, which manages the EMAC configuration using the CSL. Also manages DSP interrupts and memory allocation for packet buffers in buffer descriptors using the NDK Operating Systems Abstraction Layer (OS AL). See Section 1.2.2.
- ❑ **Generic EMAC/MDIO Chip Support Library (CSL)**, which contains the generic APIs and data structures needed to control and configure EMAC/MDIO peripherals. Also manages buffer descriptors and interrupt service routines. See Section 1.2.3.

The NIMU-specific layer in previous versions of the NDK was generic enough to be ported to different platforms with ease. However, the mini-driver was not easily portable and had to be rewritten from scratch every time it had to be ported to a new platform. This led to different flavors of the Ethernet device drivers for different platforms—thus increasing the development, maintenance, and debugging effort.

To overcome the limitations of this architecture, the architecture of Ethernet drivers in the NSPs have been reorganized to optimize for a better development and debugging experience. The Generic EMAC/MDIO Chip Support Library (CSL) component is new; it has been split apart from the Ethernet mini-driver component to better isolate portions that commonly require changes when porting.

### 1.2.1 NIMU-Specific Layer

The Network Interface Management Unit (NIMU) specific layer acts as the interface between the Ethernet driver and the NDK core stack. It provides an implementation for the APIs defined by the NIMU specification for this EMAC device. These APIs let the NDK core stack control and configure the EMAC device at runtime and transmit packets. They also enable the driver to hand any received packets back to the stack.

This layer is fairly generic and doesn't change between different platforms.

This layer's functionality and role are the same as in versions of NDK prior to v2.0.

### 1.2.2 Ethernet Mini-Driver

This layer is responsible for setting up parameters for EMAC and MDIO configuration according to system needs. It uses APIs and data structures exported by the underlying Chip Support Library (CSL) layer. It is also responsible for setting up EMAC interrupts into the DSP using data structures and APIs exposed by the "Interrupt Manager Wrapper" in the NDK OS AL.

This layer acts as the sole memory manager in the Ethernet driver. That is, it handles all memory allocations, initializations, and frees of packet buffers for use in the buffer descriptors (BDs) in the Transmit (Tx) and Receive (Rx) paths. For memory management, it again uses the data structures and APIs defined by the NDK OS AL.

For the most part, the mini-driver invokes CSL APIs for setup, Tx, and interrupt service operations. The CSL layer, however, can also invoke the mini-driver layer. The CSL layer can invoke the mini-driver registered callback functions (set up during EMAC\_open) for updating statistics and reporting errors. On receiving a packet, it can hand over the packet to be passed up the stack or for memory allocation/free of buffers in BDs.

This layer is OS agnostic, since it uses the NDK OS AL for all memory and interrupt management operations. However, this layer is device-dependent since the EMAC peripheral setup requires knowledge of the capabilities of EMAC on this platform/device and will have to be customized for each platform and for application needs. So, this layer needs to be ported and customized from one platform to another.

### 1.2.3 Generic EMAC/MDIO Chip Support Library

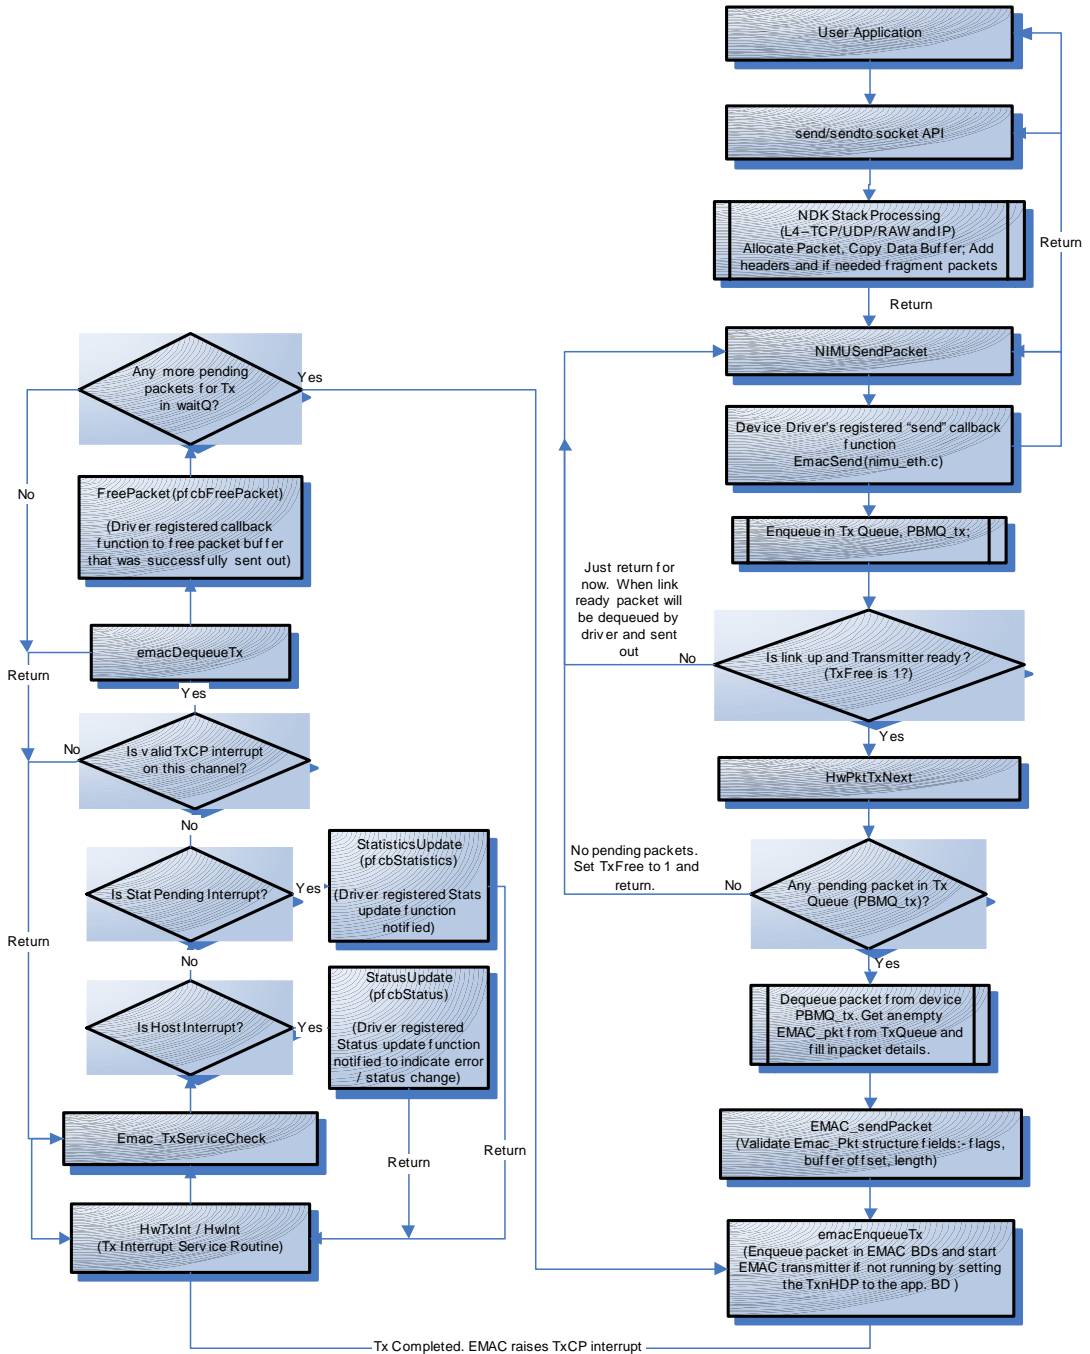
This layer enables the generic driver architecture by doing the following:

- ❑ **EMAC APIs.** It defines the data structures and interfaces (APIs) required to configure and use EMAC for transmit and receive operations.
- ❑ **MDIO and SGMII APIs.** It exposes APIs for managing the PHY-related (physical layer) configuration through the MDIO and SGMII (if the PHY is capable of gigabit speed) modules.
- ❑ **BD logic.** It implements the basic logic for CPPI Buffer Descriptor management (setup, enqueueing, and dequeueing operations).
- ❑ **ISR logic.** It contains the central logic for interrupt service routines. However, it uses the mini-driver's registered callback functions to report packet reception, statistics, errors, and obtaining or freeing a buffer for filling up a BD.

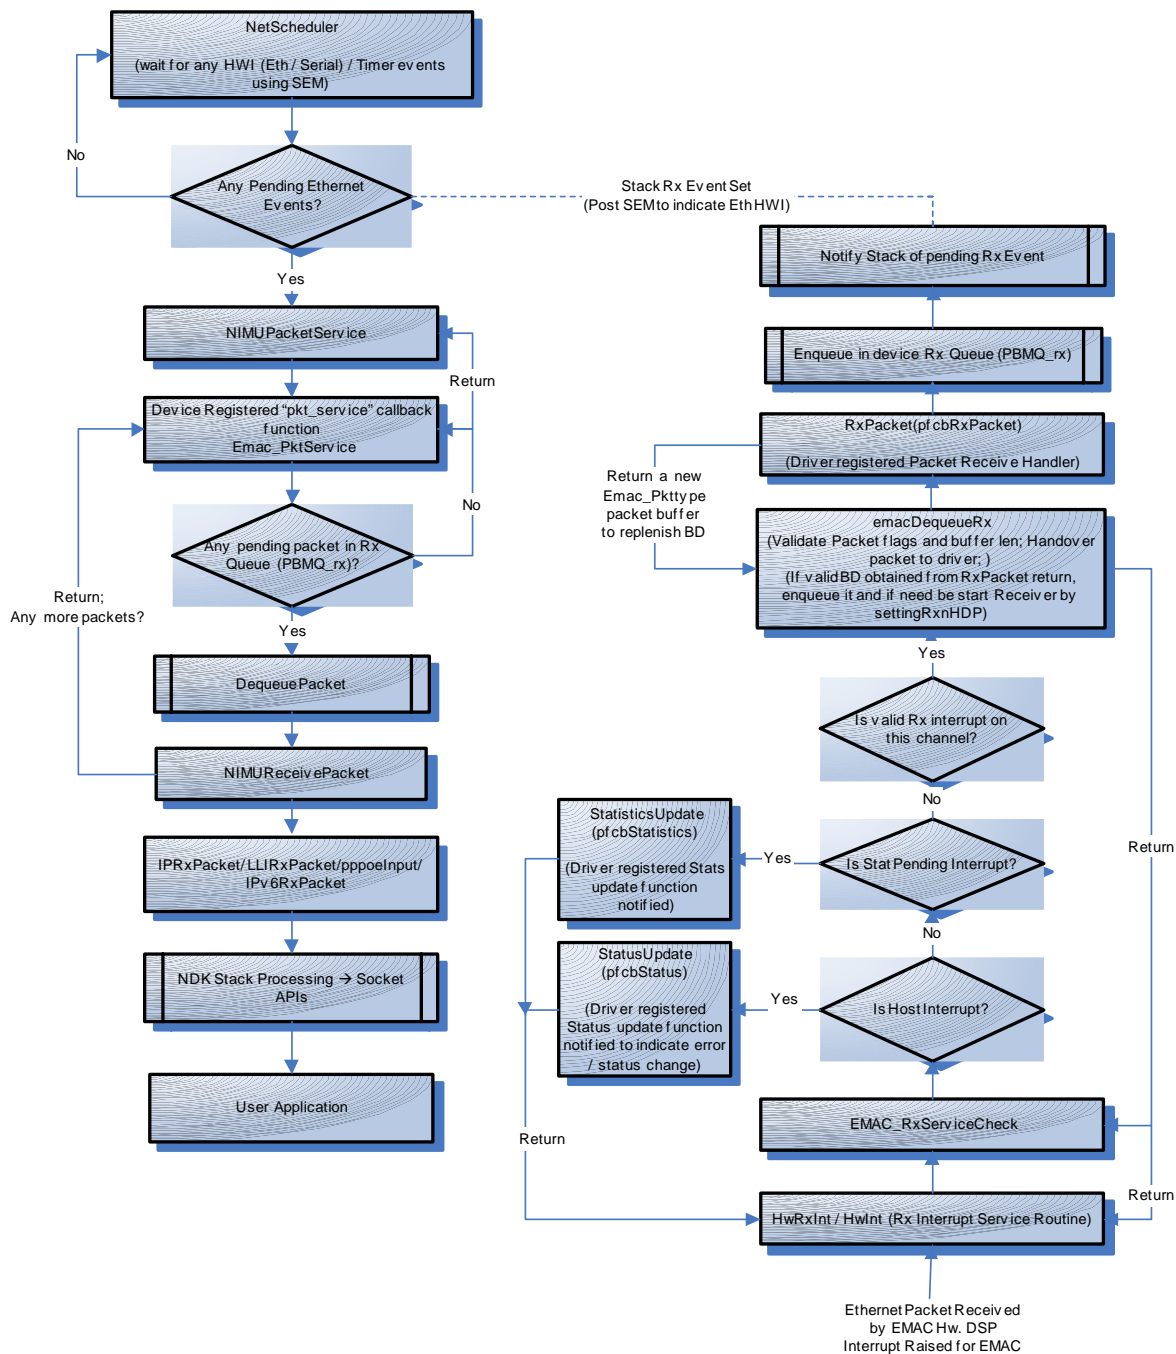
This layer is largely generic and doesn't vary much from platform to platform unless the EMAC capabilities change a whole lot. For example, the CSL for an EMAC peripheral connecting to a PHY switch would be very different from an EMAC that connects to a single PHY port. This layer is easily portable to different devices with similar capabilities.

### 1.3 Flow Charts

The transmission path for Ethernet packets is as follows:



The receive path for Ethernet packets is as follows:



## 1.4 Background

To port NDK Support Package device drivers, you should be familiar with the following constructs and concepts.

### 1.4.1 Network Control (NETCTRL) Module

The Network Control Module (NETCTRL) is at the center of the NDK and controls the interface of the HAL device drivers to the internal stack functions.

The NETCTRL module and its related APIs are described in both the *TMS320C6000 Network Developer's Kit (NDK) Software Programmer's Reference Guide* (SPRU524) and the *TMS320C6000 Network Developer's Kit (NDK) Software User's Guide* (SPRU523). To write device drivers, you must be familiar with NETCTRL. The description given in the *TMS320C6000 Network Developer's Kit (NDK) Software User's Guide* (SPRU523) is more appropriate for device driver work.

### 1.4.2 Stack Event (STKEVENT) Object

The STKEVENT object is a central component in the low-level architecture. It ties the HAL layer to the scheduler thread in the network control module (NETCTRL). The network scheduler thread waits on events from various device drivers in the system, including the Ethernet, serial, and timer drivers.

Device drivers use the STKEVENT object to inform the scheduler that an event has occurred. The STKEVENT object and its related API are described in the *TMS320C6000 Network Developer's Kit (NDK) Software Programmer's Reference Guide* (SPRU524). Device driver writers need to be familiar with STKEVENT.

### 1.4.3 Packet Buffer (PBM) Object

The PBM object is a packet buffer that is sourced and managed by the Packet Buffer Manager (PBM). The PBM is part of the OS adaptation layer (OS AL). It provides packet buffers for all packet-based devices in the system. Therefore, the serial port and Ethernet drivers both make use of this module.

The PBM module in the NDK OS AL manages packet buffers up to 3 KB in size. Any packet buffer allocation larger than 3 KB is managed by the Jumbo Packet Buffer Manager (Jumbo PBM). A default Jumbo PBM

implementation is provided in NDK OS AL; this implementation might need customization according to the application needs and system's memory constraints.

The PBM object, its related API, and the Jumbo PBM API are described in the *TMS320C6000 Network Developer's Kit (NDK) Software Programmer's Reference Guide* (SPRU524). The *TMS320C6000 Network Developer's Kit (NDK) Software User's Guide* (SPRU523) also includes a section on adapting the PBM to a particular included software.

#### 1.4.4 NDK Interrupt Manager

The NDK Interrupt Manager is a module in the NDK OS AL that abstracts out the OS (BIOS) specific APIs and data structures required for interrupt configuration and management. It exposes a simple interface to the driver writer to configure EMAC interrupts into the DSP core. Interrupt Setup (IntSetup) Object is a data structure defined by this module.

Depending on the system specification, there can be a single or multiple system event/interrupt numbers defined for the EMAC module's Transmit (Tx) and Receive (Rx) events. Also based on the system specification, one could register a single Interrupt Service Routine (ISR) for both Tx and Rx events or register separate ISRs for each event.

The following NDK Interrupt Manager APIs can be used by the Ethernet driver in setting up the interrupts:

- ❑ Interrupt\_add
- ❑ Interrupt\_delete
- ❑ Interrupt\_enable
- ❑ Interrupt\_disable

Please see the sample Ethernet driver code packaged as the NDK Support Package (NSP) for any C64x+ device for an illustration of interrupt configuration using NDK Interrupt Manager APIs. The NDK Interrupt Manager, along with its related API and data structures, are described in the *TMS320C6000 Network Developer's Kit (NDK) Software Programmer's Reference Guide* (SPRU524).

#### 1.4.5 Data Alignment

The NDK libraries have been built with the assumption that the IP header in a data packet is 16-bit aligned. In other words, the first byte of the IP packet (the version/length field) must start on an even 16-bit boundary. In any fixed-length header protocol, this requirement can be met by

backing off any odd byte header size, and adding it to the header padding specified to the stack. For Ethernet and peer-to-peer protocol (PPP), the only requirement is that the Ethernet or PPP packet not start on an odd byte boundary.

In addition, some drivers in the NDK are set up to have a 22-byte header. This is the header size of a PPPoE packet when sent using a 14-byte Ethernet header. When all arriving packets use the 22-byte header, it guarantees that they can be routed to any egress device with a header requirement up to that size. For Ethernet operation, this requires that a packet has 8 bytes of pre-pad to make its total header size 22 bytes.

The value of this pre-pad is #defined as PKT\_PREPAD in the Ethernet driver include files.

## 1.5 API Overview

The various APIs exposed by the three main layers—the NIMU-specific layer, mini-driver, and generic EMAC/MDIO CSL layer—can be classified based on their functionality into the following categories:

- ❑ **Initialization and Shutdown APIs.** These APIs are called during Ethernet device start up to initialize the EMAC environment or during shutdown to bring down the Ethernet controller and its subsystems.
- ❑ **Configuration APIs.** These APIs are called to get/set the EMAC configuration. The configuration APIs are generally useful in setting the following parameters:
  - multicast configuration
  - receive filters on the Ethernet device
- ❑ **Transmit APIs.** These APIs provide a well-defined interface for the NDK stack to pass down any available Ethernet packets onto the wire using the Ethernet driver.
- ❑ **Receive APIs.** These APIs provide a well-defined interface for the driver to pass up an Ethernet packet to the NDK stack and into an application.
- ❑ **Polling APIs.** These APIs provide an interface for the NDK core stack to monitor the status of the Ethernet link on a periodic basis and to perform any necessary configuration of the EMAC depending on a change of state, if any.



The following table groups the APIs defined by each of the Ethernet driver layers under one of these five categories.

*Table 1-2. API Mapping between the Ethernet driver layers*

<b>API Category</b>	<b>NIMU Layer</b>	<b>Mini-Driver Layer</b>	<b>CSL Layer</b>
<b>Initialization</b>	EmacInit	HwPktInit	--none--
	EmacStart	HwPktOpen	EMAC_open / MDIO_open
<b>Shutdown</b>	EmacStop	HwPktClose, HwPktShutdown	EMAC_close / MDIO_close
<b>Configuration</b>	Emacioctl	HwPktSetRx	EMAC_setReceiveFilter, EMAC_getReceiveFilter, EMAC_setMulticast, EMAC_getStatus, EMAC_getStatistics, EMAC_enumerate
<b>Transmit</b>	EmacSend	HwPktTxNext	EMAC_sendPacket, EMAC_TxServiceCheck (Tx ISR)
<b>Receive</b>	EmacPktService	HwInt / HwRxInt (depends on whether the EMAC has separate system events mapped into DSP for Rx/Tx or just one for both)	EMAC_RxServiceCheck (Rx ISR)
<b>Polling</b>	EmacPoll	_HwPktPoll	EMAC_TimerTick, MDIO_timerTick, MDIO_getStatus

The following chapters discuss each of the layers and APIs in detail.



# NIMU Layer

---

---

---

This chapter describes Network Interface Management Unit (NIMU) layer API.

<b>Topic</b>	<b>Page</b>
<b>2.1 Overview of the NIMU Layer . . . . .</b>	<b>2-2</b>
<b>2.2 NIMU APIs . . . . .</b>	<b>2-2</b>

## 2.1 Overview of the NIMU Layer

The Network Interface Management Unit (NIMU) layer interfaces with the NDK core stack. It enables the stack to control the device at runtime. This layer is platform-independent and is easily portable across various platforms.

## 2.2 NIMU APIs

Driver writers need to implement APIs as follows to make their driver NIMU-compliant:

- 1) Register a driver Init callback function with the core NDK NIMU layer by populating the function in the NIMUDeviceTable.
- 2) Allocate and initialize the NETIF\_DEVICE structure for this device with the appropriate parameters and callback functions defined for the following NIMU-defined APIs:
  - start
  - stop
  - poll
  - send
  - pkt\_service
  - ioctl
  - add\_header
- 3) Invoke the NIMURegister API to register this device with the NDK core's NIMU layer for further management.
- 4) Finally, implement all the callback functions as per the NIMU architecture guidelines and the API descriptions described in "Network Interface Management Unit" section of the *TMS320C6000 Network Developer's Kit (NDK) Software Programmer's Reference Guide* (SPRU524).

Please see the nimu\_eth.c file in the sample Ethernet driver code packaged as NDK Support Package (NSP) for any C64x+ device for an example NIMU API implementation.

# Ethernet Mini-Driver Layer

---

---

---

This chapter describes Ethernet mini-driver layer interface.

<b>Topic</b>	<b>Page</b>
<b>3.1 Overview</b> .....	<b>3-2</b>
<b>3.2 Data Structures</b> .....	<b>3-2</b>
<b>3.3 Ethernet Mini-Driver APIs</b> .....	<b>3-4</b>
<b>3.4 Configuration Variables</b> .....	<b>3-7</b>

## 3.1 Overview

The Ethernet mini-driver layer in the new driver architecture is responsible for setting up the EMAC subsystem configuration. It exposes various APIs to the NIMU layer through which the NDK stack can configure, control, transmit, and receive packets using the Ethernet controller. It sets up configuration for the EMAC, MDIO, and SGMII (if the physical layer is capable of gigabit speed) modules and acts like glue between the NIMU-specific layer and the low level EMAC configuration layer—that is, the Chip Support Library (CSL) layer—for those modules.

This layer is platform-dependent. Driver writers will need to know the PHY and EMAC capabilities and interrupt definitions for the specific system and will need to configure the Ethernet module accordingly.

## 3.2 Data Structures

Device configuration information is stored in a private device instance structure called "PDINFO" that is used to communicate the device configuration between the NIMU and the mini-driver layers.

```
typedef struct _pdinfo
{
    uint        PhysIdx;        /* physical index of device */
    HANDLE      hEther;        /* handle to logical driver */
    STKEVENT_Handle hEvent;    /* semaphore handle */
    UINT8       bMacAddr[6];    /* MAC address */
    uint        Filter;        /* current RX filter */
    uint        MCastCnt;      /* current MCast addr count */
    UINT8       bMCast[6*PKT_MAX_MCAST]; /* multicast list */
    uint        TxFree;        /* transmitter "free" flag */
    PBMQ        PBMQ_tx;       /* Tx queue */

#ifdef _INCLUDE_NIMU_CODE
    PBMQ        PBMQ_rx;       /* Rx queue */
#endif

} PDINFO;
```

The following list describes the structure items in more detail:

- **PhysIdx.** Physical Index of this device ( $\geq 0$ ). The PhysIdx may range from 0 to n-1. Care should be taken to ensure that the physical index of a device is unique if multiple instances of devices exist in the system. This attribute is an auxiliary field that can be used by the NIMU and mini-driver to communicate any data at run-time.

For example, the physical index can be used to hold the EMAC channel number on which packets using this device should be transmitted, and the mini-driver can be changed to use this info when transmitting the packet.

- ❑ **hEther.** This field is no longer being used after the switch to NIMU style drivers in NDK 2.0.
- ❑ **hEvent.** The handle to the semaphore object shared by the NDK stack and the driver to communicate pending network Rx events. This handle is used with the `STKEVENT_signal()` function to signal the NDK stack that a packet has been received and enqueued by the driver for hand off to the NDK Ethernet stack.
- ❑ **bMacAddr.** The Mac (Hardware) address of this interface. This is set to a default value by the NIMU layer. The default value can be overridden by the mini-driver with a value received from the EEPROM during device open.
- ❑ **Filter.** The current receive filter setting, which indicates which types of packets are accepted. The receive filter determines how the packet device should filter incoming packets. This field is set by the NIMU layer/stack and used by the mini-driver to program the EMAC. Legal values include:
  - `ETH_PKTFLT_NOHING`. no packets
  - `ETH_PKTFLT_DIRECT`. only directed Ethernet
  - `ETH_PKTFLT_BROADCAST`. directed plus Ethernet broadcast
  - `ETH_PKTFLT_MULTICAST`. directed, broadcast, and selected Ethernet multicast
  - `ETH_PKTFLT_ALLMULTICAST`. directed, broadcast, and all multicast
  - `ETH_PKTFLT_ALL`. All packets
- ❑ **MCastCnt.** Number of multicast addresses installed.
- ❑ **bMCast.** Multicast address list. This field is a byte array of consecutive 6-byte multicast MAC addresses. The number of valid addresses is stored in the `MCastCnt` field. The multicast address list determines what multicast addresses (if any) the MAC is allowed to receive. The multicast list is configured by the application.
- ❑ **TxFree.** Transmitter free flag. The `TxFree` flag is used by NIMU layer to determine if a new packet can be sent immediately by the mini-driver, or if it should be placed on the transmit pending queue for later. If the flag is not zero, the mini-driver function `HwPktTxNext()` is

called when a new packet is queued for transmission. This flag is maintained by the mini-driver.

- ❑ **PBMQ\_tx.** Transmit pending queue. The transmit pending queue holds all the packets waiting to be sent on the Ethernet device. The mini-driver pulls PBM packet buffers off this queue in its HwPktTxNext() function and posts them to the Ethernet MAC for transmit. Once the packet has been transmitted, the packet buffer is freed by the mini-driver calling PBM\_free(). There is one Tx queue for each PKT device.
- ❑ **PBMQ\_rx.** Receive queue. All packets received by the EMAC and handed over to the mini-driver are enqueued to the Rx queue. The mini-driver also signals the NDK stack of the pending receive packet that needs to be serviced in this queue. When the NDK scheduler thread next runs and finds this pending event to service, it invokes the NIMU layer EmacPktService function, which dequeues any pending packets on this queue and hands it over to the stack for further processing. There is one Rx queue for each PKT device.

### 3.3 Ethernet Mini-Driver APIs

The following APIs are exported by the Ethernet mini-driver layer:

- ❑ HwPktInit
- ❑ HwPktOpen
- ❑ HwPktClose
- ❑ HwPktShutdown
- ❑ HwPktSetRx
- ❑ HwPktTxNext
- ❑ \_HwPktPoll

As described in Section 1.5, the APIs exposed by this layer can be conveniently grouped according to their functionality into the following categories:

- 1) **Initialization.** HwPktInit, HwPktOpen
- 2) **Shutdown.** HwPktClose, HwPktShutdown
- 3) **Configuration.** HwPktSetRx
- 4) **Transmit.** HwPktTxNext
- 5) **Receive.** HwInt, HwRxInt
- 6) **Polling.** \_HwPktPoll



### 3.3.1 HwPktInit — Initialize Packet Driver Environment

<b>Syntax</b>	<code>uint HwPktInit();</code>
<b>Parameters</b>	None
<b>Return Value</b>	The number of Ethernet devices initialized. 0 indicates an error. All other positive values are considered success.
<b>Description</b>	This function is called to initialize the mini-driver environment and enumerate the number of devices in the system. A device instance may be opened for each device represented in the return count. If the function returns zero, no devices are supported.

### 3.3.2 HwPktOpen — Open Ethernet Device Instance

<b>Syntax</b>	<code>uint HwPktOpen (PDINFO *pi);</code>
<b>Parameters</b>	pi - Pointer to Ethernet device instance structure.
<b>Return Value</b>	Returns 0 on success and a positive value to indicate an error.
<b>Description</b>	<p>This function is called to open a packet device instance. When HwPktOpen is called, the PDINFO structure is assumed to be valid. This function sets up the EMAC configuration and invokes the CSL layer's EMAC_open function to configure the EMAC peripheral. As part of the configuration passed to EMAC_open, the driver sets up the required callback functions that the CSL layer in turn invokes to allocate/free packet buffers, update statistics or status, and to hand over received packets.</p> <p>This function is also responsible for setting up the interrupts and any other PHY related configuration to ready it for Tx/Rx operations.</p>

### 3.3.3 HwPktClose — Close Ethernet Device and Disable Interrupts

<b>Syntax</b>	<code>void HwPktClose (PDINFO *pi);</code>
<b>Parameters</b>	pi - Pointer to Ethernet device instance structure.
<b>Return Value</b>	None.
<b>Description</b>	This function is called to close a packet device instance. When called, this function invokes the CSL layer EMAC_close function to disable EMAC Tx/Rx operations and free up any enqueued packets. This function also disables the EMAC interrupts.

### 3.3.4 HwPktSetRx — Configure the Ethernet Receive Filter Settings

<b>Syntax</b>	<code>void HwPktSetRx( PDINFO *pi );</code>
<b>Parameters</b>	pi - Pointer to Ethernet device instance structure.
<b>Return Value</b>	None
<b>Description</b>	This function is called when the values contained in the PDINFO instance structure for the Rx filter or multicast list are altered. The mini-driver calculates hash values based on the new settings if multicast lists are maintained through hash tables on this platform, and updates the EMAC settings by calling the CSL layer's EMAC_setReceiveFilter API.

### 3.3.5 HwPktIoctl — Execute Driver-Specific IOCTL Commands

<b>Syntax</b>	<code>uint HwPktIoctl(PDINFO *pi, uint cmd, void *arg);</code>
<b>Parameters</b>	pi - Pointer to Ethernet packet device instance structure.  cmd - Device-specific command.  arg - Pointer to command specific argument.
<b>Return Value</b>	Returns 1 on success and 0 on error.
<b>Description</b>	This function is called to execute a driver-specific IOCTL command. Not all Ethernet drivers support this API.

### 3.3.6 HwPktTxNext — Transmit Next Buffer in the Transmit Queue

<b>Syntax</b>	<code>void HwPktTxNext( PDINFO *pi );</code>
<b>Parameters</b>	pi - Pointer to Ethernet packet device instance structure.
<b>Return Value</b>	None
<b>Description</b>	This function is called to indicate that a packet buffer has been queued in the transmit pending queue contained in the device instance structure and the NIMU layer believes the transmitter to be free. This function dequeues any pending packets in the transmit queue of this device, allocates a EMAC_Pkt structure (a data structure understood by the CSL layer) and fills in the packet details and invokes the CSL layer function EMAC_sendPacket to finally transmit the packet.

### 3.3.7 `_HwPktPoll` — Mini-Driver Polling Function

<b>Syntax</b>	<pre>void _HwPktPoll( PDINFO *pi, uint fTimerTick );</pre>
<b>Parameters</b>	<p><code>pi</code> - Pointer to Ethernet packet device instance structure.</p> <p><code>fTimerTick</code> - Flag indicating whether this function has been called because the 100 ms timer expired or if it was called by some other function randomly.</p>
<b>Return Value</b>	None
<b>Description</b>	<p>This function is called by the NIMU layer at least every 100 ms, but calls can come faster when there is network activity. The mini-driver is not required to perform any operation in this function, but it can be used to check for device lockup conditions. When the call is made due to the 100 ms time tick, the <code>fTimerTick</code> calling parameter is set.</p> <p>Note that this function is not called in kernel mode (hence, the underscore in the name). This is the only mini-driver function called from outside kernel mode (to support polling drivers).</p>

## 3.4 Configuration Variables

The following configuration variables are defined by the Ethernet mini-driver layer to control various features:

- ❑ **EXTERNAL\_MEMORY.** Enable this flag to compile the code to support the cache cleaning and synchronization required when the packet buffer memory is allocated from external memory.
- ❑ **EXTMEM.** Define this bit mask to indicate the external memory address location for this platform.
- ❑ **PKT\_MAX.** Use this constant to control the number of "EMAC\_Pkt" type packet buffers that are allocated and initialized on Receive and Transmit paths respectively at this layer to optimize the data paths. During EMAC start up, in the `HwPktOpen()` function, buffers of type "EMAC\_Pkt" structure are allocated and enqueued to a free queue/receive queue. Packets from this "RxQueue" are used to replenish the CSL layer with buffers for its BDs. Similarly, for the Tx path a queue of such EMAC\_Pkt initialized structures are held. A packet buffer from the "TxQueue" is dequeued and used in filling up the NDK packet buffer details before being handed over to the CSL layer. This constant controls the number of such replenishing buffers at this layer.

This constant can be fine-tuned during performance tuning to suit the application's needs. For example, increasing this constant helps in cases where the NDK stack or application is transmitting packets at a faster rate than the EMAC hardware. In this case, the packets are buffered up here at the mini-driver and get transmitted at the next suitable opportunity. But, it's important to note that this constant needs to be tuned according to the memory available in the system. The smallest number this can be set to is 8.

- ❑ **PKT\_PREPAD.** The number of bytes to reserve before the Ethernet header for any additional headers like PPP. This is typically defined to be 8 to include the PPP header.
- ❑ **RAM\_MCAST.** Define this configuration variable as 1 if the EMAC on this device supports RAM-based multicast lists. That is, if the EMAC is capable of storing multicast addresses in RAM and has defined appropriate registers to store them.
- ❑ **HASH\_MCAST.** Enable this or define this as 1 if the EMAC on this device is capable of maintaining the multicast address list using hash tables.
- ❑ **PKT\_MAX\_MCAST.** This constant defines the maximum number of multicast addresses that can be configured and supported by the EMAC peripheral on this device. This is typically set to 31.

# Generic EMAC/MDIO CSL Layer

---

---

---

This chapter describes the EMAC/MDIO CSL layer interface.

<b>Topic</b>	<b>Page</b>
<b>4.1 Overview</b> .....	<b>4-2</b>
<b>4.2 CSL Data Structures</b> .....	<b>4-2</b>
<b>4.3 EMAC APIs</b> .....	<b>4-2</b>
<b>4.4 Callback Functions</b> .....	<b>4-3</b>

## 4.1 Overview

The EMAC/MDIO CSL layer defines data structures and APIs that enable the driver to configure the EMAC hardware and send and receive packets.

This CSL layer is fairly generic and can be ported easily across different platforms so long as the EMAC hardware specification don't vary a lot. For example, the CSL for an EMAC with switch capabilities would be very different from the CSL for an EMAC with support for a single PHY. This layer abstracts out all the EMAC/MDIO register layer configuration details from the higher layers and makes them easier to write and understand.

## 4.2 CSL Data Structures

The CSL layer exports various data structures to enable configuration of EMAC, MDIO, and other Ethernet associated modules. Discussing all the data structures is beyond the scope of this document. The definitions can be viewed from the code or by obtaining a doxygen output of the code.

## 4.3 EMAC APIs

The following APIs are exported by the CSL EMAC layer:

- ❑ EMAC\_enumerate
- ❑ EMAC\_open
- ❑ EMAC\_close
- ❑ EMAC\_setReceiveFilter
- ❑ EMAC\_getReceiveFilter
- ❑ EMAC\_setMulticast
- ❑ EMAC\_getStatus
- ❑ EMAC\_getStatistics
- ❑ EMAC\_sendPacket
- ❑ EMAC\_RxServiceCheck
- ❑ EMAC\_TxServiceCheck
- ❑ EMAC\_TimerTick

As described in Section 1.5, the APIs exposed by this layer can be conveniently grouped according to their functionality into the following categories:

- 1) **Initialization.** EMAC\_open
- 2) **Shutdown.** EMAC\_close
- 3) **Configuration.** EMAC\_setReceiveFilter, EMAC\_getReceiveFilter, EMAC\_setMulticast, EMAC\_getStatus, EMAC\_getStatistics, EMAC\_enumerate
- 4) **Transmit.** EMAC\_sendPacket, EMAC\_TxServiceCheck (Tx ISR)
- 5) **Receive.** EMAC\_RxServiceCheck (Rx ISR)
- 6) **Polling.** EMAC\_TimerTick

The Ethernet mini-driver layer can invoke CSL APIs to perform any configuration or interrupt related processing only after opening and setting up the EMAC peripheral successfully using the "EMAC\_open" API. All the error codes, macros, and constants used are defined in the header files included with the source code and can be found in the "inc" directory.

## 4.4 Callback Functions

The CSL layer doesn't perform any OS specific operations such as memory allocation, free, initialization, copy etc. Instead, this layer defines the required callback functions in the "EMAC\_Config" data structure and mandates that the driver implement these functions and register them with the driver during the "EMAC\_open" call. The callback functions that need to be implemented by the driver and their description are describe in the subsections that follow.

See the Ethernet driver code for sample implementations of these functions.

### 4.4.1 pfcBGetPacket

This function is called by the CSL layer when it needs an empty packet buffer to replenish a receive EMAC Buffer Descriptor (BD) in the EMAC RAM. This function needs to implement logic to allocate an EMAC packet (of type "EMAC\_Pkt") and to initialize the buffers and offsets appropriately for use by the CSL layer. This function is typically called during EMAC initialization to initialize the Receive BDs or can be called during a receive interrupt servicing to re-fill any empty BDs.

#### **4.4.2 pfcFreePacket**

This function is called by the CSL layer to free the memory allocated for an EMAC packet (of type "EMAC\_Pkt") and any buffers held within it. This function is typically called during EMAC close, when an error occurs, or during a Transmit complete interrupt handling for cleaning up the associated buffers.

#### **4.4.3 pfcRxPacket**

This function is the driver-registered receive handler for all Ethernet packets received and validated by the EMAC and handed over to the CSL layer when a receive interrupt occurs.

This function is required to save the packet buffer received to hand it over to the stack for further processing. At that point, it is the responsibility of the driver/stack to free the packet buffer. This function is also required to return a new EMAC packet buffer in return to replenish the BD just serviced.

#### **4.4.4 pfcStatus**

This function is called by the CSL to notify the driver of a status change or the occurrence of an error during EMAC processing (HOSTPEND interrupt).

#### **4.4.5 pfcStatistics**

This function is called by the CSL to update the driver with the latest snapshot of statistics (STATPEND interrupt).



## A

- acronyms 1-2
- add\_header function 2-2
- alignment 1-9
- APIs 1-2
  - CSL layer 1-11, 4-2
  - mini-driver 1-11, 3-4
  - NIMU layer 1-11, 2-2
  - overview 1-10
- architecture 1-3

## B

- BD 1-2
- bMacAddr field 3-3
- bMCast field 3-3
- Buffer Descriptor 1-2

## C

- callback functions 1-5, 4-3
- Chip Support Library 1-2
- configuration
  - APIs 1-10
  - device 3-2
  - EMAC 4-2
  - EMAC subsystem 3-2
  - MDIO 4-2
  - variables 3-7
- CSL layer 1-2, 4-1
  - APIs 1-11
  - architecture overview 1-3

## D

- data alignment 1-9
- data flow 1-6
- data structures 4-2
- doxygen output 4-2
- DSP 1-2

## E

- EMAC 1-2
- EMAC APIs 1-5
- EMAC/MDIO CSL layer 4-1
  - architecture overview 1-3
  - description 1-5
- EMAC\_close function 4-2
- EMAC\_enumerate function 4-2
- EMAC\_getReceiveFilter function 4-2
- EMAC\_getStatistics function 4-2
- EMAC\_getStatus function 4-2
- EMAC\_open function 4-2
- EMAC\_RxServiceCheck function 4-2
- EMAC\_sendPacket function 4-2
- EMAC\_setMulticast function 4-2
- EMAC\_setReceiveFilter function 4-2
- EMAC\_TimerTick function 4-2
- EMAC\_TxServiceCheck function 4-2
- error handling 4-4
- Ethernet Medium Access Protocol 1-2
- Ethernet mini-driver 3-1
  - architecture overview 1-3
  - description 1-4
- EXTERNAL\_MEMORY constant 3-7
- EXTMEM constant 3-7

## F

- Filter field 3-3
- flow chart 1-6

## H

- HAL layer 1-8
- HASH\_MCAST constant 3-8
- header size 1-9
- hEther field 3-3
- hEvent field 3-3
- HOSTPEND interrupt 4-4
- HwPktClose function 3-4, 3-5
- HwPktInit function 3-4, 3-5
- HwPktIoctl function 3-6

HwPktOpen function 3-4, 3-5  
 \_HwPktPoll function 3-4, 3-7  
 HwPktSetRx function 3-4, 3-6  
 HwPktShutdown function 3-4  
 HwPktTxNext function 3-4, 3-6

## I

Init callback function 2-2  
 initialization APIs 1-10  
 Interrupt Manager 1-9  
 Interrupt\_add function 1-9  
 Interrupt\_delete function 1-9  
 Interrupt\_disable function 1-9  
 Interrupt\_enable function 1-9  
 ioctl function 2-2  
 ISRs 1-9

## J

Jumbo Packet Buffer Manager 1-8

## L

layers 1-3  
 LL 1-2  
 logic  
   buffer descriptors 1-5  
   ISRs 1-5  
 Low Level Packet Driver 1-2

## M

Mac address 3-3  
 Management Data Input/Output Interface 1-2  
 MCastCnt field 3-3  
 MDIO 1-2  
 MDIO APIs 1-5  
 MDIO layer 4-1  
 memory manager 1-4  
 mini-driver 3-1  
   APIs 1-11  
   description 1-4  
 multicast addresses 3-3, 3-8

## N

NDK 1-2  
 NDK core stack 1-4  
 NDK Interrupt Manager 1-9  
 NDK Support Package 1-2

NETCTRL module 1-8  
 NETIF\_DEVICE structure 2-2  
 Network Control Module 1-8  
 Network Developer's Kit 1-2  
 Network Interface Management Unit 1-2  
 NIMU 1-2  
 NIMU layer 2-1  
   APIs 1-11, 2-2  
   architecture overview 1-3  
   description 1-4  
 nimu\_eth.c file 2-2  
 NIMUDeviceTable structure 2-2  
 NIMURegister function 2-2  
 NSP 1-2

## O

Operating Systems Abstraction Layer 1-2  
 OS AL 1-2

## P

Packet Buffer object 1-8  
 packet buffers 3-7  
   allocating 4-3  
   freeing 4-4  
   receiving 4-4  
 packet flow 1-6  
 PBM object 1-8  
 PBMQ\_rx field 3-4  
 PBMQ\_tx field 3-4  
 PDINFO structure 3-2  
 Peer-to-Peer Protocol 1-10  
 performance 3-8  
 pfcFreePacket function 4-4  
 pfcGetPacket function 4-3  
 pfcRxPacket function 4-4  
 pfcStatistics function 4-4  
 pfcStatus function 4-4  
 PhysIdx field 3-2  
 PKT\_MAX constant 3-7  
 PKT\_MAX\_MCAST constant 3-8  
 PKT\_PREPAD constant 1-10, 3-8  
 pkt\_service function 2-2  
 poll function 2-2  
 polling APIs 1-10  
 PPP 1-10

## R

RAM\_MCAST constant 3-8  
 receive APIs 1-10  
 receive path 1-7

Rx 1-2  
  packet flow 1-7

## S

sample code 2-2  
semaphore object 3-3  
send function 2-2  
Serial Gigabit Media Independent Interface 1-2  
SGMII 1-2  
SGMII APIs 1-5  
shutdown APIs 1-10  
Stack Event object 1-8  
start function 2-2

STATPEND interrupt 4-4  
STKEVENT object 1-8  
STKEVENT\_signal function 3-3  
stop function 2-2

## T

transmission path 1-6  
transmit APIs 1-10  
tuning 3-8  
Tx 1-2  
  packet flow 1-6  
TxFree field 3-3

