

# Intel<sup>®</sup> IXP400 Software

## Programmer's Guide

---

| *July 2006*



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY RELATING TO SALE AND/OR USE OF INTEL PRODUCTS, INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT, OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

MPEG is an international standard for video compression/decompression promoted by ISO. Implementations of MPEG CODECs, or MPEG enabled platforms may require licenses from various entities, including Intel Corporation.

This document and the software described in it are furnished under license and may only be used or copied in accordance with the terms of the license. The information in this document is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by Intel Corporation. Intel Corporation assumes no responsibility or liability for any errors or inaccuracies that may appear in this document or any software that may be provided in association with this document. Except as permitted by such license, no part of this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the express written consent of Intel Corporation.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web site at <http://www.intel.com>.

Intel, the Intel logo, and Intel XScale are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

\*Other names and brands may be claimed as the property of others.

Copyright © Intel Corporation 2006. All Rights Reserved.



# Contents

<b>1.0</b>	<b>Introduction</b>	21
1.1	Versions Supported by this Document	21
1.2	Hardware Supported by this Release	21
1.3	Intended Audience	21
1.4	How to Use this Document	21
1.5	About the Processors	22
1.6	Related Documents	24
1.7	Acronyms	25
<b>2.0</b>	<b>Software Architecture Overview</b>	31
2.1	High-Level Overview	31
2.2	Deliverable Model	32
2.3	Operating System Support	33
2.4	Development Tools	33
2.5	Access Library Source Code Documentation	33
2.6	Release Directory Structure	33
2.7	Threading and Locking Policy	35
2.8	Polled and Interrupt Operation	36
2.9	Statistics and MIBs	36
2.10	Global Dependency Chart	36
<b>3.0</b>	<b>Buffer Management</b>	37
3.1	What's New	37
3.2	Overview	37
3.3	IX_OSAL_MBUF Structure	40
3.3.1	IX_OSAL_MBUF Structure and Macros	40
3.4	Mapping of ix_osdep_buf to Shared Structure	45
3.5	ix_osdep_buf Structure	46
3.6	Mapping to OS Native Buffer Types	48
3.6.1	VxWorks* M_BLK Buffer	48
3.6.2	Linux* skbuff Buffer	49
3.7	Intel® IXP400 Software Caching Strategy	50
3.7.1	Tx Path	51
3.7.2	Rx Path	52
3.7.3	Caching Strategy Summary	52
<b>4.0</b>	<b>Access-Layer Components:</b>	
	<b>ATM Driver Access (IxAtmdAcc) API</b>	53
4.1	What's New	53
4.2	Overview	53
4.3	IxAtmdAcc Component Features	53
4.4	ATM Background	55
4.4.1	Quality of Service	56
4.4.2	Adaptation Layers	56
4.5	Configuration Services	59
4.5.1	UTOPIA Port-Configuration Service	60
4.5.2	ATM Traffic-Shaping Services	60
4.5.3	VC-Configuration Services	61
4.5.4	Uninitialize ATM Driver	61
4.6	Transmission Services	62
4.6.1	Scheduled Transmission	63
4.6.1.1	Schedule Table Description	64
4.6.2	Transmission Triggers (Tx-Low Notification)	65



- 4.6.2.1 Transmit-Done Processing ..... 66
- 4.6.2.2 Transmit Disconnect ..... 68
- 4.6.3 Receive Services ..... 69
  - 4.6.3.1 Receive Triggers (Rx-Free-Low Notification) ..... 71
  - 4.6.3.2 Receive Processing ..... 71
  - 4.6.3.3 Receive Disconnect ..... 73
- 4.6.4 Buffer Management ..... 74
  - 4.6.4.1 Buffer Allocation ..... 74
  - 4.6.4.2 Buffer Contents ..... 74
  - 4.6.4.3 Buffer Size Constraints ..... 75
  - 4.6.4.4 Buffer-Chaining Constraints ..... 75
  - 4.6.4.5 Starvation and Throttling ..... 75
- 4.6.5 Error Handling ..... 76
  - 4.6.5.1 API-Usage Errors ..... 76
  - 4.6.5.2 Real-Time Errors ..... 76
- 5.0 Access-Layer Components:**
  - ATM Manager (IxAtmm) API ..... 77**
    - 5.1 What's New ..... 77
    - 5.2 IxAtmm Overview ..... 77
    - 5.3 IxAtmm Component Features ..... 77
    - 5.4 ixAtmmAcc API ..... 78
    - 5.5 UTOPIA Level-2 Port Initialization ..... 79
    - 5.6 ATM-Port Management Service Model ..... 79
    - 5.7 Tx/Rx Control Configuration ..... 82
    - 5.8 Dependencies ..... 84
    - 5.9 Error Handling ..... 84
    - 5.10 Management Interfaces ..... 84
    - 5.11 Memory Requirements ..... 84
    - 5.12 Performance ..... 85
- 6.0 Access-Layer Components:**
  - ATM Transmit Scheduler (IxAtmSch) API ..... 87**
    - 6.1 What's New ..... 87
    - 6.2 Overview ..... 87
    - 6.3 IxAtmSch Component Features ..... 88
    - 6.4 IxAtmSch API ..... 88
    - 6.5 Traffic Types ..... 89
    - 6.6 Connection Admission Control (CAC) Function ..... 89
    - 6.7 VC Oversubscription and Priority Feature ..... 90
    - 6.8 Scheduling and Traffic Shaping ..... 91
      - 6.8.1 Schedule Table ..... 91
        - 6.8.1.1 Minimum Cells Value (minCellsToSchedule) ..... 92
        - 6.8.1.2 Maximum Cells Value (maxCells) ..... 92
      - 6.8.2 Schedule Service Model ..... 92
      - 6.8.3 Timing and Idle Cells ..... 93
    - 6.9 Dependencies ..... 93
    - 6.10 Error Handling ..... 93
    - 6.11 Memory Requirements ..... 94
      - 6.11.1 Code Size ..... 94
      - 6.11.2 Data Memory ..... 94
    - 6.12 Performance ..... 94
      - 6.12.1 Latency ..... 94
- 7.0 Access-Layer Components:**
  - Security (IxCryptoAcc) API ..... 95**
    - 7.1 What's New ..... 95



7.2	Overview .....	95
7.3	Internet Security — Background Information.....	97
7.3.1	Encryption and Message Authentication.....	97
7.3.2	Key Management .....	98
7.4	Security Architecture .....	98
7.4.1	IxCryptoAcc Interfaces.....	98
7.4.1.1	Intel XScale® Processor Software.....	99
7.4.1.2	Offloading to Hardware.....	99
7.4.1.3	API Usage.....	99
7.4.2	Basic API Flow .....	99
7.4.3	Context Registration and the Cryptographic Context Database.....	101
7.4.4	Buffer and Queue Management .....	104
7.4.5	Using Key Management Service.....	104
7.4.5.1	Use of Random Numbers .....	104
7.4.5.2	Use of Hashing .....	105
7.4.5.3	Use of Large Number Arithmetic .....	105
7.4.6	Config and Utility Functionality .....	105
7.4.7	Memory Requirements .....	107
7.4.8	Dependencies .....	107
7.4.9	Error Handling .....	109
7.4.10	Endianness.....	109
7.4.11	Import and Export of Cryptographic Technology .....	109
7.5	IPSec Services .....	109
7.5.1	IPSec Background and Implementation .....	109
7.5.2	IPSec Packet Formats .....	110
7.5.2.1	Reference ESP Dataflow.....	111
7.5.2.2	Reference AH Dataflow .....	112
7.5.3	Hardware Support for IPSec Services .....	113
7.5.4	IPSec API Call Flow.....	113
7.5.5	Special API Use Cases.....	115
7.5.5.1	HMAC with Key Size Greater Than 64 Bytes.....	115
7.5.5.2	Performing CCM (AES CTR-Mode Encryption and AES CBC-MAC Authentication) for IPSec .....	115
7.5.6	IPSec Assumptions, Dependencies, and Limitations.....	118
7.6	WEP Services.....	118
7.6.1	WEP Background and Implementation.....	118
7.6.2	Hardware Support for WEP Services.....	119
7.6.3	WEP API Call Flow .....	120
7.7	SSL and TLS Protocol Usage Models.....	122
7.8	Supported Encryption and Authentication Algorithms.....	123
7.8.1	Encryption Algorithms .....	123
7.8.2	Cipher Modes.....	124
7.8.2.1	Electronic Code Book (ECB).....	124
7.8.2.2	Cipher Block Chaining (CBC) .....	124
7.8.2.3	Counter Mode (CTR).....	124
7.8.2.4	Counter-Mode Encryption with CBC-MAC Authentication (CCM) for CCMP in 802.11i .....	124
7.8.3	Authentication Algorithms.....	125
7.9	Support for Large Number Arithmetic.....	125
7.9.1	Large Number Arithmetic .....	125
<b>8.0</b>	<b>Access-Layer Components:</b>	
	<b>DMA Access Driver (IxDmaAcc) API .....</b>	<b>127</b>
8.1	What's New .....	127
8.2	Overview .....	127
8.3	Features .....	127



- 8.4 Assumptions ..... 127
- 8.5 Fast Data Transfer for Little Endian Systems..... 128
- 8.6 Builds ..... 128
- 8.7 Dependencies..... 128
- 8.8 DMA Access-Layer API ..... 129
  - 8.8.1 IxDmaAccDescriptorManager..... 131
- 8.9 Parameters Description..... 131
  - 8.9.1 Source Address ..... 131
  - 8.9.2 Destination Address ..... 131
  - 8.9.3 Transfer Mode ..... 132
  - 8.9.4 Transfer Width ..... 132
  - 8.9.5 Addressing Modes ..... 132
  - 8.9.6 Transfer Length ..... 132
  - 8.9.7 Supported Modes ..... 133
- 8.10 Data Flow ..... 135
- 8.11 Control Flow..... 135
  - 8.11.1 DMA Initialization..... 136
  - 8.11.2 DMA Configuration and Data Transfer ..... 137
- 8.12 Restrictions of the DMA Transfer ..... 139
- 8.13 Error Handling..... 140
- 9.0 Access-Layer Components:**
  - Ethernet Access (IxEthAcc) API ..... 141**
    - 9.1 What's New ..... 141
    - 9.2 New APIs..... 141
    - 9.3 IxEthAcc Overview..... 142
    - 9.4 Ethernet Access Layer and its Interface to Other Components: Architectural Overview ..... 142
      - 9.4.1 Role of the Ethernet NPE..... 143
      - 9.4.2 Queue Manager ..... 144
      - 9.4.3 Learning/Filtering ..... 144
      - 9.4.4 MAC/PHY Configuration ..... 145
    - 9.5 Ethernet Access Layers: Component Features..... 145
    - 9.6 Data Plane..... 146
      - 9.6.1 Port Initialization ..... 146
      - 9.6.2 Ethernet Frame Transmission ..... 147
        - 9.6.2.1 Transmission Flow ..... 147
        - 9.6.2.2 Transmit Buffer Management and Priority ..... 148
        - 9.6.2.3 Using Chained IX\_OSAL\_MBUFs for Transmission / Buffer Sizing.... 150
      - 9.6.3 Ethernet Frame Reception..... 150
        - 9.6.3.1 Receive Flow ..... 153
        - 9.6.3.2 Receive Queue Interrupt Disable/Enable..... 153
        - 9.6.3.3 Receive Buffer Management and Priority..... 153
        - 9.6.3.4 Additional Receive Path Information..... 157
      - 9.6.4 Data-Plane Endianness ..... 158
      - 9.6.5 Maximum Ethernet Frame Size ..... 158
    - 9.7 Control Path..... 158
      - 9.7.1 Ethernet MAC Control..... 159
        - 9.7.1.1 MAC Duplex Settings ..... 160
        - 9.7.1.2 MII I/O ..... 160
        - 9.7.1.3 Frame Check Sequence ..... 160
        - 9.7.1.4 Frame Padding..... 160
        - 9.7.1.5 MAC Filtering ..... 161
        - 9.7.1.6 802.3x Flow Control..... 161
        - 9.7.1.7 NPE Loopback ..... 162
        - 9.7.1.8 Emergency Security Port Shutdown ..... 162



9.7.1.9	Soft-error Handling .....	162
9.8	Initialization .....	162
9.9	Uninitialization .....	163
9.10	Shared Data Structures .....	163
9.11	Management Information .....	167
9.12	Ethernet and HSS Channelized services co-existence.....	169
9.12.1	Queue Manager Queues .....	169
9.12.2	Live lock prevention scheme in HSS and Ethernet co-existence.....	170
<b>10.0</b>	<b>Access-Layer Components:</b>	
	<b>Ethernet Database (IxEthDB) .....</b>	<b>171</b>
10.1	Overview .....	171
10.2	What's New .....	171
10.3	New APIs .....	171
10.4	IxEthDB Functional Behavior .....	172
10.4.1	Feature set.....	172
10.4.2	Additional Database Features .....	173
10.4.2.1	User-Defined Field .....	173
10.4.2.2	Database Clear .....	174
10.4.3	MAC Address Learning and Filtering .....	174
10.4.3.1	Learning and Filtering .....	174
10.4.3.2	Other MAC Learning/Filtering Usage Models.....	176
10.4.3.3	Learning/Filtering General Characteristics.....	176
10.4.4	Frame Size Filtering.....	179
10.4.4.1	Filtering Example Based Upon Maximum Frame Size.....	179
10.4.5	Source MAC Address Firewall .....	180
10.4.6	IPv4 and IPv6 Payload Detection .....	181
10.4.7	802.1Q VLAN .....	182
10.4.7.1	Background – VLAN Data in Ethernet Frames.....	182
10.4.7.2	Database Records Associated With VLAN IDs .....	183
10.4.7.3	Acceptable Frame Type Filtering .....	183
10.4.7.4	Ingress Tagging and Tag Removal.....	183
10.4.7.5	Port-Based VLAN Membership Filtering.....	184
10.4.7.6	Port and VLAN-Based Egress Tagging and Tag Removal .....	185
10.4.7.7	Port ID Extraction .....	187
10.4.8	802.1Q User Priority / QoS Support .....	187
10.4.8.1	Priority Aware Transmission .....	187
10.4.8.2	Receive Priority Queuing.....	188
10.4.8.3	Priority to Traffic Class Mapping .....	189
10.4.9	802.3 / 802.11 Frame Conversion .....	190
10.4.9.1	Background — 802.3 and 802.11 Frame Formats .....	190
10.4.9.2	Destination Port ID Indication for a Forwarding Frame .....	192
10.4.9.3	VLAN Support for 802.11 Frames .....	193
10.4.9.4	How the 802.3 / 802.11 Frame Conversion Feature Works.....	193
10.4.9.5	Pad Field Addition/Removal for 802.11 Frames .....	195
10.4.9.6	802.3 <-> 802.11 API Details .....	195
10.4.10	Spanning Tree Protocol Port Settings .....	197
10.4.11	Soft-error Handling.....	198
10.5	IxEthDB .....	198
10.5.1	Enabling/Disabling EthDB .....	198
10.5.2	Initialization .....	198
10.5.3	Dependencies .....	199
10.5.4	Dependencies on IxEthAcc Configuration .....	199
10.5.4.1	Promiscuous-Mode Requirement .....	200
10.5.4.2	FCS Appending.....	200



- 11.0 Access-Layer Components:**
  - Ethernet PHY (IxEthMii) API** ..... 201
  - 11.1 What's New ..... 201
  - 11.2 Overview ..... 201
  - 11.3 Features ..... 201
  - 11.4 Supported PHYs ..... 201
  - 11.5 Dependencies ..... 202
- 12.0 Access-Layer Components:**
  - Feature Control (IxFeatureCtrl) API** ..... 203
  - 12.1 What's New ..... 203
  - 12.2 Overview ..... 203
  - 12.3 Hardware Feature Control ..... 203
    - 12.3.1 Using the Product ID-Related Functions ..... 205
    - 12.3.2 Using the Feature Control Register Functions ..... 205
  - 12.4 Software Configuration ..... 205
  - 12.5 Dependencies ..... 206
- 13.0 Access-Layer Components:**
  - HSS-Access (IxHssAcc) API** ..... 207
  - 13.1 What's New ..... 207
  - 13.2 Overview ..... 207
    - 13.2.1 Coexistence of HSS Services in NPE-A ..... 208
    - 13.2.2 Coexistence of HSS services and ATM services in NPE-A ..... 208
    - 13.2.3 Coexistence of HSS Channelized services and Ethernet services in NPE-A ..... 210
  - 13.3 IxHssAcc API Overview ..... 210
    - 13.3.1 IxHssAcc Interfaces ..... 210
    - 13.3.2 Basic API Flow ..... 211
    - 13.3.3 HSS and HDLC Theory and Coprocessor Operation ..... 212
    - 13.3.4 Packetized Data Stream ..... 215
      - 13.3.4.1 56-Kbps, Packetized Raw Mode ..... 215
    - 13.3.5 High-Level API Call Flow ..... 215
    - 13.3.6 Dependencies ..... 216
    - 13.3.7 Key Assumptions ..... 217
    - 13.3.8 Error Handling ..... 218
  - 13.4 HSS Port Initialization Details ..... 218
  - 13.5 HSS Channelized Operation ..... 219
    - 13.5.1 Channelized Connect and Enable ..... 220
      - 13.5.1.1 Normal Mode ..... 220
      - 13.5.1.2 Bypass Mode ..... 222
    - 13.5.2 Channelized Tx/Rx Methods ..... 225
      - 13.5.2.1 Interrupt Mode via CallBacks ..... 225
      - 13.5.2.2 Polling Mode ..... 226
    - 13.5.3 Channelized Disconnect ..... 228
  - 13.6 HSS Packetized Operation ..... 228
    - 13.6.1 Packetized Connect and Enable ..... 228
    - 13.6.2 Packetized Tx ..... 230
    - 13.6.3 Packetized Rx ..... 232
    - 13.6.4 Packetized Disconnect ..... 235
  - 13.7 Buffer Allocation Data-Flow Overview ..... 235
    - 13.7.1 Data Flow in Packetized Service ..... 235
    - 13.7.2 Data Flow in Channelized Service ..... 237
- 14.0 Access-Layer Components:**
  - NPE-Downloader (IxNpeDI) API** ..... 241
  - 14.1 What's New ..... 241



14.2	Overview .....	241
14.3	Microcode Images .....	241
14.4	Standard Usage Example.....	242
14.5	Custom Usage Example.....	248
14.6	IxNpeDI Uninitialization .....	248
14.7	Deprecated APIs .....	249
<b>15.0</b>	<b>Access-Layer Components:</b>	
	<b>NPE Message Handler (IxNpeMh) API .....</b>	<b>251</b>
15.1	What's New .....	251
15.2	Overview .....	251
15.3	Initializing the IxNpeMh .....	252
	15.3.1 Interrupt-Driven Operation .....	252
	15.3.2 Polled Operation.....	252
15.4	Uninitializing IxNpeMh .....	252
15.5	NPE Parity Error Handling .....	253
15.6	Sending Messages from an Intel XScale® Processor Software Client to an NPE.....	253
	15.6.1 Sending an NPE Message.....	253
	15.6.2 Sending an NPE Message with Response.....	254
15.7	Receiving Unsolicited Messages from an NPE to Client Software .....	255
15.8	Dependencies .....	257
15.9	Error Handling .....	257
<b>16.0</b>	<b>Access-Layer Components:</b>	
	<b>Parity Error Notifier (IxParityENAcc) API .....</b>	<b>259</b>
16.1	What's New .....	259
16.2	Introduction .....	259
	16.2.1 Background.....	259
	16.2.2 Parity and ECC Capabilities in the Intel® IXP45X and Intel® IXP46X Product Line of Network Processors .....	260
	16.2.2.1 Network Processing Engines .....	260
	16.2.2.2 Switching Coprocessor in NPE B (SWCP) .....	261
	16.2.2.3 AHB Queue Manager (AQM) .....	261
	16.2.2.4 DDR SDRAM Memory Controller Unit (MCU).....	261
	16.2.2.5 Expansion Bus Controller .....	261
	16.2.2.6 PCI Controller.....	261
	16.2.2.7 Secondary Effects of Parity Interrupts.....	262
	16.2.3 Interrupt Prioritization.....	262
16.3	IxParityENAcc API Details .....	263
	16.3.1 Features .....	263
	16.3.2 Dependencies .....	263
16.4	IxParityENAcc API Usage Scenarios.....	264
	16.4.1 Summary Parity Error Notification Scenario.....	265
	16.4.2 Summary Parity Error Recovery Scenario.....	267
	16.4.3 Summary Parity Error Prevention Scenario.....	267
	16.4.4 Parity Error Notification Detailed Scenarios .....	267
<b>17.0</b>	<b>Access-Layer Components:</b>	
	<b>Error Handler (ixErrHdlAcc) API .....</b>	<b>273</b>
17.1	What's New .....	273
17.2	ixErrhdlAcc Overview .....	273
17.3	Architectural Overview .....	274
17.4	Functional Details for NPE Error and AQM Parity Error.....	274
17.5	Dependencies .....	275
17.6	Software Interfaces .....	276
17.7	Intel® IXP400 Software Enabling of Soft Error Detection and Handling on the Intel® IXP4XX Development Platforms.....	277



- 17.7.1 Soft Error Detection, Handling Configuration and Initialization Pseudo Code.. 277
- 17.7.2 Initializing the Intel® IXP400 Software Soft Error Modules..... 279
- 17.7.3 Detection and handling of Soft Error..... 280
- 17.7.4 Unloading the ixErrHdlAcc..... 284
- 17.7.5 Ethernet Device Driver Modifications for Soft Error Detection and Handling... 285
  - 17.7.5.1 Resolving NPE IRQ Sharing Conflicts ..... 285
  - 17.7.5.2 QM Queue Dispatcher Binding Setup ..... 285
- 17.8 Functional Soft Error Recovery and Handling of AQM SRAM Parity Error and NPE on an Ethernet Application..... 286
  - 17.8.1 NPE Soft Error Recovery ..... 286
    - 17.8.1.1 Background..... 286
    - 17.8.1.2 Intel® IXP400 Software Access Layer Components Error Handling.. 286
    - 17.8.1.3 Scalability ..... 287
    - 17.8.1.4 Multiple Event Errors ..... 287
    - 17.8.1.5 Memory Leakage Prevention Methods..... 287
- 18.0 Access-Layer Components:**
  - Queue Manager (IxQMgr) API ..... 289**
  - 18.1 What's New ..... 289
  - 18.2 Overview ..... 289
  - 18.3 Features and Hardware Interface ..... 290
  - 18.4 IxQMgr Initialization and Uninitialization ..... 291
  - 18.5 Queue Configuration ..... 291
  - 18.6 Queue Identifiers ..... 292
  - 18.7 Configuration Values ..... 292
  - 18.8 Dispatcher ..... 292
  - 18.9 Dispatcher Modes ..... 293
  - 18.10 Livelock Prevention ..... 296
  - 18.11 Threading ..... 298
  - 18.12 Dependencies..... 298
  - 18.13 NPE Parity Error Handling ..... 298
- 19.0 Access-Layer Components:**
  - Synchronous Serial Port (IxSspAcc) API ..... 301**
  - 19.1 What's New ..... 301
  - 19.2 Introduction..... 301
  - 19.3 IxSspAcc API Details ..... 301
    - 19.3.1 Features..... 301
    - 19.3.2 Dependencies..... 302
  - 19.4 IxSspAcc API Usage Models..... 302
    - 19.4.1 Initialization and General Data Model..... 302
    - 19.4.2 Interrupt Mode ..... 303
    - 19.4.3 Polling Mode ..... 305
- 20.0 Access-Layer Components:**
  - Time Sync (IxTimeSyncAcc) API ..... 307**
  - 20.1 What's New ..... 307
  - 20.2 Introduction..... 307
    - 20.2.1 IEEE 1588 PTP Protocol Overview..... 308
    - 20.2.2 IEEE 1588 Hardware Assist Block ..... 309
    - 20.2.3 IxTimeSyncAcc..... 312
    - 20.2.4 IEEE 1588 PTP Client Application..... 312
  - 20.3 IxTimeSyncAcc API Details..... 312
    - 20.3.1 Features..... 312
    - 20.3.2 Dependencies..... 313
    - 20.3.3 Error Handling..... 313
  - 20.4 IxTimeSyncAcc API Usage Scenarios ..... 314



20.4.1	Polling for Transmit and Receive Timestamps .....	314
20.4.2	Interrupt Mode Operations.....	314
20.4.3	Polled Mode Operations .....	315
<b>21.0</b>	<b>Access-Layer Components:</b>	
	<b>UART-Access (IxUARTAcc) API .....</b>	<b>317</b>
21.1	What's New .....	317
21.2	Overview .....	317
21.3	Interface Description .....	317
21.4	UART / OS Dependencies .....	318
21.4.1	FIFO Versus Polled Mode .....	318
21.5	Dependencies .....	319
<b>22.0</b>	<b>Access-Layer Components:</b>	
	<b>USB Access (ixUSB) API .....</b>	<b>321</b>
22.1	What's New .....	321
22.2	Overview .....	321
22.3	USB Controller Background.....	321
22.3.1	Packet Formats .....	322
22.3.2	Transaction Formats .....	323
22.4	ixUSB API Interfaces.....	326
22.4.1	ixUSB Setup Requests.....	326
22.4.1.1	Configuration .....	328
22.4.1.2	Frame Synchronization .....	329
22.4.2	ixUSB Send and Receive Requests .....	329
22.4.3	ixUSB Endpoint Stall Feature .....	329
22.4.4	ixUSB Error Handling .....	331
22.5	USB Data Flow .....	331
22.6	USB Dependencies .....	331
<b>23.0</b>	<b>Codelets .....</b>	<b>333</b>
23.1	What's New .....	333
23.2	Overview .....	333
23.3	ATM Codelet (IxAtmCodelet) .....	333
23.4	Crypto Access Codelet (IxCryptoAccCodelet) .....	333
23.5	DMA Access Codelet (IxDmaAccCodelet).....	334
23.6	Ethernet AAL-5 Codelet (IxEthAal5App).....	334
23.7	Ethernet Access Codelet (IxEthAccCodelet) .....	334
23.8	HSS Access Codelet (IxHssAccCodelet).....	335
23.9	Parity Error Notifier Codelet (IxParityENAccCodelet) .....	335
23.10	Performance Profiling Codelet (IxPerfProfAccCodelet) .....	335
23.11	Time Sync Codelet (IxTimeSyncAccCodelet).....	335
23.12	USB RNDIS Codelet (IxUSBRNDIS) .....	336
<b>24.0</b>	<b>Operating System</b>	
	<b>Abstraction Layer (OSAL) .....</b>	<b>337</b>
24.1	What's New .....	337
24.2	New APIs .....	337
24.3	Overview .....	337
24.4	OS-Independent Core Module .....	339
24.5	OS-Dependent Module .....	339
24.6	Optional Modules.....	339
24.6.1	Buffer Translation Module .....	340
24.7	OSAL Library Structure .....	340
24.8	OSAL Modules and Related Interfaces .....	343
24.8.1	Core Module .....	343
24.8.2	Buffer Management Module .....	346



- 24.8.3 I/O Memory and Endianness Support Module ..... 346
- 24.9 Supporting a New OS ..... 349
- 24.10 Supporting New Platforms ..... 349
  - 24.10.1 Module Specific Requirements ..... 350
  - 24.10.2 General Purpose Requirements ..... 351
- 24.11 Testing Strategy ..... 352
- 25.0 ADSL Driver ..... 355**
  - 25.1 What's New ..... 355
  - 25.2 Device Support ..... 355
  - 25.3 ADSL Driver Overview ..... 355
    - 25.3.1 Controlling STMicroelectronics\* ADSL Modem Chipset Through CTRL-E ..... 356
  - 25.4 ADSL API ..... 356
  - 25.5 ADSL Line Open/Close Overview ..... 356
  - 25.6 Limitations and Constraints ..... 357
- 26.0 I<sup>2</sup>C Driver (IxI2cDrv) ..... 359**
  - 26.1 What's New ..... 359
  - 26.2 Introduction ..... 359
  - 26.3 I<sup>2</sup>C Driver API Details ..... 359
    - 26.3.1 Features ..... 359
    - 26.3.2 Dependencies ..... 360
    - 26.3.3 Error Handling ..... 361
      - 26.3.3.1 Arbitration Loss Error ..... 361
      - 26.3.3.2 Bus Error ..... 362
  - 26.4 I<sup>2</sup>C Driver API Usage Models ..... 362
    - 26.4.1 Initialization and General Data Model ..... 362
    - 26.4.2 Example Sequence Flows for Slave Mode ..... 364
    - 26.4.3 I<sup>2</sup>C Using GPIO Versus Dedicated I<sup>2</sup>C Hardware ..... 367
- 27.0 Endianness in Intel® IXP400 Software v2.3 ..... 369**
  - 27.1 What's New ..... 369
  - 27.2 Overview ..... 369
  - 27.3 The Basics of Endianness ..... 369
    - 27.3.1 The Nature of Endianness: Hardware or Software? ..... 370
    - 27.3.2 Endianness When Memory is Shared ..... 370
  - 27.4 Software Considerations and Implications ..... 370
    - 27.4.1 Coding Pitfalls — Little Endian/Big Endian ..... 371
      - 27.4.1.1 Casting a Pointer Between Types of Different Sizes ..... 371
      - 27.4.1.2 Network Stacks and Protocols ..... 372
      - 27.4.1.3 Shared Data Example: LE Re-Ordering Data for BE Network Traffic ..... 372
    - 27.4.2 Best Practices in Coding of Endian-Independence ..... 373
    - 27.4.3 Macro Examples: Endian Conversion ..... 373
      - 27.4.3.1 Macro Source Code ..... 373
  - 27.5 Endianness Features of the Intel® IXP4XX Product Line of Network Processors ..... 374
    - 27.5.1 Supporting Little Endian Mode ..... 376
    - 27.5.2 Reasons for Choosing a Particular LE Coherency Mode ..... 376
    - 27.5.3 Silicon Endianness Controls ..... 378
      - 27.5.3.1 Hardware Switches ..... 378
      - 27.5.3.2 Intel XScale® Processor Endianness Mode ..... 379
      - 27.5.3.3 Little Endian Data Coherence Enable/Disable ..... 380
      - 27.5.3.4 MMU P-Attribute Bit ..... 381
      - 27.5.3.5 PCI Bus Swap ..... 381
      - 27.5.3.6 Summary of Silicon Controls ..... 381
    - 27.5.4 Silicon Versions ..... 381
  - 27.6 Little endian Strategy in Intel® IXP400 Software and Associated BSPs ..... 382
    - 27.6.1 APB Peripherals ..... 383



27.6.2	IXP400 Software Little Endian Strategy on APB .....	383
27.6.3	AHB Memory-Mapped Registers .....	384
27.6.4	Intel® IXP400 Software Core Components .....	384
27.6.4.1	Queue Manager — IxQMgr .....	384
27.6.4.2	NPE Downloader — IxNpeDI .....	385
27.6.4.3	NPE Message Handler — IxNpeMh .....	385
27.6.4.4	Ethernet Access Component — IxEthAcc .....	385
27.6.4.5	ATM and HSS .....	390
27.6.5	PCI .....	390
27.6.6	Intel® IXP400 SoftwareOS Abstraction .....	390
27.6.7	VxWorks* Considerations .....	391
27.6.8	Software Versions .....	393

## Figures

1	Intel® IXP46X Product Line Network Processor Block Diagram .....	23
2	Intel® IXP400 Software v2.3 Architecture Block Diagram .....	32
3	Global Dependencies .....	36
4	Intel® IXP400 Software v2.3 Buffer Flow .....	38
5	IX_OSAL_MBUF User Interface .....	39
6	IX_OSAL_MBUF Structure .....	40
7	OSAL IX_OSAL_MBUF Structure and Macros .....	41
8	API User Interface to IX_OSAL_MBUF .....	42
9	Pool Management Fields .....	43
10	IX_OSAL_MBUF: ix_osdep_buf Structure .....	44
11	IX_OSAL_MBUF: ix_ctrl Structure .....	44
12	IX_OSAL_MBUF: NPE Shared Structure .....	45
13	Internal Mapping of ix_osdep_buf to the Shared NPE Structure .....	46
14	AAL5 PDU Transmission for a Scheduled Port .....	63
15	IxAtmdAccScheduleTable Structure and Order Of ATM Cell .....	65
16	Tx Done Recycling — Using a Threshold Level .....	67
17	Tx Done Recycling — Using a Polling Mechanism .....	68
18	Tx Disconnect .....	69
19	Rx Using a Threshold Level .....	71
20	RX Using a Polling Mechanism .....	72
21	Rx Disconnect .....	73
22	Services Provided by Ixatmm .....	81
23	Configuration of Traffic Control Mechanism .....	83
24	Component Dependencies of IxAtmm .....	84
25	Multiple VCs for Each Port, Multiplexed onto Single Line by the ATM Scheduler .....	91
26	Translation of IxAtmScheduleTable Structure to ATM Tx Cell Ordering .....	92
27	Basic IxCryptoAcc API Flow .....	101
28	IxCryptoAcc API Call Process Flow for CCD Updates .....	103
29	IxCryptoAcc Component Dependencies .....	108
30	IxCryptoAcc, NPE and IPSec Stack Scope .....	109
31	Relationship Between IPSec Protocol and Algorithms .....	110
32	ESP Packet Structure .....	111
33	Authentication Header .....	111
34	ESP Data Flow .....	112
35	AH Data Flow .....	113
36	IPSec API Call Flow .....	114
37	CCM Operation Flow .....	116
38	CCM Operation on Data Packet .....	116
39	AES CBC Encryption For MIC .....	117
40	AES CTR Encryption For Payload and MIC .....	117
41	WEP Frame with Request Parameters .....	119



42	WEP Perform API Call Flow.....	121
43	ixDmaAcc Dependencies.....	129
44	IxDmaAcc Component Overview.....	130
45	IxDmaAcc Control Flow.....	136
46	IxDMAcc Initialization.....	137
47	DMA Transfer Operation.....	138
48	Ethernet Component Overview.....	143
49	Ethernet Access Layers Block Diagram.....	146
50	Ethernet Transmit Frame API Overview.....	147
51	Ethernet Transmit Frame Data Buffer Flow.....	149
52	Ethernet Receive Frame Overview.....	152
53	Ethernet Receive Plane Data Buffer Flow.....	155
54	IxEthAcc and Secondary Components.....	159
55	Example Network Diagram for MAC Address Learning and Filtering with Two Ports.....	175
56	Egress VLAN Control Path for Untagged Frames.....	186
57	QoS on Receive for 802.1Q Tagged Frames.....	188
58	QoS on Receive for Untagged Frames.....	189
59	AP-STA and AP-AP Modes.....	191
60	Passing Only the Two Types of BPDUs.....	197
61	HSS/HDLC Access Overview.....	212
62	T1 Tx Signal Format.....	214
63	IxHssAcc Component Dependencies.....	217
64	Normal and Bypass Mode Illustration.....	220
65	Channelized Connect For Normal Mode.....	222
66	Channelized Connect For Bypass Mode.....	224
67	Channelized Transmit and Receive.....	227
68	Packetized Connect.....	229
69	Packetized Transmit.....	232
70	Packetized Receive.....	234
71	HSS Packetized Receive Buffering.....	236
72	HSS Packetized Transmit Buffering.....	237
73	HSS Channelized Receive Operation.....	239
74	HSS Channelized Transmit Operation.....	240
75	Message from Intel XScale® Processor Software Client to an NPE.....	254
76	Message with Response from Intel XScale® Processor Software Client to an NPE.....	255
77	Receiving Unsolicited Messages from NPE to Software Client.....	256
78	ixNpeMh Component Dependencies.....	257
79	IxParityENAcc Dependency Diagram.....	264
80	Parity Error Notification Sequence.....	265
81	Data Abort with No Parity Error.....	268
82	Parity Error with No Data Abort.....	269
83	Data Abort followed by Unrelated Parity Error Notification.....	269
84	Unrelated Parity Error Followed by Data Abort.....	270
85	Data Abort Caused by Parity Error.....	270
86	Parity Error Notification Followed by Related Data Abort.....	271
87	Data Abort with both Related and Unrelated Parity Errors.....	271
88	Architectural View of ixErrHdlAcc Component.....	274
89	Layered Block Diagram Depicting the Dependencies of Intel® IXP400 Software ParityENAcc Access Component with Notification and Soft Error Recovery.....	275
90	Interface Architecture.....	276
91	Flow Chart Depicting the Initialization and Configuration of the Soft Error Detection and Handling.....	278
92	A Basic Software Architecture Showing the Soft Error Detection and Handling Implementation using the Intel® IXP400 Software Access Libraries.....	280
93	Data Flow Diagram Depicting the Start of an Interrupt Error as it Recovers.....	283



94	AQM Hardware Block .....	290
95	Dispatcher in Context of an Interrupt.....	295
96	Dispatcher in Context of a Polling Mechanism.....	296
97	IxSspAcc Dependencies .....	302
98	Interrupt Scenario.....	304
99	Polling Scenario .....	306
100	IxTimeSyncAcc Component Dependencies .....	308
101	Block Diagram of Intel® IXP46X Network Processor .....	310
102	Polling for Timestamps of Sync or Delay_Req.....	314
103	Interrupt Servicing of Target Time Reached Condition .....	315
104	Polling for Auxiliary Snapshot Values .....	315
105	UART Services Models.....	319
106	USBSetupPacket .....	327
107	STALL on IN Transactions.....	330
108	STALL on OUT Transactions.....	330
109	USB Dependencies .....	332
110	OSAL Architecture .....	338
111	OSAL Directory Structure .....	342
112	OSAL Intel® IXP4XX Processor Variant Directory Structure.....	343
113	Requirements for the Routine TimeStamp.....	351
114	STMicroelectronics* ADSL Chipset on the Intel® IXDP425 / IXCDP1100 Development Platform.....	356
115	Example of ADSL Line Open Call Sequence .....	357
116	I <sup>2</sup> C Driver Dependencies.....	361
117	Sequence Flow Diagram for Slave Receive / General Call in Interrupt Mode.....	364
118	Sequence Flow Diagram for Slave Transmit in Interrupt Mode .....	365
119	Sequence Flow Diagram for Slave Receive in Polling Mode .....	366
120	Sequence Flow Diagram for Slave Transmit in Polling Mode.....	367
121	Endianness in Big Endian-Only Software Release .....	375
122	Intel® IXP4XX Product Line of Network Processors Endianness Controls.....	379
123	Ethernet Frame (Big Endian) .....	386
124	One Half-Word-Aligned Ethernet Frame (LE Address Coherent) .....	387
125	Intel XScale® Processor Read of IP Header (LE Data Coherent) .....	388
126	VxWorks* Data Coherent Swap Code.....	392

## Tables

1	Acronym Listing.....	25
2	Internal ix_osdep_buf Field Format .....	47
3	ix_osdep_buf Field Details.....	47
4	ix_osdep_buf to M_BLK Mapping.....	49
5	Buffer Translation Functions .....	50
6	Acronyms .....	55
7	ATM Cell Header .....	56
8	Negotiable Traffic and QoS Parameters .....	56
9	AAL5 CPCS PDU .....	57
10	OAM Reserved VPI/VCI and Their Function .....	58
11	OAM Payload Format .....	58
12	OAM Type and Function Type.....	58
13	List of Configuration Routines .....	59
14	List of Transmission Routines.....	62
15	List of Receive Routines .....	70
16	IX_OSAL_MBUF Fields Required for Transmission .....	74
17	IX_OSAL_MBUF Fields of Available Buffers for Reception .....	74
18	IX_OSAL_MBUF Fields Modified During Reception .....	75



19	Real-Time Errors .....	76
20	List of Routines .....	78
21	List of Routines in the Component .....	89
22	Supported Traffic Types and Traffic Parameters .....	89
23	IxAtmSch Data Memory Usage .....	94
24	IxCryptoAcc API .....	98
25	Difference Between the NPE and PKE-Based Hash Routine.....	106
26	IxCryptoAcc Data Memory Usage.....	107
27	Supported Encryption Algorithms.....	123
28	Supported Authentication Algorithms .....	125
29	DMA Modes Supported for Addressing Mode of Incremental Source Address and Incremental Destination Address .....	133
30	DMA Modes Supported for Addressing Mode of Incremental Source Address and Fixed Destination Address.....	134
31	DMA Modes Supported for Addressing Mode of Fixed Source Address and Incremental Destination Address .....	135
32	IPv6/IPv4 Payload Detection .....	157
33	IX_OSAL_MBUF Structure Format .....	164
34	ixp_ne_flags Field Format .....	164
35	IX_OSAL_MBUF Header Definitions for the Ethernet Subsystem .....	164
36	IX_OSAL_MBUF "Port ID" Field Format .....	166
37	IX_OSAL_MBUF "Port ID" Field Values .....	167
38	ixp_ne_flags.link_prot Field Values .....	167
39	Managed Objects for Ethernet Receive.....	168
40	Managed Objects for Ethernet Transmit .....	169
41	IxEthDB Feature Set .....	172
42	Possible IP Types.....	181
43	Untagged MAC Frame Format.....	182
44	VLAN Tagged MAC Frame Format.....	182
45	VLAN Tag Format .....	182
46	Egress VLAN Tagging/Untagging Behavior Matrix .....	187
47	Default Priority to Traffic Class Mapping.....	189
48	IEEE 802.11 Frame Format .....	190
49	IEEE802.11 Frame Control (FC) Field Format.....	191
50	STA Frame Format.....	192
51	802.3 to 802.11 Header Conversion Rules .....	194
52	802.11 to 802.3 Header Conversion Rules .....	195
53	PHYs Supported by IxEthMii .....	202
54	Product ID Values.....	203
55	Feature Control Register Values.....	204
56	HSS Tx Clock Output frequencies and PPM Error .....	213
57	HSS TX Clock Output Frequencies and Associated Jitter Characterization.....	213
58	Jitter Definitions .....	214
59	HSS Frame Output Characterization .....	214
60	NPE-A Images.....	243
61	NPE-B Images.....	245
62	NPE-C Images.....	246
63	Parity Error Interrupts.....	262
64	Parity Capabilities Supported by IxParityENAcc .....	263
65	Parity Error Interrupt Deassertion Conditions .....	266
66	Interface Identity and Types .....	277
67	AQM Configuration Attributes .....	292
68	Default IEEE 1588 Hardware Assist Block States upon Hardware/Software Reset.....	311
69	IN, OUT, and SETUP Token Packet Format .....	322
70	SOF Token Packet Format.....	322



71	Data Packet Format.....	323
72	Handshake Packet Format.....	323
73	Bulk Transaction Formats.....	324
74	Isochronous Transaction Formats.....	324
75	Control Transaction Formats, Set-Up Stage.....	325
76	Control Transaction Formats.....	325
77	Interrupt Transaction Formats.....	325
78	API interfaces Available for Access Layer.....	326
79	Host-Device Request Summary.....	328
80	Detailed Error Codes.....	331
81	OSAL Core Interface.....	344
82	OSAL Buffer Management Interface.....	346
83	OSAL I/O Memory and Endianness Interface.....	348
84	Intel XScale® Processor Little Endian Writes in Address Coherent Mode and NPE Reads to/from SDRAM.....	377
85	Intel XScale® Processor Little Endian Writes in Data Coherent Mode and NPE Reads to/from SDRAM.....	378
86	Endian Hardware Summary.....	381
87	Intel® IXP400 Software Macros.....	391
88	Endian Conversion Macros.....	391
89	Intel® IXP400 Software Versions.....	393



## Revision History

Date	Revision	Description
May 2006	009	<p>Updated guide for IXP400 Software Version 2.3. New features are described in the following chapters:            New chapter added --            Chapter 17.0, "Access-Layer Components: Error Handler (IxErrHdlAcc) API"            Updates made to following chapters:            Chapter 4.0, "Access-Layer Components: ATM Driver Access (IxAtmdAcc) API"            Chapter 7.0, "Access-Layer Components: Security (IxCryptoAcc) API"            Chapter 9.0, "Access-Layer Components: Ethernet Access (IxEthAcc) API"            Chapter 10.0, "Access-Layer Components: Ethernet Database (IxEthDB)"            Chapter 13.0, "Access-Layer Components: HSS-Access (IxHssAcc) API"            Chapter 14.0, "Access-Layer Components: NPE-Downloader (IxNpeDI) API"            Chapter 15.0, "Access-Layer Components: NPE Message Handler (IxNpeMh) API"            Chapter 16.0, "Access-Layer Components: Parity Error Notifier (IxParityENAcc) API"            Chapter 18.0, "Access-Layer Components: Queue Manager (IxQMgr) API"</p>
November 2005	008	<p>Intel® IXP42X product line stepping A0 and IXC1100 control plane processor are no longer supported.            Updated guide for IXP400 Software Version 2.1. New features are described in the following chapters:</p> <ul style="list-style-type: none"> <li>• Chapter 7.0, "Access-Layer Components: Security (IxCryptoAcc) API"</li> <li>• Chapter 6.0, "Access-Layer Components: ATM Transmit Scheduler (IxAtmSch) API"</li> <li>• Chapter 9.0, "Access-Layer Components: Ethernet Access (IxEthAcc) API"</li> <li>• Chapter 10.0, "Access-Layer Components: Ethernet Database (IxEthDB)"</li> <li>• Chapter 13.0, "Access-Layer Components: HSS-Access (IxHssAcc) API"</li> <li>• Chapter 14.0, "Access-Layer Components: NPE-Downloader (IxNpeDI) API"</li> <li>• Chapter 16.0, "Access-Layer Components: Parity Error Notifier (IxParityENAcc) API"</li> <li>• Chapter 23.0, "Codelets"</li> </ul> <p>General enhancements and updates were made to the following chapters:</p> <ul style="list-style-type: none"> <li>• Chapter 3.0, "Buffer Management"</li> <li>• Chapter 4.0, "Access-Layer Components: ATM Driver Access (IxAtmdAcc) API"</li> <li>• Chapter 18.0, "Access-Layer Components: Queue Manager (IxQMgr) API"</li> <li>• Chapter 24.0, "Operating System Abstraction Layer (OSAL)"</li> <li>• Chapter 27.0, "Endianness in Intel® IXP400 Software v2.3"</li> </ul> <p>Change bars indicate areas of change.</p>



Date	Revision	Description
April 2005	007	Updated guide for IXP400 Software Version 2.0. Added: <ul style="list-style-type: none"> <li>• Chapter 16, "Access-Layer Components: Parity Error Notifier (IxParityENAcc) API"</li> <li>• Chapter 19, "Access-Layer Components: Synchronous Serial Port (IxSspAcc) API"</li> <li>• Chapter 20, "Access-Layer Components: Time Sync (IxTimeSyncAcc) API"</li> <li>• Chapter 26, "I2C Driver (IxI2cDrv)"</li> </ul> Removed: Access-Layer Components: Fast-Path Access (IxFpathAcc) API Change bars indicate areas of change.
November 2004	006	Updated guide for IXP400 Software Version 1.5. Added Chapter 24, "Endianness in Intel® IXP400 Software v1.5", and revised: <ul style="list-style-type: none"> <li>• Chapter 3, "Buffer Management"</li> <li>• Chapter 9, "Access-Layer Components: Ethernet Access (IxEthAcc) API"</li> <li>• Chapter 10, "Access-Layer Components: Ethernet Database (IxEthDB) API"</li> <li>• Chapter 18, "Access-Layer Components: Queue Manager (IxQMgr) API"</li> <li>• Chapter 22, "Operating System Abstraction Layer (OSAL)"</li> </ul> Change bars indicate areas of change.
December 2003	005	Updated manual for IXP400 Software Version 1.4. Removed API documentation (in a separate reference).
September 2003	004	Made two minor corrections.
August 2003	003	Updated manual for IXP400 Software Version 1.3.
February 2003	002	Removed "Intel Confidential" classification.
February 2003	001	Initial release of document.





## 1.0 Introduction

---

This chapter contains important information to help you understand and use the Intel® IXP400 Software v2.3 release.

### 1.1 Versions Supported by this Document

This programmer's guide is intended to be used in conjunction with software release 2.3. Refer to the accompanying release notes for the latest information regarding proper documentation sources, and so forth.

Previous versions of the programmer's guide (for earlier IXP400 software releases) can be found at the following Web site:

<http://developer.intel.com/design/network/products/npfamily/docs/ixp4xx.htm>

To identify your version of software:

1. Open the file `ixp400_xscale_sw/src/include/IxVersionId.h`.
2. Check the value of `IX_VERSION_ID`.

### 1.2 Hardware Supported by this Release

The Intel® IXP400 Software v2.3 release supports all Intel® IXP4XX Product Line of Network Processors. Processor capabilities differ between processor product lines or processor variants. Not all capabilities of the processor may be supported by this software release.

*Note:* The Intel® IXP42X product line stepping A0 and IXC1100 control plane processor are no longer supported.

### 1.3 Intended Audience

This document describes the software release 2.3 architecture and is intended for software developers and architects employing the Intel® IXP4XX product line processors. The document describes each software component's functionality, dependencies between the components, and presents the common design policies used by each component.

### 1.4 How to Use this Document

This programmer's guide is organized as follows:



Chapters	Description
Chapters 1 and 2	Introduces the Intel® IXP400 Software v2.3 and the supported processors, including an overview of the software architecture.
Chapters 4 through 22	Provide functional descriptions of the various access-layer components.
Chapter 3 and 25	Describe the memory buffer management and operating system abstraction layers, needed for a more in-depth architectural understanding of the software.
Chapter 24 and 26–28	Describe codelets (example applications), ADSL driver, I <sup>2</sup> C driver, and endianness.

For the developer interested in a limited number of specific features of the software release 2.3, a recommended reading procedure would be:

1. Read Chapters 1 through 3 to get a general knowledge of the products' software and hardware architecture.
2. Read the chapters on the specific access-layer component(s) of interest.

**Note:** Many of the access-layer components have dependencies on other components — particularly on IxNpeDI and IxQmgr. For that reason, developers also should review those chapters.

3. Review the codelet descriptions in [Chapter 23.0](#) and their respective source code for those codelets that offer features of interest.
4. Refer to the API source code and source code documentation found in the software release documents folder as necessary.

## 1.5 About the Processors

Next-generation networking solutions must meet the growing demands of users for high-performance data, voice, and networked multimedia products. Manufacturers of networking equipment must develop new products under stringent time-to-market deadlines and deliver products whose software can be easily upgraded. The IXP4XX product line processors family is designed to meet the needs of broadband and embedded networking products such as high-end residential gateways; small to medium enterprise (SME) routers, switches, security devices; DSLAMs (Digital Subscriber Line Access Multiplexers) for multi-dwelling units (MxU); wireless access points; industrial control systems; and networked printers.

The IXP4XX product line processors deliver wire-speed performance and sufficient “processing headroom” for manufacturers to add a variety of rich software services to support their applications. These are highly integrated network processors that support multiple WAN and LAN technologies, giving customers a common architecture for multiple applications. With its development platforms, a choice of operating systems, and a broad range of development tools, the processor family is supported by a complete development environment for faster time-to-market. This network processor family offers the choice of multiple clock speeds at 266, 400, 533 and 667 MHz, with both commercial (0° to 70° C) and extended (-40° to 85° C) temperature options.

The IXP4XX product line processors have a unique distributed processing architecture that features the performance of the Intel XScale® Processor and up to three Network Processor Engines (NPEs). The combination of the four high-performance processors provides tremendous processing power and enables wire-speed performance at both the LAN and WAN ports. The three NPEs are designed to offload many computationally intensive data plane operations from the Intel XScale® Processor. This provides ample “processing headroom” on the Intel XScale® Processor for developers to add differentiating product features.

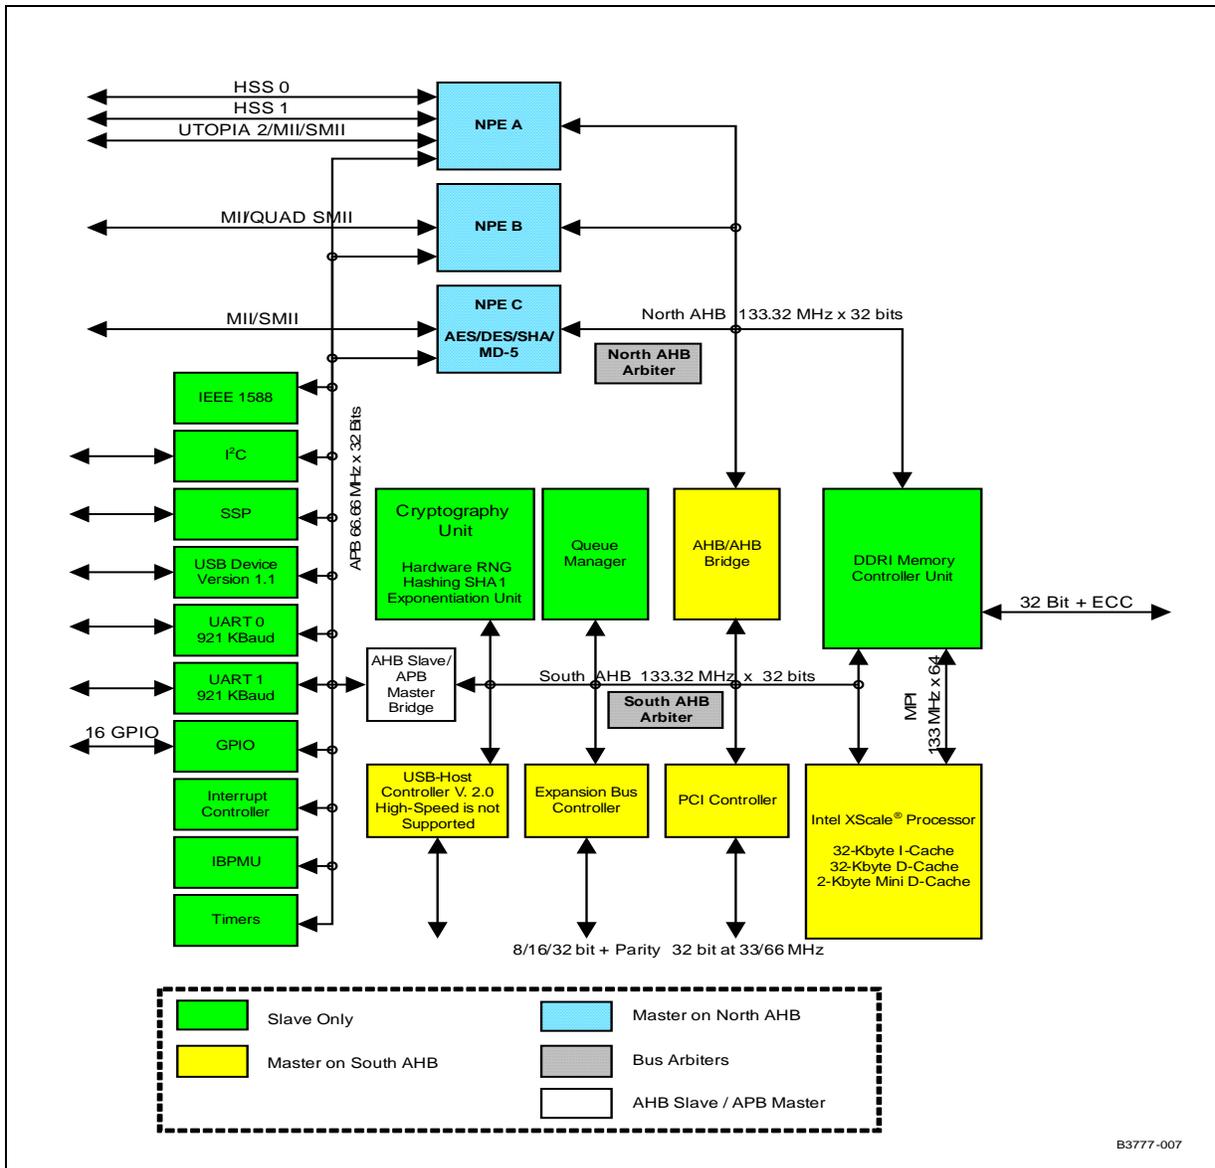


For a list of IXP42X product line features, see the *Intel® IXP42X Product Line of Network Processors and IXC1100 Control Plane Processor Datasheet*.

For a list of IXP46X product line features, see the *Intel® IXP45X and Intel® IXP46X Product Line of Network Processors Datasheet*.

For a quick reference, see [Figure 1 on page 23](#) for a block diagram of the Intel® IXP465 Network Processor.

**Figure 1. Intel® IXP46X Product Line Network Processor Block Diagram**



Software development is eased by the extensive Intel XScale® Processor tools environment that includes compilers, debuggers, operating systems, models, support services from third party vendors, and fully documented evaluation hardware platforms



and kits. The compiler, assembler, and linker support specific optimizations designed for the Intel XScale® technology, the ARM\* instruction set v.5TE, and Intel DSP extensions.

## 1.6 Related Documents

Users of this document should always refer to the associated **Software Release Notes** for the specific release. Additional Intel documents listed below are available from your field representative or from the following Web site:

<http://www.intel.com/design/network/products/npfamily/docs/ixp4xx.htm>

Document Title	Document #
Intel® IXP400 Software Specification Update	273795
<a href="http://www.intel.com/intelpress/sum_ixp4.htm">http://www.intel.com/intelpress/sum_ixp4.htm</a> , <i>Designing Embedded Networking Applications</i> (an Intel Press book)	ISBN 0-9743649-3-2
Intel® IXP42X Product Line of Network Processors and IXC1100 Control Plane Processor Developer's Manual	252480
Intel® IXP42X Product Line of Network Processors and IXC1100 Control Plane Processor Datasheet	252479
Intel® IXP45X and Intel® IXP46X Product Line of Network Processors Datasheet	306261
Intel® IXP45X and Intel® IXP46X Product Line of Network Processors Developer's Manual	306262
Intel® IXP4XX Product Line of Network Processors Specification Update	306428
Intel® IXDP425 / IXCDP1100 Development Platform Specification Update	253527
Intel® IXDP465 Development Platform Specification Update	306509
ARM* Architecture Version 5TE Specification	ARM DDI 0100E (ISBN 0 201 737191)
PCI Local Bus Specification, Revision 2.2	–
Universal Serial Bus Specification, Revision 1.1	–
UTOPIA Level 2 Specification, Revision 1.0	–
IEEE 802.3 Specification	–
IEEE 1149.1 Specification	–
IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems (IEEE Std. 1588™ - 2002)	–
ARM Ltd., <i>AMBA Specification</i> , Rev. 2.0, May 1999	–
<a href="http://www.pcisig.com/reflector/msg01668.html">http://www.pcisig.com/reflector/msg01668.html</a> , a discussion on a PCI bridge between little and big endian devices.	–



## 1.7 Acronyms

Table 1. Acronym Listing

Acronym	Description
AAL	ATM Adaptation Layer
ABR	Available Bit Rate
ACK	Acknowledge Packet
ADSL	Asymmetric Digital Subscriber Line
AES	Advanced Encryption Standard
AH	Authentication Header (RFC 2402)
AHB	Advanced High-Performance Bus
AL	Adaptation Layer
AP	Access Permission
APB	Advanced Peripheral Bus
API	Application Programming Interface
AQM	AHB Queue Manager
ARC4	Alleged RC4
ATM	Asynchronous Transfer Mode
ATU-C	ADSL Termination Unit — Central Office
ATU-R	ADSL Termination Unit — Remote
BE	Big endian
BSD	Berkeley Software Distribution
BSP	Board Support Package
CAC	Connection Admission Control
CAS	Channel Associated Signaling
CBC	Cipher Block Chaining
CBR	Constant Bit Rate
CCD	Cryptographic Context Database
CCM	Counter mode encryption with CBC-MAC authentication
CDVT	Cell Delay Variation Tolerance
CFB	Cipher FeedBack
CPCS	Common Part Convergence Sublayer
CPE	Customer Premise Equipment
CPU	Central Processing Unit
CRC	Cyclic Redundancy Check
CTR	Counter Mode
DDR	Double Data Rate
DES	Data Encryption Standard
DMT	Discrete Multi-Tone
DOI	Domain of Interpretation
DSL	Digital Subscriber Line
DSP	Digital Signal Processor
DUE	Detected unrecoverable error



Table 1. Acronym Listing (Continued)

Acronym	Description
E	Empty
E1	Euro 1 trunk line (2.048 Mbps)
ECB	Electronic Code Book
ECC	Error Correction Code
EISA	Extended ISA
ERP	Endpoint Request Packet
ESP	Encapsulation Security Payload (RFC2406)
Eth0	Ethernet NPE A
Eth1	Ethernet NPE B
F	Full
FCS	Frame Check Sequence
FIFO	First In First Out
FIQ	Fast Interrupt Request
FRAD	Frame Relay Access Device
FRF	Frame Relay Forum
FXO	Foreign Exchange Office
FXS	Foreign Exchange Subscriber
G.SHDSL	ITU G series specification for symmetric High Bit Rate Digital Subscriber Line
GCI	General Circuit Interface
GE	Gigabit Ethernet
GFR	Guaranteed Frame Rate
GPIO	General Purpose Input/Output
HDLC	High-Level Data Link Control
HDSL2	High Bit-Rate Digital Subscriber Line version 2
HEC	Header Error Check
HLD	High Level Design
HMAC	Hashed Message Authentication Code
HPI	Host Port Interface
HPNA	Home Phone Network Alliance
HSS	High Speed Serial
HSSI	High Speed Serial Interface
HW	Hardware
IAD	Integrated Access Device
IBPMU	Internal Bus PMU
ICV	Integrity Check Value
IKE	Internet Key Exchange
IMA	Inverse Multiplexing over ATM
IP	Internet Protocol
IPsec	Internet Protocol Security
IRQ	Interrupt Request



Table 1. Acronym Listing (Continued)

Acronym	Description
ISA	Industry Standard Architecture
ISR	Interrupt Service Routine
ISR	Interrupt Sub-Routine
IV	Initialization Vector
IX_OSAL_MBUF	BSD 4.4-like mbuf implementation for software release 2.3. Referred to as IX_MBUF, IXP_BUF and IX_OSAL_MBUF interchangeably.
IX_MBUF	BSD 4.4-like mbuf implementation for software release 2.3. Referred to as IX_MBUF, IXP_BUF and IX_OSAL_MBUF interchangeably.
IXA	Internet Exchange Architecture
IXP	Internet Exchange Processor
IXP_BUF	BSD 4.4-like mbuf implementation for software release 2.3. Referred to as IX_MBUF, IXP_BUF and IX_OSAL_MBUF interchangeably.
LAN	Local Area Network
LE	Little endian
LSB	Least Significant Bit
MAC	Media Access Control
MAC	Message Authentication Code (in SSL or TLS)
MBS	Maximum Burst Size
MCR	Minimum Cell Rate
MCU	Memory Controller Unit
MD5	Message Digest 5
MFS	Maximum Frame Size
MIB	Management Information Base
MII	Media-Independent Interface
MLPPP	Multi-Link Point-to-Point Protocol
MMU	Memory Management Unit
MPHY	Multi PHY
MPI	Memory Port Interface
MSB	Most Significant Bit
MVIP	Multi-Vendor Integration Protocol
MxU	Multi-dwelling Unit
NAK	Not-Acknowledge Packet
NAPT	Network Address Port Translation
NAT	Network Address Translation
NE	Nearly Empty
NF	Nearly Full
NOTE	Not Empty
NOTF	Not Full
NOTNE	Not Nearly Empty
NOTNF	Not Nearly Full



Table 1. Acronym Listing (Continued)

Acronym	Description
NPE	Network Processing Engine
OC3	Optical Carrier - 3
OF	Overflow
OFB	Output FeedBack
OS	Operating System
OSAL	Operating System Abstraction Layer
PBX	Private Branch Exchange
PCI	Peripheral Control Interconnect
PCI	Peripheral Component Interface
PCR	Peak Cell Rate
PDU	Protocol Data Unit
PHY	Physical Layer Interface
PID	Packet Identifier
PMU	Performance Monitoring Unit
PRE	Preamble Packet
PTP	Precision Time Protocol
QM or QMgr	Queue Manager
rt-VBR	Real Time Variable Bit Rate
Rx	Receive
SA	Security Association
SAR	Segmentation and Re-assembly
SCR	Sustainable Cell Rate
SDRAM	Synchronous Dynamic Random Access Memory
SDSL	Symmetric Digital Subscriber Line
SDU	Service Data Unit
SERC	Soft-error recovery component (ixErrHdlAcc component)
SHA1	Secure Hash Algorithm 1
SIO	Standard I/O (input/output)
SIP	Session Initiation Protocol
SMII	Serial Media-Independent Interface
SNMP	Simple Network Management Protocol
SOF	Start of Frame
SPHY	Single PHY
SSL	Secure Socket Layer
SSP	Synchronous Serial Port
SVC	Switched Virtual Connection
SWCP	Switching Coprocessor
TCD	Target Controller Driver
TCI	Transmission Control Interface
TCP	Transmission Control Protocol



Table 1. Acronym Listing (Continued)

Acronym	Description
TDM	Time Division Multiplexing
TLB	Translation Lookaside Buffer
TLS	Transport Level Security
ToS	Type of Service
Tx	Transmit
UBR	Unspecified Bit Rate
UDC	Universal Serial Bus Device Controller
UF	Underflow
USB	Universal Serial Bus
UTOPIA	Universal Test and Operation PHY Interface for ATM
VBR	Variable Bit Rate
VC	Virtual Connection
VCC	Virtual Circuit Connection
VCI	Virtual Circuit Identifier
VDSL	Very High Speed Digital Subscriber Line
VLAN TCI	VLAN Tag Control Information
VLAN TPID	VLAN Tag Protocol ID
VoDSL	Voice over Digital Subscriber Line
VoFR	Voice over Frame Relay
VoIP	Voice over Internet Protocol
VPC	Virtual Path Connection
VPI	Virtual Path Identifier
VPN	Virtual Private Network
WAN	Wide Area Network
WEP	Wired Equivalent Privacy
Xcycle	Idle-Cycle Counter Utilities
xDSL	Any Digital Subscriber Line
XOR	Exclusive OR

§ §





## 2.0 Software Architecture Overview

---

### 2.1 High-Level Overview

The primary design principles of the Intel® IXP400 Software v2.3 architecture are to enable the supported processors' hardware in a manner which allows maximum flexibility. Intel® IXP400 Software v2.3 consists of a collection of software components specific to the Intel® IXP4XX Product Line of Network Processors and their supported development and reference boards.

This section discusses the software architecture of this product, as shown in [Figure 2 on page 32](#).

The **NPE microcode** consists of one or more loadable and executable NPE instruction files that implement the NPE functionality behind the software release 2.3 library. The NPEs are RISC processors embedded in the main processor that are surrounded by multiple coprocessor components. The coprocessors provide specific hardware services (for example, Ethernet processing and MAC interfaces, cryptographic processing, and so forth). The NPE instruction files are incorporated into the software release 2.3 library at build time (or at run-time for Linux\*). The library includes a NPE downloader component that provides NPE code version selection and downloading services. A variety of NPE microcode images are provided, enabling different combinations of services.

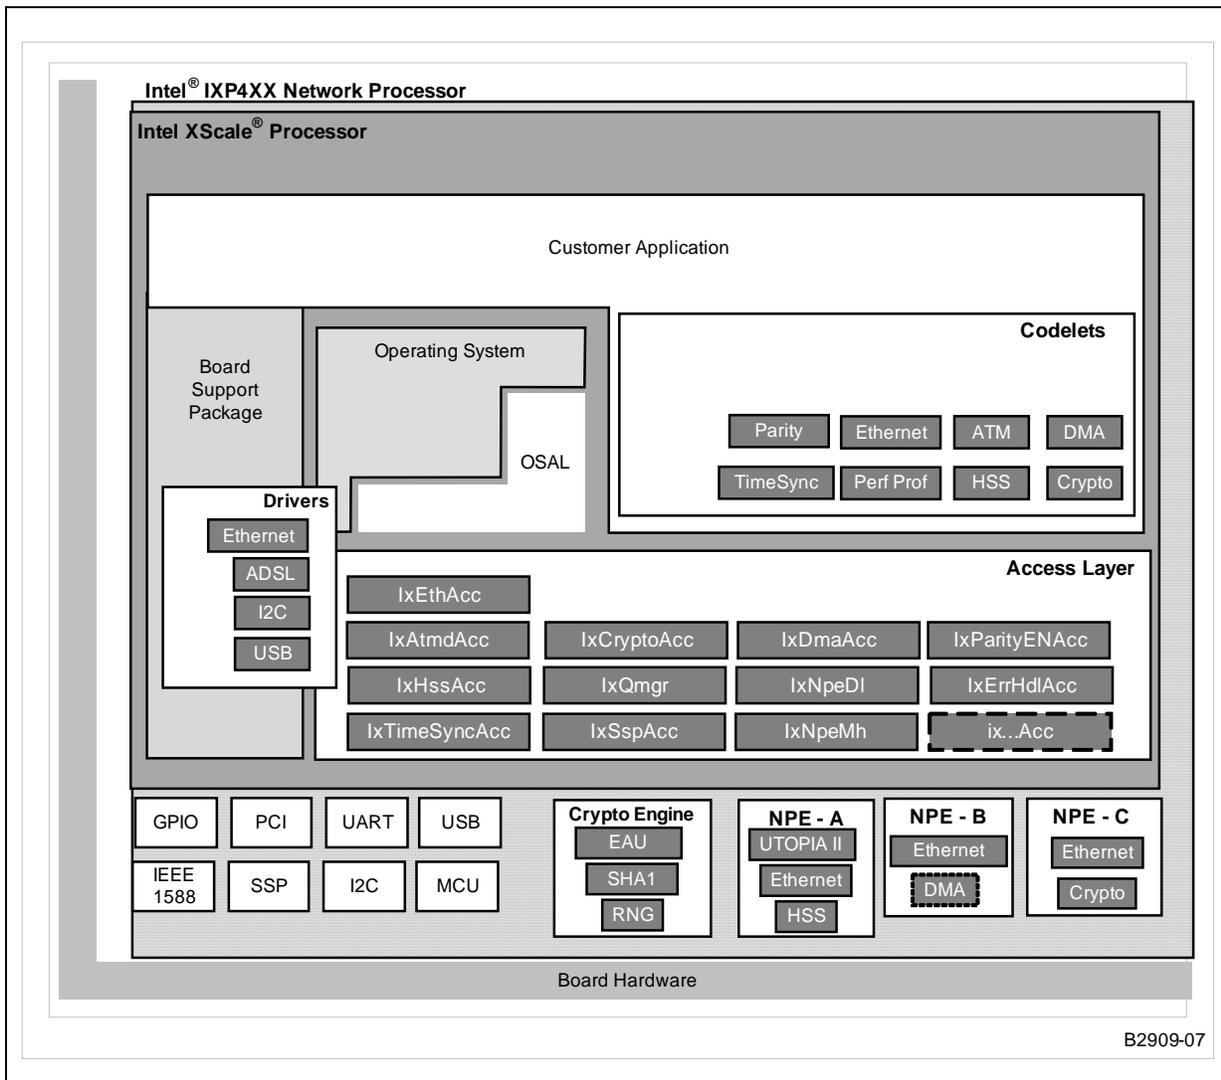
The **Access Layer** provides a software interface which gives customer code access to the underlying capabilities of the supported processors. This layer is made up of a set of software components (access-layer components), which clients can use to configure, control and communicate with the hardware. Specifically, most access-layer components provide an API interface to specific NPE-hosted hardware capabilities, such as AAL 0 and AAL 5 on UTOPIA, Cryptography, Ethernet, HSS, or DMA. The remaining access-layer components provide an API interface to peripherals on the processors (for example, UART and USB) or features of the Intel XScale® processor (for example, Product ID Registers or Performance Monitoring Unit).

The example **Codelets** are narrowly focused example applications that show how to use many of the services or functions provided by the Intel XScale® processor library and the underlying hardware. Many codelets are organized by hardware port type and typically exercise some Layer-2 functionality on that port, such as: AAL 5 PDU Transmit / Receive over UTOPIA, Channelized or HDLC Transmit / Receive over HSS, Ethernet frame Transmit / Receive.

The **Operating System Abstraction Layer (OSAL)** defines a portable interface for operating system services. The access-layer components and the codelets use the OS services as abstracted in this module.

**Device Driver** modules translate the generic Operating System specific device interface commands to the Access Layer software APIs. Some device driver modules are provided by the OS vendors' Board Support Packages. Others may be provided in conjunction with the software release 2.3.

Figure 2. Intel® IXP400 Software v2.3 Architecture Block Diagram



## 2.2 Deliverable Model

Intel® IXP400 Software v2.3 consists of these elements:

- Intel® IXP400 Software v2.3 access-layer components and OSAL layer
- Complete documentation and source code for software release 2.3 components
- NPE microcode images
- Example codelets

**Note:**

The software releases do not include tools to develop NPE software. The supplied NPE functionality is accessible through the access-layer APIs provided by the software release 2.3 library. The NPE microcode is provided as a ".c" file that must be compiled with the access-layer library. NPE microcode is compatible only with the specific access layer for which it is provided.



## 2.3 Operating System Support

The Intel XScale® technology offers a broad range of tools together with support for two widely adopted operating systems. The software release 2.3 supports VxWorks\* and the standard Linux 2.6 kernel. MontaVista\* software will provide the support for Linux\*. Support for other operating systems may be available. For further information, visit the following Internet site:

<http://developer.intel.com/design/network/products/npfamily/ixp425.htm>

The software release 2.3's software library is OS-independent in that all components are written in ANSI-C with no direct calls to any OS library function that is not covered by ANSI-C. A thin abstraction layer is provided for some OS services (timers, mutexes, semaphores, and thread management), which can be readily modified to support additional operating systems. This enables the devices to be compatible with multiple operating systems and allows customers the flexibility to port the IXP4XX product line processors to their OS of choice.

## 2.4 Development Tools

The Intel XScale® technology offers a broad range of tools together with support for two widely adopted operating systems. Developers have a wide choice of third-party tools including compilers, linkers, debuggers and board-support packages (BSPs). Tools include Wind River\* WorkBench\* 2.4 for the VxWorks\* 6.2 real-time operating system, Wind River's\* PLATFORM for Network Equipment\* and the complete GNU\* Linux\* development suite.

Refer to the release notes accompanying the software for information on specific OS support.

## 2.5 Access Library Source Code Documentation

The access library source code uses a commenting style that supports the Doxygen\* tool for use in creating source code documentation. Doxygen is an open-source tool that reads appropriately commented source code and produces hyperlinked documentation of the APIs suitable for online browsing (HTML).

The documentation output is typically multiple HTML files, but Doxygen can be configured to produce LaTeX\*, RTF (Rich Text Format\*), PostScript, hyperlinked PDF, compressed HTML, and Unix\* man pages. Doxygen is available for Linux\*, Windows\* and other operating systems.

For more information, use the following Web URL:

<http://www.doxygen.org>

The software release 2.3 compressed file contains the HTML source code documentation at `ixp400_xscale_sw\doc\index.html`. This output is suitable for **online** browsing. For a **printable** reference, see the Adobe\* Portable Document Format (PDF) file contained in the compressed software-download file.

## 2.6 Release Directory Structure

The software release 2.3 includes the following directory structure:

```
\---ixp_osal
    +---doc (API References in HTML and PDF format)
    +---include
```



```
+---os
+---src
\---ixp400_xscale_sw
+---buildUtils (setting environment vars. in VxWorks* and Linux*)
+---doc (API Reference in HTML and PDF format)
  \---src (contains access-layer and codelet source code)
    +---adsl (separate package)
    +---atmdAcc
    +---atmm
    +---atmsch
    +---codelets (sub-directory for codelet source)
      | +---atm
      | +---cryptoAcc (for crypto version only)
      | +---dmaAcc
      | +---ethAal5App
      | +---ethAcc
      | +---hssAcc
      | +---parityENAcc
      | +---perfProfAcc
      | +---timeSyncAcc
      | \---usb (separate package)
      |   +---drivers
      |   +---include
    +---cryptoAcc (for crypto version only)
  \---dmaAcc
    | \---errHdlAcc
    |   +---include
  \---ethAcc
    | +---include
  \---ethDB
```



```

|   +---include
+---ethMii
+---featureCtrl
\---hssAcc
|   +---include
+---i2c
+---include (header location for top-level public modules)
\---npeDl
|   +---include
\---npeMh
|   +---include
+---osLinux (Linux* specific operations for loading NPE microcode)
+---osServices (v1.4 backwards compatibility)
+---ossl (v1.4 backwards compatibility)
+---parityENAcc
+---perfProfAcc
+---qmgr
+---sspAcc
+---timeSyncAcc
\---uartAcc
|   +---include
\---usb
|   +---include

```

## 2.7 Threading and Locking Policy

The software release 2.3 access-layer does not implement processes or threads. The architecture assumes execution within a preemptive multi-tasking environment with the existence of multiple-client threads and uses common, real-time OS functions — such as semaphores, task locking, and interrupt control — to protect critical data and procedure sequencing. These functions are not provided directly by the OS, but by the OS abstraction components.

## 2.8 Polled and Interrupt Operation

It is possible to use access-layer components by running the Queue Manager in a polled mode or in an interrupt driven mode of operation. A customer's application code may be invoked by registering with the callback mechanisms provided in the access-layer components. Access-layer components do not autonomously bind themselves to interrupts but generally may be dispatched by an interrupt service routine that is bound to the Queue Manager interrupts. Or, a timer-based task may periodically check the queue manager status and dispatch the access-layer components that are registered to specific queues. Refer to [Chapter 18.0](#) for additional information.

All data path interfaces are executable in the context of both IRQ and FIQ interrupts, though not all operating systems may take advantage of FIQ interrupts in their default configuration.

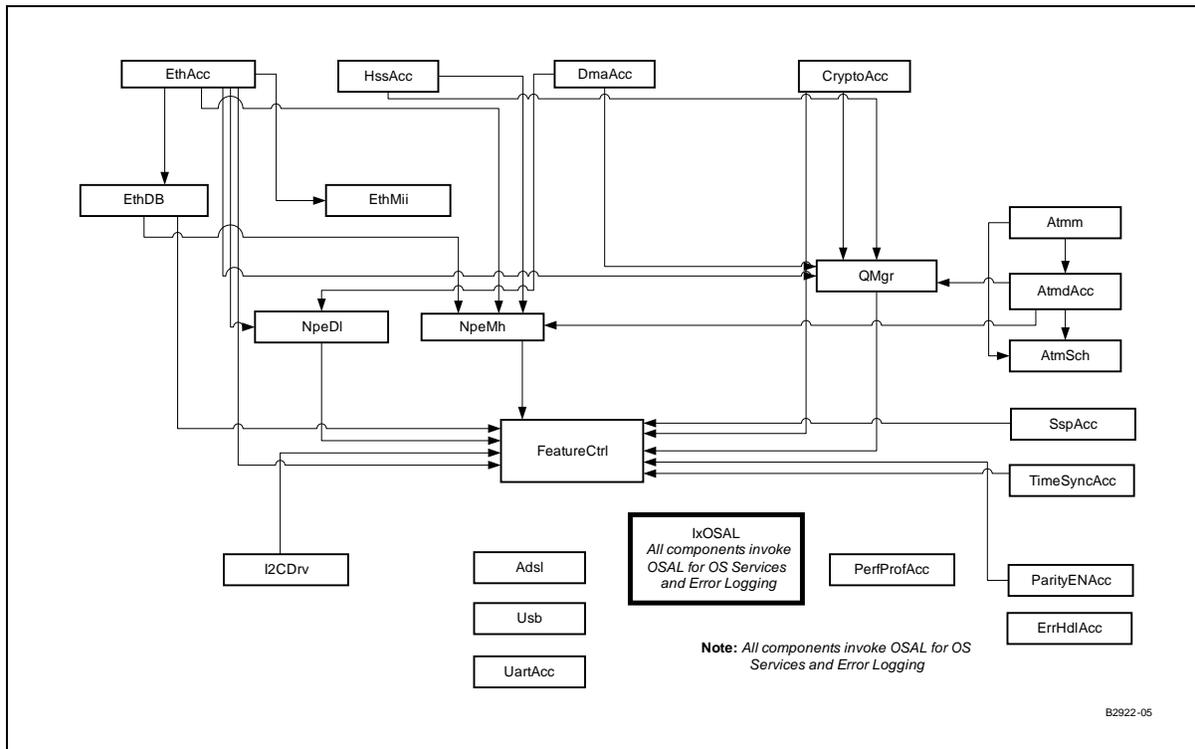
## 2.9 Statistics and MIBs

The software release 2.3 access-layer components only maintain statistics that access-layer clients cannot collect of their own accord. The access-layer components do not provide management interfaces (MIBs). Access-layer clients can use the statistics provided to implement their own MIBs.

## 2.10 Global Dependency Chart

Figure 3 shows the interdependencies for the major APIs discussed in this document.

Figure 3. Global Dependencies





## 3.0 Buffer Management

---

This chapter describes the data buffer system used in Intel® IXP400 Software v2.3, and includes definitions of the IXP400 software internal memory buffers, cache management strategies, and other related information.

### 3.1 What's New

All references to the buffer IXP\_BUF have been changed to IX\_OSAL\_MBUF.

### 3.2 Overview

Buffer management is the general principle of how and where network data buffers are allocated and freed in the entire system. Network data buffers, whose formats are known to all involved components, need to flow between access-layer components.

As shown in [Figure 4 on page 38](#), the IXP400 software access-layer follows a simple buffer-management principle: All buffers used between access-layer component and clients above the access-layer component must be allocated and freed by the clients, that is, in this case, the operating system driver. The client passes a buffer to an access-layer component for various purposes (generally, Tx and Rx), and the access-layer component returns the buffer to the client when the requested job is completed. The access-layer component's Operating System Abstraction Layer module provides the mapping of the OS buffer header fields to the IXP buffer format. Clients can also implement their own utilities to convert their buffers to the IX\_OSAL\_MBUF format and vice-versa. Depending upon the service requested, the NPE modifies the IX\_OSAL\_MBUF's shared structure and hands the buffer back to the access-layer component.

[Figure 4](#) shows different stages where the different fields in the IX\_OSAL\_MBUF buffer is updated at transmit and receive time. Follow the numeric sequence for actions taken during transmission and alphabetic sequence for actions taken during reception.

Step1, A: Both transmission and reception starts with the allocation of an IX\_OSAL\_MBUF and setting the data pointer field to point to the OS\_BUFFER supplied by the client.

Step 2, B: IX\_OSAL\_MBUF:ix\_ne sent to the Queue Manager

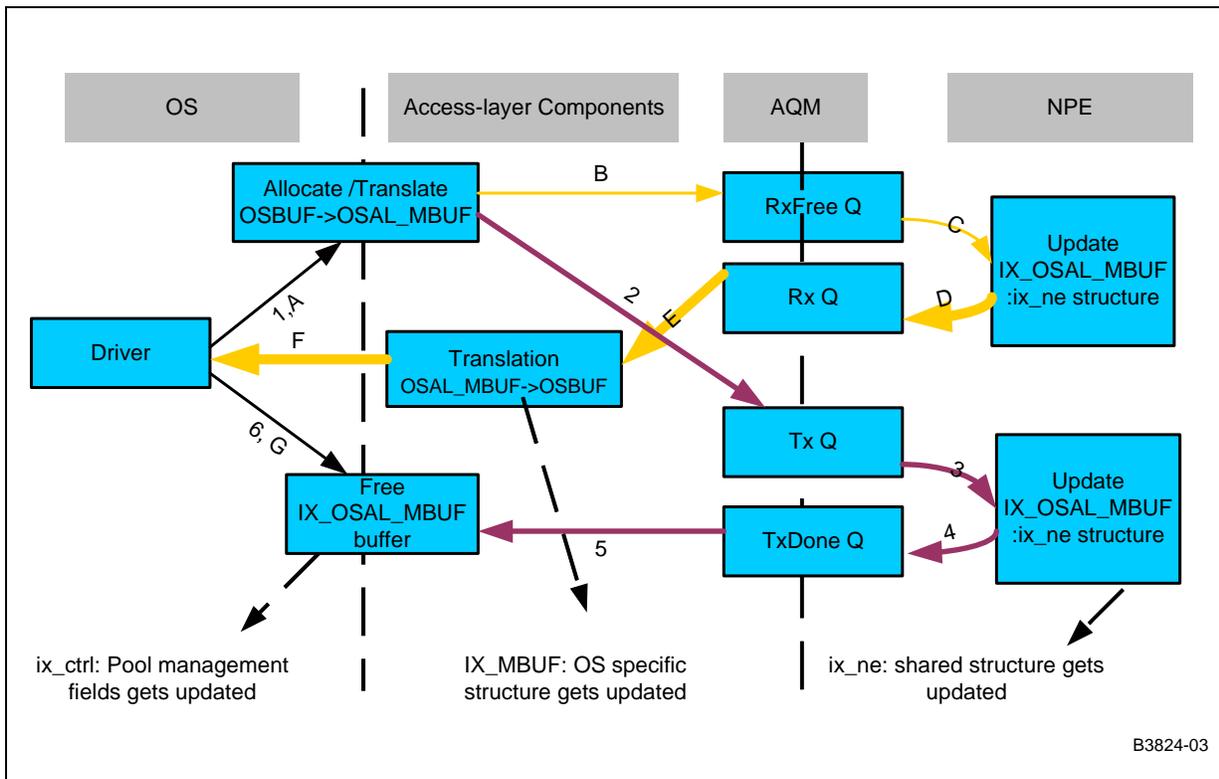
Step 3,4, C,D: IX\_OSAL\_MBUF:ix\_ne handled by NPE

Step 5, E: Access component receives IX\_OSAL\_MBUF:ix\_ne from Queue Manager

Step F: Client extracts the OS\_BUFFER data from IX\_OSAL\_MBUF

Step 6, G: At the end the transmission and reception IX\_OSAL\_MBUF is reused or released to the free pool

Figure 4. Intel® IXP400 Software v2.3 Buffer Flow



The access-layer component may call a client-registered callback function to return the buffer, or may put the buffer back on a free queue for the client to poll. The access-layer components utilize similar buffer management techniques when communicating with the NPEs.

The network data buffers and their formats (as well as management of the buffers), must be 'familiar' to all components so that the buffers can efficiently flow in the system. The IXP400 software uses two internal buffer formats for all network data:

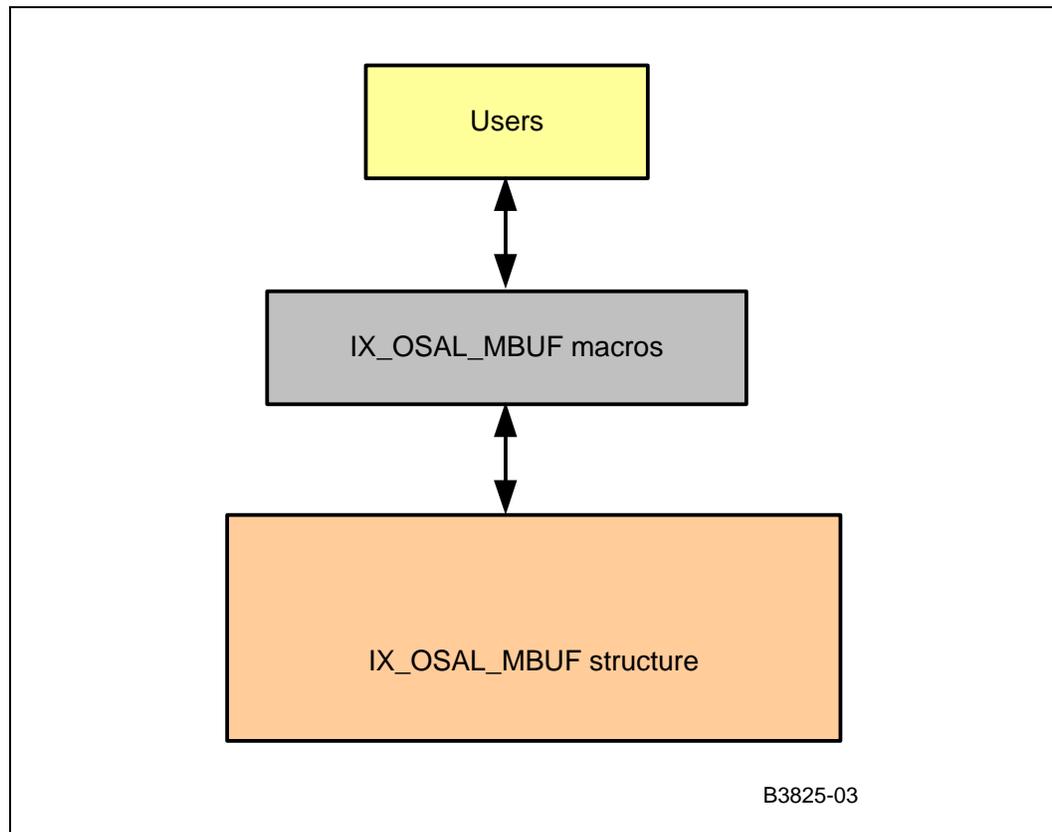
- IX\_OSAL\_MBUF
- raw buffer

These two formats are compatible with the IXP400 software's access-layer components and NPEs.

### IX\_OSAL\_MBUF

The IX\_OSAL\_MBUF is the software release 2.3 defined buffer format used by the access-layer components. As shown in Figure 5, the Operating System Abstraction Layer of software release 2.3 provides the users with macros to read and write the IX\_OSAL\_MBUF fields of the IX\_OSAL\_MBUF buffer. Intel® IXP400 Software v2.3 users are expected to use the IX\_OSAL\_MBUF\_xxx macros provided with the API to access the IX\_OSAL\_MBUF fields.

Figure 5. IX\_OSAL\_MBUF User Interface



The usual fields to be updated between the user and the IX\_OSAL\_MBUF fields depend on the access-layer component, but most of the software release 2.3 API requires the use of following fields:

- IX\_DATA
- IX\_MLEN
- IX\_PKT\_LEN
- IX\_NEXT\_BUFFER\_IN\_PKT\_PTR (in case of chained buffers)

### Raw Buffers

Raw buffer format is a contiguous section of memory represented in one of two ways. One way to pass raw buffers between two access-layer components is through an agreement to circularly access the same piece of raw buffer. One access-layer component circularly writes to the buffer while the other access-layer component circularly reads from the buffer. The buffer length and alignment are parts of the agreement. At run-time, another communication channel is needed to synchronize the read pointer and write pointers between the two components.

The other way to pass raw buffers between two components is through passing a pointer to the buffer between the components. If all buffers are the same size and that size is fixed, the length can be made known during configuration. Otherwise, another communication channel in run-time is needed to tell the length of the buffer. The raw

buffer component is typically used for circuit-switched network data (that is, TDM-based). The access-layer component IxHssAcc channelized service uses raw buffers. Refer to [Section 13.7.2](#) for additional information on raw buffers.

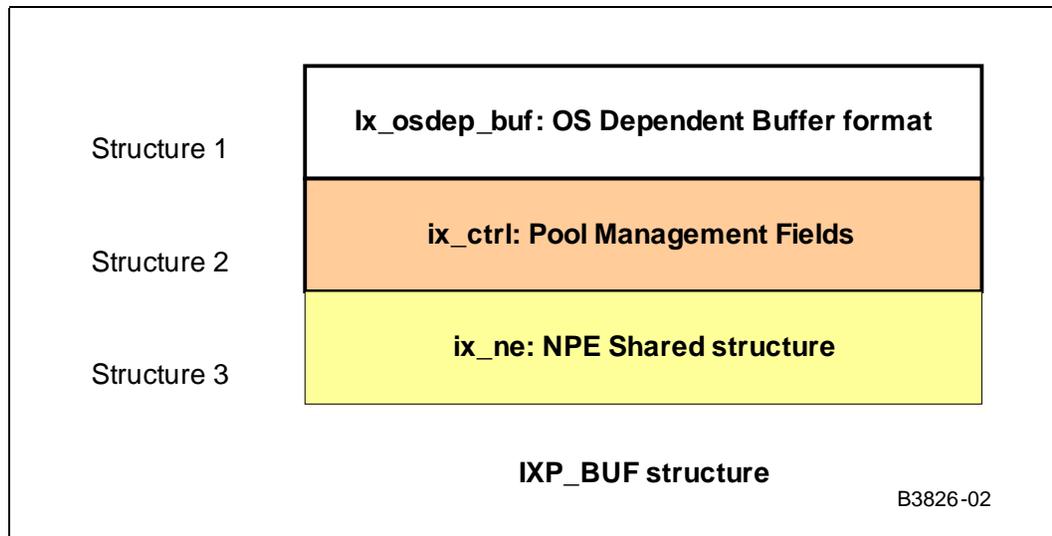
*Note:* Intel® IXP400 Software v2.3 provides OSAL macros, which can be used to allocate memory for raw buffers as a substitute to allocating IX\_OSAL\_MBUF from the pool.

### 3.3 IX\_OSAL\_MBUF Structure

As shown in [Figure 6](#), IX\_OSAL\_MBUF is comprised of the following three main structures, and each structure is comprised of eight entries four bytes long.

1. The first structure consists of an eight word fields some of which are between the OS driver / API users and the access-layer components.
2. The second structure consists of internal fields used by the pool manager, which is provided by the OSAL component.
3. The third structure is the NPE Shared structure that is composed of common header fields and NPE service specific fields. Depending upon the access-component usage, some of the service specific fields such as VLAN tags may be available for the user through use of macros.

**Figure 6. IX\_OSAL\_MBUF Structure**

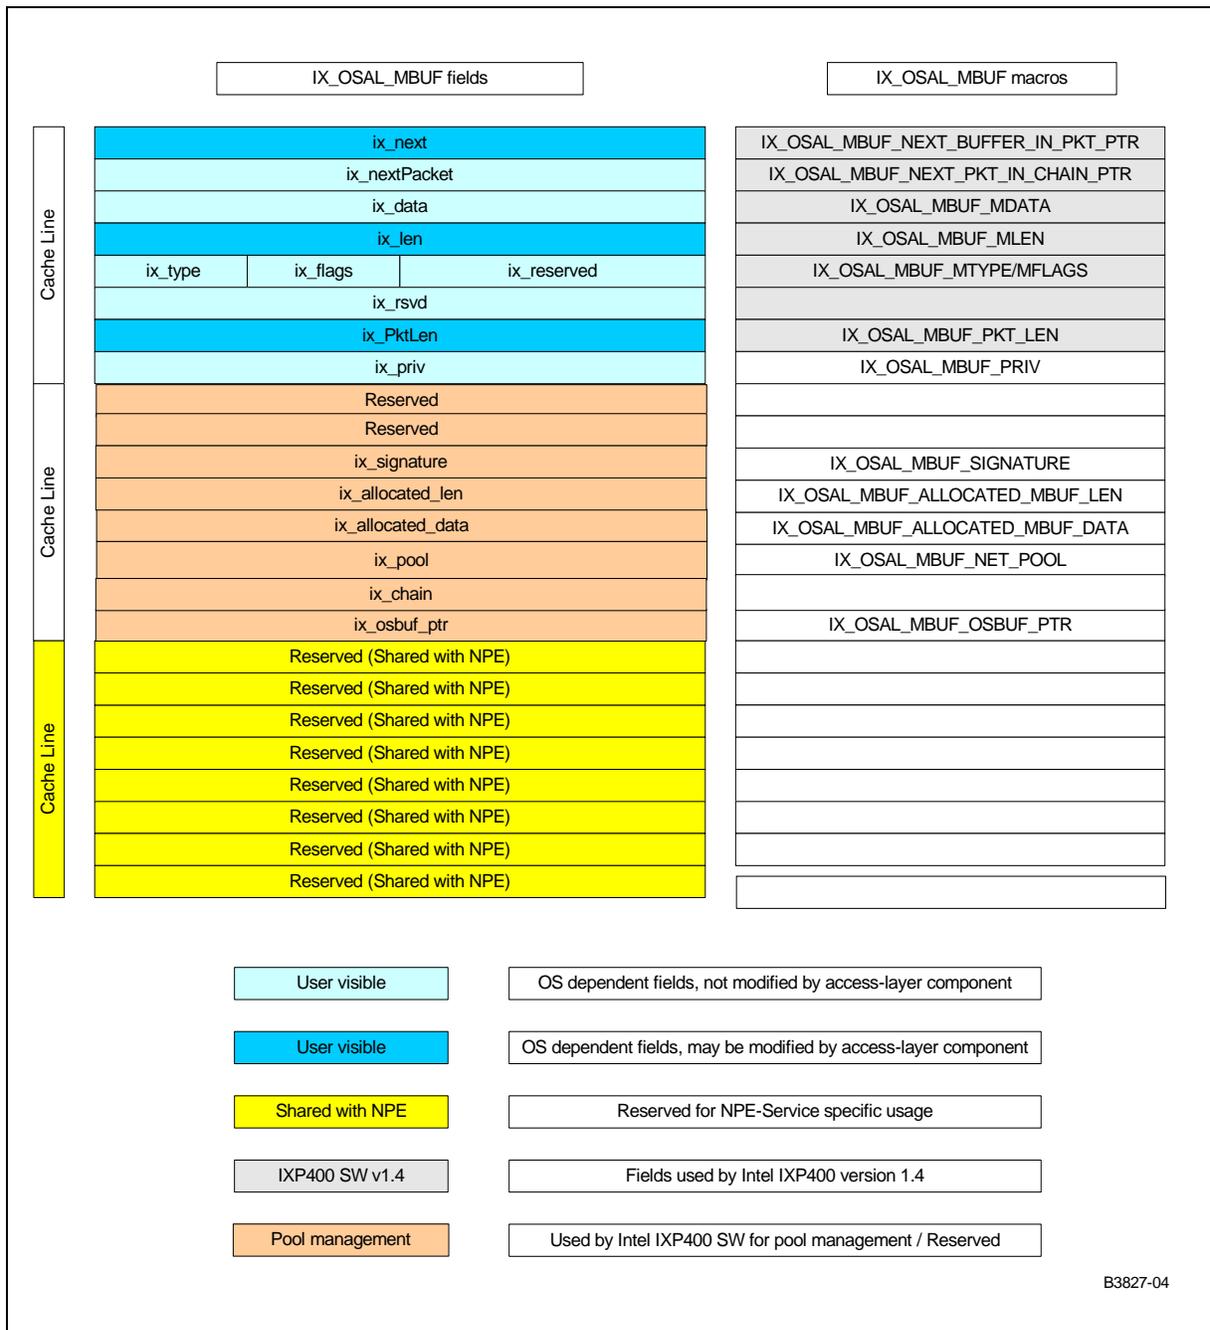


#### 3.3.1 IX\_OSAL\_MBUF Structure and Macros

Users are expected to use the following IX\_OSAL\_MBUF macros provided to access IX\_OSAL\_MBUF subfields. [Figure 7](#) shows macros defined by the OSAL layer component to be used to access the IX\_OSAL\_MBUF fields.

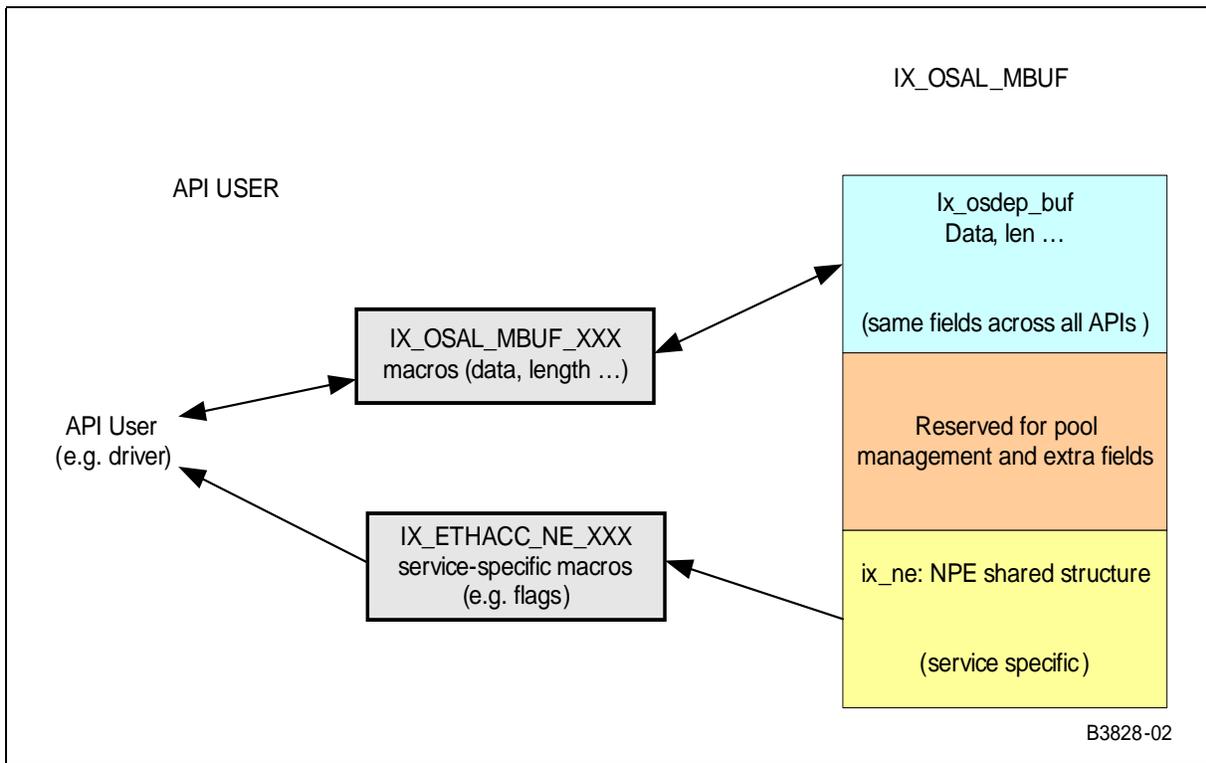


Figure 7. OSAL IX\_OSAL\_MBUF Structure and Macros



Depending upon the usage model, different software components use the structures to update the internal fields of the IX\_OSAL\_MBUF structure. Figure 8 shows a typical interface for the API users or operating system drivers to the IX\_OSAL\_MBUF fields. Depending upon the access-layer components in use the API user may or may not use the service-specific macros to read the NPE-shared structure of the IX\_OSAL\_MBUF fields. Reading of the MAC address or a VLAN tag for a quick classification is an example of NPE-shared structure use.

Figure 8. API User Interface to IX\_OSAL\_MBUF



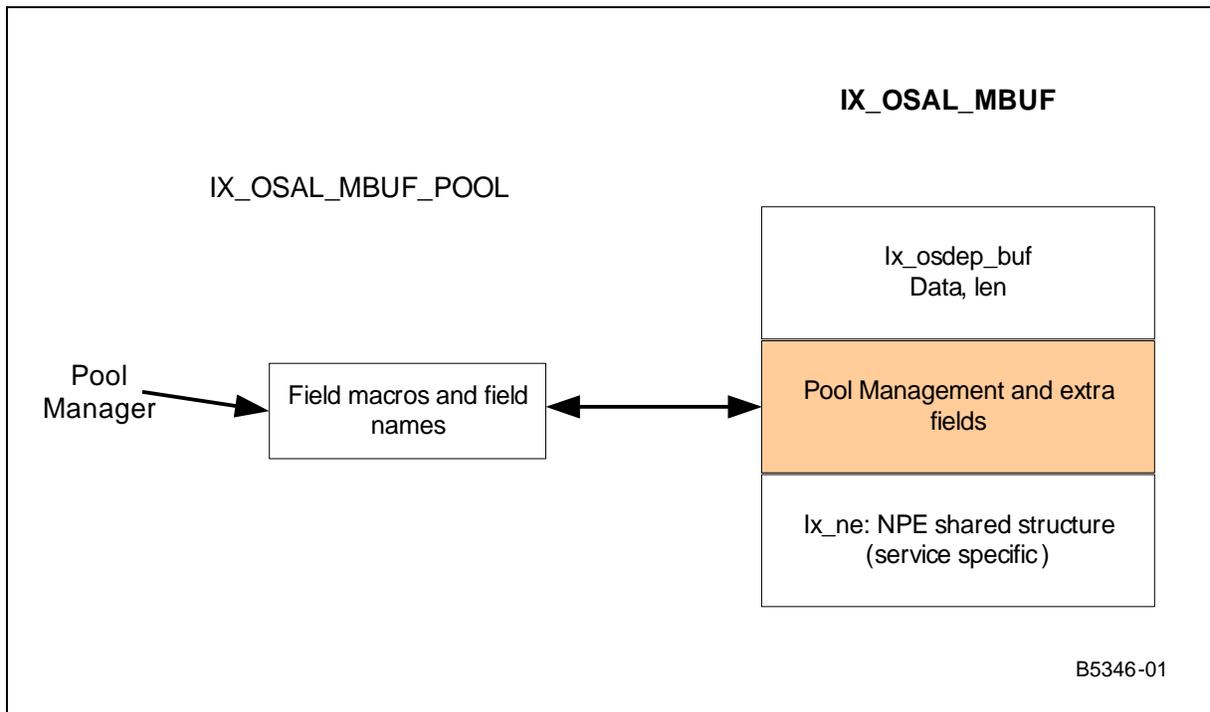
The access-layer components adapt to the endianness as defined by the Intel XScale® processor. The access-layer components can perform reads and write to the ix\_osdep\_buf fields as well as the NPE-shared structure. The service-specific fields to be updated in the NPE-shared structure may vary depending upon access-component needs.

Figure 9 shows the interface between the OSAL pool management module and the pool management fields used for pool maintenance. The pool management field also stores the os\_buf\_ptr field, which is used by the access-layer to retrieve the original pointer to the OS buffer and is set at the time of pool allocation.

**Note:** The ix\_os\_buf\_ptr field is a native OS Buffer pointer to mBlk in VxWorks\* and sk\_buff in Linux\*.



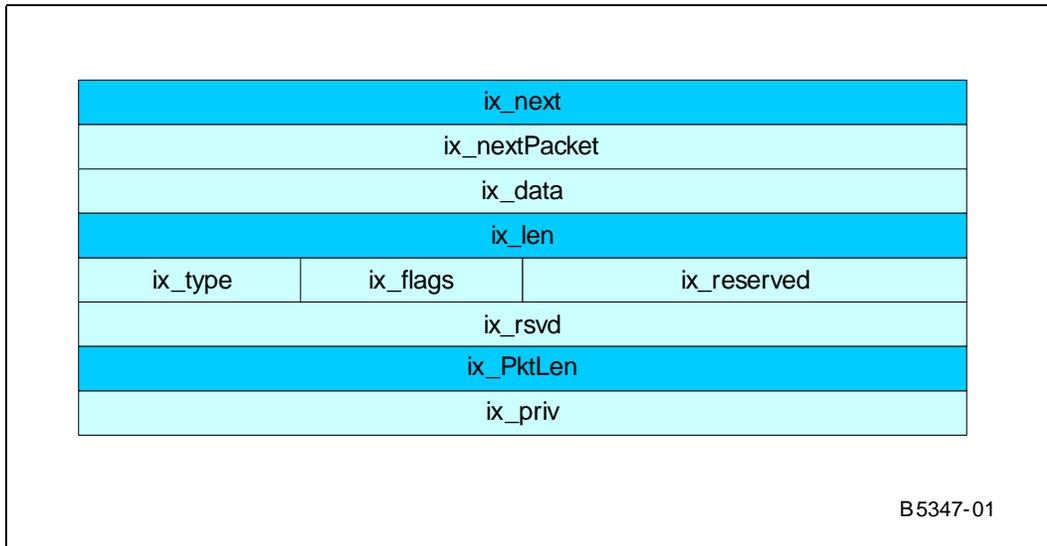
Figure 9. Pool Management Fields



#### `ix_osdep_buf`: OS-Dependent Buffer Format

As shown in Figure 10, the `ix_osdep_buf` information follows a format originally defined in Berkeley Software Distribution (BSD) TCP/IP code distribution to preserve the backward compatibility with previous Intel® IXP400 Software releases. The OSAL layer provides translation functions to map the OS-dependent buffer format to the `ix_osdep_buf` format for Linux\* and VxWorks\* operating systems. This simplifies the buffer management without sacrificing functionality and flexibility.

Figure 10. IX\_OSAL\_MBUF: ix\_osdep\_buf Structure



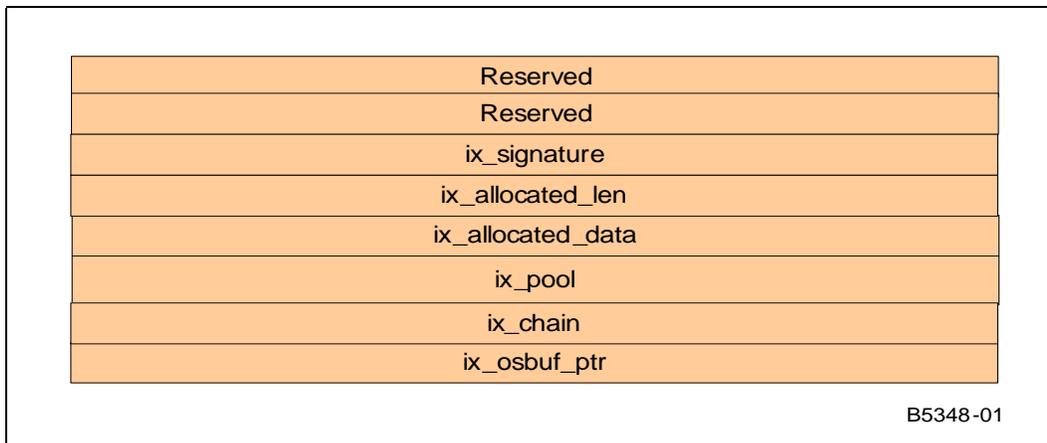
Linux\* utilizes memory structures called skbuffs. The user allocates IX\_OSAL\_MBUF and sets the data payload pointer to the skbuff payload pointer. An os\_buf\_ptr field inside the ixp\_ctrl structure (defined below) of the IX\_OSAL\_MBUF is used to save the actual skbuff pointer. In this manner, the OS buffers are not freed directly by the IXP400 software.

The IXP400 software IX\_OSAL\_MBUF to skbuff mapping is a 'zero-copy' implementation. There is no copy/performance penalty in using Linux\* skbuffs. Other proprietary buffer schemes could also be implemented with the IXP400 software using the mbuf-to-skbuff implementation as an example.

**ix\_ctrl: Intel® IXP400 Software Internal Pool Management Fields**

As shown in Figure 11, the ix\_ctrl fields are set and used by the IX\_OSAL\_MBUF pool manager provided by the OSAL component. Some of the fields can be used for specific purposes for different operating systems, for example, signature verification fields is used in Linux\* when NDEBUD is enabled. The reserved field may be used in VxWorks\* to support IPv6 format.

Figure 11. IX\_OSAL\_MBUF: ix\_ctrl Structure



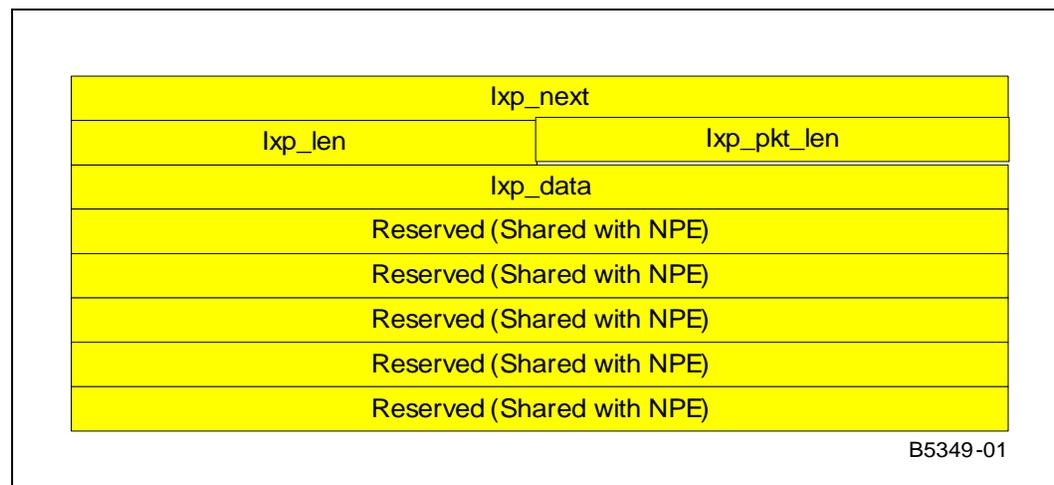


**ix\_ne: IXP400 NPE Shared Structure**

As shown in Figure 12, this structure is provided by the Intel XScale® processor to the NPE. Depending upon the access-layer component usage, some of these fields may be visible to the user through use of macros and also may be altered by the NPE. The lower five words of this structure are defined according to the needs of NPE microcode; therefore, different NPE images may have different structure for this part. The upper three words follows the same structure across all the NPE images.

*Note:* Users should not make any assumptions about usage of the service-specific fields in this NPE-shared structure. The fields are for internal NPE usage only.

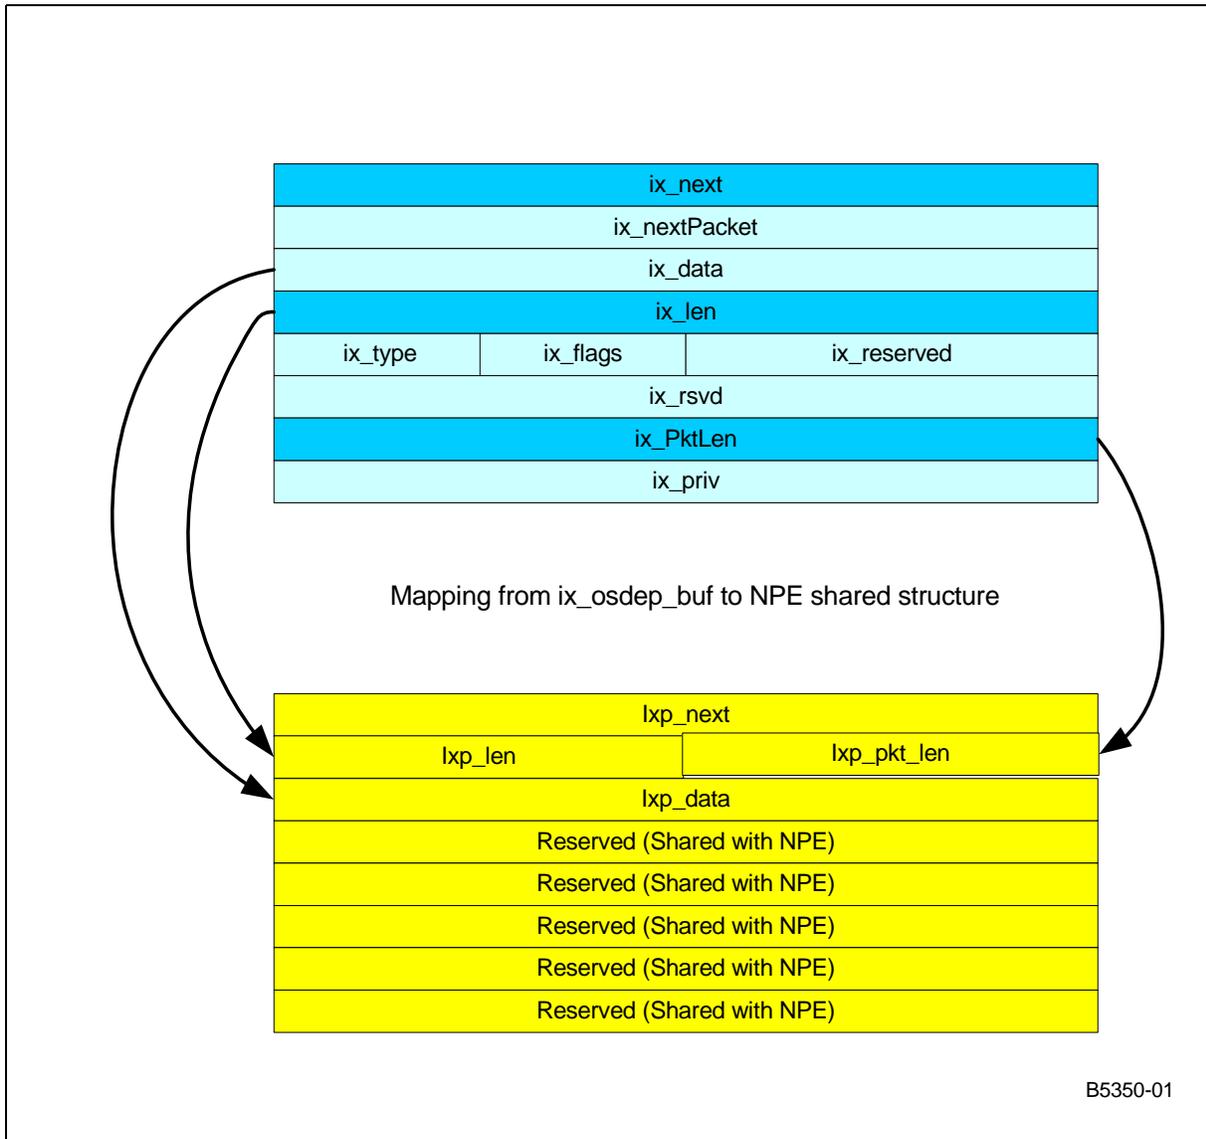
**Figure 12. IX\_OSAL\_MBUF: NPE Shared Structure**



**3.4 Mapping of ix\_osdep\_buf to Shared Structure**

Figure 13 shows an example case on how the ix\_osdep\_buf headers are internally mapped to the NPE shared structure as in the case of the Ethernet and Crypto access-layer components only. The ix\_osdep\_buf standard buffer format is used throughout the access-layer code. In order to minimize overhead in reading the whole buffer control structure from the memory to the NPE while performing NPE-specific services, the pointer to the NPE shared structure is passed to the NPE for processing the data instead of the buffer descriptor pointer itself. Therefore, for the access-layer components, only the required information (such as next buffer pointer, buffer data pointer, buffer length and packet length) from the buffer control structure is copied into NPE shared structure. Depending upon the endianness, the IXP400 software internally swaps the buffers of packetised data and the headers between the upper software layers and the NPEs for the Ethernet and the Crypto access-layer components. It is important to note that NPE shared buffer format used by the IXP400 software is hard-coded in the NPE microcode. It is not possible to change this shared buffer format.

Figure 13. Internal Mapping of ix\_osdep\_buf to the Shared NPE Structure



### 3.5 ix\_osdep\_buf Structure

Table 2 and Table 3 present ix\_osdep\_buf structure format and details.



Table 2. Internal ix\_osdep\_buf Field Format

	0	1	2	3
0	ix_next (IX_OSAL_MBUF_NEXT_BUFFER_IN_PKT_PTR)			
4	ix_nextPacket (IX_OSAL_MBUF_NEXT_PKT_IN_CHAIN_PTR)			
8	ix_data (IX_OSAL_MBUF_MDATA)			
12	ix_len (IX_OSAL_MBUF_MLEN)			
16	ix_type	ix_flags	ix_reserved	
20	ix_rsvd			
24	ix_pktlen			
28	ix_priv(Reserved)			

A set of macros are provided for the IXP400 software to access each of the fields in the buffer structure. Each macro takes a single parameter – a pointer to the buffer itself. Each macro returns the value stored in the field. More detail on the field, their usage, and the macros are detailed in the table below.

*Note:* The data pointer IX\_OSAL\_MBUF\_MDATA could be aligned on a 16 bit boundary to help align an IP header on a 32 bit boundary.

Table 3. ix\_osdep\_buf Field Details (Sheet 1 of 2)

Field / MACRO	Purpose	Used by Access-Layer?
<b>IX_OSAL_MBUF_NEXT_BUFFER_IN_PKT_PTR</b> Parameter type: <i>ix_osdep_buf *</i> Return type: <i>ix_osdep_buf *</i> Description: Returns a 32-bit pointer to the next buffer in the packet	32-bit pointer to the next buffer in a chain (linked list) of buffers. NULL entry marks end of chain.	Yes, where buffer chaining is supported.
<b>IX_OSAL_MBUF_NEXT_PKT_IN_CHAIN_PTR</b> Parameter type: <i>ix_osdep_buf *</i> Return type: <i>ix_osdep_buf *</i> Description: Returns a 32-bit pointer to the first buffer in the next packet in the packet chain	32-bit pointer to the next packet in a chain (linked list) of packets. NULL entry marks end of chain. Each packet in the chain may consist of a chain of buffers.	No. Packet chaining is not supported by IXP400 software.
<b>IX_OSAL_MBUF_MDATA</b> Parameter type: <i>ix_osdep_buf *</i> Return type: <i>char *</i> Description: Returns a pointer to the first byte of the buffer data	32-bit pointer to the data section of a buffer. The data section typically contains the payload of a network buffer.	Yes. But does not get modified by the access-layer. Presently true for Ethernet and Security component.
<b>IX_OSAL_MBUF_MLEN</b> Parameter type: <i>ix_osdep_buf *</i> Return type: <i>int</i> Description: Returns the number of octets of valid data in the data section of the buffer	Lengths (octets) of valid data in the data section of the buffer.	Yes.
<b>IX_OSAL_MBUF_TYPE</b> Parameter type: <i>ix_osdep_buf *</i> Return type: <i>unsigned char</i> Description: Returns the type field of the buffer	Buffer type	Yes, by some components.



Table 3. ix\_osdep\_buf Field Details (Sheet 2 of 2)

Field / MACRO	Purpose	Used by Access-Layer?
<b>IX_OSAL_MBUF_FLAGS</b> Parameter type: <i>ix_osdep_buf *</i> Return type: <i>unsigned char</i> Description: Returns the flags field of the buffer	Buffer flags.	Yes, by some components.
Reserved	Reserved field, used to preserve 32-bit word alignment.	No.
<b>IX_OSAL_MBUF_PKT_LEN</b> Parameter type: <i>ix_osdep_buf *</i> Return type: <i>unsigned int</i> Description: Returns the length of the packet (typically stored in the first buffer of the packet only)	Total length (octets) of the data sections of all buffers in a chain of buffers (packet). Typically set only in the first buffer in the chain (packet).	Yes, where buffer chaining is supported.
Reserved	Used by VxWorks*	No.

### 3.6 Mapping to OS Native Buffer Types

OSAL provides buffer-translation macros for users to translate OS-specific buffer formats to OSAL IXP buffer format and vice versa. The mapping of OS buffer fields to the IXP400 software buffer format is usually done in the OS specific driver component. However, for ease of users the OSAL component provides generic macros for VxWorks\*, and Linux\* operating system that does the translation. Depending upon the build, the OSAL component will translate the macros to its OS-specific implementation. The general syntax for using these macros is as follows:

- IX\_OSAL\_CONVERT\_OSBUF\_TO\_IXPBUF(osBufPtr,ixpBufPtr)
- IX\_OSAL\_CONVERT\_IXPBUF\_TO\_OS\_BUF(ixpBufPtr,osBufPtr)

These macros are intended to replace Linux\* skbuf and VxWorks\* mbuf conversions. Users can also define their own conversion utilities in their package to translate their buffers to IXP buffers (IX\_OSAL\_MBUF).

#### 3.6.1 VxWorks\* M\_BLK Buffer

The first structure ix\_osdep\_buf of the IX\_OSAL\_MBUF buffer format is compatible with VxWorks\* M\_BLK structure. It is also intended to provide a backward compatibility to previous Intel® IXP400 Software releases. For this reason, when compiled for VxWorks\*, the ix\_osdep\_buf buffer format is compatible directly as an M\_BLK buffer. The IXP400 software does not make use of all the fields defined by the M\_BLK buffer. The macros listed in Table 4 are used by the IXP400 software to access the correct fields within the M\_BLK structure.

The M\_BLK structure is defined in the global VxWorks\* header file “netBufLib.h”.

Note that the M\_BLK structure contains many fields that are not used by the IXP400 software. These fields are ignored and are not modified by the IXP400 software.

M\_BLK buffers support two levels of buffer chaining:

- *buffer chaining* — Each buffer can be chained together to form a packet. This is achieved using the **IX\_OSAL\_MBUF\_NEXT\_BUFFER\_IN\_PKT\_PTR** equivalent field in the M\_BLK. This is supported and required by the IXP400 software.
- *packet chaining* — Each packet can consist of a chain of one or more buffers. Packets can also be chained together (to form a chain of chains). **This is not used by the IXP400 software.** The **IX\_OSAL\_MBUF\_NEXT\_PKT\_IN\_CHAIN\_PTR**



equivalent field of the M\_BLK buffer structure is used for this purpose. Most IXP400 software components will ignore this field.

*Note:* The VxWorks\* netMbuf pool library functions will not be supported to allocate and free the IX\_OSAL\_MBUF buffers.

Table 4 shows the field mapping between the ix\_osdep\_buf and the M\_BLK buffer structures through OSAL macros.

**Table 4.** ix\_osdep\_buf to M\_BLK Mapping

ix_osdep_buf	M_BLK
IX_OSAL_MBUF_NEXT_BUFFER_IN_PKT_PTR	mBlkHdr.mNext
IX_OSAL_MBUF_NEXT_PKT_IN_CHAIN_PTR	mBlkHdr.mNextPkt
IX_OSAL_MBUF_MDATA	mBlkHdr.mData
IX_OSAL_MBUF_MLEN	mBlkHdr.mLen
IX_OSAL_MBUF_TYPE	mBlkHdr.mType
IX_OSAL_MBUF_FLAGS	mBlkHdr.mFlags
IX_OSAL_reserved	mBlkHdr.reserved
IX_OSAL_MBUF_NET_POOL	mBlkPktHdr.rcvif
IX_OSAL_MBUF_PKT_LEN	mBlkPktHdr.len
priv	pCIBlk

### 3.6.2 Linux\* skbuff Buffer

The buffer format native to the Linux\* OS is the “skbuff” buffer structure, which is significantly different from the ix\_osdep\_buf buffer format used by the IXP400 software.

The Linux\* skbuf structure is attached to the os\_buf\_ptr field during transmit or receive and is detached during TxDone. The user must allocate an IX\_OSAL\_MBUF header, make a call to a translational function and pass the IX\_OSAL\_MBUF buffer to the IXP400 software release. The translation functions enter all the required fields from the OS buffers to respective fields in the first structure, that is, the ix\_osdep\_buf structure within the IX\_OSAL\_MBUF structure. The translation of fields from the ix\_osdep\_buf structure into the NPE shared structure is accomplished by the OSAL component on Transmit and Receive Replenish. On TxDone the user may recycle the IX\_OSAL\_MBUF back to the IX\_OSAL\_MBUF\_POOL or to an internal data structure.

The OSAL layer provides buffer translation macros for users to translate OS-specific buffer formats to IX\_OSAL\_MBUF buffer format and vice versa.

It works on the following principles:

- Each IX\_OSAL\_MBUF is mapped to an skbuff (1:1 mapping)
- The **os\_buf\_ptr** field of the ix\_ctrl structure is used to store a pointer to the corresponding skbuff.
- The **ix\_data** pointer field of the ix\_osdep\_buf structure within the IX\_OSAL\_MBUF structure is set to point to the **data** field of the corresponding skbuff through use of the IX\_OSAL\_MBUF\_MDATA macro.
- The **ix\_len** and **ix\_pkt\_len** fields of the ix\_osdep\_buf structure within the IX\_OSAL\_MBUF structure is set to the length of the skbuff data section (the **len** field in the skbuff structure) through use of the IX\_OSAL\_MBUF\_PKT\_LEN and IX\_OSAL\_MBUF\_MLEN macros.



The prototype for this function is shown in [Table 5](#).

**Table 5. Buffer Translation Functions**

<ul style="list-style-type: none"><li>• <code>IX_OSAL_CONVERT_OSBUF_TO_IXPBUF(osBufPtr,ixpBufPtr)</code></li></ul> <p>The following fields of <code>ix_osdep_buf</code> within the <code>IX_OSAL_MBUF</code> structure will get updated:</p> <ul style="list-style-type: none"><li>– <code>ix_len</code></li><li>– <code>ix_pktlen</code></li><li>– <code>ix_data</code></li><li>– <code>ix_ctrl.os_buf_ptr</code></li></ul> <ul style="list-style-type: none"><li>• <code>IX_OSAL_CONVERT_IXPBUF_TO_OS_BUF(ixpBufPtr)</code></li></ul> <p>The following fields will get updated in the <code>skbuffer</code></p> <ul style="list-style-type: none"><li>– <code>(skb)osBufPtr = ix_ctrl.os_buf_ptr</code></li><li>– <code>skb-&gt;data = IX_OSAL_MBUF_MDATA(ixMbufPtr)</code></li><li>– <code>skb-&gt;len = IX_OSAL_MBUF_MLEN(ixMbufPtr)</code></li><li>– <code>skb-&gt;len = IX_OSAL_MBUF_PKT_LEN(ixMbufPtr)</code></li></ul>
--

The suggested usage model of this function is:

- Allocate a pool of `IX_OSAL_MBUF` buffer headers. Do not allocate data sections for these buffers.
- When passing a buffer from higher-level software (for example, OS network stack) to the IXP400 software, attach the `skbuff` to an `IX_OSAL_MBUF` using the translation function.
- When receiving an `IX_OSAL_MBUF` passed from the IXP400 software to higher-level software, use the translation function to retrieve a pointer to the `skbuff` that was attached to the `IX_OSAL_MBUF`, and use that `skbuff` with the OS network stack to process the data.

The Intel® IXP400 Software Linux\* Ethernet Device driver (“`ixp425_eth.c`”), which is included in the IXP400 software distribution in form of a patch, contains an example of this suggested usage model.

### 3.7 Intel® IXP400 Software Caching Strategy

The general caching strategy in the IXP400 software architecture is that the software (include Intel XScale® processor-based code and NPE microcode) only concerns itself with the parts of a buffer which it modifies. For all other parts of the buffer, the user (higher-level software) is entirely responsible.

`IX_OSAL_MBUF` buffers typically contain a header section and a data section. The header section contains fields that can be used and modified by the IXP400 software and the NPEs. Examples of such fields are:

- pointer to the data section of the `IX_OSAL_MBUF`
- length of the data section of the `mbuf`
- pointer to the next `mbuf` in a chain of `mbufs`
- buffer type field



- buffer flags field

As a general rule, IXP400 software concerns itself only with IX\_OSAL\_MBUF headers, and assumes that the user (that is, higher-level software) will handle the data section of buffer.

The use of cached memory for IX\_OSAL\_MBUF buffer is strongly encouraged, as it will result in a performance gain as the buffer data is accessed many times up through the higher layers of the operating system's network stack. However, use of cached memory has some implications that need to be considered when used for buffers passed through the IXP400 software Access-Layer.

The code that executes on Intel XScale® processor accesses the buffer memory via the cache in the Intel XScale® processor MMU. However, the NPEs bypass the cache and access this external SDRAM memory directly. This has different implications for buffers transmitted from Intel XScale® processor to NPE (Tx path), and for buffers received from NPE to Intel XScale® processor (Rx path).

### 3.7.1 Tx Path

If a buffer in cached memory has been altered by Intel XScale® processor code, the change will exist in the cached copy of the IX\_OSAL\_MBUF, but may not be written to memory yet. In order to ensure that the memory is up-to-date, the portion of cache containing the altered data must be *flushed*.

The cache flushing strategy uses the following general guidelines:

- The “user” is responsible for flushing the data section of the IX\_OSAL\_MBUF. Only those portions of the data section which have been altered by the Intel XScale® processor code need to be flushed. This must be done **before** submitting an IX\_OSAL\_MBUF to the IXP400 software for transmission via the component APIs (for example, ixEthAccPortTxFrameSubmit()).
- The IXP400 software is responsible for writing and flushing the ix\_ne shared section of the buffer header. This must be done before submitting an IX\_OSAL\_MBUF to the NPE. Communication to the NPEs is generally performed by access-layer components by sending IX\_OSAL\_MBUF headers through the IxQMgr queues.

Since flushing portions of the cache is an expensive operation in terms of CPU cycles, it is not advisable to flush both the header **and** data sections of each IX\_OSAL\_MBUF. To minimize the performance impact of cache-flushing, the IXP400 software only flushes sections that it modifies (the IX\_OSAL\_MBUF header) and leaves the flushing of the data section as the responsibility of the user. You can minimize the performance impact by flushing only what you must.

#### Tx Cache Flushing Example

In the case of an Ethernet bridging system, only the user can determine that it is not necessary to flush any part of the packet payload. In a routing environment, the stack can determine that only the beginning of the mbuf may need to be flushed (for example, if the TTL field of the IP header is changed). Additionally, with the VxWorks\* OS, mbufs can be from cached memory or uncached memory. Only the user knows which buffers need to be flushed or invalidated and which buffers do not.

When the NPE has transmitted the data in a buffer, it will return the buffer back to the Intel XScale® processor. In most cases, the cache copy is still valid because the NPE will not modify the contents of the buffer on transmission. Therefore, as a general rule, the IXP400 software does not invalidate the cached copy of IX\_OSAL\_MBUF used for transmission after they are returned by the NPE.



### 3.7.2 Rx Path

If a buffer has been altered by an NPE, the change will exist in memory but the copy of the buffer in Intel XScale® processor cache may not be up-to-date. We need to ensure that the cached copy is up-to-date by invalidating the portion of cache that contains the copy of the altered buffer data.

The strategy for dealing with data received by the NPEs uses the following general guidelines:

- The “user” is responsible for invalidating the data section of the IX\_OSAL\_MBUF. Again, only the user knows which portions of the data section to access. In some instances, the user may be required to submit free IX\_OSAL\_MBUFs to be used to hold received data (for example, ixEthAccPortRxFreeReplenish()). It is strongly recommended that the cache location holding the data portion of the free IX\_OSAL\_MBUFs be invalidated before submitting them via the API.
- The IXP400 software is responsible for writing and flushing the ix\_ne shared section of the buffer header. The IXP400 software may modify the header of the IX\_OSAL\_MBUF before passing it to the NPE, hence the need to flush and then invalidate the header section of the IX\_OSAL\_MBUF. This should be done before submitting an IX\_OSAL\_MBUF to the NPE for reception (via IxQMgr queues).

*Note:*

In some cases, the Access-Layer will flush the header section of the IX\_OSAL\_MBUF before submitting the IX\_OSAL\_MBUF to the NPE, and will invalidate the header section after receiving it back from the NPE with data. This approach is also acceptable; however, the approach listed above is considered more efficient and more robust.

As in the flushing operations listed in the previous section, invalidating portions of the cache is an expensive operation in terms of CPU cycles. To minimize the performance impact of cache-invalidating, the IXP400 software only invalidates sections that it modifies (the IX\_OSAL\_MBUF header) and leaves the invalidating of the data section as the responsibility of the user. The user can minimize the performance impact by invalidating only what is necessary. When recycling IX\_OSAL\_MBUFs, only the user knows what was the previous use of the IX\_OSAL\_MBUF and the parts of payload that may need to be invalidated.

### 3.7.3 Caching Strategy Summary

Before the NPE reads the memory, ensure that the memory is up-to-date by flushing cached copies of any parts of the buffer memory modified by the Intel XScale® processor.

After the NPE modifies the memory, ensure that the Intel XScale® processor MMU cache is up-to-date by invalidating cached copies of any parts of the buffer memory that the Intel XScale® processor will need to read. It is more robust to invalidate before the NPE gets a chance to write to the SDRAM.

OS-independent macros are provided for both flushing (IX\_OSAL\_CACHE\_FLUSH) and invalidating (IX\_OSAL\_CACHE\_INVALIDATE). For more information, refer to the header file `ixp_osal/include/IxOsal.h`.

§ §



## 4.0 Access-Layer Components: ATM Driver Access (IxAtmdAcc) API

---

This chapter describes the Intel® IXP400 Software v2.3's "ATM Driver-Access" access-layer component.

### 4.1 What's New

NPE A supports the co-existence of ATM MPORT (4 ports) and 16 HSS Channelized voice channels (including 2 HSS Bypass pairs). The NPE image ID for this configuration is **IX\_NPEDL\_NPEIMAGE\_NPEA\_HSSO\_ATM\_MPHY\_4\_PORT**. To include this image with the object code, compile `ixp400` or `libixp400` with the option `"IX_NPE_HSS_MPHY4PORT=1"`. For example, in Linux\*, do

```
make ixp400 IX_NPE_HSS_MPHY4PORT=1
```

and in VxWorks\* do

```
make libixp400 IX_NPE_HSS_MPHY4PORT=1
```

### 4.2 Overview

The ATM access-driver component is the `IxAtmdAcc` software component and provides a unified interface to AAL transmit and receive hardware. The software release 2.3 supports AAL 5, AAL 0, and OAM. This component provides an abstraction to the Intel® IXP4XX Product Line of Network Processors' ATM cell-processing hardware. It is designed to support ATM transmit and receive services for multiple ports and VCs.

This chapter describes the configuration, control, and transmit/receive flow of ATM PDU data through the `IxAtmdAcc` component.

The general principle of improving performance by avoiding unnecessary copying of data is adhered to in this component. The BSD-based buffering scheme is used.

Since AAL 0 is conceptually a raw cell service, the concept of an AAL-0 PDU can be somewhat misleading. In the context of software release 2.3, an AAL-0 PDU is defined as containing an integral number of 48-byte (cell payload only) or 52-byte (cell payload and cell header without HEC field) cells.

### 4.3 IxAtmdAcc Component Features

The services offered by the `ixAtmdAcc` component are:

- Support for the configuration and activation of up to 12 ports on the UTOPIA Level-2 interface.
- Support for the following transmission services on a particular port and VC (PDUs may consist of single or chained `IXP_OSAL_BUFFS`):
  - AAL-5 CPCS with fully formed PDU. AAL-5 CRC calculation is performed by NPE hardware.
  - AAL-0-48 PDU containing an integral number of 48-byte cells.
  - AAL-0-52 PDU containing an integral number of 52-byte cells.



- OAM PDU containing an integral number of 52-byte OAM cells.
- Support for the following reception services on a particular port and VC (PDUs may consist of single or chained IXP\_OSAL\_BUFS):
  - AAL-5 CPCS with fully formed PDU with error detection for CRC errors, priority queuing, and corrupt-packet delivery.
  - AAL-0-48 PDU containing an integral number of 48-byte cells.
  - AAL-0-52 PDU containing an integral number of 52-byte cells.
  - OAM PDU containing an integral number of 52-byte OAM cells with good HEC and CRC10.
- Support for ATM traffic shaping
  - Scheduler registration: Allows registration of ATM traffic-shaping entities on a per-ATM-port basis. A registered scheduler must be capable of accepting per-VC-cell demand notifications from AtmdAcc.
  - Transmission control: Allows ATM traffic-shaping entities to determine when cells are sent and the number of cells sent from each VC at a time.
- Supports the ability to set or view the CLP for AAL-5 CPCS SARed PDUs.
- The component does not process cell headers for AAL-0-52/OAM. Macros have been defined for getting and setting CLP, PTI and GFC in cell headers.
- Allows the client to determine the port on which the PDU was received, for all client service types.
- Supports up to
  - 32 virtual channels for AAL-0/AAL-5 transmit services. Also, support up to 32 channels for AAL-0/AAL-5 Receive services. One client per channel is supported.
  - one dedicated OAM transmit channel (OAM-VC) per port. This channel supports transmission of OAM cells on any VC.
  - one dedicated OAM receive channel (OAM-VC) for all ports. This channel supports reception of OAM cells from any port on any VC.
- Support for coexistence of ATM MPORT (4 ports) and 16 HSS Channelized voice channels (including 2 pairs for HSS Bypass service).

**Note:** To compile the object code with this feature, compile `ixp400` or `libIxp400` with the option “`IX_NPE_HSS_MPHY4PORT=1`”. For example, in Linux\*, do  
`make ixp400 IX_NPE_HSS_MPHY4PORT=1`  
and in VxWorks\* do  
`make libIxp400 IX_NPE_HSS_MPHY4PORT=1`
- Provides an interface to retrieve statistics. These statistics include the number of cells received, the number of cells receive with an incorrect cell size, the number of cells containing parity errors, the number of cells containing HEC errors, and the number of idle cells received.
- Provides an interface to use a threshold mechanism, which allows the client actions to be driven by events or a polling mechanism, through which the client decides where and when to invoke the functions of the interface.
- Supports the possibility to be used in a complete polling environment, or in a complete interrupt environment, or a mixture of both. This is done by providing the control over the Rx and TxDone dispatch functions, transmit and replenish functions. The user may trigger them from interrupts, or poll them, or both, assuming the user provides an exclusion mechanism when needed.

The `ixAtmdAcc` component communicates with the NPEs' ATM-over-UTOPIA component through entries placed on Queue Manager queues, `IXP_OSAL_MBUFs`, and associated descriptors — located in external memory and through the message bus interface.



## 4.4 ATM Background

Table 6 provides a list of acronyms used this chapter.

**Table 6. Acronyms**

Acronym	Description
ATM	Asynchronous transfer mode
HEC	Header Error Control
GFC	Generic flow control
PTI	Payload type identifier
CLP	Cell loss priority
AAL	ATM Adaptation layer
VPI	Virtual path identifier
VCI	Virtual channel identifier
SAR	Segmentation and reassembly
CAC	Call admission control
CBR	QoS class: Constant bit rate
UBR	QoS class: Unspecified bit rate
ABR	QoS class: Available Bit Rate
rt-VBR	QoS class: Real time variable bit rate
nrt-VBR	QoS class: on-real time variable bit rate
PCR	Traffic: Peak cell rate
SCR	Traffic: Sustained cell rate
MBS	Traffic: Maximum burst size
CDVT	Traffic: CDV tolerance
peak to peak CDV	QoS: peak to peak Cell delay variation
MaxCTD	QoS: Cell transfer delay
CLR	QoS: Cell loss ratio

ATM is a connection-oriented data transfer mode based on fast packet switching using 53 byte sized packets called cells. Before sending cells that carry user data, a virtual connection between source and destination has to be established. All packets of a connection follow the same path within the network. During the connection setup, each switch generates an entry in the Virtual Path Identifier (VPI) / Virtual Channel Identifier (VCI) translation table. This enables the switch to move an incoming packet from its VP / VC to corresponding outgoing VP / VC. As an advantage, this kind of routing requires a smaller header. Only a locally valid address (for example, VPI / VCI) must be carried in the packet.

The 53-byte ATM cell consist of 5-byte header and 48-byte data. As shown in [Table 7](#), the ATM header contains information about destination, type, and priority of the cell.

ATM Networks are thought to transmit data with varying characteristics. Various applications need different Qualities of Service (QoS). Some applications, such as telephony, may be very sensitive to delay, but rather insensitive to loss, whereas others (for example, compressed video) are quite sensitive to loss.



**Table 7. ATM Cell Header**

	8	7	6	5	4	3	2	1
Byte1	GFC				VPI			
Byte2	VPI				VCI			
Byte3	VCI							
Byte4	VCI				PTI			CLP
Byte5	HEC							

### 4.4.1 Quality of Service

The ATM Forum specifies several Quality of Service (QoS) categories:

- CBR (Constant Bit Rate)
- rt-VBR (real-time Variable Bit Rate)
- nrt-VBR (non-real-time Variable Bit Rate)
- ABR (Available Bit Rate)
- UBR (Unspecified Bit Rate)

Table 8 shows the negotiated parameters for any QoS category.

**Table 8. Negotiable Traffic and QoS Parameters**

	CBR	rt-VBR	nrt-VBR	UBR	ABR
Traffic Parameters					
PCR & CDVT	Specified				
SCR, MBS, CDVT	NA	Specified		NA	
MCR	NA			Specified	
QoS Parameters					
Peak to Peak CDV	Specified		Unspecified		
MaxCTD	Specified		Unspecified		
CLR	Specified			Specified	Network Specific

### 4.4.2 Adaptation Layers

The ATM Adaptation Layer (AAL) relays ATM cells between ATM Layer and higher layer. When relaying information received from the higher layers, it segments the data into ATM cells. When relaying information received from the ATM Layer, it must reassemble the payloads into a format the higher layers can understand. This operation is called Segmentation and Reassembly (SAR), and is the main task of AAL. Various AALs were defined in supporting various traffic or service expected to be used.

The service classes and the corresponding types of AALs are as follows:



**Class A** - Constant Bit Rate (CBR) service: AAL1 supports a connection-oriented service in which the bit rate is constant. Examples of this service include 64 Kbit/sec. voice, fixed-rate uncompressed video and leased lines for private data networks.

**Class B** - Variable Bit Rate (VBR) service: AAL2 supports a connection-oriented service in which the bit rate is variable but requires a bounded delay for delivery. Examples of this service include compressed packetized voice or video. The requirement on bounded delay for delivery is necessary for the receiver to reconstruct the original uncompressed voice or video.

**Class C** - Connection-oriented data service: For connection-oriented file transfer and in general data network applications where a connection is set up before data is transferred, this type of service has variable bit rate and does not require bounded delay for delivery. Two AAL protocols that were defined to support this service class have been merged into a single type and are called AAL3/4. But with its high complexity, the AAL5 protocol is often used to support this class of service.

**Class D** - Connectionless data service: Examples of this service include datagram traffic and in general, data network applications where no connection is set up before data is transferred. AAL3/4 or AAL5 can be used to support this class of service.

**Operation Administration and Maintenance (OA&M)** - OA&M is defined for supervision, testing, and performance monitoring. It uses loop-back for maintenance and ITU TS standard CMIP, and is organized into five hierarchical levels: Virtual Channel (F5 - Between VC endpoints), Virtual Path (F4- Between VP endpoints), Transmission Path (F3- Between elements that perform assembling, disassembling of payload, header, or control), Digital Section (F2 Between section end-points, performs frame synchronization) and Regenerator Section (F1- Between regeneration sections).

**AAL0 PDU:**

AAL0 payload consists of 48 bytes without special field, is also referred to as raw cells.

AAL0 payload with 52 bytes consists of 48 byte data and 7 byte header without the HEC.

**AAL5 CPCS PDU:**

AAL5 is a simple and efficient AAL (SEAL - Simple and Efficient Adaptation Layer), and is the one used most for data traffic. It has no per-cell length nor per-cell CRC fields. AAL 5 is the most widely used ATM Adaptation Layer Protocol. The CPCS PDU is broken up into a 48-byte ATM cell payload. See [Table 9](#) for the format of a CPCS PDU.

**Table 9. AAL5 CPCS PDU**

1-65535 bytes	0-47	1	1	2	4 bytes
PDU Payload	PAD	UU	CPI	LI	CRC-32

PDU payload - Variable length user information field. Note information comes before any length indication.

PAD - Padding used to cell align the trailer, and is between 0 and 47 bytes long.

UU - CPCS user-to-user indication to transfer one byte of user information.

CPI - Common Part Indication

LI - Length indicator.



**OA&M cells:**

These use pre-defined (reserved) VPI/VCI numbers. See [Table 10](#) and [Table 11](#) on [page 58](#) for the reserved vpi/vci pairs and payload format, respectively.

**Table 10. OAM Reserved VPI/VCI and Their Function**

VPI	VCI	Function
0	0	Idle cells
0	1	Meta signalling (default)
non-zero	1	Meta signalling
0	2	General broadcast signalling (default)
non-zero	2	General broadcast signalling
0	5	Point-to-point signalling (default)
non-zero	5	Point-to-point signalling
non-zero	3	Segment OAM F4 flow cell
non-zero	4	End-to-End OAM F4 flow cell
non-zero	6	RM-VPC cells for rate management
0	15	SPANS
0	16	ILMI
0	18	PNNI Signalling

**Table 11. OAM Payload Format**

4 bits	4 bits	45 bytes	6 bits	10 bits
OAM Type	Function Type	Function Spec	Reserve	CRC-10

**OAM Type / Function Type:**

The possible values for OAM type and function type are listed in [Table 12](#).

**Table 12. OAM Type and Function Type**

OAM Type	Function Type
Fault Mgmnt (0001)	Alarm Indication Signals (0000)
	Far End Receive Failure (0001)
	OAM Cell Loopback (1000)
	Continuity Check (0100)
Performance Mgmnt (0010)	Forward Monitoring (0000)
	Backward Reporting (0001)
	Monitoring & Reporting (0010)
Activation/ Deactivation (1000)	Performance Monitoring (0000)
	Continuity Check (0001)

**OAM F4** cells operate at the VP level. They use the same VPI as the user cells; however, they use two different reserved VCIs, as follows:



VCI=3 Segment OAM F4 cells.

VCI=4 End-end OAM F4 cells.

**OAM F5** cells operate at the VC level. They use the same VPI and VCI as the user cells. To distinguish between data and OAM cells, the PTI field is used as follows:

PTI=100 (4) Segment OAM F5 cells processed by the next segment.

PTI=101 (5) End-to-end OAM F5 cells, and are only processed by end stations terminating an ATM link.

CRC-10 - Cyclic redundancy check calculated over the SAR header.  $G(x) = x^{10} + x^9 + x^5 + x^4 + x + 1$

## 4.5 Configuration Services

IxAtmdAcc supports three configuration services:

- UTOPIA port configuration
- ATM traffic shaping
- VC configuration

Table 13 presents a list of configuration routines.

**Table 13. List of Configuration Routines**

Name	Function
ixAtmdAccInit	Initialize the component
ixAtmdAccUninit	Uninitialize the component
ixAtmdAccUtopiaConfigSet	Download the configuration structure to the UTOPIA interface
ixAtmdAccUtopiaConfigReset	Reset the configuration structure to the UTOPIA interface
ixAtmdAccUtopiaStatusGet	Get the UTOPIA interface configuration
ixAtmdPortEnable	Enable a PHY logical port
ixAtmdPortDisable	Disable a PHY logical port
ixAtmdAccTxVcTryConnecct	Connect to a AAL PDU transmit service for a particular port/VPI/VCI
ixAtmdAccTxVcTryDisconnecct	Disconnect to a AAL PDU transmit service for a particular port/VPI/VCI
ixAtmdAccPortTxScheduledModeEnable	Put the port into a scheduled mode
ixAtmdAccPortTxScheduledModeDisable	Put the port into a unscheduled mode
ixAtmdAccPortRxVcConnect	Connect to a AAL PDU receive service for a particular port/VPI/VCI
ixAtmdAccPortRxVcDisconnect	Disconnect to a AAL PDU receive service for a particular port/VPI/VCI

### 4.5.1 UTOPIA Port-Configuration Service

The UTOPIA interface is the Intel® IXP4XX product line processors' interface by which ATM cells are sent to and received from external PHYs. In order to configure the UTOPIA interface, IxAtmdAcc provides an interface that allows a configuration structure to be sent to and/or retrieved from the UTOPIA interface.



IxAtmdAcc provides the interface to configure the hardware and enable/disable traffic on a per-port basis.

### 4.5.2 ATM Traffic-Shaping Services

An ATM scheduling entity provides a mechanism where VC traffic on a port is shaped in accordance with its traffic parameters. IxAtmdAcc does not itself provide such a traffic-shaping service, but can be used in conjunction with external scheduling services.

The scheduler registration interface allows registration of ATM traffic-shaping entities on a per-port basis. These entities, or proxies thereof, are expected to support the following callbacks on their API:

- A function to exchange VC identifiers. A VC identifier identifies a port, VPI, VCI, and is usually specific to ATM layer interface. IxAtmdAcc has identifier known as a `connId` and the scheduling entity is expected to have its own identifier known as a `vcId`. This callback also serves to allow the scheduling entity to acknowledge the presence of VC.
- A function to update the scheduling entity on a per VC basis with the number of ATM cells. This function is used every time the user submits a new PDU for transmission.
- A function to clear the cell count related to a particular VC. This function is used during a disconnect to stop the scheduling services for a VC.

No locking or mutual exclusion is provided by the IxAtmdAcc component over these registered functions.

The transmission-control API expects to be called with an updated transmit schedule table on a regular basis for each port. This table contains the overall number of cells, the number of idle cells to transmit, and — for each VC — the number of cells to transmit to the designated ATM port.

The ATM Scheduler can be different for each logical port and the choice of the ATM scheduler is a client decision. ATM scheduler registrations should be done before enabling traffic on the corresponding port. Once registered, a scheduler cannot be unregistered. If no ATM scheduler is registered for one port, transmission for this port is done immediately.

### 4.5.3 VC-Configuration Services

IxAtmdAcc provides an interface for registering VCs in both Tx and Rx directions. The ATM VC is identified by a logical PHY port, an ATM VPI, and an ATM VCI. The total number of ATM AAL-5 or AAL-0 VCs supported — on all ports and in both directions — is 32. IxAtmdAcc supports up to 32 Rx channels, and up to 32 Tx channels on all ports. For AAL-5 and AAL-0, the number of logical clients supported per-VC is one.

In addition to the 32 VCs mentioned above, one dedicated OAM transmit VC per port and one dedicated OAM receive VC are supported. These dedicated OAM VCs behave like an **OAM interface** for the OAM client, and are used to carry OAM cells for any VPI/VCI (even if that VPI/VCI is one of the 32 connected for AAL services).

In the TX direction, the client has to register the ATM traffic characteristics to the ATM scheduler before invoking the IxAtmdAcc connect function. The `ixAtmdAccTxVcConnect` function performs the following actions:

- Check that the port has been configured and is UP. Otherwise, the fail status is returned to the caller.



- Check whether the ATM VC is already in use in an other TX connection. Otherwise, the fail status is returned to the caller.
- Check the registration of this VC to the registered ATM Scheduler.
- Bind the VC with the scheduler associated with this port.
- Register the transmit done callback where the transmitted buffers (used to store the AAL5 PDU) get recycled.
- Register the notification callback and the hardware asks for more data to transmit.
- Allocate a connection ID and return it to the client.

In the RX directions, the ixAtmdAccRxVcConnect steps involve the following actions:

- Check the PHY port is enabled.
- Check for ATM VC already in use in an other RX connection.
- Register the callback and the received buffers (contains AAL5 PDU) get pushed into the client's protocol stack.
- Register the notification callback and the hardware asks for more available buffers.
- Allocate a connection ID and return it to the client.

When connecting, a connection ID is allocated and must be used to identify the VC, in all calls to the API. The connection IDs for Receive and Transmit, on the same ATM VC, are different.

The client has the choice of using a threshold mechanism provided by IxAtmdAcc or polling the various resources. When using the threshold mechanism, the client must register a callback function and supply a threshold level. As a general rule, when configuring threshold values for various services, the lower the threshold value is, the higher the interrupt rate is.

#### 4.5.4 Uninitialize ATM Driver

IxAtmdAcc provides an interface for deallocating/destroying the allocated buffers or mutexes done during the initialization. This function must be called to free the allocated buffers. The uninit function internally calls various other uninitialization functions, based on requirement.

## 4.6 Transmission Services

The IxAtmdAcc transmit service currently supports AAL 5, AAL 0-48, AAL 0-52, and OAM only and operates in scheduled mode.

In the scheduled mode, the fully formed AAL5 PDU stored in buffers are accepted and internally queued in IxAtmdAcc until they are scheduled for transmission by a scheduling entity. The scheduling entity determines the number cells to be transmitted from a buffer at a time, this allows cells from various VCs to be interleaved on the wire.

IxAtmdAcc accepts outbound ATM payload data for a particular VC from its client in the form of chained IX\_OSAL\_MBUF. An IX\_OSAL\_MBUF chain represents a PDU and can contain 1-65535 payload octets. A PDU is however a multiple of 48 octets when padding and the AAL5 trailer are included.

The submission rate of buffers for transmission should be based on the traffic contract for the particular VC and is not known to IxAtmdAcc. There are a maximum number of buffers that IxAtmdAcc can hold at a time and a maximum number of buffers that the underlying hardware can hold before and during transmission. This maximum is guaranteed to facilitate the port rate saturation at 64byte packets.



Under the ATM Scheduler control (Scheduled Mode), IxAtmdAcc interpret the schedule table, build and send requests to the underlying hardware, to be segmented into 48 byte SAR payloads and transmitted with ATM cell headers over the UTOPIA bus.

Once the transmission is complete, IxAtmdAcc passes back the IX\_OSAL\_MBUFs to its client (on a per connection basis). The client can free them and return them to the IX\_OSAL\_MBUF pool. The preferred option is to reuse the buffers during the next traffic. Processing of transmit done buffers from IxAtmdAcc is controlled by the client.

Transmit Done is a system-wide entity and provides a service to multiple ports. A system using multiple ports with various transmit activity results in latency effects for low-activity ports. The user must tune the number of buffers needed to service a low rate port or channel if the overall user application involves a port configured with a VC supporting a very different traffic rate. This tuning is at the clients discretion and hence is beyond the scope of this document

In the case of OAM, a PDU containing OAM cells for any port, VPI, or VCI must be submitted for transmission on the dedicated OAM-VC for that port. This is true regardless of whether an AAL-5/AAL-0-48/AAL-0-52 transmit service connection exists for the given VPI or VCI. The dedicated OAM-VC is scheduled just like any other VC.

Table 14 presents a list of transmission routines.

Table 14. List of Transmission Routines

Name	Function
ixAtmdAccTxVcConnect	Connect to a AAL PDU transmit service for a particular port/vpi/ vci
ixAtmdAccTxVcPduSubmit	Submit a PDU for transmission on connection
ixAtmdAccTxVcTryDisconnect	Disconnect from AAL PDU transmit service for a particular port/VPI/VCI
ixAtmdAccPortTxFreeEntriesQuery	Get the number of scheduled cells pending transmission to the hardware
ixAtmdAccPortTxCallbackRegister	Configure the Tx port threshold value and register a callback to handle threshold notifications
ixAtmdAccPortTxScheduledModeEnable	Put the port into scheduled mode
ixAtmdAccPortTxProcess	Transmit queue cells to the HW based on the supplied queue table
ixAtmdAccTxDoneDispatch	Process a number of pending transmit done buffers from the hardware
ixAtmdAccPortTxDoneLevelQuery	Query the current number of transmit IX_OSAL_MBUFs reading for recycling from the hardware
ixAtmdAccTxDoneDispatcherRegister	Configure the TxDone stream threshold value and register a callback to handle threshold notifications
ixAtmdAccTxDoneDispatcherUnregister	Unregister a callback to handle threshold notifications
ixAtmdAccPortTxCallbackUnregister	This function unregisters a callback to handle threshold notifications

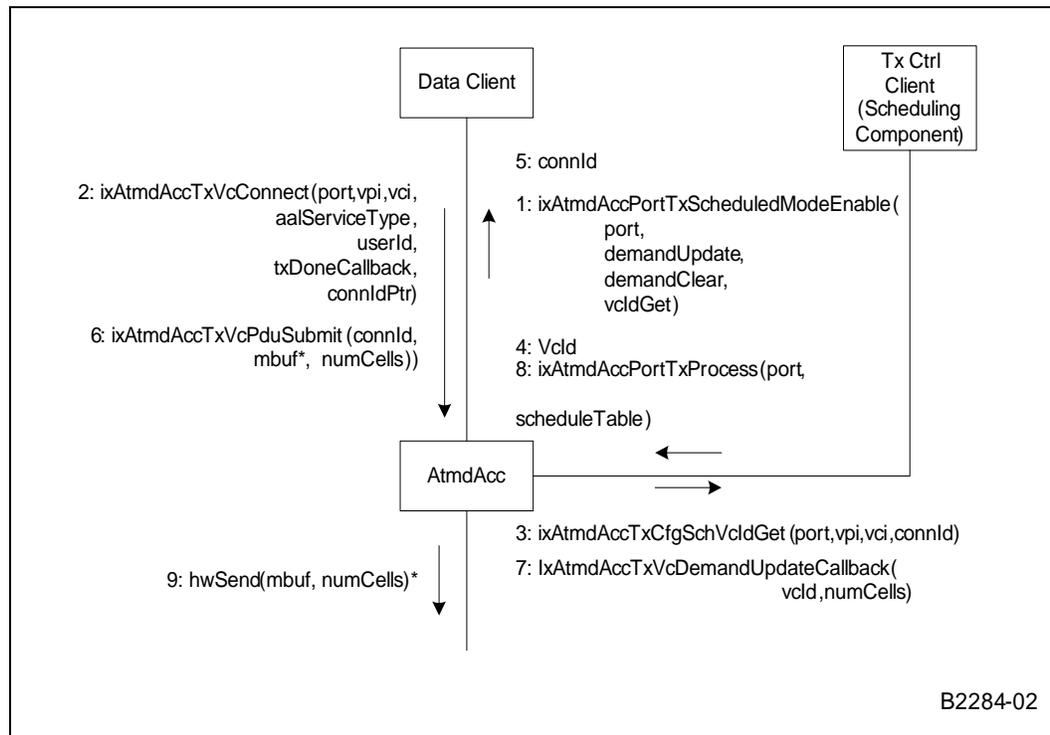
### 4.6.1 Scheduled Transmission

The scheduling entity controls the VC, from which cells are transmitted and when they are transmitted. Buffers on each VC are always sent in the sequence they are submitted to IxAtmdAcc. However, cells from various VCs can be interleaved.

Figure 14 on page 63 shows VC connection and buffer transmission for a scheduled port.



Figure 14. AAL5 PDU Transmission for a Scheduled Port



1. A control client wants to use an ATM traffic shaping entity that controls the transmission of cells on a particular port, ensuring VCs on that port conform to their traffic descriptor values. The client therefore calls `ixAtmdAccPortTxScheduledModeEnable ()` passing the port and some callback functions as parameters. `IxAtmdAcc` has no client connections active for that port and accepts to scheduler registration.
2. Later a data client wants to use the `IxAtmdAcc` AAL5 transmit service for a VC on the same port, and therefore calls `ixAtmdAccTxVcConnect()`.
3. `IxAtmdAcc` invokes the callback function `IxAtmdAccTxCfgSchVclIdGet ()` that was registered for the port, and supplies the `AtmScheduler` VC identifier as a parameter to request permission to send traffic on this VC.
4. The shaping entity acknowledges the validity of the VC, stores the `IxAtmdAcc` connection ID and issues a `VclId` to `IxAtmdAcc`.
5. `IxAtmdAcc` accepts the connection request from the data client and returns a connection ID to be used by the client in further `IxAtmdAcc` API calls for that VC.
6. Some time later the data client has a fully formed AAL5 PDU in an `IX_OSAL_MBUF` ready for transmission on the AAL5 service (The CRC in the AAL5 trailer is not pre-calculated). The client calls `ixAtmdAccTxVcPduSubmit()` passing the `IX_OSAL_MBUF` and numbers of cells contained in the chained `IX_OSAL_MBUF` as parameters.
7. `IxAtmdAcc` ensures the connection is valid and submits new demand in cells to the shaping entity by calling `ixAtmdAccTxVcDemandUpdate` callback function. The shaping entity accepts the demand and `IxAtmdAcc` internally enqueues the `IX_OSAL_MBUFs` for later transmission.
8. The traffic shaping entity decides at certain time, by its own timer mechanism or by using the **Tx Low Notification** service provided by `IxAtmdAcc` component for this



port, that cells should be transmitted on the port based on the demand it has previously obtained from AtmdAcc. It creates a transmit schedule table and passes it to the IxAtmdAcc by calling ixAtmdAccPortTxProcess().

9. IxAtmdAcc takes the schedule, interprets it, and sends scheduled cells to the hardware. In the case of hardware queue being full (only possible if the **Tx Low Notification** service is not used), the ixAtmdAccPortTxProcess call returns an overloaded status, so that the traffic shaping entity can retry this again later.

#### 4.6.1.1 Schedule Table Description

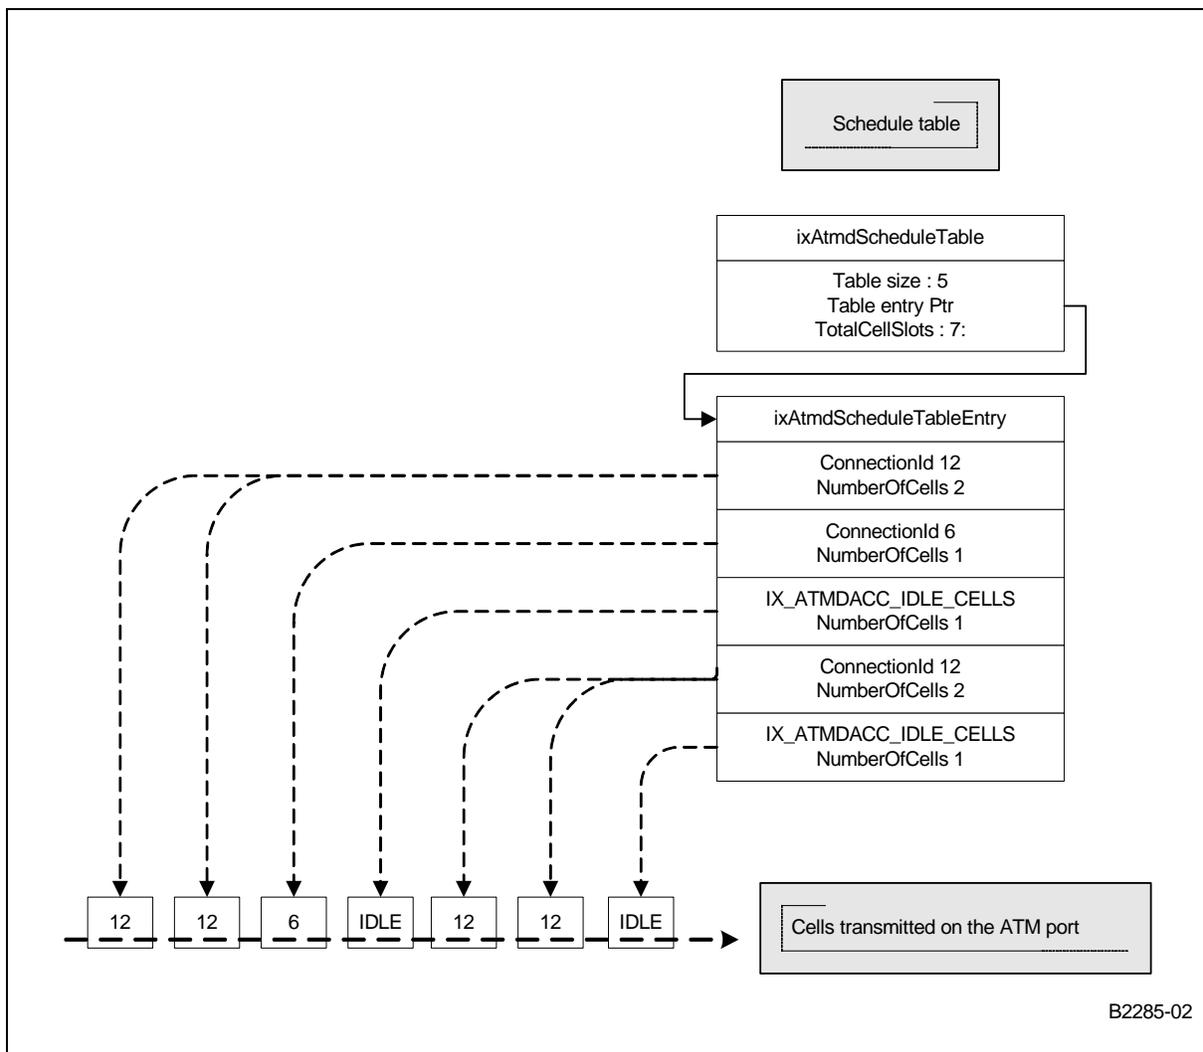
IxAtmdAcc uses a schedule table when transmitting cell information to the hardware. This schedule table drives the traffic on one port.

The schedule table is composed of an array of table entries, and each specifies a ConnectionID and the number of cells (up to 16) to transmit from that VC. Idle cells are inserted in the table with the ConnectionID identifier set to IX\_ATMDACC\_IDLE\_CELLS.

[Figure 15 on page 65](#) shows how this table is translated into an ordered sequence of cells transmitted to one ATM port.



Figure 15. IxAtmdAccScheduleTable Structure and Order Of ATM Cell



#### 4.6.2 Transmission Triggers (Tx-Low Notification)

In Scheduled Mode, the rate and exact point at which the `ixAtmdAccTxProcess()` interface should be called by the shaping entity is at the clients discretion, and hence is beyond the scope of this document. However, `ixAtmdAcc` transmit service does provide a Tx Low Notification service, which can be configured to execute a client supplied notification callback when the number of cells not yet transmitted by the hardware reaches a certain low level. The service only supports a single client per port and the default threshold value is 0, where there are no cells to be transmitted.

The choice to implement a threshold based on a number of cells is driven by the following facts:



A Tx rate is expressed in cells per second, not in PDUs per second. Expressing a rate in PDU units, which can have any length (from a few bytes up to many kilobytes), is valid only as a long-term average and does not take into account the instant rate. The more granularity in the unit, the better the throughput can be configured.

When many channels need to share the same hardware interface, the natural division of big packets (shaping) is in cells. Breaking big packets into cells allows the user to tune the transmit rate and txLow notification rate efficiently. It also allows txLow notification rate to be set to a predictable and constant value.

A threshold expressed in cells allows the user to manage the maximum latency of the system. The latency can be estimated to be (threshold / cell rate). For example, if the cell rate is 1800 cells / sec, and the threshold is 4, the latency is  $4 / 1800 = 2.2$  milliseconds. The user can decide the way to build the schedule table. If the schedule table comprises of 5 cells, the user is notified after exactly 5 cells are sent (5.5 milliseconds).

#### 4.6.2.1 Transmit-Done Processing

When buffers have been sent on a port, they are placed in a single, transmit-complete stream, which is common to all ports. IxAtmdAcc does not autonomously process this stream — the client, instead, deciding when and how many buffers are processed.

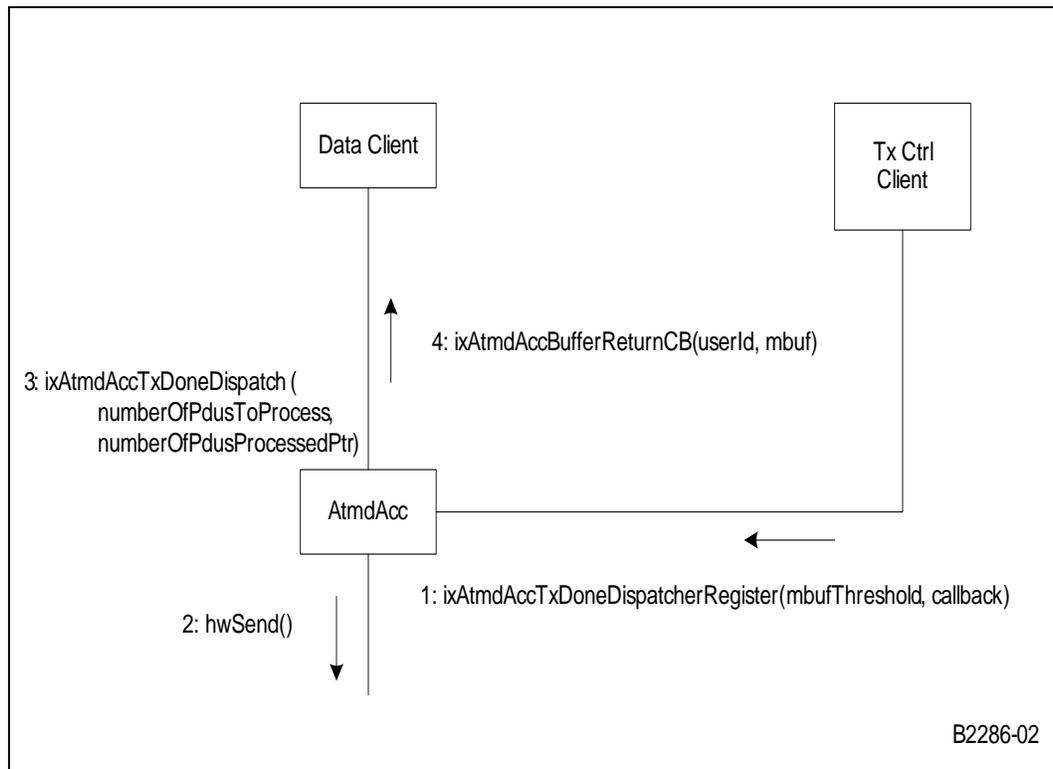
Processing primarily involves handing back ownership of buffers to clients. The rate at which this is done must be sufficient to ensure that client-buffer starvation does not occur. The details of the exact rate at which this must be done is implementation-dependent and not within the scope of this document. Because the Tx-Done resource is a system-wide resource, it is important to note that failing to poll it causes transmission to be suspended on all ports.

##### **Transmit Done — Based on a Threshold Level**

IxAtmdAcc does provide a notification service whereby a client can choose to be notified when the number of outstanding PDUs in the transmit done stream has reached a configurable threshold, as shown in [Figure 16](#).



Figure 16. Tx Done Recycling — Using a Threshold Level



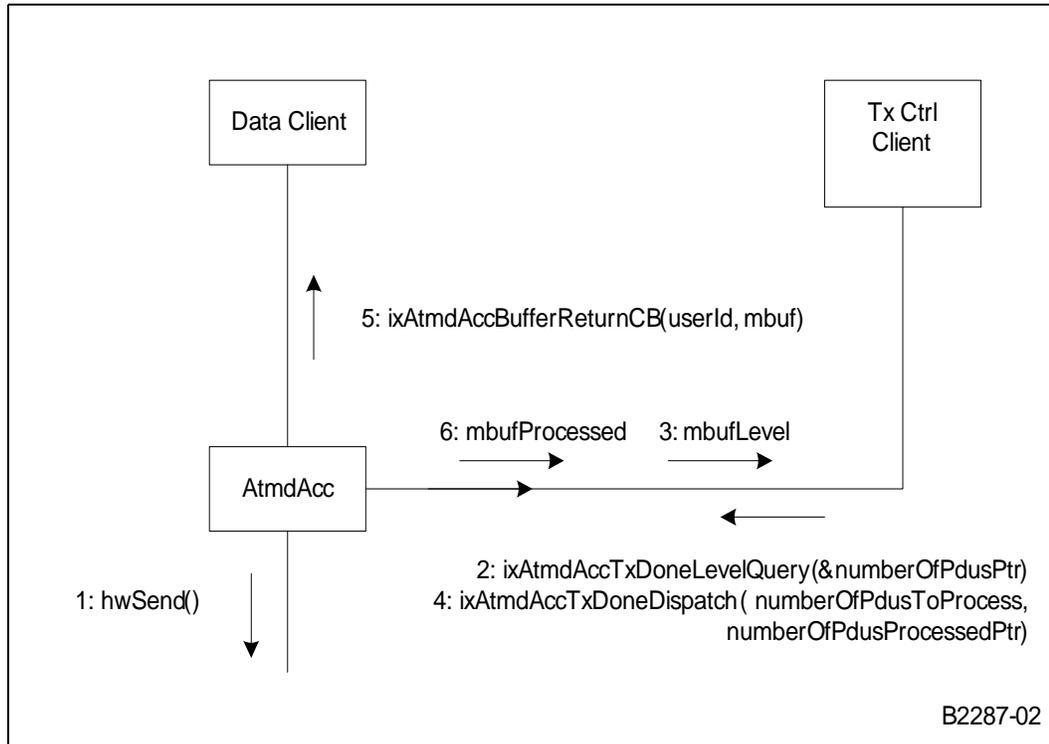
1. The control client wants to use the threshold services to process the transmitted PDUs. He calls the `ixAtmdAccTxDoneDispatcherRegister()` function to set a buffer threshold level and register a callback. `IxAtmdAcc` provides the function `ixAtmdAccTxDoneDispatch()` to be used by the control client. This function itself can be used directly as the callback. `IxAtmdAccTxDoneDispatcherRegister` allows the client to register its own callback. From this callback (where he can use an algorithm to decide the number of PDUs to service, depending on system load or any user constraint), the user has to call the `IxAtmdAccTxDoneDispatch()` function
2. Sometime earlier, the data client sent data to transmit. Cells are sent over the UTOPIA interface and the complete PDUs are available.
3. At a certain point in time, the threshold level of available PDUs is reached, the control client's callback is invoked by `IxAtmdAcc`. In response to this callback, the control client calls `ixAtmdAccTxDoneDispatch()`. This function gets the transmitted PDUs and retrieve the `connId` associated with this PDU.
4. Based on `connId`, `ixAtmdAccTxDoneDispatch` identified the data client to whom this buffer belongs. The corresponding data client's `TxDoneCallback` function, as registered during a `TxVcConnect`, is invoked with the `IX_OSAL_MBUF`. This `TxDoneCallback` function is likely to free or recycle the chained `IX_OSAL_MBUFs`.

*Note:* Even if threshold level is used, there is a need for polling to ensure that the `TxDone` queue gets cleared out when idle/low traffic or when the VCs are being shut down. Currently, the `IxAtmm` component has a thread that ensures the `TxDone` queue gets cleared at a certain interval time.

### Transmit Done — Based on Polling Mechanism

A polling mechanism can be used instead of the threshold service to trigger the recycling of the transmitted PDUs, as shown in Figure 17.

Figure 17. Tx Done Recycling — Using a Polling Mechanism



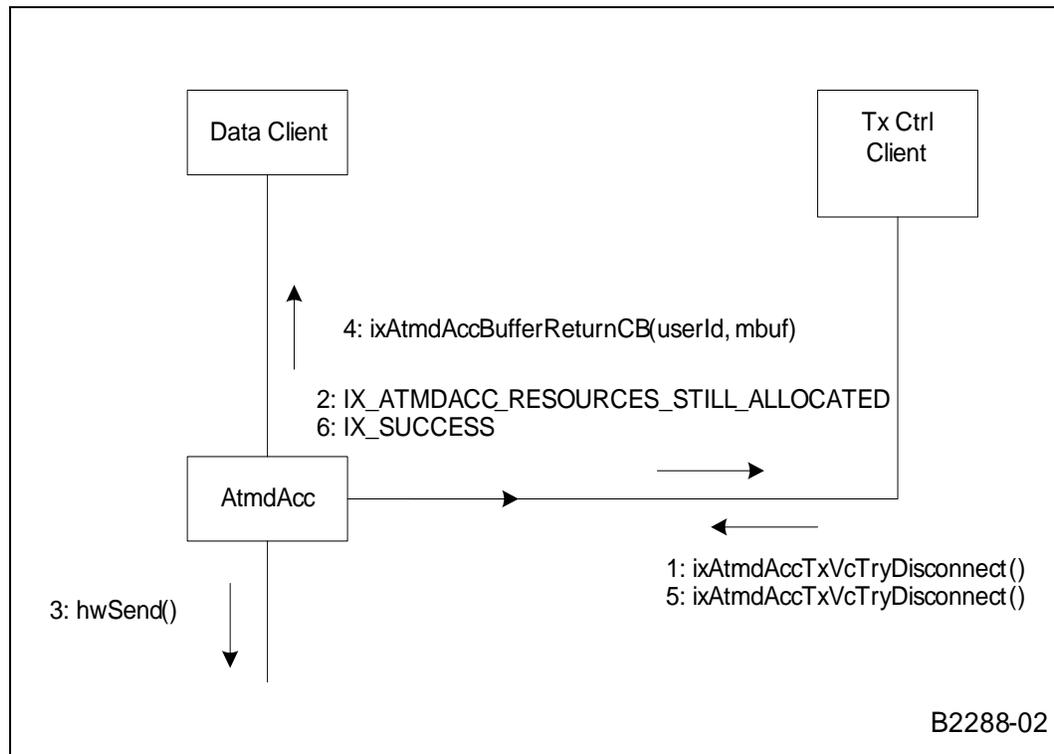
1. Sometime earlier, the data client sent data to transmit. Cells are sent over the UTOPIA interface and the PDUs are available.
2. A control client does not want to use the threshold services to process the transmitted PDUs. Therefore, the client invokes the `ixAtmdAccTxDoneLevelQuery()` function to get the current number of PDUs already transmitted.
3. The control client invokes `ixAtmdAccTxDoneDispatch()` with PDUs to do additional processing.
4. `ixAtmdAccTxDoneDispatch()` uses `connId` to identify the client to which this buffers belongs. The corresponding client's `TxDoneCallback()` function, as registered during a `TxVcConnect`, is invoked with the `IX_OSAL_MBUF` parameter. This `TxDoneCallback()` function may free or recycle the chained `IX_OSAL_MBUFs`.
5. The client gets the number of buffer processed to the control client. This number may be different of the number requested in the case of multiple instances of the `ixAtmdAccTxDoneDispatch()` function are used at the same time.

#### 4.6.2.2 Transmit Disconnect

Before a client disconnects from a VC, all resources must have been recycled, as shown in Figure 18. This is done by calling the `ixAtmdAccTxVcTryDisconnect()` function until all PDUs are transmitted by the hardware and all buffers are sent back to the client.



Figure 18. Tx Disconnect



1. The Data Client sends the last PDUs and the Control Client wants to disconnect the VC. `ixAtmdAccTxVcTryDisconnect()` invalidates further attempts to transmit more PDUs. Any call to `ixAtmdAccTxVcPduSubmit()` fails for this VC.
2. If there are resources still in use, the `ixAtmdAccTxVcTryDisconnect()` functions returns `IX_ATMDACC_RESOURCES_STILL_ALLOCATED`. This means that the hardware didn't finish transmitting all of the cells and there are `IX_OSAL_MBUFs` pending for transmission, or `IX_OSAL_MBUFs` in the `TxDone` stream.
3. Remaining packets are transmitted. (no new traffic is accepted through `ixAtmdAccTxVcPduSubmit()`)
4. The client waits a certain delay, depending on the TX rate for this VC, and asks again to disconnect the VC.
5. There are no resources still in use, the `ixAtmdAccTxVcTryDisconnect()` functions returns `IX_SUCCESS`. This means that the hardware did finish to transmit all of the cells and there are no `IX_OSAL_MBUFs` pending for transmission or in the `txDone` stream.

### 4.6.3 Receive Services

`IxAtmdAcc` processes inbound AAL payload data for individual VCs, received in `IXP_OSAL_MBUFs`. In the case of AAL 5, `IXP_OSAL_MBUFs` may be chained. In the case of AAL 0-48/52/OAM, chaining of `IXP_OSAL_MBUFs` is not supported. In the case of OAM, an `ix_IXP_OSAL_MBUF` contains only a single cell.

In the case of AAL 0, Rx cells are accumulated into an `IXP_OSAL_MBUF` under supervision of an Rx timer. The `IXP_OSAL_MBUF` is passed to the client when the `IXP_OSAL_MBUF` is passed to the client or when the `IXP_OSAL_MBUF` is filled or when the timer expires. The Rx timer is implemented by the NPE A.



In order to receive a PDU, the client layer must allocate IXP\_OSAL\_MBUFs and pass their ownership to the IxAtmdAcc component. This process is known as replenishment. Such buffers are filled out with cell payload. Complete PDUs are passed to the client. In the case of AAL 5, an indication about the validity of the PDU — and the validity of the AAL-5 CRC — is passed to the client.

In the case of AAL 0, PDU completion occurs when an IXP\_OSAL\_MBUF is filled, or is controlled by a timer expiration. The client is able to determine this by the fact that the IXP\_OSAL\_MBUF is not completely filled, in the case that completion was due to a timer expiring.

Refer to the API for details about the AAL-0 timer.

IxAtmdAcc supports prioritization of inbound traffic queuing by providing two separate receive streams. The algorithms and tuning required to service these streams are different, so management of latency and other priority constraints, on receive VCs, is allowed. As an example, one stream can be used for critical-time traffic (such as voice) and the other stream for data traffic.

The streams can be serviced in two ways:

- Setting a threshold level (when there is data available)
- Polling mechanism

Both mechanisms pass buffers to the client through a callback. Once the client is finished processing the buffer, it can ask to replenish the channel with available buffers or free the buffer back directly to the buffer pool.

Table 15 lists the receive routines.

**Table 15. List of Receive Routines**

Name	Function
ixAtmdAccRxDispatcherRegister	Register a notification callback to be invoked when there is at least one entry on a particular Rx queue
ixAtmdAccRxVcConnect	Connect to a AAL PDU receive service for a particular port/VPI/VCI and service type
ixAtmdAccRxVcFreeReplenish	Provide free IXP_OSAL_MBUFs for data reception on a connection
ixAtmdAccRxVcFreeLowCallbackregister	Configure the Rx Free threshold value and register a callback to handle threshold notifications
ixAtmdAccRxVcFreeEntriesQuery	Get the number of rx_mbufs the system can accept to replenish the rx reception mechanism on a particular channel
ixAtmdAccRxDispatch	Connect function that executes Rx processing for a particular Rx stream
ixAtmdAccRxLevelQuery	Query the number of entries in a particular Rx queue
ixAtmdAccRxVcEnable	Start the Rx service on a VC
ixAtmdAccRxVcDisable	Stop the Rx service on a VC
ixAtmdAccRxVcTryDisconnect	Disconnect a VC channel from the Rx service
ixAtmdAccRxDispatcherUnregister	Unregister a notification callback to be invoked when there is at least one entry on a particular Rx queue



### 4.6.3.1 Receive Triggers (Rx-Free-Low Notification)

IxAtmdAcc receive service does provide a Rx-free-low notification service that can be configured to execute a client supplied notification callback when the number of available buffers reaches a certain low level. The service is supported on a per-VC basis and the maximum threshold level is QSize for RxFree (set to 32 during initialization) unchained IXP\_OSAL\_MBUFs.

### 4.6.3.2 Receive Processing

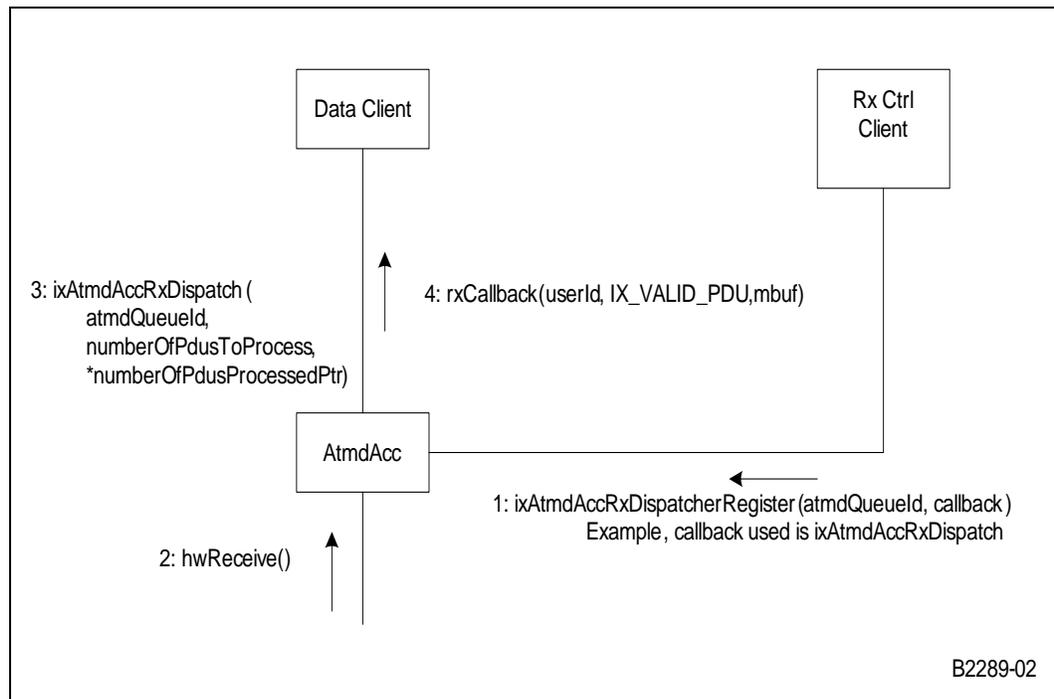
When buffers have been received on a port, they are placed in one of two Rx streams common to the VCs sharing this resource as decided by the client when establishing a connection. IxAtmdAcc does not autonomously process this stream, but instead the client decides when and how many buffers are processed.

Processing primarily involves handing back ownership of buffers to clients. The rate at which this is done must be sufficient to ensure that client requirements in terms of latency are met. The details of the exact rate at which this must be done is implementation-dependent and not within the scope of this document.

#### Receive — Based on a Threshold Level

IxAtmdAcc provides a notification service where a client can choose to be notified when incoming PDUs are ready in a receive stream as shown in Figure 19.

Figure 19. Rx Using a Threshold Level



1. The client invokes ixAtmdAccRxDispatcherRegister () to register a callback function. IxAtmdAcc provides the ixAtmdAccRxDispatch() function to be used by this callback. This function itself can be used directly as the callback or the ixAtmdAccRxDispatcherRegister allows the client to register its own callback. From this callback (where he can use an algorithm to decide the number of IXP\_OSAL\_MBUFs to service, depending on system load or any user constraint), the user has to call the IxAtmdAccRxDispatch() function. The atmdQueueId is the

receive queue type for receive traffic. There are two queues: IX\_ATM\_RX\_A and IX\_ATM\_RX\_B.

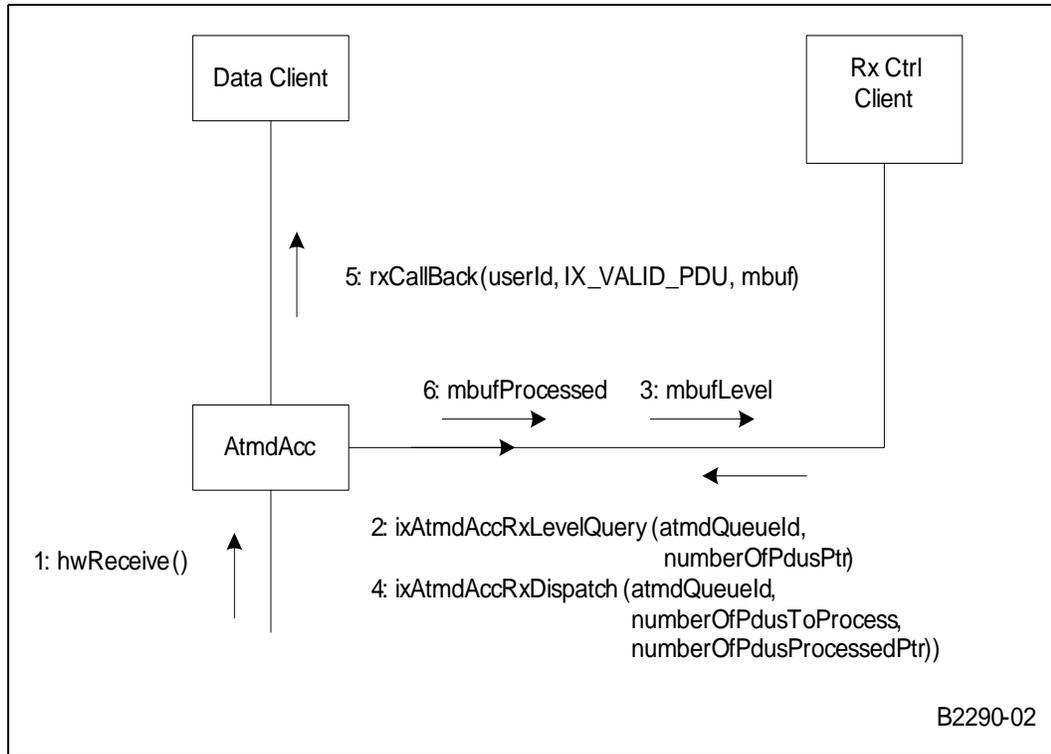
2. Cells are received over the UTOPIA interface and there is a PDU available.
3. When a complete PDU is received, the callback is invoked and the function `ixAtmdAccRxDispatch()` runs. This function iterates through the received buffers and retrieve the `connId` associated with each buffer.
4. `ixAtmdAccRxDoneDispatch()` uses `connId` to identify the client to which this buffers belongs. The corresponding client's `RxCallback` function (as registered during a `RxVcConnect`) is invoked with the first `IX_OSAL_MBUF` of a PDU. This `RxCallback` function is likely to push the received information to the protocol stack, and then to free or recycle the `IX_OSAL_MBUFs`. The `RxCallback` is invoked once per PDU. If there are many PDUs related to the same VC, the `RxCallback` is called many times.

*Note:* Even if threshold level is used, there is a need to use a polling to ensure that the receive queue gets cleared out when idle/low traffic or when the VCs are being shut down. Currently, the `IxAtmm` component has a thread that ensures the receive queue gets cleared at a certain interval time.

**Received — Based on a Polling Mechanism**

A polling mechanism can also be used to collect received buffers as shown in [Figure 20](#).

**Figure 20. RX Using a Polling Mechanism**



1. Cells are received over the UTOPIA interface and a complete PDU is available.
2. The client queries the current number of PDUs already received using the `ixAtmdAccRxLevelQuery()` function.
3. The control client instructs `IxAtmdAcc` to process a certain number of PDUs from one of the streams by calling `ixAtmdAccTxDoneDispatch()`.

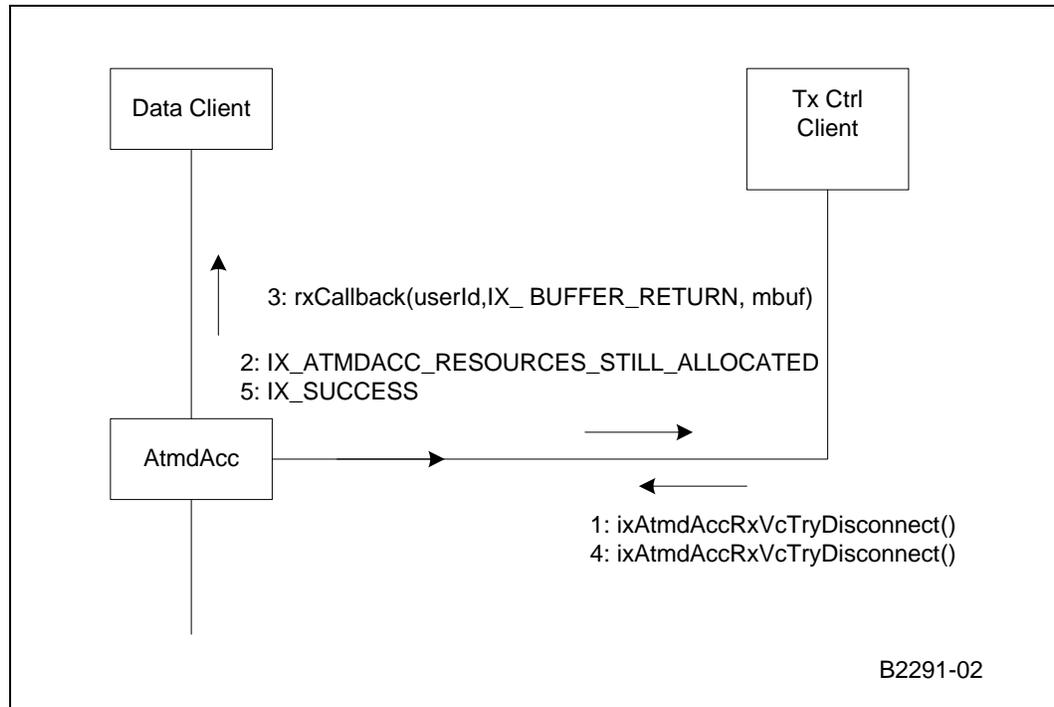


4. IxAtmdAcc gets the requested number of PDUs from the underlying hardware. Based on connId, ixAtmdAccRxDispatch() identifies the data clients to which the buffers belong. The corresponding data client's RxCallback functions — as registered during a ixAtmdAccRxVcConnect — is invoked with the first IXP\_OSAL\_MBUF a PDU.
5. This RxCallback function is likely to push the received information to the protocol stack, and then to free or recycle the IXP\_OSAL\_MBUFs. The RxCallback is invoked once per PDU. If there are many PDUs related to the same VC, the RxCallback is called many times.
6. IxAtmdAcc returns the number of PDUs processed.

#### 4.6.3.3 Receive Disconnect

Before a client disconnects from a VC, all resources must have been recycled as shown in Figure 21 on page 73.

Figure 21. Rx Disconnect



1. The control client wants to disconnect the VC. IxAtmdAccRxVcTryDisconnect() tell IxAtmdAcc to discard any rx traffic and — if resources are still in use — the IxAtmdAccRxVcTryDisconnect() function returns IX\_ATMDACC\_RESOURCES\_STILL\_ALLOCATED.
2. Reception of remaining traffic is discarded.
3. The client waits a certain delay — depending on the Rx drain rate for this VC — and asks again to disconnect the VC. If resources are still in use, the IxAtmdAccRxVcTryDisconnect() function returns IX\_ATMDACC\_RESOURCES\_STILL\_ALLOCATED.
4. Because there are no resources still in use, the IxAtmdAccRxVcDisconnect() function returns IX\_SUCCESS. This means that there are no resources or IXP\_OSAL\_MBUFs pending for reception or in the rxFree queue for this VC.



#### 4.6.4 Buffer Management

The IxAtmdAcc Interface is based on IXP\_OSAL\_MBUFs. The component addressing space for physical memory is limited to 28 bits. Therefore IXP\_OSAL\_MBUF headers should be located in the first 256 Mbytes of physical memory.

##### 4.6.4.1 Buffer Allocation

IXP\_OSAL\_MBUFs used by IxAtmdAcc are allocated and released by the client through the appropriate operating-system functions. During the disconnect steps, pending buffers are released by the IxAtmdAcc component using the callback functions provided by the client, on a per-VC basis.

##### 4.6.4.2 Buffer Contents

For performance reasons, the data pointed to by an IXP\_OSAL\_MBUF is not accessed by the IxAtmdAcc component.

The IXP\_OSAL\_MBUF fields required for transmission are described in Table 16. These fields are not changed during the Tx process.

Table 16. IXP\_OSAL\_MBUF Fields Required for Transmission

IXP_OSAL_MBUF_NEXT_BUFFER_IN_PKT_PTR	Required when IXP_OSAL_MBUFs are chained to build a PDU. In the last IXP_OSAL_MBUF of a PDU, this field value has to be 0
IXP_OSAL_MBUF_NEXT_PKT_IN_CHAIN_PTR	Not used
IXP_OSAL_MBUF_MDATA	Required. This field should point to the part of PDU data
IXP_OSAL_MBUF_MLEN	Required. This field is the length of data pointed to by ix_data
IXP_OSAL_MBUF_MTYPE	Not used
IXP_OSAL_MBUF_MFLAGS	Not used
IXP_OSAL_MBUF_PKT_LEN	Required in the first IXP_OSAL_MBUF of a chained PDU. This is the total length of the PDU.

The fields of available IXP\_OSAL\_MBUFs are described in Table 17. They are set by the client and must provide available buffers to IxAtmdAcc Rx service.

Table 17. IXP\_OSAL\_MBUF Fields of Available Buffers for Reception

IXP_OSAL_MBUF_NEXT_BUFFER_IN_PKT_PTR	This field value has to be 0. Buffer chaining is not supported when providing available buffers.
IXP_OSAL_MBUF_NEXT_PKT_IN_CHAIN_PTR	Not used
IXP_OSAL_MBUF_MDATA	This field is the pointer to PDU data
IXP_OSAL_MBUF_MLEN	This field is the length of data pointed to by ix_data
IXP_OSAL_MBUF_MTYPE	Not used
IXP_OSAL_MBUF_MFLAGS	Not used
IXP_OSAL_MBUF_PKT_LEN	Set to 0

The IXP\_OSAL\_MBUF fields in received buffers are set during traffic reception, and are described in Table 18.



**Table 18. IX\_OSAL\_MBUF Fields Modified During Reception**

IX_OSAL_MBUF_NEXT_BUFFER_IN_PKT_PTR	Modified when IX_OSAL_MBUFs are chained to build a PDU and it points to the next IX_OSAL_MBUF. In the last IX_OSAL_MBUF of a PDU, this field value has to be 0.
IX_OSAL_MBUF_NEXT_PKT_IN_CHAIN_PTR	Not used
IX_OSAL_MBUF_MDATA	This field is the pointer to PDU data
IX_OSAL_MBUF_MLEN	Modified. This field is the length of data pointed to by mh_data
IX_OSAL_MBUF_MTYPE	Not used
IX_OSAL_MBUF_MFLAGS	Not used
IX_OSAL_MBUF_PKT_LEN	Not used

#### 4.6.4.3 Buffer Size Constraints

Any IXP\_OSAL\_MBUF size can be transmitted, but a full PDU *must* be a multiple of a cell size (48/52 bytes, depending on AAL type). Similarly, the system can receive and chain IXP\_OSAL\_MBUFs that are a multiple of a cell size.

When receiving and transmitting AAL PDUs, the overall packet length is indicated in the first IXP\_OSAL\_MBUF header. For AAL 5, this length includes the AAL-5 PDU padding and trailer.

Buffers with an incorrect size are rejected by IxAtmdAcc functions.

#### 4.6.4.4 Buffer-Chaining Constraints

IX\_OSAL\_MBUFs can be chained to build PDUs up to 64 kilobytes of data + overhead. To submit a PDU for transmission, the client must supply a chained IX\_OSAL\_MBUF. When receiving a PDU, the client gets a chained IX\_OSAL\_MBUF. However, when replenishing of buffers to the Rx queue occurs (by using the provided interface), the receive queue contains the un-chained IX\_OSAL\_MBUFs. Similarly, Tx Done queue also contains un-chained IX\_OSAL\_MBUFs.

#### 4.6.4.5 Starvation and Throttling

Starvation of resources may occur when the number of IX\_OSAL\_MBUFs in the system and the threshold levels are not set appropriately. If the txDone threshold is set to 32, then IX\_OSAL\_MBUFs are held until 32 of them are internally stored. Then they all are released and the txDone user callback is called nearly 32 times. The user may decide to use these buffers to replenish a different subsystem (inside IxAtmdAcc or outside IxAtmdAcc). The size of the various resources has to be carefully monitored depending the application running (for example, replenishing the rxFree mechanism of a subsystem with the txDone mechanism from an other sub-system)

At a system-wide level, this makes the system run with alternation of starvation followed by burst periods. As a general rule, resources starvation can be used to throttle the traffic rate, but the mechanism to stop starvation should not be triggered by a burst.

There is a trade-off between the number of buffers, the maximum acceptable burst levels, the CPU bandwidth and the traffic average rate.



### 4.6.5 Error Handling

#### 4.6.5.1 API-Usage Errors

The AtmdAcc component detects the following misuse of the API:

- Inappropriate use of connection IDs
- Incorrect parameters
- Mismatches in the order of the function call — for example, using start() after disconnect()
- Use of resources already allocated for an other VC — for example, port/VPI/VCI

Error codes are reported as the return value of a function API.

The AAL client is responsible for using its own reporting mechanism and for taking the appropriate action to correct the problem.

#### 4.6.5.2 Real-Time Errors

Errors may occur during real-time traffic. Table 19 shows the various possible errors and the way to resolve them.

Table 19. Real-Time Errors

Cause	Consequences and Side Effects	Corrective Action
Rx-free queue underflow	<ul style="list-style-type: none"> <li>• System is not able to store the inbound traffic, and gets dropped.</li> <li>• AAL-5 CRC errors</li> <li>• PDU length invalid</li> <li>• Cells missing</li> <li>• PDUs missing</li> </ul>	<ul style="list-style-type: none"> <li>• Use the replenish function more often</li> <li>• Use more and bigger IXP_OSAL_MBUFs</li> </ul>
Tx-Done overflow	The hardware is blocked because the Tx-done queue is full.	<ul style="list-style-type: none"> <li>• Poll the TxDone queue more often.</li> <li>• Change the TxDone threshold.</li> </ul>
IxAtmdAccPduSubmit() reports IX_ATMD_OVERLOADED	System is unable to transmit a PDU.	<ul style="list-style-type: none"> <li>• Increase the scheduler-transmit speed.</li> <li>• Slow down the submitted traffic.</li> </ul>
Rx overflow	<ul style="list-style-type: none"> <li>• Inbound traffic is dropped.</li> <li>• AAL-5 CRC errors</li> <li>• PDU length invalid</li> </ul>	Poll the Rx streams more often.





## 5.0 Access-Layer Components: ATM Manager (IxAtmm) API

---

This chapter describes the Intel® IXP400 Software v2.3's "ATM Manager API" access-layer component.

IxAtmm is an example software release 2.3 component. The phrase "Atmm" stands for "ATM Management."

The chapter describes the following details of ixAtmm:

- Functionality and services
- Interfaces to use these services
- Conditions and constraints for using the services
- Dependency on other software release 2.3 components
- Performance and resource usage

### 5.1 What's New

The following new routines have been added:

**IxAtmmUninit()** destroys the mutex object initialized by IxAtmmInit, and unregisters the VC Callback. It also clears all the tables initialized by ixAtmmInit.

**ixAtmmPortUninitialize()** uninitializes the data path of that port.

### 5.2 IxAtmm Overview

The software release 2.3's IxAtmm component is a demonstration ATM configuration and management component intended as a "point of access" for clients to the ATM layer of the Intel® IXP4XX Product Line of Network Processors.

This component, supplied only as a demonstration, encapsulates the configuration of ATM components in one unit. It can be modified or replaced by the client as required.

### 5.3 IxAtmm Component Features

The ixAtmm component is an ATM-port, virtual-connection (VC), and VC-access manager. It does not provide support for ATM OAM services and it does not directly move any ATM data.

IxAtmm services include:

- Configuring and tracking the usage of the (physical) ATM ports on Intel® IXP4XX Product Line of Network Processors.

In software release 2.3, up to 12 parallel logical ports are supported over UTOPIA Level 2. IxAtmm configures the UTOPIA device for a port configuration supplied by the client.



- Initializing the IxAtmSch ATM Scheduler component for each active port. IxAtmm assumes that the client will supply initial upstream port rates once the capacity of each port is established.
- Ensuring traffic shaping is performed for each registered port. IxAtmm acts as transmission control for a port by ensuring cell demand is communicated to the IxAtmSch ATM Scheduler from IxAtmdAcc and cell transmission schedules produced by IxAtmSch are supplied at a sufficient rate to IxAtmdAcc component.
- Determining the policy for processing transmission buffers recycled from the hardware. In the software release 2.3, the component ensures this processing is done on an event-driven basis. That is, a notification of threshold number of outstanding recycled buffers trigger processing of the recycled buffers.
- Controlling the processing of receive buffers via IxAtmdAcc. IxAtmdAcc supports two incoming Rx buffer streams termed high- and low-priority streams.
  - The high-priority stream is serviced in an event-driven manner. For example, as soon a buffer is available in the stream, it is serviced.
  - The low-priority stream is serviced on a timer basis.
- Allowing clients to register VCCs (Virtual Channel Connections) on all serving ATM ports for transmitting and/or receiving ATM cells. IxAtmm checks the validity (type of service, traffic descriptor, and so forth) of the registration request and rejects any request that presents invalid traffic parameters. IxAtmm does not have the capability to signal, negotiate, and obtain network admission of a connection. The client will make certain that the network has already admitted the requested connection before registering a connection with IxAtmm. IxAtmm also may reject a connection registration that exceeds the port capacity on a first-come-first-serve basis, regardless of whether the connection has already been admitted by the network.
- Enabling query for the ATM port and registered VCC information on the port.
- Allowing the client to modify the port rate of any registered port after initialization.
- IxAtmmUninit checks if the Atmm component is initialized and if found true destroys the mutex object initialized by IxAtmmInit, and unregisters the VC Callback. It also clears all the tables initialized by ixAtmmInit.
- IxAtmmPortUninitialize is done only while UTOPIA is initialized. It also validates for input parameter and if found correct uninitializes the data path of that port.

## 5.4 ixAtmmAcc API

Table 20 lists the routines available in this component.

Table 20. List of Routines

Name	Function
ixAtmmInit	Initialize the IxAtmm software component
ixAtmmUninit	Uninitialize the IxAtmm software component
ixAtmmUtopiaInit	Initialize the UTOPIA Level-2 ATM coprocessor for the specified number of physical ports.
ixAtmmUtopiaUninit	Uninitialize the UTOPIA Level-2 ATM coprocessor.
ixAtmmPortInitialize	Activate the registered ATM port with IxAtmm

**Table 20. List of Routines**

Name	Function
ixAtmmPortUninitialize	Uninitializes the Data Path. Unintializes the specified port.
ixAtmmPortModify	Change the existing port rate (expressed in bits/second) on an established ATM port
ixAtmmPortQuery	Request details on currently registered transmit and receive rates for an ATM port
ixAtmmPortEnable	Enable transmit for an ATM port. At initialization, all the ports are disabled.
ixAtmmPortDisable	Disable transmit for an ATM port.
ixAtmmVcRegister	Register an ATM Virtual Connection on the specified ATM port.
ixAtmmVcDeregister	Deregister a VC from the system
ixAtmmVcQuery	Supplies information about an active VC on a particular port
ixAtmmVcIdQuery	Supplies information about an active VC on a particular port
ixAtmmVcChangeCallbackRegister	Supply a function to IxAtmm which is called to notify the client if a new VC is registered with IxAtmm or an existing VC is removed.
ixAtmmVcChangeCallbackDeregister	Deregister a previously supplied callback function
ixAtmmUtopiaStatusShow	Display UTOPIA status counters
ixAtmmUtopiaCfgShow	Display UTOPIA information (config registers and status registers).

## 5.5 UTOPIA Level-2 Port Initialization

IxAtmm is responsible for the initial configuration of the Intel® IXP4XX Product Line of Network Processors' UTOPIA Level-2 device. This is performed through a user interface that facilitates specification of UTOPIA-specific parameters to the IxAtmm component.

IxAtmm supports up to 12 logical ports over the UTOPIA interface.

The data required for each port to configure the NPE UTOPIA coprocessor is the 5-bit address of the Transmit and Receive PHY interfaces on the UTOPIA bus.

The NPE UTOPIA coprocessor can also be initialized in loop-back mode. Loop-back is only supported, however, in a single port configuration.

All other UTOPIA configuration parameters are configured to a static state by the IxAtmm and are not configurable through the functional interface of this component. Clients that wish a greater level of control over the UTOPIA device should modify and recompile the IxAtmm component with the new static configuration. Alternately, they can use the interface provided by the IxAtmdAcc component.

## 5.6 ATM-Port Management Service Model

IxAtmm can be considered an "ATM-port management authority." It does not directly perform data movement, although it does control the ordering of cell transmission through the supply of ATM cell-scheduling information to the lower levels.

IxAtmm manages the usage of registered ATM ports and allows or disallows a VC to be established on these ports — depending on existing active-traffic contracts and the current upstream port rate.



Once a connection is established, a client can begin to use it. The client makes data transfer requests directly to corresponding AAL layer through the IxAtmdAcc component. The AAL layer passes the request to the Intel® IXP4XX Product Line of Network Processors through the appropriate hardware layers, under direction from IxAtmm.

The IxAtmm service model consists of two basic concepts:

- ATM port
- VC/VCC (virtual channel/virtual channel connection) connections that are established over this port

A VC is a virtual channel through a port. A VC is *unidirectional* and is associated with a unique VPI/VCI value. Two VCs — in opposite direction on the same port — can share the same VPI/VCI value. A VCC is an end-to-end connection through linked VCs, from the local ATM port to another device across the ATM network.

Initially, a port is “bare” or “empty.” A VC must be attached (registered) to a port. Registration means, “to let IxAtmm know that — from now on — the VC can be considered usable on this port.”

IxAtmm is not responsible for signaling and obtaining admission from the network for a VCC. A client must use other means, where necessary, to obtain network admission of a VCC. A client specifies to IxAtmm the traffic descriptor for the requested VCC. IxAtmm will accept or deny this request based only on the port rate available and the current usage of the port by VCCs already registered with the system. This CAC functionality is provided by the IxAtmSch component.

IxAtmm presumes that the client has already negotiated — or will negotiate — admission of the VCC with the network.



Figure 22. Services Provided by Ixatmm

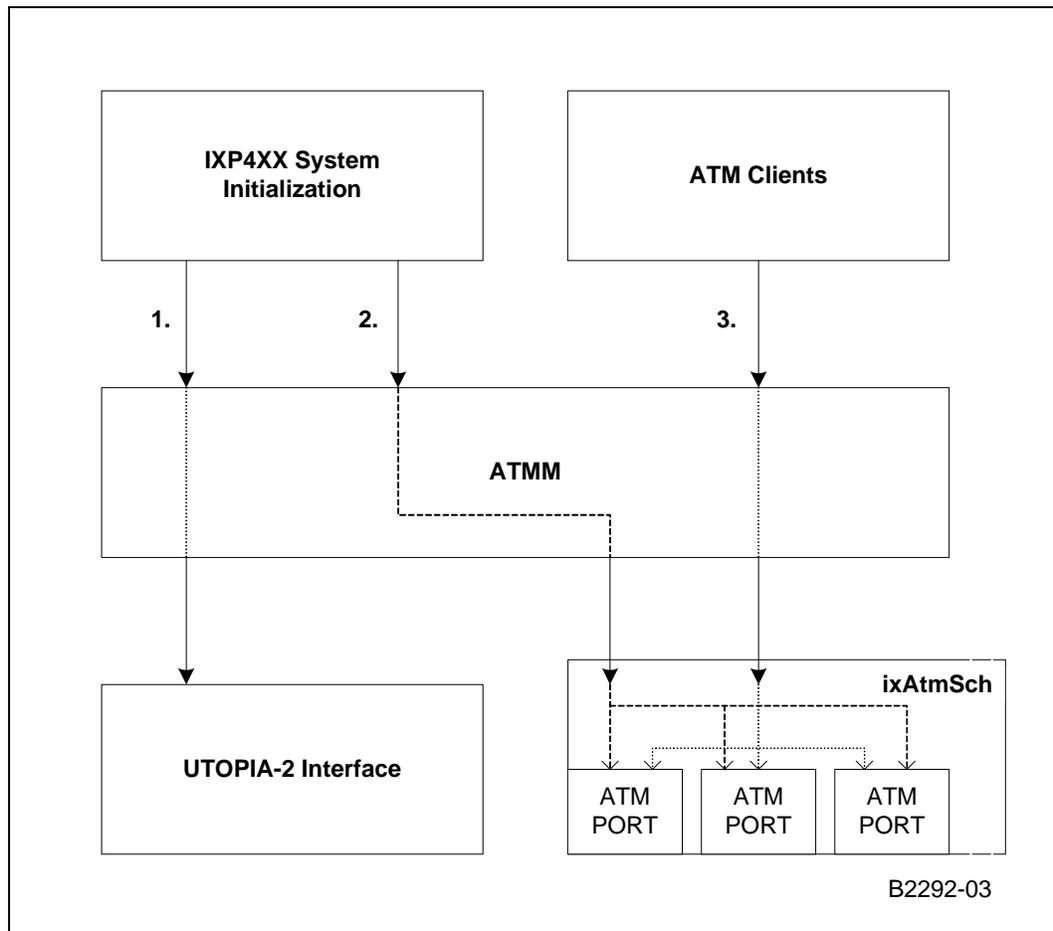


Figure 22 shows the main services provided by the IxAtmm component. In this diagram, the three services outlined are:

1. Intel® IXP4XX Product Line of Network Processors system-initialization routine will invoke an IxAtmm interface function to initialize the UTOPIA Level-2 device for all active ATM ports in the system. This function call is only performed once, encompassing the hardware configuration of all ports in a single call to the interface.
2. Once the link is established for each active port and the line rates are known to the system, IxAtmm is informed of the upstream and downstream rate for each port. The upstream rate is required by the ATM scheduler component in order to provide traffic shaping and admission services on the port. The port rates must be registered with IxAtmm before any VCs may be registered. In addition, once the scheduling component is configured, it is bound to IxAtmdAcc. This ensures shaped transmission of cells on the port.
3. Once the port rate has been registered, the client may register VCs on the established ports. Upstream and downstream VCs must be registered separately. The client is assumed to have negotiated any required network access for these VCs before calling IxAtmm. IxAtmm may refuse to register upstream VCs — the ATM scheduler's admission refusal being based on port capacity.



Once IxAtmm has allowed a VC, any future transmit and receive request on that VC will not pass through IxAtmm. Instead, they go through corresponding AAL layer directly to the Intel® IXP4XX Product Line of Network Processors' hardware.

Further calls to IxAtmdAcc must be made by the client following registration with IxAtmm to fully enable data traffic on a VC.

IxAtmm does *not* support the registration of Virtual Path Connections (VPCs). Registration and traffic shaping is performed by IxAtmm and IxAtmSch on the VC/VCC level only.

## 5.7 Tx/Rx Control Configuration

The IxAtmm application is responsible for the configuration of the mechanism by which the lower-layer services will drive transmit and receive of traffic to and from the IXP4XX product line processor' hardware. This configuration is achieved through the IxAtmdAcc component interface.

Configuration of these services is performed when the first active port is registered with IxAtmm. See [Figure 23 on page 83](#)

IxAtmm will configure IxAtmdAcc for the following traffic events:

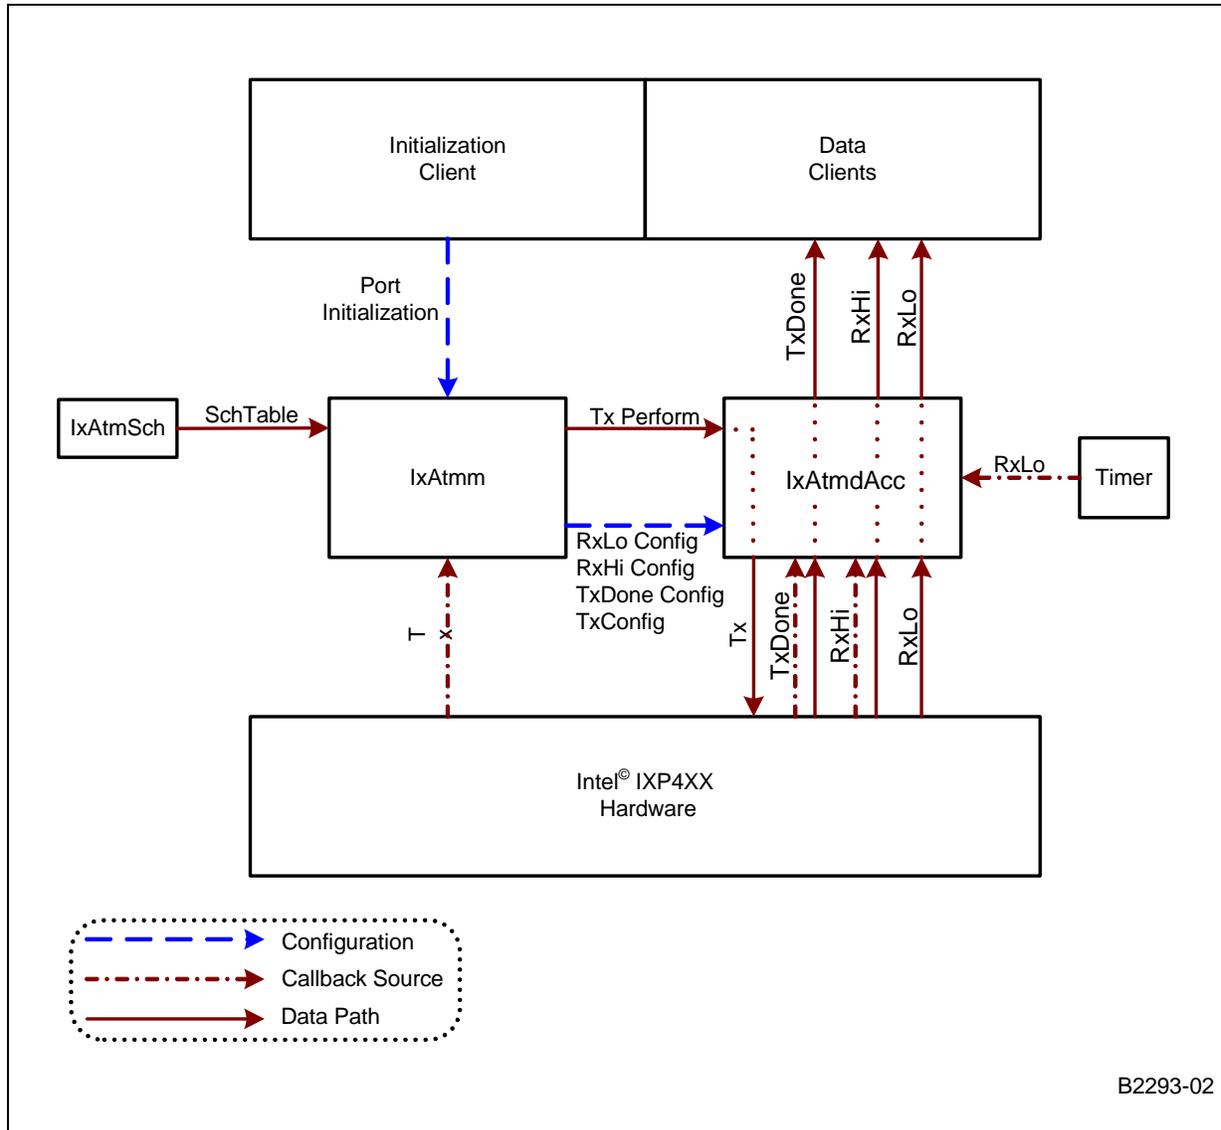
- **Transmit Required** — The Intel® IXP4XX Product Line of Network Processors' hardware requires more cells to be scheduled for transmission on a particular port. IxAtmm will implement a callback function that is registered as a target for the low-queue notification callback with IxAtmdAcc. When invoked, this function will generate a transmit schedule table for the port through the IxAtmSch component and pass this table to the IxAtmdAcc interface to cause more cells to be transmitted to the hardware, according to the generated schedule table.
- **Transmit Done** — When all data from a particular buffer has been transmitted, it is necessary for the Intel® IXP4XX Product Line of Network Processors' hardware to return the buffer to the relevant client. IxAtmm will configure the Intel® IXP4XX Product Line of Network Processors such that the processing of these buffers is performed whenever there are a specific number of buffers ready to be processed. IxAtmm will configure the system such that the default IxAtmdAcc interface returns these buffers to the appropriate clients and are then invoked automatically.
- **High-Priority Receive** — Data received on the any high-priority receive channel (such as voice traffic) is required to be supplied to the client in a timely manner. IxAtmm will configure the IxAtmdAcc component to process the receipt of data on high-priority channels using a low threshold value on the number of received data packets. The default IxAtmdAcc receive processing interface is invoked whenever the number of data packets received by the Intel® IXP4XX Product Line of Network Processors reaches the supplied threshold. These packets will then be dispatched to the relevant clients by the IxAtmdAcc component.
- **Low-Priority Receive** — Data received on low-priority receive channels (for example, data traffic) is not as urgent for delivery as the high-priority data and is, therefore, expected to be tolerant of some latency when being processed by the system. IxAtmm will configure the Intel® IXP4XX Product Line of Network Processors such that the receive processing of low-priority data is handled according to a timer. This will cause the processing of this data to occur at regular time intervals, each time returning all pending low-priority data to the appropriate clients.

The IxAtmm component is responsible only for the configuration of this mechanism. Where possible the targets of threshold and timer callbacks are the default interfaces for the relevant processing mechanism, as supplied by IxAtmdAcc. The exception is the processing of cell transmission, which is driven by an IxAtmm callback interface that passes ATM scheduling information to the IxAtmdAcc component, as required to drive



the transmit function. As a result, all data buffers in the system — once configured — will pass directly through IxAtmdAcc to the appropriate clients. No data traffic will pass through the IxAtmm component at any stage.

**Figure 23. Configuration of Traffic Control Mechanism**



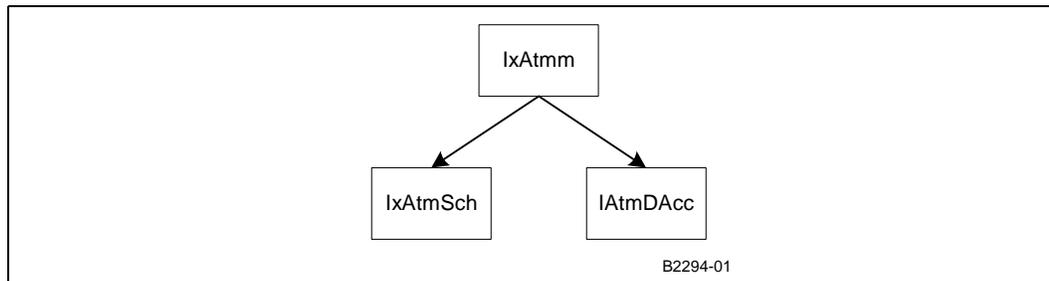
Only transmit traffic — which has already been queued by the client with IxAtmdAcc when the request for more traffic is made — is scheduled and sent to the hardware. (That is, no callback to the data client is made in the context of the transmit processing.) IxAtmdAcc makes IxAtmSch aware of the existence of this pending traffic when it is queued by the client through the use of a previously registered callback interface.

The supply of empty buffers to the hardware — for use in the receive direction — is the responsibility of the individual client on each active VC. As a result, the target callback for this event on each VC is outside of the visibility of the IxAtmm component, being

part of the client logic. It is the responsibility of each client, therefore, to ensure that the supply mechanism of free buffers for receive processing is configured correctly before traffic may begin passing on the system.

## 5.8 Dependencies

Figure 24. Component Dependencies of IxAtmm



IxAtmm configures the Intel® IXP4XX Product Line of Network Processors' UTOPIA Level-2 device through an interface provided by the IxAtmdAcc component.

IxAtmm is also responsible for configuring VC registrations with the IxAtmSch demo ATM scheduler component and relaying CAC decisions to the client in the event of VC registration failure.

IxAtmm is responsible for port traffic shaping by conveying traffic and scheduling information between the ATM scheduler component and the cell transmission control interface provided by the IxAtmdAcc component.

## 5.9 Error Handling

IxAtmm returns an error type to the user when the client is expected to handle the error. Internal errors is reported using the Intel® IXP4XX Product Line of Network Processors' OSAL error reporting mechanism.

The established state of the IxAtmm component (registered ports, VCs, and so forth) is not affected by the occurrence of any error.

## 5.10 Management Interfaces

No management interfaces are supported by the IxAtmm component. If a management interface is required for the ATM layer, the IxAtmm is the logical place for this interface to be implemented, as the component is intended to provide an abstract public interface to the non-data path ATM functions.

## 5.11 Memory Requirements

IxAtmm code is approximately 26 Kbytes in size.

IxAtmm data memory requirement — under peak cell-traffic load — is approximately 20 Kbytes.



## 5.12 Performance

The IxAtmm does not operate on the data path of the Intel® IXP4XX Product Line of Network Processors. Because it is primarily concerned with registration and deregistration of port and VC data, IxAtmm is typically executed during system initialization.

§ §





## 6.0 Access-Layer Components: ATM Transmit Scheduler (IxAtmSch) API

---

This chapter describes the Intel® IXP400 Software v2.3's "ATM Transmit Scheduler" (IxAtmSch) access-layer component.

### 6.1 What's New

The following features are supported:

- rt-VBR QoS class
- Oversubscription mode of operation
- Prioritization of cells in oversubscription mode

A section on design constraints and limitations on ATM high speed rate has been added.

The following two routines have been added to ixAtmSch API:

**ixAtmSchUninit (void)** This function is used to uninitialized the ixAtmSch component.

**ixAtmSchPortModelUninitialize ()** This function shall be called to uninitialized an ATM port.

### 6.2 Overview

IxAtmSch is an "example" software release 2.3 component, an ATM scheduler component supporting ATM transmit services on Intel® IXP4XX Product Line of Network Processors.

This chapter discusses the following IxAtmSch component details:

- Functionality and services
- Interfaces to use the services
- Conditions and constraints for using the services
- Component dependencies on other software release 2.3 components
- Component performance and resource usage estimates

IxAtmSch is a simplified scheduler with limited capabilities. See [Table 22 on page 89](#) for details of scheduler capabilities.

The IxAtmSch API is specifically designed to be compatible with the IxAtmdAcc transmission-control interface. However, if a client decides to replace this scheduler implementation, they are urged to reuse the API presented on this component.

IxAtmSch conforms to interface definitions for the Intel® IXP4XX Product Line of Network Processors' ATM transmission-control schedulers.



## 6.3 IxAtmSch Component Features

The IxAtmSch component is provided as a demonstration ATM scheduler for use in the processor's ATM transmit. It provides two basic services for managing transmission on ATM ports:

- Outbound (transmission) virtual connection admission control on serving ATM ports
- Schedule table to the ATM transmit function that will contain information for ATM cell scheduling and shaping

IxAtmSch implements a fully operational ATM traffic scheduler for use in the processor's ATM software stack. It is possible (within the complete software release 2.3 architecture) to replace this scheduler with one of a different design. If replaced, this component still is valuable as a model of the interfaces that the replacement scheduler requires to be compatible with the software release 2.3 ATM stack. IxAtmSch complies with the type interfaces for an software release 2.3 compatible ATM scheduler as defined by the IxAtmdAcc software component.

The IxAtmSch service model consists of two basic concepts: ATM port and VCC. Instead of dealing with these real hardware and software entities in the processor and software stack, IxAtmSch models them. Because of this, there is no limit to how many ATM ports it can model and schedule — given enough run-time computational resources.

IxAtmSch does not currently model or schedule Virtual Paths (VPs) or support any VC aggregation capability.

In order to use IxAtmSch services, a client first must ask IxAtmSch to establish the model for an ATM port. Virtual connections then can be attached to the port.

IxAtmSch models the virtual connections and controls the admission of a virtual connection, based on the port model and required traffic parameters. IxAtmSch schedules and shapes the outbound traffic for all VCs on the ATM port. IxAtmSch generates a scheduling table detailing a list of VCs and number of cells of each to transmit in a particular order.

The IxAtmSch component's two basic services are related. If a VC is admitted on the ATM port, IxAtmSch is committed to schedule all outbound cells for that VC, so that they are conforming to the traffic descriptor. The scheduler does not reject cells for transmission as long as the transmitting user(s) (applications) do not over-submit. Conflict may happen on the ATM port because multiple VCs are established to transmit on the port.

If a scheduling commitment cannot be met for a particular VC, it is not be admitted. The IxAtmSch component admits a VC based only on the port capacity, current-port usage, and required-traffic parameters.

The current resource requirements are for a maximum of 12 ports and a total of 32 VCs across all ports. This may increase in the future.

Oversubscription feature is supported

rt-VBR QoS class is supported

## 6.4 IxAtmSch API

The routines in the API are listed in [Table 21](#).



**Table 21. List of Routines in the Component**

Name	Function
ixAtmSchInit	Initialize the ixAtmSch component.
ixAtmSchUninit	Uninitialize the ixAtmSch component
ixAtmSchPortModelInitialize	Initialize an ATM port
ixAtmSchPortModelUninitialize	Uninitialize an ATM port
ixAtmSchPortRateModify	Modify the portRate on a previously initialized port,
ixAtmSchVcModelSetup	Set up an upstream (transmitting) virtual connection model (VC) on the specified ATM port
ixAtmSchVcConnIdSet	Set the vcUserConnId for a VC on the specified ATM port
ixAtmSchVcModelRemove	Remove a previously established VC on a particular port.
ixAtmSchVcQueueUpdate	Notify IxAtmSch that the user of a VC has submitted cells for transmission
ixAtmSchVcQueueClear	Remove all currently queued cells from a registered VC
ixAtmSchTableUpdate	Update of the schedule table for a particular ATM port.
ixAtmSchShow	Print statistics on the current and accumulated state of VCs and traffic
ixAtmSchStatsClear	Reset all counter statistics in the ATM scheduler to zero

## 6.5 Traffic Types

Table 22 shows the ATM service categories supported in the software release 2.3 scheduler model.

**Table 22. Supported Traffic Types and Traffic Parameters**

Traffic Type	Supported	Num VCs	CDVT	PCR	SCR	MBS
rt-VBR	Yes	Up to 32 VCs	No	Yes	Yes	Yes
nrt-VBR	Yes	Up to 32 VCs	No	Yes	Yes	Yes
UBR	Yes	Up to 32 VCs	No	No	No	No
CBR	Yes	Up to 32 VCs	No	Yes	No	No
<b>Note:</b>						
1. The traffic type rt-VBR does not support any of its QoS parameters (CDVT, peak-to-peak CDV, MaxCTD and CLR). Hence, rt-VBR and nrt-VBR are equivalent traffic except in oversubscription mode.						
2. rt-VBR and CBR support is provided with IXP400 software version 2.1 onwards. nrt-VBR and UBR support is provided in IXP400 software version 1.0 onwards.						

## 6.6 Connection Admission Control (CAC) Function

IxAtmSch makes outbound virtual connection admission decisions based a simple ATM port reference model. Only one parameter is needed to establish the model: outbound (upstream) port rate R, in terms of (53 bytes) ATM cells per second.

IxAtmSch assumes that the “real-world” ATM port is a continuous pipe that draws the ATM cells at the constant cell rate. IxAtmSch does not rely on a hardware clock to get the timing. Its timing information is derived from the port rate. It assumes  $T = 1/R$  seconds pass for sending every ATM cell.

IxAtmSch determines if a new (modeled) VC admission request on any ATM port is acceptable using the following information supplied by its client:



- Outbound port rate
- Required traffic parameters for the new VC
- Traffic parameters of existing VCs on that port

IxAtmSch works on a first-come-first-served basis. For example, if three existing CBR VCs on the ATM port each use one-fourth of the port's capacity ( $PCR = R/4$ ), the fourth CBR VCC asking for 1/3 of the port capacity ( $PCR = R/3$ ) is rejected. IxAtmSch issues a globally unique VCC ID for each accepted VCC.

For rt-VBR and nrt-VBR VCs — where the SCR and PCR values are different — only the SCR value is used to determine the required capacity for the VC. This is based on the principle that, over a long term, the required capacity of the VC is equal to the SCR value, even if the VC may burst at rates above that rate for short periods.

Upon a successful registration via the CAC function, each VC is issued a port-unique identifier value. This value is a positive integer. This value is used to identify the VC to IxAtmSch during any subsequent calls. The combination of port and VC ID values will uniquely identify any VC in the processor device to the IxAtmSch component.

## 6.7 VC Oversubscription and Priority Feature

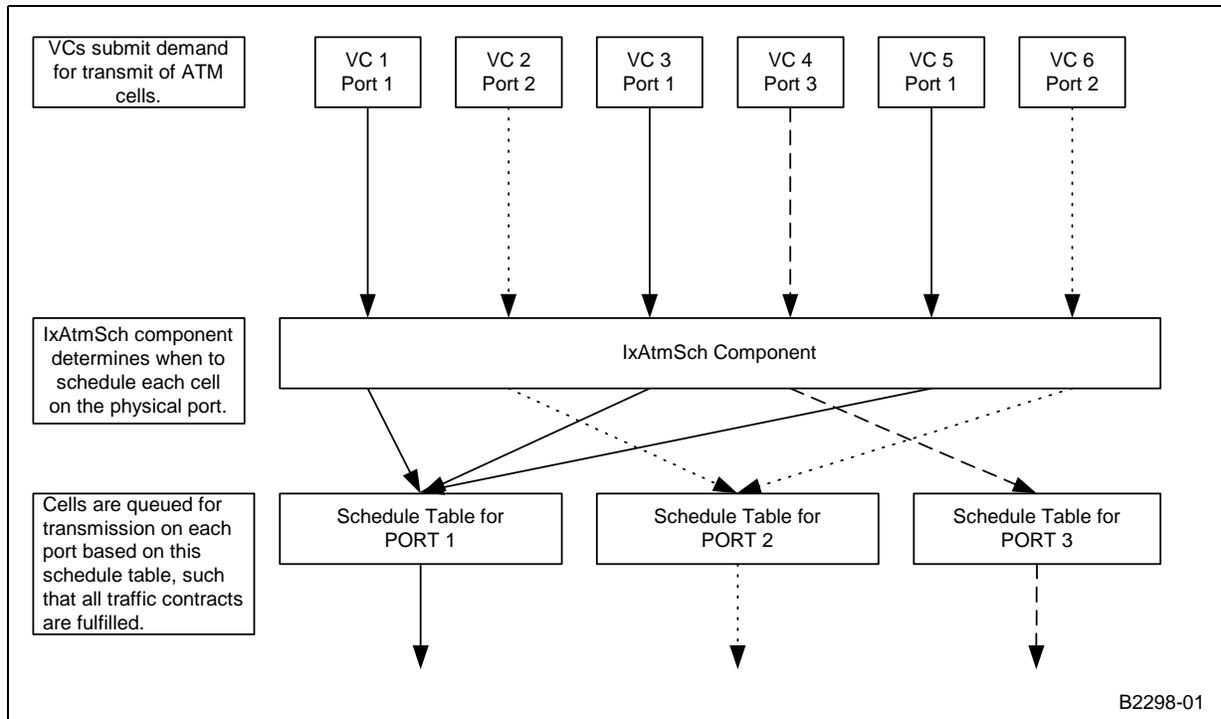
The oversubscription feature is activated when the total allocated bandwidth for CBR, rt-VBR and nrt-VBR services is greater than the ATM port rate. This situation could occur if the port rate is reduced using `ixAtmSchPortRateModify()`. For example, assume that the initial port rate  $R = 800$  kbps and three CBR VCs are utilizing the full bandwidth. Now, if the port rate is reduced to 400 kbps using `ixAtmSchPortRateModify()` then the VCs are oversubscribed their bandwidth requirement is greater than 400 kbps.

When the oversubscription occurs, IxAtmSch component will use the priority feature. The priority scheduling is a feature where the CBR has the highest priority for cell scheduling followed by rt-VBR and nrt-VBR. UBR has the lowest priority. At each priority level, IxAtmSch component will ensure that the VC has cells to be scheduled and, most importantly, it is eligible for cell scheduling. This eligibility is based on whether the VC is compliant for cell scheduling.



## 6.8 Scheduling and Traffic Shaping

Figure 25. Multiple VCs for Each Port, Multiplexed onto Single Line by the ATM Scheduler



B2298-01

### 6.8.1 Schedule Table

Once an ATM port is modeled and VCs are admitted on it, the client can request IxAtmSch to publish the schedule table that indicates how the cells — on all modeled VCs over the port — is interleaved and transmitted.

IxAtmSch publishes a scheduling table each time its scheduling function is called by a client for a particular port. The schedule table data structure returned specifies an ordering on which cells should be transmitted from each VCs on the port for a forthcoming period. The client is expected to requests a table for a port when the transmit queue is low on that port.

The number of cells that are scheduled by each call to the scheduling function will vary depending on the traffic conditions. The schedule table contains an element, `totalCellSlots`, which specifies how many cell slots are scheduled in this table returned, including idle cells.

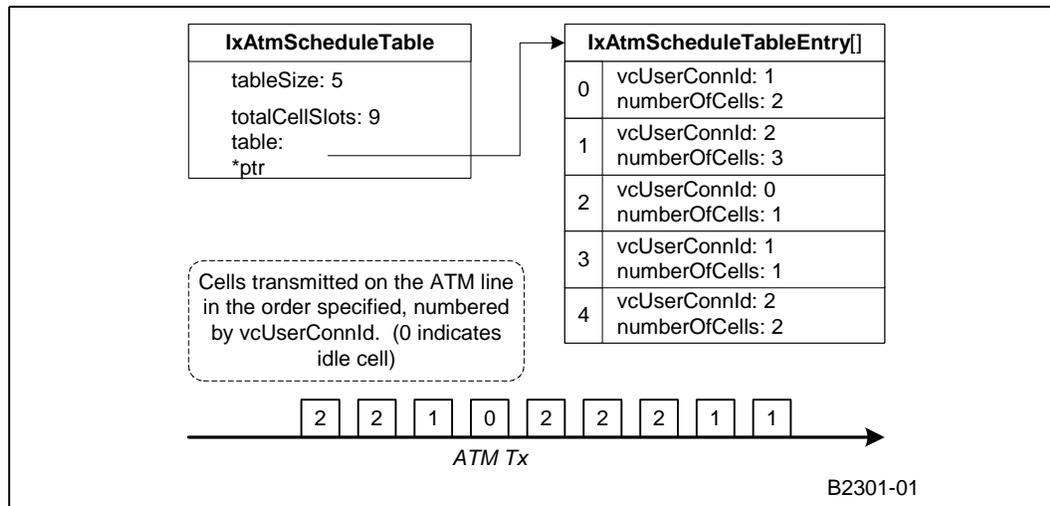
When the client calls the schedule function, the scheduler assumes that all previously scheduled cells on this port have been transmitted and that it may overwrite the previous schedule table with the new table. The client, therefore, must not be dependent on the integrity of the previous table when a request is made for a new schedule table. Additionally, the client should ensure that the current schedule table has been processed by the transmit mechanism before it requests for a new table.

The schedule table is composed of an array of table entries, each of which specifies a VC ID and a number of cells to transmit from that VC. The scheduler explicitly inserts idle cells into the table, where necessary, to fulfill the traffic contract of the VCs registered in the system. Idle cells are inserted in the table with the VC identifier set to 0.

The exact format of the schedule table is defined in `IxAtmTypes.h`.

Figure 26 shows how this table is translated into an ordered sequence of cells transmitted to the ATM port.

**Figure 26. Translation of IxAtmScheduleTable Structure to ATM Tx Cell Ordering**



### 6.8.1.1 Minimum Cells Value (minCellsToSchedule)

When a port model is created the minimum number of cells (`minCellstoSchedule`) that the scheduler should schedule per table is specified. Therefore, as long as there is at least one cell available to schedule the scheduler will guarantee to generate a table containing a minimum `totalCellSlots` value of `minCellsToSchedule`. If the number of outstanding cells available for scheduling is less than `minCellsToSchedule`, idle cells are scheduled to make up the difference. This value is setup once per port and cannot be modified.

*Note:* The `minCellstoSchedule` facility is provided to simplify the transmission control code in the case where queue threshold values are used to drive scheduling. The threshold value in cells can be matched to the `minCellsToSchedule` so that scheduler is always guaranteed to schedule enough cells to fill the Tx Q above its threshold value.

### 6.8.1.2 Maximum Cells Value (maxCells)

The maximum number of cells that the scheduler produces in a table can be limited by the `maxCells` parameter. This can be controllable on a table by table basis. The actual number of cells scheduled is the lesser of `maxCells` and `minCellsToSchedule`.

## 6.8.2 Schedule Service Model

`IxAtmSch` provides schedule service through two functional interfaces: "VC queue update" and "Schedule table update."



The client calls the VC queue update interface whenever the user of the VC submits cells for transmission. The structure of the VC queue update interface is compatible with the requirements of the IxAtmdAcc component.

The client calls the schedule-table-update interface whenever it needs a new table. Internally, IxAtmSch maintains a transmit queue for each VC.

IxAtmSch also provides a “VC queue clear” interface for use when the client wishes to cancel pending demand on a particular VC. This interface is useful, for example, when the client wishes to remove a VC from the system.

### 6.8.3 Timing and Idle Cells

IxAtmSch does not rely on a hardware clock for timing. Instead, the component derives timing information from the supplied port transmit rate for each modeled ATM port. IxAtmSch assumes that  $T = 1/R$  seconds pass for sending every ATM cell. IxAtmSch also assumes that all cells scheduled in a schedule table are transmitted immediately following the cells previously scheduled by the scheduler on that port. (No cells — other than those scheduled by IxAtmSch — are being transmitted on the port.)

The client is responsible for calling “update table” in the following timely fashion, if the demand is always there. Suppose the “update table” calls for a port corresponds to time spending  $T(1), T(2), \dots$ , where one  $T(n)$  is the time needed to transmit cells scheduled in the  $n$ 'th updated table. Then, if the demand is always there, the client must call the  $n$ 'th “update table” before  $T(1)+T(2)+\dots+T(n-1)$  has passed, assuming the client's first such call is at time 0. This can be easily achieved by making sure that port transmission is never empty when the demand is continuously pouring in.

When all registered VC transmit queues are exhausted, an empty schedule table is returned by the `ixAtmSchTableUpdate` interface. It is assumed that the client will instruct the lower layers to transmit idle cells until new cells are submitted for transmit on a registered VC. IxAtmSch is not aware of the number of idle cells transmitted in this situation and will reset its internal clock to its starting configuration when new cells are queued.

A further interface is provided to allow the client to update the transmit port rate of an ATM port which has already been registered with the IxAtmSch device, and may have established VCs with pending transmit demand. This interface is provided to cater for the event of line-rate drift, as can occur on transmit medium.

In the event that the new port rate is insufficient to support all established VC transmit contracts, IxAtmSch will refuse to perform this modification. The client is expected to explicitly remove or modify some established VC in this event, such that all established contracts can be maintained and then resubmit the request to modify the ATM port transmit rate.

## 6.9 Dependencies

The IxAtmSch component has an idealized local view of the system and is not dependent on any other software release 2.3 component.

Some function interfaces supplied by the software release 2.3 component adhere to structure requirements specified by the IxAtmdAcc component. However, no explicit dependency exists between the IxAtmSch component and the IxAtmdAcc component.

## 6.10 Error Handling

IxAtmSch returns an error type to the user when the client is expected to handle the error. Internal errors is reported using standard processor error-reporting techniques.



## 6.11 Memory Requirements

Memory estimates have been sub-divided into two main areas: performance critical and not performance critical.

### 6.11.1 Code Size

The ixAtmSch code size is approximately 35 Kbytes.

### 6.11.2 Data Memory

There are a maximum of 32 VCs per port and 12 ports supported by the IxAtmSch component. These multipliers are used in [Table 23](#).

**Table 23. IxAtmSch Data Memory Usage**

Parameter	Per VC Data (bytes)	Per Port Data + 32 VC data (bytes)	Total (bytes)
Performance-Critical Data	36	$44 + (32 * 36) = 1,196$	14,352
Non-Critical Data	40	$12 + (40 * 32) = 1292$	15,504
Total	76	2,488	29,856

## 6.12 Performance

The key performance measure for the IxAtmSch component is the rate at which it can generate the schedule table, measured by time per cell. The rate at which queue updates are performed is also important. As this second situation will happen less frequently, however — because a great many cells may be queued in one call to the update function — it is of secondary importance.

The remaining functionality provided by the IxAtmSch is infrequent in nature, being used to initialize or modify the configuration of the component. This situation is not performance-critical as it does not affect the data path of the Intel® IXP42X product line.

### 6.12.1 Latency

The transmit latency introduced by the IxAtmSch component into the overall transmit path of the processor is zero under normal operating conditions. This is due to the fact that — when traffic is queued for transmission — scheduling is performed in advance of the cell slots on the physical line becoming available to transmit the cells that are queued.

§ §



## 7.0 Access-Layer Components: Security (IxCryptoAcc) API

---

This chapter describes the security architecture and the IxCryptoAcc access-layer component that provides access to the security features in Intel® IXP400 Software v2.3.

The IxCryptoAcc access-layer component (also referred to as IxCryptoAcc API) is the “Security API” that in most cases offloads encryption, authentication and key management services used in security applications such as IPSec, SSL, TLS and WEP to the coprocessors in the NPE and PKE engine. The offloading of cryptographic functions frees Intel XScale® Processor cycles to implement additional non-crypto features in the security application. The use of coprocessors may also accelerate the execution of cryptographic functions.

### 7.1 What’s New

The following existing routine in IxCryptoAcc API is deprecated and is removed from IXP400 software:

**ixCryptoAccCryptoServiceStop:** The programmer should use the **ixCryptoAccUninit** instead.

### 7.2 Overview

The IxCryptoAcc component provides the following capabilities:

- Support the following cryptographic operation mode:
  - Encryption only
  - Decryption only
  - Authentication Calculation Only
  - Authentication Check Only
  - Encryption followed by Authentication Calculation
  - Authentication Check followed by Decryption.
- Support the following cryptographic algorithms:
  - DES (64 bit block cipher size, 64 bit key)
  - Triple DES (64 bit block cipher size, 3 keys - 64 bit each, hence total key size is 192)
  - AES (128 bit block cipher size, key sizes - 128, 192, 256 bit).
  - ARC4 (8 bit cipher size, input key size of 128 bits only). This cipher algorithm can only be used with no authentication or along with WEP CRC authentication.
- Support the following mode of operation for encryption and decryption:
  - ECB



- CBC
- CTR (for AES algorithm ONLY).  
Notes: For CTR mode, client shall construct the CTR counter block (NONCE + IV + counter) and pass the counter block to the access component as IV.
- Single pass AES-CCM encryption and authentication as specified in 802.11i security specification.  
**Note:** Mode of operations listed above is only applicable to block cipher and not applicable to stream cipher.
- Support the following hashing algorithms:
  - SHA-1 (512 bit data block size, 160 bits message digest size).
  - MD5 (512 bit data block size, 128 bits message digest size).
- Support the following authentication algorithms:
  - HMAC-SHA-1 (512 bit data block size, key size: 20 bytes - 64 bytes, up to 160 bits Message Authentication Code size).
  - HMAC-MD5 (512 bit data block size, key size: 16 bytes - 64 bytes, up to 128 bits Message Authentication Code size).  
**Note:** For key size greater than  $L$  bytes (where  $L$  is the length of the message digest produced by the hashing algorithms as stated as above, different hashing algorithm has different message digest length), the authentication key must be hashed to become shorter key before using it as authentication key for HMAC authentication. For key size less than  $L$ , the authentication key must be padded with 0s to become  $L$  bytes of key by the client before using it as the authentication key for HMAC authentication. The client can specify the MAC length required when registering the cryptographic context. The maximum lengths supported are stated above. The cryptographic component will truncate the generated digest as required. For example a client requiring compatibility with IPsec could specify a 96 bit MAC length using the HMAC-SHA-1 implementation
  - Generation and verification of ICV using the standard 32 bit -CRC polynomial as defined in IEEE-802.11, clause 8.2 - Wired Equivalency Privacy (WEP).
  - Support hashing of authentication key or hash data to become  $L$  bytes ( $L = 20$  for SHA-1 and 16 for MD5) of authentication key or digest on NPE. Minimum of data length is 1 bytes and maximum is 65399 bytes. Padding is taken care by the access component.
- Support a maximum of 10,000 crypto contexts simultaneously. The maximum crypto context can be modified depending on the performance requirements. In order to avoid memory wastage, the default maximum crypto context is set to 1000. The supplied and generated keys during cryptographic process are zeroed when the crypto context is destroyed.
- Report operation failure to client.
- Support reentrancy with the assumptions below:
  - It is client's responsibility to ensure there are no pending requests for the crypto context before un-registering it.
  - Different clients should NOT hold the same crypto context. All the requests for same crypto context should be initiated from same client (thread). In other words, the same client could submit multiple requests for the same crypto context.
- Support implicit cryptographic synchronization only.
- Support SHA-1 hashing using SHA unit on PKE Crypto Engine. Minimum of data length is 1 bytes and maximum is 65399 bytes. Padding is taken care by the access component.



- Support pseudo-random number generation, length unit of pseudo-random number is in words (multiple of 32-bit).
- Support following exponential arithmetic operations.
  - Blinding feature for Exponent/ Modulo arithmetic is enabled as default. See [“Use of Large Number Arithmetic”](#) on page 105 for details.

## 7.3 Internet Security — Background Information

The widely used Internet security standards such as IPSec, SSL, TLS and WEP uses encryption, message authentication and authenticates the identity of the other party.

Encryption converts an unmodified plaintext into a modified ciphertext using a well defined cipher algorithm. A cipher algorithm uses a secret key to covert the plaintext to ciphertext of same length. The Sender encrypts plaintext and transmits the ciphertext over the Internet to the Receiver. The Receiver decrypts the ciphertext into plaintext using a complimentary algorithm and the same or a paired key. This process of encryption and decryption makes it difficult to eavesdrop on the communication between the Sender and Receiver.

Message authentication (Data Integrity) prevents tampering of the data in transit. In its simplest form, the Sender hashes the plaintext using well defined algorithm and produces a Message Digest. The Message Digest is encrypted using another secret key and transmitted to the Receiver along with the ciphertext. The receiver decrypts the Message Digest and compares it with the Message Digest computed by the receiver. If they match then no tampering has taken place.

Encryption and Message authentication does not protect the Receiver from communicating with a fake Sender. This is avoided by the Sender transmitting his Digital Certificate to the Receiver. A Digital Certificate contains the identity of the Sender and is guaranteed by a Certification Authority (CA).

### 7.3.1 Encryption and Message Authentication

Modern cryptography is based on confidentiality, data integrity, authentication of the identity of the communicating entity and making cryptoanalysis difficult. All the above features uses Symmetric and Asymmetric keys. Keys are large integer values with size from 40 bits to 2048 bits or larger. Cryptography based on symmetric keys uses a shared secret between communication entities. Whereas, cryptography based on asymmetric keys assigns a pair of keys — one private and another public — to each communicating entity. Confidential messages sent using the public key can only be deciphered using the matching private key and vice versa. Creating cipher text using a public key is called encryption and using a private key is called the Signing.

Encryption and Message Authentication can be implemented using symmetric or asymmetric keys. Although use of asymmetric keys mitigates the problem of secure key exchange, longer asymmetric keys are required to provide the same level of security as shorter symmetric keys. Also, asymmetric key algorithms are more CPU intensive than symmetric key algorithms. Hence, symmetric keys are widely used for bulk encryption and message authentication.

The security standards IPSec, SSL and TLS use block cipher encryption algorithms DES, 3DES, AES operating in various modes ECB, CBC and CTR. The various modes provide stronger and/or improved performance. The same standards use message authentication algorithms such as HMAC-SHA-1 and HMAC-MD5. The security standard WEP uses stream cipher encryption using ARC4 and ICV computations.



### 7.3.2 Key Management

Cryptoanalysis is the practice of codebreaking and typically involves using a computer in finding the secret keys used by communicating entities. Prolonged use of the same secret symmetric keys provides large amounts of cipher text that aids successful cryptoanalysis. Even if the keys could be changed easily, the manual/third-party-generation of keys and logistics of managing keys for a large or growing network is impractical.

Internet security standards such as IPSec, SSL and TSL use RSA encryption, Diffie-Hellman key exchange protocol to share or independently derive master secret. The master secret is used to generate symmetric keys used in bulk encryption and message authentication. The Internet security standards also use digital certificates signed by CAs (Certifying Authority) to authenticate the pairing of public key with the communicating entity's identity such as domain name and other data such as Diffie-Hellman parameters stored in the certificate.

## 7.4 Security Architecture

This section describes the overall security architecture for the cryptographic services provided by the IxCryptoAcc API and its use in the Internet security applications.

### 7.4.1 IxCryptoAcc Interfaces

The IxCryptoAcc API provides generic crypto, authentication, and key management services to all security applications. In addition, there are configuration routines and utilities to handle special requirements. In Linux\*, the IxCryptoAcc API are available as Kernel mode API.

Table 24 lists the interface routines available in IxCryptoAcc API.

Table 24. IxCryptoAcc API

Routine Name	Function	Type
ixCryptoAccConfig	Selects Crypto interface to initialize	Config
ixCryptoAccInit	Initializes ixCryptoAcc access component	Config
ixCryptoAccCtxRegister	Registers Crypto context	Config
ixCryptoAccCtxUnregister	Unregisters Crypto context	Config
ixCryptoAccCryptoServiceStop	Unregisters and stops all Crypto services (deprecated, use ixCryptoAccUninit instead)	Config
ixCryptoAccUninit	Unregisters crypto contexts and stops all crypto services	Config
ixCryptoAccAuthCryptPerfom	Performs Crypto and/or authentication	Block cipher
ixCryptoAccNpeWepPerform	Performs Crypto and/or ICV calculation for WEP using NPE resources	Stream Cipher
ixCryptoAccXscaleWepPerform	Performs Crypto and/or ICV calculation for WEP using Intel XScale® Processor WEP engine	Stream Cipher
ixCryptoAccCipherKeyUpdate	Updates key without changing context	Block/ Stream Cipher
ixCryptoAccPkeHashPerform	Generates Hash key digest	Key Management
ixCryptoAccPkePseudoRandomNumberGet	Generates pseudo random number	Key Management
ixCryptoAccPkeEauExpConfig	Configures the exponent option for modular exponential operations	Key Management
ixCryptoAccPkeEauPerform	Computes large number arithmetic operations	Key Management
ixCryptoAccShow	Prints statistics and status	Utility

**Table 24. IxCryptoAcc API**

Routine Name	Function	Type
ixCryptoAccShowWithId	Prints statistics, status and crypto contexts	Utility
ixCryptoAccHashKeyGenerate	Generates hash digest using NPE Hashing coprocessor. Also a generic hashing function.	Utility
ixCryptoAccHashPerform	Alias for ixCryptoAccHashKeyGenerate	Utility

The API utilizes a number of other access-layer components, including the Intel XScale® Processor as well as hardware-based crypto functionality available on the NPEs and PKE Crypto engine. Figure 27 on page 101 shows the high-level architecture of IxCryptoAcc.

#### 7.4.1.1 Intel XScale® Processor Software

The Intel XScale® Processor WEP Engine is a software-based “engine” for performing ARC4 and WEP ICV calculations used by WEP clients. While this differs from the model of NPE-based crypto functionality, it provides additionally design flexibility for products that require NPE A to perform non-crypto operations.

#### 7.4.1.2 Offloading to Hardware

IxQMgr is another access-layer component that interfaces to the hardware-based AHB Queue Manager (AQM). The AQM is SRAM memory used to store pointers to data in SDRAM memory, which is accessible by both the Intel XScale® Processor and the NPEs. These items are the mechanism by which data is transferred between IxCryptoAcc and the NPEs. Separate hardware queues are used for both IPSec and WEP services.

The security application can offload computations for block cipher algorithms (DES, 3DES, AES operating in various modes ECB, CBC and CTR) and corresponding message authentication (HMAC-SHA-1 and HMAC-MD5) to the coprocessors in NPE C. Similarly, the computations for stream cipher ARC4 and ICV computation can be offloaded to the NPE A processor and coprocessors. Note that ARC4 algorithm is executed by the NPE A core and not by a coprocessor. The ICV algorithm uses AAL coprocessor to compute CRC32.

The RSA, DSS and Diffie-Hellman algorithms and protocols use random numbers, hashing and large number arithmetic extensively. These functions can be offloaded by Intel XScale® Processor to the SHA-1, RNG and EAU coprocessors in the PKE Crypto engine. Unlike the crypto support in NPEs, Intel XScale® Processor directly reads from and writes to the registers of the PKE Crypto engine coprocessors.

#### 7.4.1.3 API Usage

In this chapter IPSec and WEP are used as examples that makes use of cryptoAcc access-layer API to perform the authentication and encryption operations. In this software release, the IxCryptoAccCodelet is provided as an example of application software.

### 7.4.2 Basic API Flow

This section describes a high-level flow of the IxCryptoAcc API. A more detailed example of API usage is provided in a subsequent section.

#### Encryption/ Decryption and authentication using NPE

The flow of the API is similar for both IPSec and WEP services:

- The client application initializes the IxCryptoAcc access component.



- The client application defines the cryptographic context consisting of crypto algorithm and its mode to be used, direction of data and pointers to the callback routines.
- Packets for encryption, decryption and authentication are prepared by the client and passed to one of the “Perform” API routines along with the cryptographic context defined in the previous step.
- The IxCryptoAcc API invokes ixQMgr that instructs the relevant NPE to gather the data and crypto context from the SDRAM.
- The NPE performs requested encryption/decryption and/or authentication functions using the appropriate coprocessors and stores the results in SDRAM.
- Finally the previously registered callback function is executed.

### **Encryption/Decryption and authentication using Intel XScale® Processor WEP engine**

- The client application initializes the IxCryptoAcc access component.
- The client application defines the cryptographic context consisting of crypto algorithm and its mode to be used, direction of data and pointers to the callback routines.
- Packets for encryption, decryption and authentication are prepared by the client and passed to the synchronous “Perform” API routine along with the cryptographic context defined in the previous step.
- The Intel XScale® Processor performs requested encryption/decryption and/or authentication function using the optimized software and returns with the results. Note that a call to the Intel XScale® Processor WEP engine is a synchronous call and no callback function is involved.

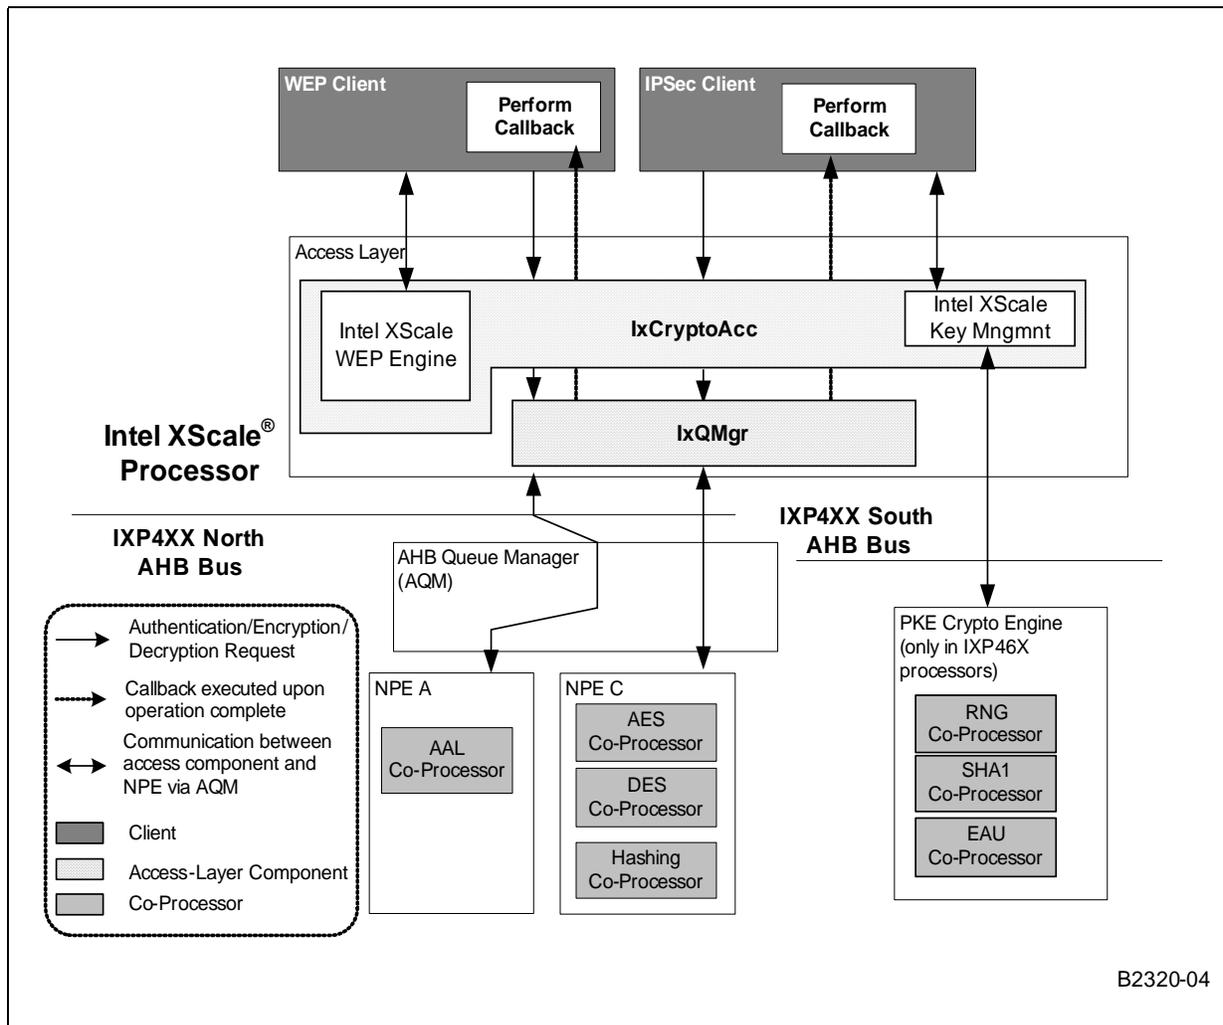
### **Key management function using PKE Crypto Engine**

- The client application initializes the IxCryptoAcc access component.
- The relevant “Pke.Perform” API is called with relevant parameters.
- The Intel XScale® Processor performs requested function using the appropriate coprocessors and stores the results in SDRAM.
- Pseudo Random number generation API is a synchronous call that returns with a result. The Hashing and Exponent arithmetic functions invoke the handler bound to its interrupt.

The basic API flow described above is shown in [Figure 27](#).



Figure 27. Basic IxCryptoAcc API Flow



### 7.4.3 Context Registration and the Cryptographic Context Database

The IxCryptoAcc access component supports up to 1,000 simultaneous security association (SA) tunnels. While the term SA is well-known in the context of IPSec services, the IxCryptoAcc component defines these security associations more generically, as they can be used for WEP services as well. Note that Key Management function does not use Context Registration and Queue Manager features.

Depending upon the application's requirements, the maximum active tunnels supported by IxCryptoAcc access-layer component can be changed by the client. The number of active tunnels will not have any impact on the performance, but will have an impact on the memory needed to keep the crypto context information. The memory requirement will depend on the number of tunnels.

Each cryptographic "connection" is defined by registering it as a cryptographic context containing information such as algorithms, keys, and modes. Each of these connections is given an ID during the context registration process and stored in the Cryptographic



Context Database. The information stored in the CCD is stored in a structure detailed below, and is used by the NPE or Intel XScale® Processor WEP Engine to determine the specific details of how to perform the cryptographic processing on submitted data.

The context-registration process creates the structures within the CCD, but the crypto context for each connection must be previously defined in an IxCryptoAccCtx structure. The IxCryptoAccCtx structure contains the following information:

- The type of operation for this context. For example, encrypt, decrypt, authenticate, encrypt and authenticate, and so forth
- Cipher parameters, such as algorithm, mode, and key length
- Authentication parameters, such as algorithm, digest length, and hash length
- In-place versus non-in-place operation. In-place operation means once the crypto processing of the source data is completed, the resulting data is placed onto the same IX\_OSAL\_MBUF as it was read from.

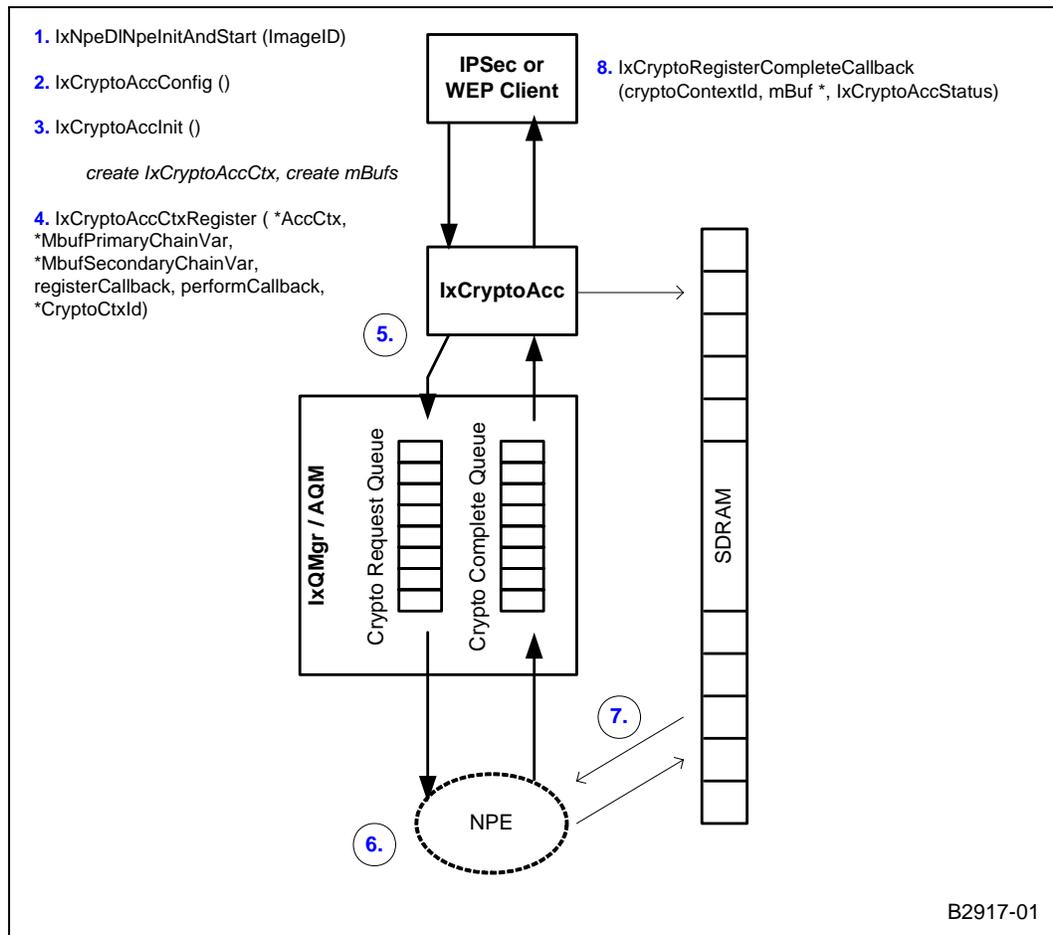
When the client performs calls the `ixCryptoAccCtxRegister()` function, the following data must be provided or received:

- The client provides a pointer to the crypto context (for example., SA definition) being registered.
- The client is required to allocate two IX\_OSAL\_MBUFs for the access component to compute the primary and secondary chaining variables for SHA-1/MD5.
- The client must register two callbacks. One callback is executed upon the completion of the registration function, the second is executed each time a cryptographic procedure (“perform” functions) has completed on the NPE for this context. There is one exception for the perform callback function, noted in section “[ixCryptoAccXscaleWepPerform\(\)](#)” on page 120.
- The function returns a context ID upon successful registration in the CCD.

[Figure 28 on page 103](#) shows the IxCryptoAcc API call process flow that occurs when registering security associations within the CCD. This process is identical for both IPsec and WEP services except in situations where NPE crypto functionality will not be used, such as when using WEP services using only the Intel XScale® Processor WEP engine. For more detailed information on this usage model see “[ixCryptoAccXscaleWepPerform\(\)](#)” on page 120.



Figure 28. IxCryptoAcc API Call Process Flow for CCD Updates



1. The proper NPE microcode images must be downloaded to the NPEs and initialized, if applicable.
2. IxCryptoAcc must be configured appropriately according to the NPEs and services that is utilized. By default, IxCryptoAccConfig() configured the component for using NPE C and enabled the Intel XScale® Processor WEP engine.
3. IxCryptoAcc must be initialized. At this point the client application should define the crypto context to be registered, as well as create the buffers for the initial chaining variables.
4. The crypto context must be registered using the IxCryptoAccCtxRegister() function.
5. The IxCryptoAcc API will write the crypto context structure to SDRAM. If NPE-based crypto functionality is being used then IxCryptoAcc will use IxQMigr to place a descriptor for the crypto context being registered into the Crypto Request Queue.
6. The NPE will read the descriptor on the Crypto Ready Queue, generate any reverse keys required, and generate the initial chaining variable if required.
7. The NPE or Intel XScale® Processor WEP Engine writes the resulting data in the Crypto Context Database residing in SDRAM. The NPE will then enqueue a descriptor onto the Crypto Complete Queue to alert the IxCryptoAcc component that registration is complete.



8. IxCryptoAcc will return a context Id to the client application upon successful context registration, and will call the Register Complete callback function.

#### 7.4.4 Buffer and Queue Management

The IX\_OSAL\_MBUF buffer format is for use between the IxCryptoAcc access component and the client when doing encryption and message authentication. Key management functions don't use IX\_OSAL\_MBUF buffers. All buffers used between the IxCryptoAcc access component and clients are allocated and freed by the clients. The client will allocate the IX\_OSAL\_MBUFs and the buffers is passed to IxCryptoAcc. The CryptoAcc access-layer component will allocate memory for the CCD. The client passes a buffer to IxCryptoAcc when it requests NPE based crypto services, and the IxCryptoAcc component returns the buffer to the client when the requested job is done.

The component assumes that the allocated IX\_OSAL\_MBUFs are sufficient in length and no checking has been put in place for the IX\_OSAL\_MBUF length within the IX\_OSAL\_MBUF structure. There is, however, IX\_OSAL\_MBUF checking when the code is compiled in DEBUG mode. When appending the ICV at the end of the payload, it is assumed that the IX\_OSAL\_MBUF's length is sufficient and will not cause memory segmentation. The ICV offset should be within the length of the IX\_OSAL\_MBUF.

Depending on the transfer mode in-place before returning the buffer to the client, the encrypted / decrypted payload is written into the source buffer or destination buffer. This selection of in-place versus non-in-place buffer operation may be defined for each crypto context prior to context registration.

When the AHB Queue Manager is full, the NPE crypto hardware will return IX\_CRYPTACC\_QUEUE\_FULL to the client. The client will must re-send the data to be encrypted or decrypted or authenticated after a random interval.

#### 7.4.5 Using Key Management Service

The key management service API are wrapped in an #ifdef in the software. They are available only when the build environment variables are configured for the IXP46X network processors. Use of key management services in an IXP42X product line build will result in compile time errors.

The client should check for the return status of the API because, unlike the NPE-based services, there is no queuing mechanism available. The mutex mechanism is used to ensure that only one operation can be executed at a time. If RETRY status is returned, the coprocessor is in use and the clients will must resend the request after a random interval.

The key management client uses hashing, random numbers and large number arithmetic extensively, and most of these functions could be offloaded to the PKE Crypto engine.

##### 7.4.5.1 Use of Random Numbers

Random numbers are integral to generation and update of symmetric keys, implementation of PFS (perfect forward secrecy) and protection from replay attacks. Also, random number-based cookies are included in messages to provide protection from Denial of Service attacks.

Random numbers are also used in the creation of asymmetric keys. The RSA asymmetric keys are created starting with two large prime numbers. There is no magic formula to create a prime number. Since prime numbers are random in nature, we start with a random number of desired length and then employ mathematical methods to compute the probability of the generated random number being a prime. If the probability is sufficient high then we assume the random number to be a prime number,



otherwise repeat the above process with another random number. Internet security standards such as IPSec, SSL and TSL use RSA encryption, Diffie-Hellman key exchange protocol to share or independently derive master secret. The master secret is used to generate symmetric keys used in bulk encryption and message authentication and select the one that has a high probability of being a prime. High-quality random numbers can be generated by creating a pseudo random number and randomizing the value further using SHA-1 and MD5 Hashing. The key creation client can offload pseudo random number generation and SHA-1 hashing to the RNG and SHA-1 coprocessors in the PKE Crypto engine.

The client can use `ixCryptoAccPkeRandomNumberGet()` to generate a pseudo random number of user specified length.

#### 7.4.5.2 Use of Hashing

SHA-1 hashing is used in symmetric key generation algorithms, random number generation algorithm, and Diffie-Hellman key authentication. Any SHA-1 hashing can be offloaded to the Hashing coprocessor in the PKE Crypto engine.

The client can use `ixCryptoAccPkeHashPerform()` to hash a given data.

#### 7.4.5.3 Use of Large Number Arithmetic

Asymmetric key generation, encryption, decryption, signing and signature verification algorithms employed in RSA and DSS use very large number arithmetic including addition, subtraction, multiplication, exponentiation and modular operations. The security application can offload all large number computations except division to EAU coprocessor in the PKE Crypto engine.

##### Asymmetric Key Generation

Since EAU coprocessor does not support large number division, only parts of the asymmetric key creation can be offloaded to the EAU coprocessor.

##### Timing Attacks

A timing attack is where the attacker tries to uncover private keys by analyzing the time taken to execute cryptographic algorithms. In asymmetric key algorithms, computation time for a private key operation is dependent on the key in some way. Blinding techniques can be used to remove the correlation between private key and encryption time. Blinding can be enforced by disabling Short Exponent and Fast Exponent setting and is the default behavior. Enabling Short Exponent and Fast Exponent speeds up computations for Exponent/ Modulo arithmetic.

The client can use `ixCryptoAccPkeEauExpConfig()` and `ixCryptoAccPkeEauPerform()` to configure and execute large number arithmetic.

#### 7.4.6 Config and Utility Functionality

The config and utility functionality provide the following features:

- A number of status definitions, useful for determining the cause of registration or cryptographic processing errors.
- The ability to un-register a specific crypto context from the CCD.
- Two status and statistics functions are provided. These function show information such as the number of packets returned with operation fail, number of packets encrypted/ decrypted/authenticated, the current status of the queue, whether the queue is empty or full or current queue length.
- The ability to stop the Hardware Crypto functionality and WEP engine services.



The following functions are used in specific situations that merit further explanation.

**ixCryptoAccHashKeyGenerate()**

This is a generic SHA-1 or MD5 hashing function that takes as input the specification of a basic hashing algorithm, some data and the length of the digest output. There are several useful scenarios for this function.

This function should be used in situations where an HMAC authentication key of greater than 64 bytes is required for a crypto context, and should be called prior to registering that crypto context in the CCD. An initialization vector is supplied as input.

The function can also be used by SSL client applications as part of the SSL protocol MAC generation by supplying the record protocol data as input.

The Hashing request is queued and executed in the NPE C.

**ixCryptoAccHashPerform()**

The function ixCryptoAccHashKeyGenerate has been extended to become a generic hashing function. It could be used to hash any key size or data that is greater than 0 and less than 65339 bytes. ixCryptoAccHashPerform is an alias for ixCryptoAccHashKeyGenerate.

**ixCryptoAccPkeHashPerform()**

This function will use the Hashing coprocessor in the PKE Crypto engine. This function will use more Intel XScale® Processor bandwidth compared to ixCryptoAccHashPerform because Intel XScale® Processor does memory-mapped IO to read and write to the SHA-1 coprocessor in the PKE Crypto engine.

The client can determine which Hash Perform routine is appropriate for the application. The differences between the two HashPerform functions are shown in Table 25.

**Table 25. Difference Between the NPE and PKE-Based Hash Routine**

<b>ixCryptoAccHashPerform()</b>	<b>ixCryptoAccPkeHashPerform()</b>
Requests are queued	Requests are executed immediately if free. Otherwise, the operation should be retried after random time.
Can be used to compute both SHA-1 and MD5	Can be used to compute only SHA-1
Executed in NPE C	Executed in PKE Crypto engine
Uses less Intel XScale® Processor cycles	Uses more Intel XScale® Processor cycles

**ixCryptoAccCtxCipherKeyUpdate()**

This function is called to change the key value of a previously registered context. Key change for a registered context is only supported for CCM cipher mode. This is done in order to quickly change keys for CCM mode, without going through the process of context deregistration and registration. Changes to the key lengths are not allowed for a registered context. This function should only be used if one is invoking cryptographic operations using CCM as cipher mode.

The client should make sure that there are no pending requests on the "cryptoCtxId" for the key change to happen successfully. If there are pending requests on this context the result of those operations are undefined.



For contexts registered with other modes, the client should unregister and re-register a context for the particular security association in order to change keys and other parameters.

### 7.4.7 Memory Requirements

This section shows the amount of data memory required by IxCryptoAcc for it to operate under peak call-traffic load. The IxCryptoAcc component allocates its own memory for the CCD to store the required information, and for the NPE queue descriptors required when using NPE-based crypto functionality. The total memory allocation follows this general formula:

Total Memory Allocation = (Size of NPE queue descriptor + size of additional authentication data) \* Number of descriptors + (size of crypto context) \* (number of crypto contexts).

This shows the memory requirements for 1,000 security associations, the default value set by IX\_CRYPTTO\_ACC\_MAX\_ACTIVE\_SA\_TUNNELS. This value can be increased or decreased as needed by the client.

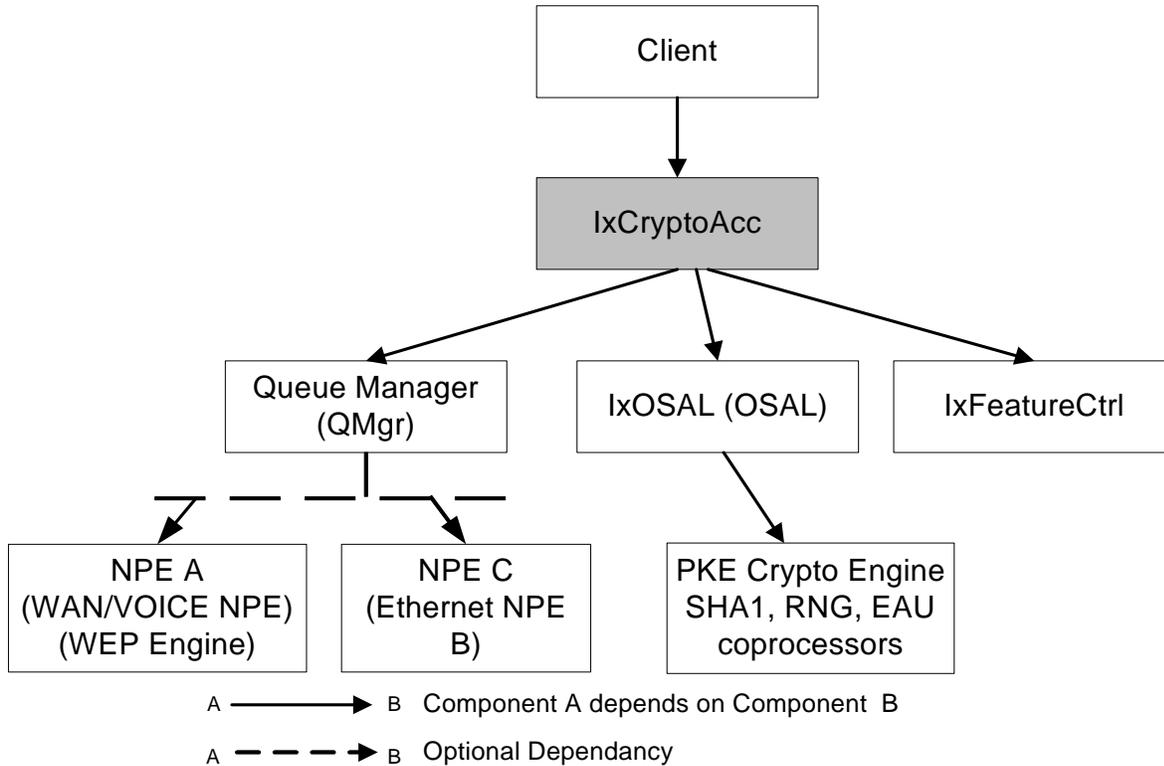
**Table 26. IxCryptoAcc Data Memory Usage**

Structure	Size in Bytes	Total Size in Bytes
NPE Queue Descriptor	96	
Additional Authentication Data	64	
Total Memory per NPE Descriptor	96+64=160	
Number of NPE Descriptors	278	
Total Memory Allocated for NPE Descriptors	160 * 278=	44,480
Crypto Context	152	
Number of Crypto Context (IX_CRYPTTO_ACC_MAX_ACTIVE_SA_TUNNELS)	1,000	
Total Memory Allocated for Crypto Contexts	152 * 1000=	152,000
Size of KeyCryptoParam Structures	256	
Total memory allocated for KeyCryptoParam Structures	104*256	26624
Total Memory Allocated by IxCryptoAcc	44480 + 152000 + 26624=	~218Kbytes

### 7.4.8 Dependencies

Figure 29 shows the component dependencies of the IxCryptoAcc component.

Figure 29. IxCryptoAcc Component Dependencies



B3835-02

Figure 29 can be summarized as follows:

- Client component will call IxCryptoAcc for cryptographic services. NPE will perform the encryption, decryption, and authentication process via IxQMgr.
- IxCryptoAcc depends on the IxQMgr component to configure and use the hardware queues to access the NPE.
- OS Abstraction Layer access-component is used for key management services. Coprocessors in the PKE Crypto engine will perform SHA-1 hashing, pseudo random number generation and large number arithmetic.
- OS Abstraction Layer access-component is used for error handling and reporting, IX\_OSAL\_MBUF handling, endianness handling, mutex handling, and for memory allocation.
- IxFeatureCtrl access-layer component is used to detect the processor capabilities at runtime, to ensure the necessary coprocessors are available for the requested cryptographic context registrations. The IxFeatureCtrl will only issue an warning and will not return any errors if it detects that the coprocessors are not available on the silicon. The client should make sure that they do not use the cryptographic features if a particular version of silicon does not support the cryptographic features.



- In situations where only the Intel XScale® Processor WEP Engine is used, the IxQMgr component is not utilized. Instead, local memory is used to pass context between the IxCryptoAcc API and the Intel XScale® Processor WEP Engine.

After the CCD has been updated, the API can then be used to perform cryptographic processing on client data, for a given crypto context. This service request functionality of the API is described in “IPSec Services” on page 109 and “WEP Services” on page 118.

### 7.4.9 Error Handling

IxCryptoAcc returns an error type to the client and the client is expected to handle the error. Internal errors is reported using an IxCryptoAcc-specific, error-handling mechanism listed in IxCryptoAccStatus.

### 7.4.10 Endianness

The mode supported by this component is both big endian and little endian.

### 7.4.11 Import and Export of Cryptographic Technology

Some of the cryptographic technologies provided by this software (such as 3DES and AES) may be subjected to both export controls from the United States and import controls worldwide. Where local regulations prohibit, some described modes of operation may be disabled.

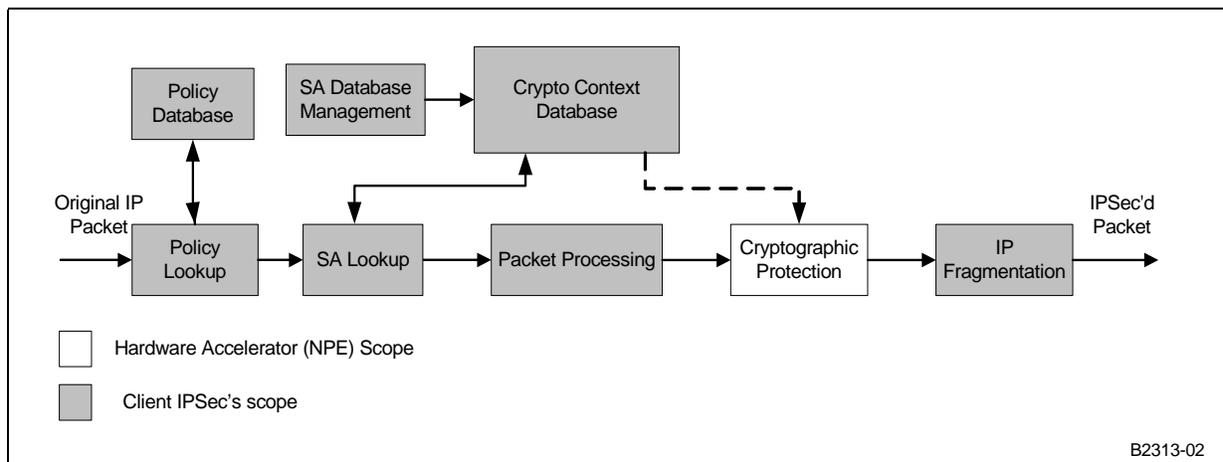
## 7.5 IPSec Services

This section describes the way that IxCryptoAcc is used in an IPSec usage model.

### 7.5.1 IPSec Background and Implementation

When deploying IPSec-related applications, the generalized architecture in Figure 30 is used. The figure shows the scope and the roles played by the NPE and the IxCryptoAcc component in an IPSec application.

Figure 30. IxCryptoAcc, NPE and IPSec Stack Scope



The IPsec protocol stack provides security for the transported packets by encrypting and authenticating the IP payload. Before an IP packet is sent out to the public network, it is processed by the IPsec application (the IxCryptoAcc and supporting components, in this scenario) to encapsulate the IP packet into the ESP or AH packet format.

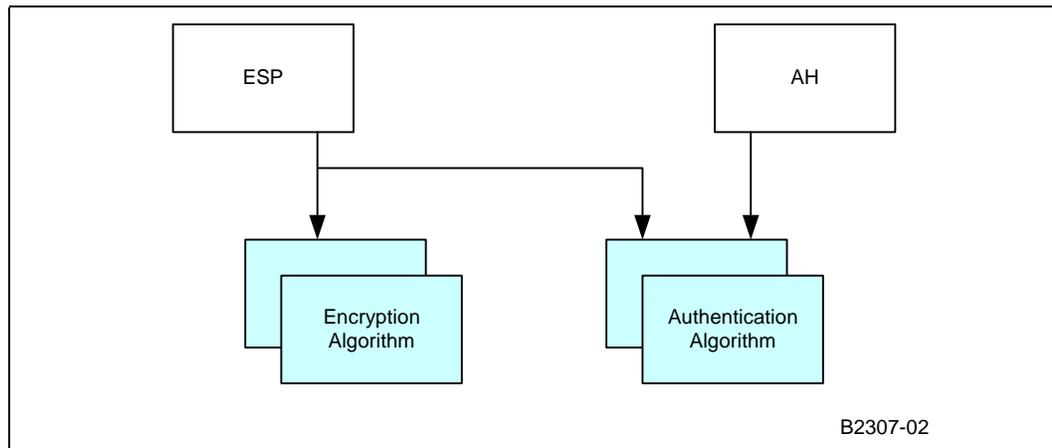
The information within the SA database that is required for the cryptographic protection is passed in via the client to the NPE (in the Cryptographic Protection Block). The client looks up the crypto context policy and SA database to determine the mode of transporting packets, the IPsec protocol (ESP or AH), and so forth. The client determines use of the transport or tunnel mode from the registered security context. The mode is transparent to the NPE crypto hardware and the ixCryptoAcc component.

The client processes the IP packet into ESP- or AH-packet format, the IP packet is padded accordingly (if ESP is chosen), and the IP header mutable fields are handled (if AH). Then, based on the SA information, the NPE executes cryptographic protection algorithms (encryption and/or authentication). This is done regardless of whether transport or tunnel mode is used.

The client sends out the protected IP packet after the cryptographic protection is applied. If the IP packet is too large in size, the client fragments the packet before sending.

Figure 31 shows the relationship of encryption and authentication algorithms within the IPsec protocol.

Figure 31. Relationship Between IPsec Protocol and Algorithms



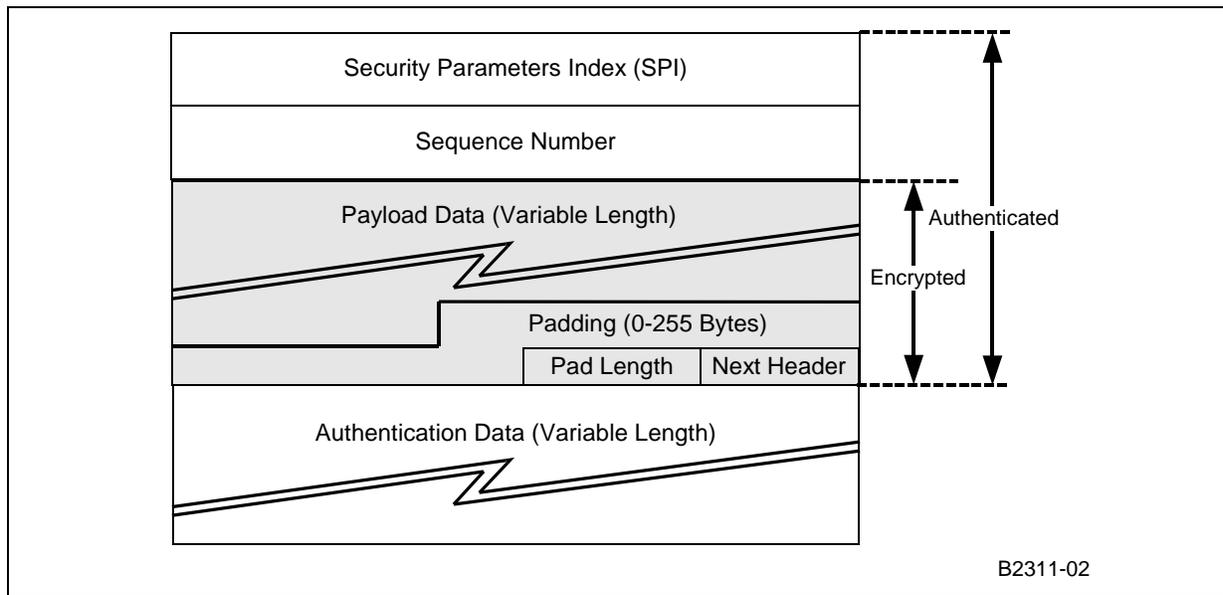
### 7.5.2 IPsec Packet Formats

IPsec standards have defined packet formats. The authentication header (AH) provides data integrity and the encapsulating security payload (ESP) provides confidentiality and data integrity. In conjunction with SHA-1 and MD5 algorithms, both AH and ESP provide data integrity. The IxCryptoAcc component supports both different modes of authentication. The ICV is calculated through SHA-1 or MD5 and inserted into the AH packet and ESP packet.

In ESP authentication mode, the ICV is appended at the end of the packet, which is after the ESP trailer if encryption is required.



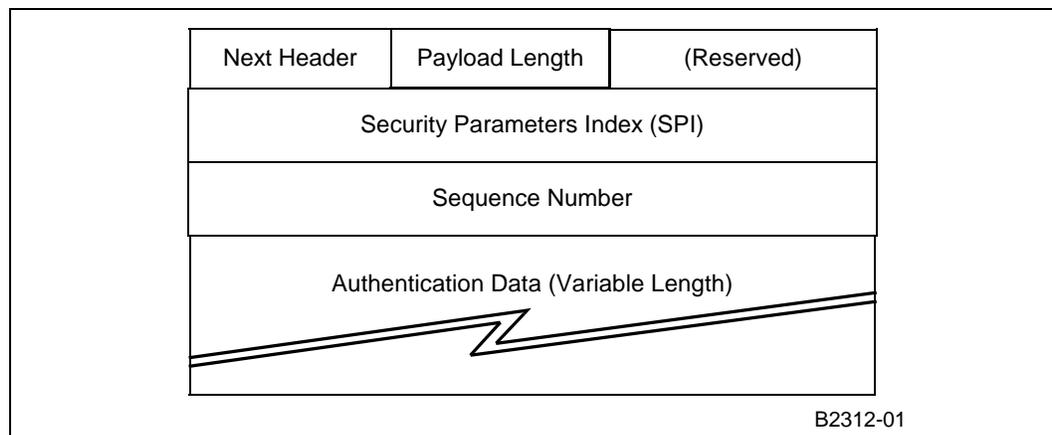
**Figure 32. ESP Packet Structure**



In AH mode, the ICV value is part of the authentication header. AH is embedded in the data to be protected. This results in AH being included for ICV calculation, which means the authentication data field (ICV value) must be cleared before executing the ICV calculation. The same applies to the ICV verification — the authentication data needing to be cleared before the ICV value is calculated and compared with the original ICV value in the packet. If the ICV values don't match, authentication is failed.

NPE determines where to insert the ICV value, based on the ICV offset specified in the perform function.

**Figure 33. Authentication Header**

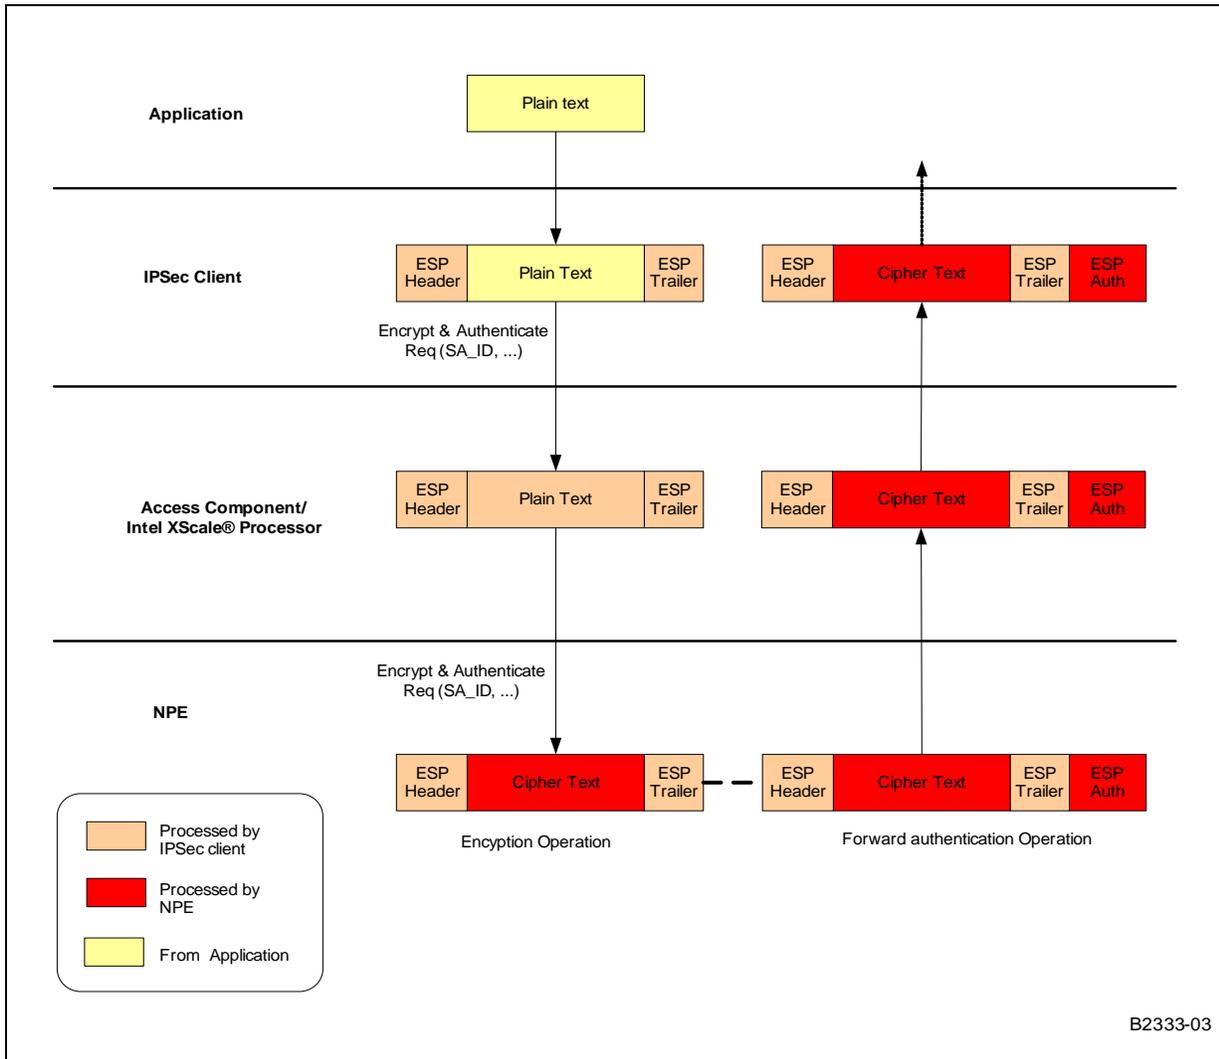


### 7.5.2.1 Reference ESP Dataflow

Figure 34 shows the example data flow for IP Security environment. Transport mode ESP is used in this example. The IP header is not indicated in the figure.

The IP header is located in front of the ESP header while plain text is the IP payload.

Figure 34. ESP Data Flow

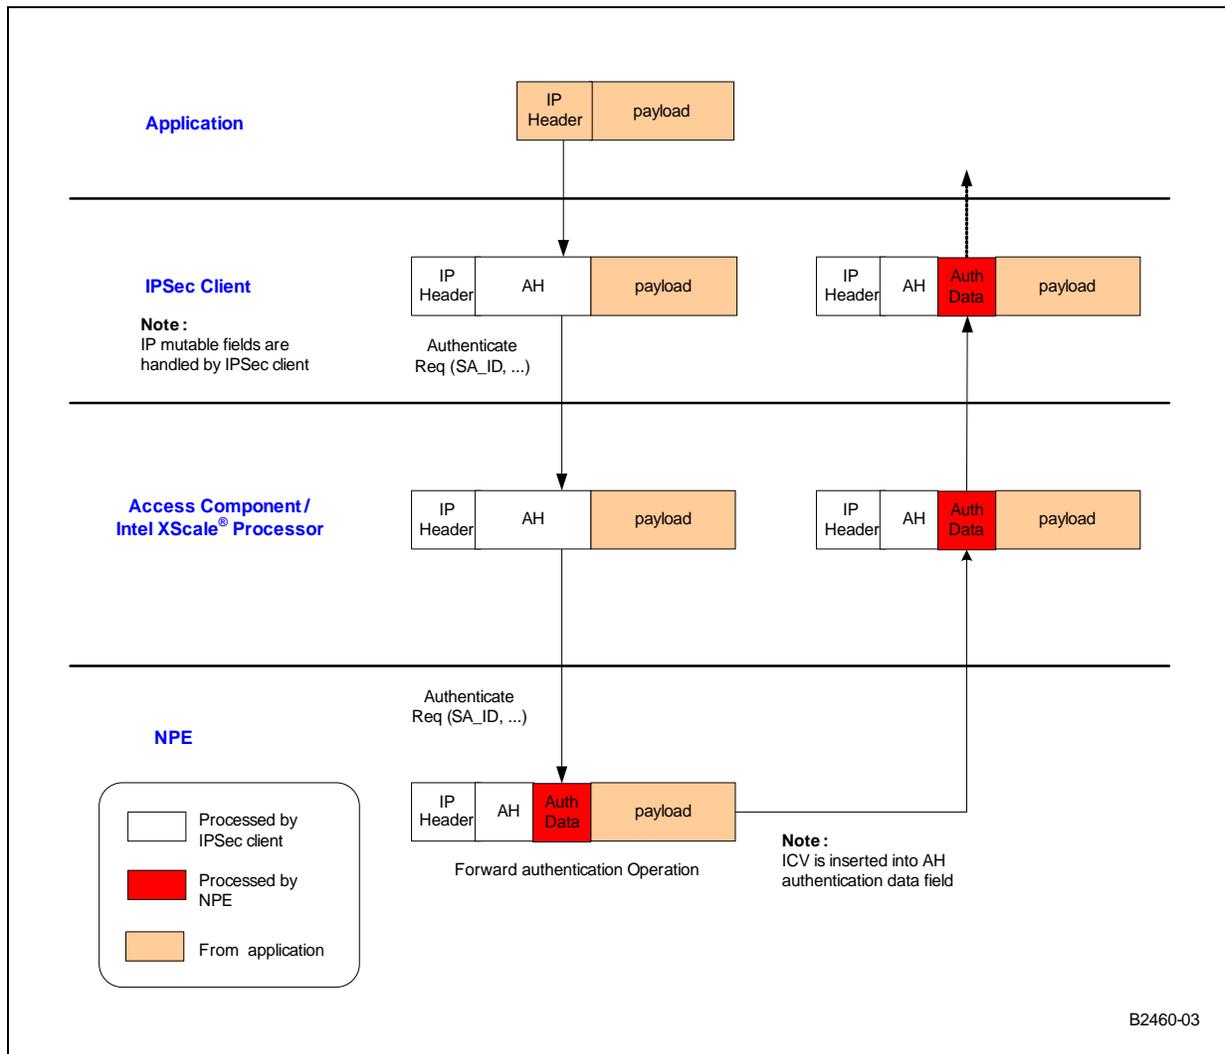


### 7.5.2.2 Reference AH Dataflow

Figure 35 shows the example data flow for IP Security environment. Transport mode AH is used in this example. IPSec client handles IP header mutable fields.



Figure 35. AH Data Flow



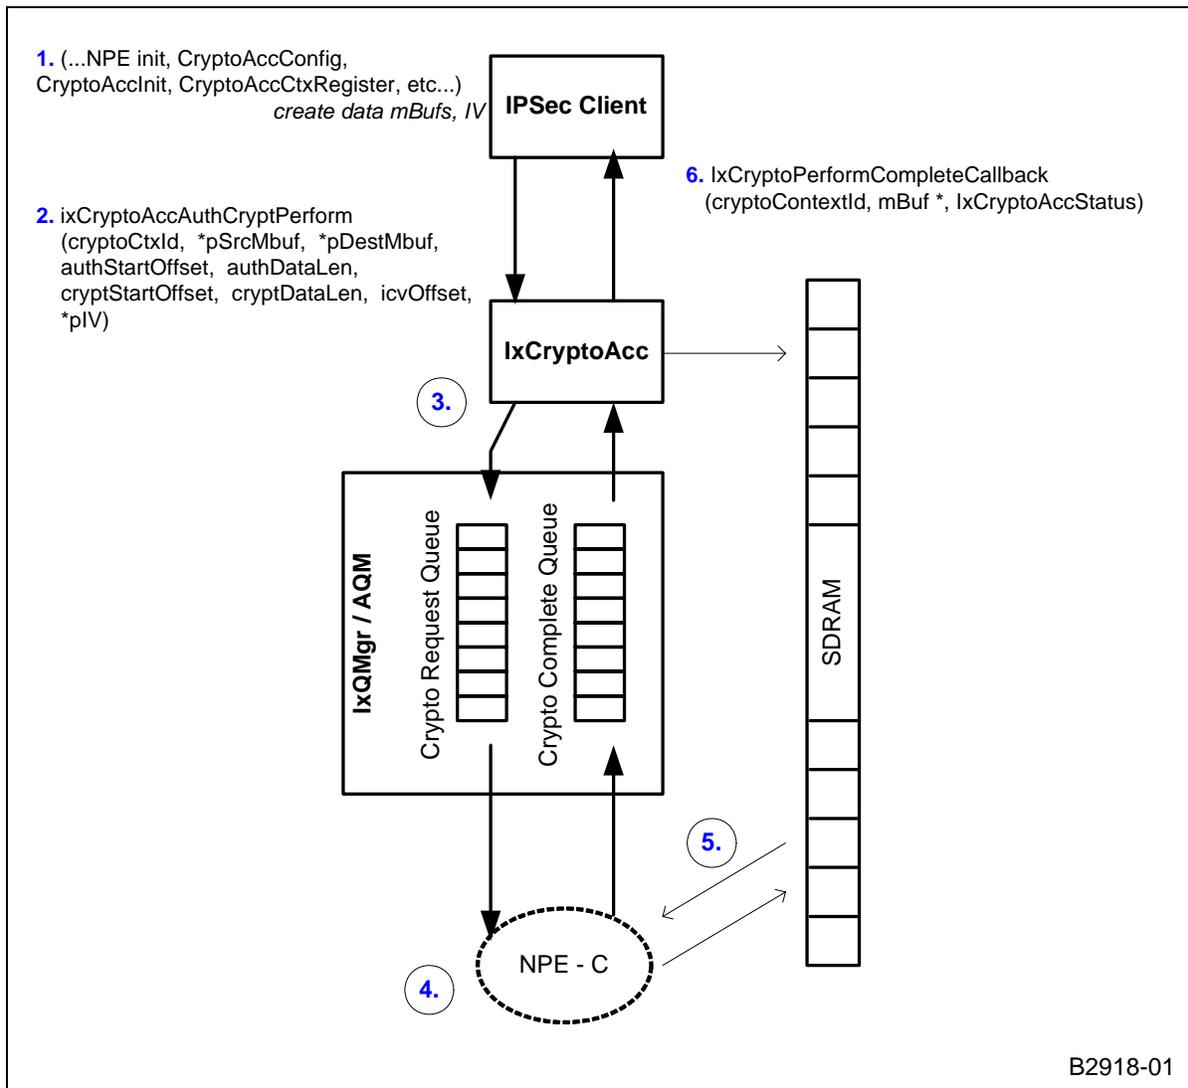
### 7.5.3 Hardware Support for IPSec Services

The IxCryptoAcc API is dependant upon hardware resources within NPE C (also known as Ethernet NPE B) in order to perform many of the cryptographic encryption, decryption, or authentication functions. Specifically, NPE C provides an AES coprocessor, DES coprocessor and a hashing coprocessor (for MD5 and SHA-1 calculations).

### 7.5.4 IPSec API Call Flow

Figure 36 on page 114 details the IxCryptoAcc API call flow that occurs when submitted data for processing using IPSec services. The process assumes that the API has been properly configured and that a crypto context has been created and registered in the CCD, as described in “Context Registration and the Cryptographic Context Database” on page 101.

Figure 36. IPSec API Call Flow



B2918-01

1. The proper NPE microcode images must have been downloaded to the NPE and initialized. Additionally, the IxCryptoAcc API must be properly configured, initialized, and the crypto context registration procedure must have completed. At this point, the client must create the IX\_OSAL\_MBUFs that will hold the target data and populate the source IX\_OSAL\_MBUF with the data to be operated on. Depending on the encryption/decryption mode being used, the client must supply an initialization vector for the AES or DES algorithm.
2. The client submits the `IxCryptoAccAuthCryptPerform()` function, supplying the crypto context ID, pointers to the source and destination buffer, offset and length of the authentication and crypto data, offset to the integrity check value, and a pointer to the initialization vector.
3. IxCryptoAcc uses IxQMgr to place a descriptor for the data into the Crypto Request Queue.
4. The NPE will read the descriptor on the Crypto Ready Queue and performs the encryption/decryption/authentication operations, as defined in the CCD for the



submitted crypto context. The NPE inserts the Integrity Checksum Value (ICV) for a forward-authentication operation and verifies the ICV for a reverse-authentication operation.

5. The NPE writes the resulting data to the destination IX\_OSAL\_MBUF in SDRAM. This may be the same IX\_OSAL\_MBUF in which the original source data was located, if the crypto context defined in-place operations. The NPE will then enqueue a descriptor onto the Crypto Complete Queue to alert the IxCryptoAcc component that the perform operation is complete.
6. IxCryptoAcc will call the registered Perform Complete callback function.

## 7.5.5 Special API Use Cases

### 7.5.5.1 HMAC with Key Size Greater Than 64 Bytes

As specified in the RFC 2104, the authentication key used in HMAC operation must be at least of L bytes length, where L = 20 bytes for SHA-1 or L = 16 bytes for MD5. Authentication key with a key length greater than or equal to 'L' and less than or equal to 64 bytes can be used directly in HMAC authentication operation. No further hashing of authentication key is needed. Thus the authentication key can be used directly in crypto context registration.

However, authentication key with key length greater than 64 bytes must be hashed to become L bytes of key size before it can be used in HMAC authentication operation. The authentication key must be hashed before calling crypto context registration API as shown in steps below:

- a. Call `ixCryptoAccHashKeyGenerate()` function and pass in the original authentication key using an IX\_OSAL\_MBUF. Also, you will need to register a callback function for when this operation is complete.
- b. Wait for callback from IxCryptoAcc.
- c. Copy generated authentication key from IX\_OSAL\_MBUF into a cryptographic context structure (`IxCryptoAccCtx`) and call `ixCryptoAccCtxRegister()` to register the crypto context for this HMAC operation.

### 7.5.5.2 Performing CCM (AES CTR-Mode Encryption and AES CBC-MAC Authentication) for IPSec

A generic CCM cipher is not supported in the software release 2.3. However, it is possible to perform AES-CCM operations in an IPSec-application style. Single-pass AES-CCM is supported for WEP Services only, as documented in "[Counter-Mode Encryption with CBC-MAC Authentication \(CCM\) for CCMP in 802.11i](#)" on page 124.

The overall strategy to accomplish the AES-CCM request involves two operations. The first operation does the AES-CBC operation to get the CBC-MAC. The second operation is to perform a AES-CTR encryption operation to encrypt the payload and create the CBC-MAC to get the MIC. Two crypto contexts are registered and two crypto perform service requests are invoked in order to complete the encryption and authentication for a packet.

[Figure 37 on page 116](#) and [Figure 38 on page 116](#) show the steps needed to encrypt and authenticate a packet in general by using CCM mode. Those steps are:

1. Use AES CBC-MAC to compute a MIC on plaintext header, and payload.  
The last cipher block from this operation becomes MIC.
2. Use AES-CTR mode to encrypt the payload with counter values 1, 2, 3, ...
3. Use AES-CTR mode to encrypt the MIC with counter value 0 (First key stream (S0) from AES-CTR operation)

Figure 37. CCM Operation Flow

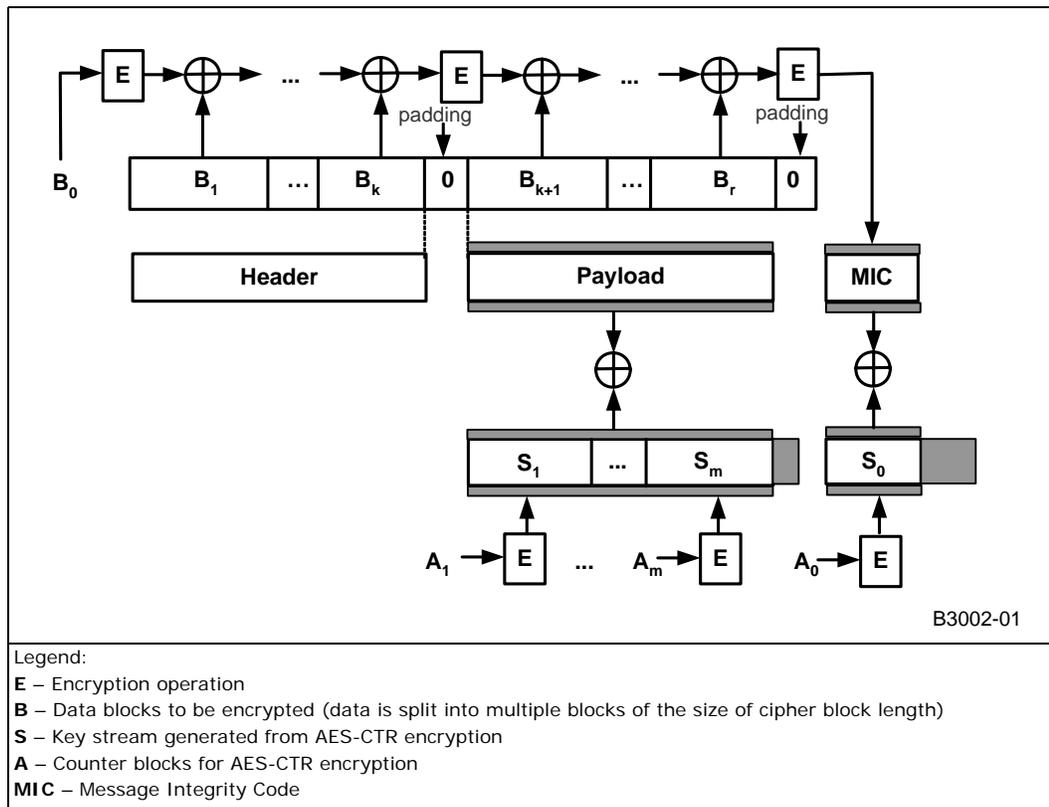
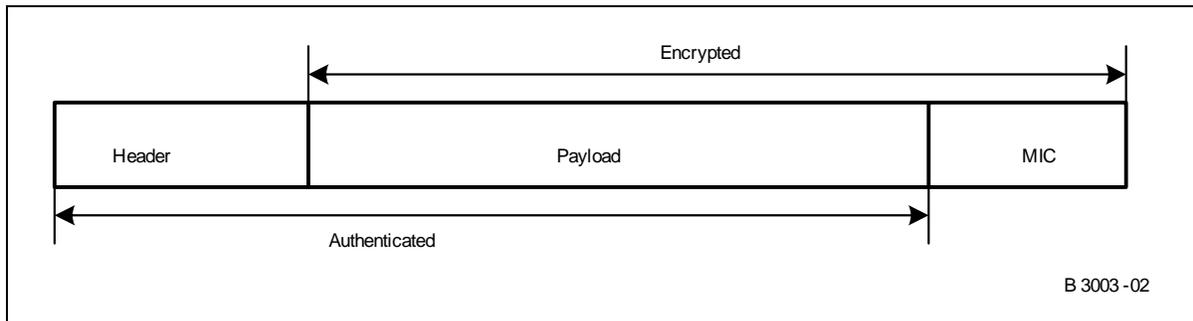


Figure 38. CCM Operation on Data Packet



The API usage for performing an IPSec-style AES-CCM operation is as follows:

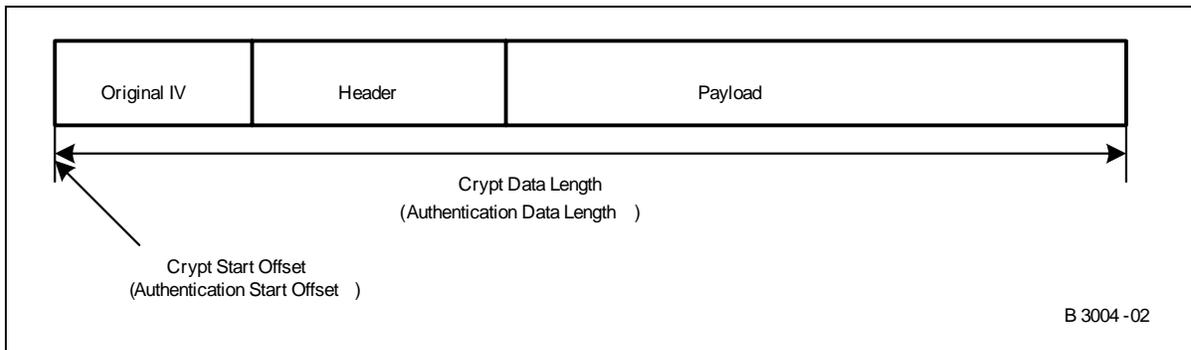
1. Register a crypto context for AES-CBC encryption (cipher context). A crypto context ID (A, in this example) is obtained in this operation. Non-in-place operation must be chosen (useDifferentSrcAndDestMbufs in IxCryptoAccCtx must set to TRUE) to avoid the original data being overwritten by encrypted data. This crypto context is used only for the purpose of authentication and generating the MIC.
2. Register another crypto context for AES-CTR encryption (cipher context). A crypto context ID (B) will also be obtained in this operation. This crypto context is used for payload and MIC encryption only.
3. After both crypto context registration for both contexts is complete, call the crypto perform API using context ID A. The IV for this packet is inserted as first block of



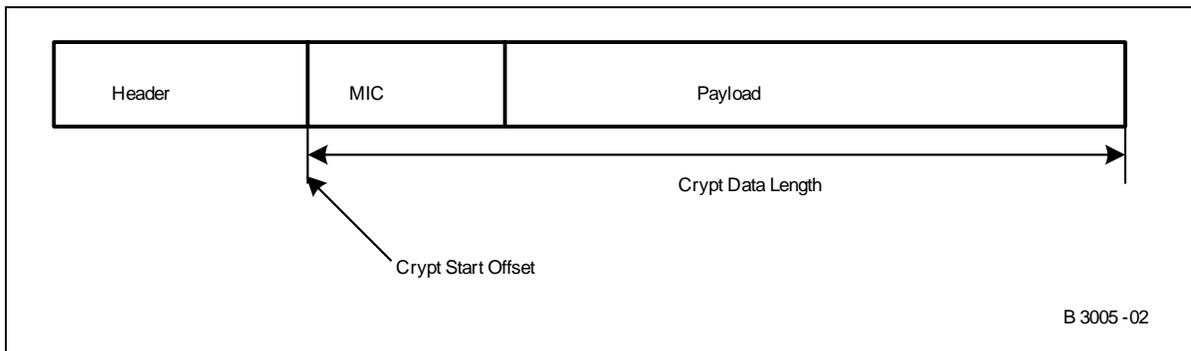
message in the packet. The input IV to the crypto perform function is set to zeroes. Crypt start offset and crypt data length parameters are set to the same values as authentication start offset and authentication data length, as shown in [Figure 39 on page 117](#). Authentication start offset and authentication data length can be ignored in the API for this operation, as this is an encryption operation only. The client should handle all the above-mentioned steps before calling the crypto perform function.

4. Wait for the operation in step 3 to complete and extract the MIC from the destination IX\_OSAL\_MBUF using the callback function.
5. Append the MIC from step 4 into the IX\_OSAL\_MBUF before the payload data.
6. Call the crypto perform function with crypto context ID B. Change the crypt start offset to point to the start offset of the MIC and change the crypt data length to include the length of MIC, as shown in [Figure 40 on page 117](#).
7. Wait for operation in step 6 to complete and move the MIC back to its original location in IX\_OSAL\_MBUF. The MIC is the final authentication data.

**Figure 39. AES CBC Encryption For MIC**



**Figure 40. AES CTR Encryption For Payload and MIC**



Since the data has to be read twice by the NPE, this two-pass mechanism will have slower throughput rate compared to the other crypto perform operations that combine encryption and authentication.

Note that memory copying is needed when performing the CCM request on a packet as mentioned above. Chained IX\_OSAL\_MBUFs could be used to avoid excessive memory copying in order to get better performance. If a single IX\_OSAL\_MBUF is used, memory copying is needed to insert MIC from AES-CBC operation into the packet, between header and payload. The payload must be moved in order to hold MIC in the packet. An



efficient method of doing this could be to split the header and payload into two different IX\_OSAL\_MBUFs. Then the MIC can be inserted after the header into the header IX\_OSAL\_MBUF for the AES CTR encryption operation.

## 7.5.6 IPsec Assumptions, Dependencies, and Limitations

- Mutable fields in IP headers should be set to a value of 0 by the client.
- The client must pad the IP datagram to be a multiple of the cipher block size, using ESP trailer for encryption (RFC 2406, explicit padding).
- The IxCryptoAcc component handles any necessary padding required during authentication operations, where the IP datagram is not a multiple of the authentication algorithm block size. The NPE pads the IP datagram to be a multiple of the block size, specified by the authentication algorithm (RFC 2402, implicit padding).
- The client must provide an initialization vector to the access component for the DES or AES algorithm, in CBC mode and CTR mode.
- IxCryptoAcc generates the primary and secondary chaining variables which are used in authentication algorithms.
- IxCryptoAcc generates the reverse keys from the keys provided for AES algorithm.

## 7.6 WEP Services

### 7.6.1 WEP Background and Implementation

The Wired Equivalent Privacy (WEP) specification is designed to provide a certain level of security to wireless 802.11 connections at the data-link level. The specification dictates the use of the ARC4 stream cipher algorithm and the use of a CRC-32 Integrity Check Value (ICV) calculation on the payload and data header.

The IxCryptoAcc API provides both the encryption/decryption and ICV calculation or verification in a single-pass implementation. The API uses two functions for performing WEP service operations, depending on the crypto component being utilized. The stream cipher features that support a WEP usage model can also be used by client applications to offload other cryptography protocols, such as SSL. Refer to [“ARC4” on page 123](#).

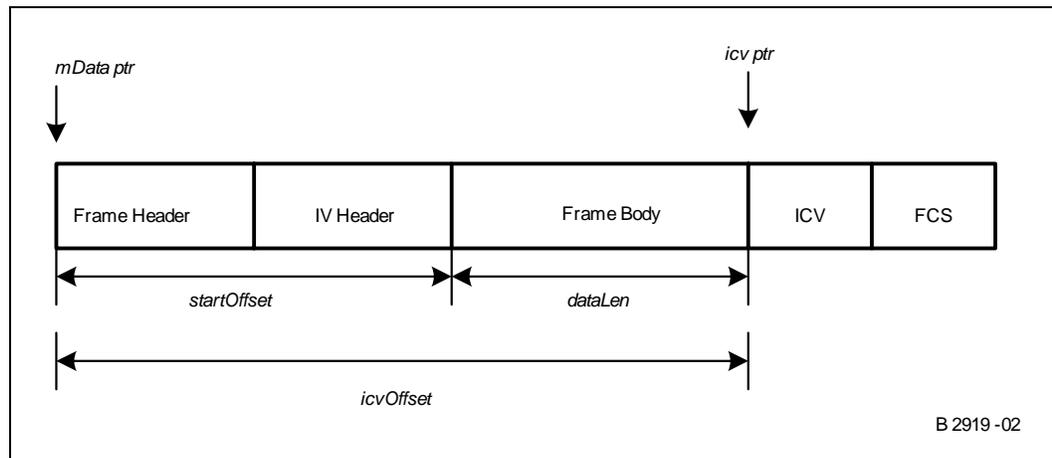
`ixCryptoAccXScaleWepPerform()` is used to submit data for WEP services using the Intel XScale® Processor-based WEP engine.

`ixCryptoAccNpeWepPerform()` is used to submit data for WEP services using NPE A processor cycles.

Both functions operate in a substantially similar manner, taking in the parameters discussed below and shown in [Figure 41](#).



Figure 41. WEP Frame with Request Parameters



- \*pSrcMbuf — a pointer to IX\_OSAL\_MBUF, which contains data to be processed. This IX\_OSAL\_MBUF structure is allocated by client. Result of this request is stored in the same IX\_OSAL\_MBUF and overwritten the original data if UseDifferentSrcAndDestMbufs flag in IxCryptoAccCtx is set to FALSE (in-place operation). Otherwise, if UseDifferentSrcAndDestMbufs flag is set to TRUE, the result is written into destination IX\_OSAL\_MBUF (non-in-place operation) and the original data in this IX\_OSAL\_MBUF will remain unchanged.
- \*pDestMbuf — Only used if UseDifferentSrcAndDestMbufs is TRUE. This is the buffer where the result is written to. This IX\_OSAL\_MBUF structure is allocated by client. The length of IX\_OSAL\_MBUF *must* be big enough to hold the result of operation. The result of operation *cannot* span into two or more different IX\_OSAL\_MBUFs, thus the IX\_OSAL\_MBUF supplied must be at least the length of expected result. The data is written back starting at startOffset in the pDestMbuf.
- startOffset — Supplied by the client to indicate the start of the payload to be decrypted/encrypted or authenticated.
- dataLen — Supplied by the client to indicate the length of the payload to be decrypted/encrypted in number of bytes.
- icvOffset — Supplied by the client to indicate the start of the ICV (Integrity Check Value) used for the authentication. This ICV field should not be split across multiple IX\_OSAL\_MBUFs in a chained IX\_OSAL\_MBUF.
- \*pKey — Pointer to IX\_CRYPTO\_ACC\_ARC4\_KEY\_128 bytes of per packet ARC4 keys. This pointer can be NULL if the request is WEP IV gen or verify only.

In the figure above, it is assumed for the sake of simplicity that mData is a contiguous buffer starting from byte 0 to the end of the FCS.

FCS is not computed or touched by the component.

## 7.6.2 Hardware Support for WEP Services

The WEP services provided in IxCryptoAcc depend on hardware-based resources for some of the cryptographic functions. This differs from the model of NPE-based crypto functionality typically found in the software release 2.3 in that the client software can select to use NPE-based crypto functionality or an Intel XScale® Processor-based software engine that both provide equivalent functionality.

These crypto components provide the following services to IxCryptoAcc:



- ARC4 (Alleged RC4) encryption / decryption
- WEP ICV generation and verification

The API provides two functions for performing WEP operations. `ixCryptoAccXScaleWepPerform()` is used to submit data for WEP services using the Intel XScale® Processor-based WEP engine. `ixCryptoAccNpeWepPerform()` is used to submit data for WEP services using NPE A processor cycles.

It is important to note that the perform requests are always executed entirely on the specified engine. However, a single crypto context may be submitted to either engine. There are some specific behavioral characteristics for each engine.

#### **ixCryptoAccNpeWepPerform()**

The NPE-based WEP perform function acts identically to the IPSec service perform functions in terms of callback behavior. During crypto context registration, a callback is specified to be executed upon completion of the perform operation. For `ixCryptoAccNpeWepPerform()`, this callback is executed asynchronously. When the NPE has completed the required processing, it will initiate the client callback.

#### **ixCryptoAccXscaleWepPerform()**

The WEP perform function using the Intel XScale® Processor WEP engine has two distinct differences from the NPE-based function.

First, `ixCryptoAccXscaleWepPerform()` operates synchronously. This is to say that once the perform function is submitted, the Intel XScale® Processor function retains the context until the perform operation is complete. The Intel XScale® Processor perform function will not execute the registered *performCallback* function. The client should initiate any local callback function on its own.

The second behavior difference is that the Intel XScale® Processor perform function does not support non-in-place memory operations. The function returns an error if the non-in-place operation is requested.

#### **NPE Microcode Images**

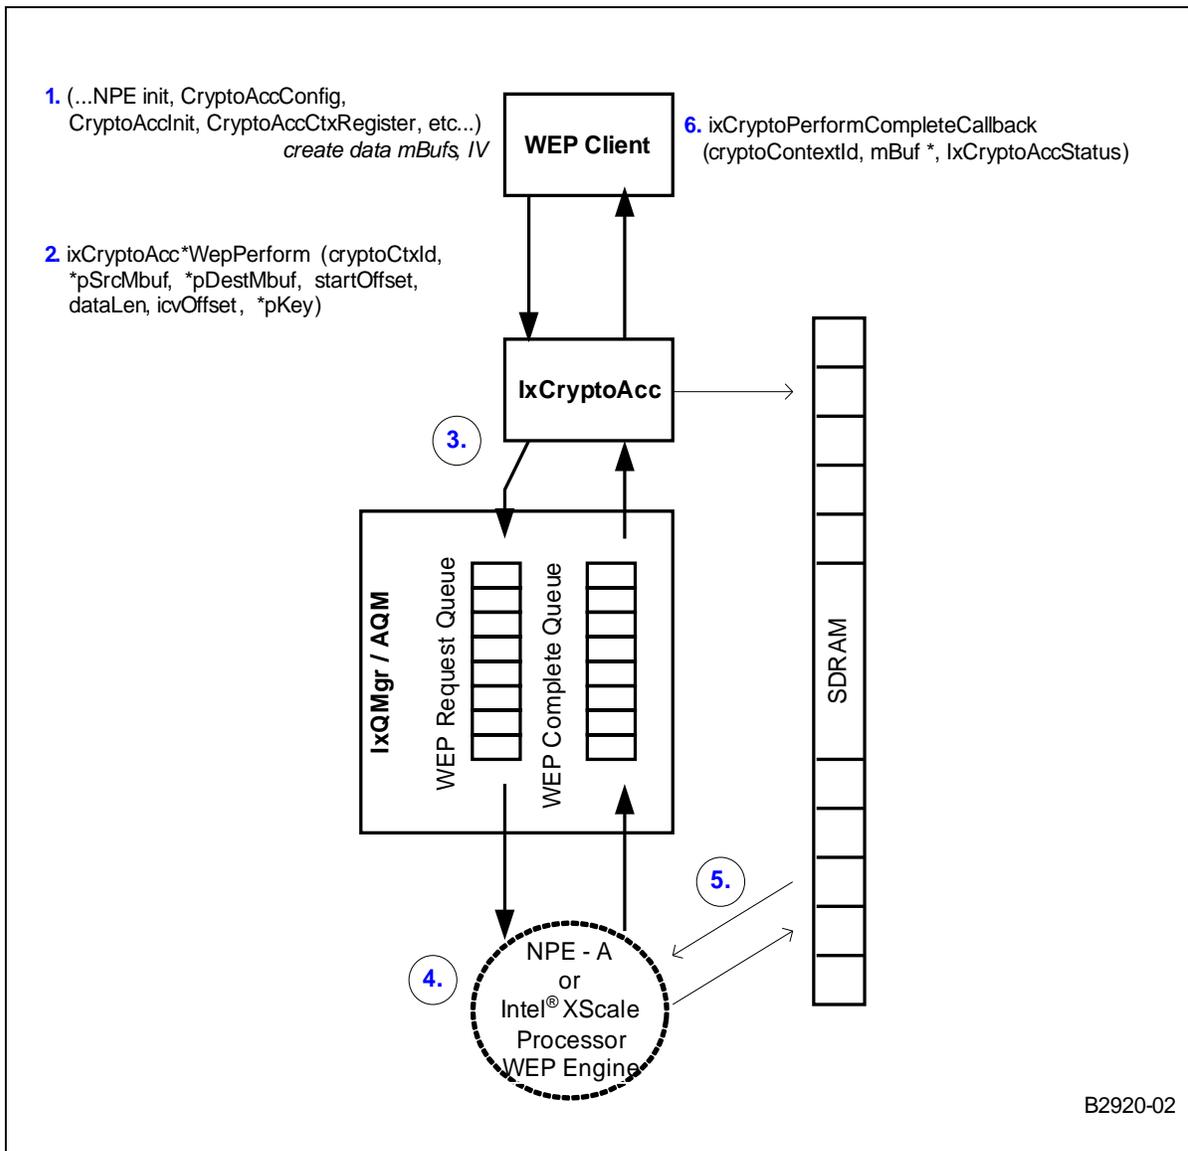
The WEP NPE image `IX_NPEDL_NPEIMAGE_NPEA_WEP` makes autonomous use of NPE A (also known as the WAN/Voice NPE) and cannot be used simultaneously with any other NPE images on NPE A. If the product design requires NPE A to be used for another purpose (DMA or ATM processing, for example), then the Intel XScale® Processor WEP engine should be used.

### **7.6.3 WEP API Call Flow**

Figure 42 details the IxCryptoAcc API call flow that occurs when submitted data for processing using WEP services. The process listed below assumes that the API has been properly configured and that a crypto context has been created and registered in the CCD, as described in [“Context Registration and the Cryptographic Context Database”](#) on page 101.



Figure 42. WEP Perform API Call Flow



B2920-02

1. The proper NPE microcode images must have been downloaded to the NPE and initialized. Additionally, the IxCryptoAcc API must be properly configured, initialized, and the crypto context registration procedure must have completed. At this point, the client must create the IX\_OSAL\_MBUFs that will hold the target data and populate the source IX\_OSAL\_MBUF with the data to be operated on. The client must supply the ARC4 key for the ARC4 algorithm.
2. The client submits the ixCryptoAccNpeWepPerform() or ixCryptoAccXscaleWepPerform() function, supplying the crypto context ID, pointers to the source and destination buffer, offset and length of the authentication and crypto data, offset to the integrity check value, and a pointer to the ARC4 key.
3. IxCryptoAcc will use IxQMgr to place a descriptor for the data into the WEP Request Queue.



4. The NPE will read the descriptor on the Crypto Request Queue and performs the encryption/decryption/authentication operations, as defined in the CCD for the submitted crypto context. The NPE will also insert or verify the WEP ICV integrity check value.
5. The NPE writes the resulting data to the destination IX\_OSAL\_MBUF in SDRAM. This may be the same IX\_OSAL\_MBUF in which the original source data was located, if the crypto context defined in-place operations. The NPE will then enqueue a descriptor onto the WEP Complete Queue to alert the IxCryptoAcc component that the perform operation is complete.
6. If the `ixCryptoAccNpeWepPerform()` function was executed in Step 2, IxCryptoAcc will call the registered Perform Complete callback function. Otherwise the client will need initiate any callback-type actions itself.

## 7.7 SSL and TLS Protocol Usage Models

SSL version 3 and TLS version 1 protocol clients can use several features provided by the IPsec and WEP services, described in earlier sections of this chapter. SSL and TLS are similar in many ways. The primary difference related to the IxCryptoAcc API is that TLS uses the HMAC (RFC 2104) hashing method for record protocol authentication. SSLv3 uses a keyed hashing mechanism for MAC generation that is similar, but not identical, to the HMAC specification.

### Authentication

SSL does not use the HMAC method of MAC generation that is provided with the IxCryptoAcc `ixCryptoAccAuthCryptPerform()` function. An SSL client can instead use `ixCryptoAccHashPerform()` for basic SHA-1 or MD-5 hashing capabilities, as part of its MAC calculation activities. Refer to “[ixCryptoAccHashKeyGenerate\(\)](#)” on page 106.

TLS clients may use the `ixCryptoAccAuthCryptPerform()` function for authentication calculation or verification crypto contexts.

### Encryption/Decryption

Both protocols can take advantage of the DES-CBC and 3DES-CBC encryption. The CipherSpec value of DES\_EDE\_CBC in the SSL and TLS protocols refers to the 3DES-CBC operation mode. Both types of clients may use the `ixCryptoAccAuthCryptPerform()` function for encrypt-only or decrypt-only contexts.

### ARC4 Steam Cipher

SSL and TLS clients may use the ARC4 cipher capabilities of the `ixCryptoAccNpeWepPerform()` and `ixCryptoAccXscaleWepPerform()` functions. Note that only 128-bit key strength is supported for contexts that do not use WEP-CRC calculation.

### Combined Mode Operations

One fundamental difference between SSL / TLS protocols and IPsec operations lies in the order of authenticate and encryption/decryption operations. SSL and TLS protocols generate the MAC prior to encryption (and verify the authentication code after decrypting the message). The IPsec ESP protocol generates its HMAC-based Integrity Check Value (ICV) on the encrypted IP packet payload (and verifies the ICV before decrypting the packet payload).



The `ixCryptoAccAuthCryptPerform()` functionality described in “IPSec Services” on page 109 offers capabilities to perform encrypt /decrypt AND authentication calculations in one submission for IPSec style clients only. This “single-pass” method does not work for SSL and TLS clients. SSL and TLS clients must register two contexts; one for encryption/decryption only and the other for authentication create / verify.

## 7.8 Supported Encryption and Authentication Algorithms

### 7.8.1 Encryption Algorithms

IxCryptoAcc supports four different ciphering algorithms

- Data Encryption Standard (DES)
- Triple DES
- Advanced Encryption Standard (AES)
- ARC4 (Alleged RC4)

Table 27 summarizes the supported cipher algorithms and the key sizes. The actual key size in DES and 3DES is less because every byte has one parity bit. The parity bit is not used in the encryption process.

**Table 27. Supported Encryption Algorithms**

Cipher Algorithm	Key Sizes (Bits)	Parity Bit (Bits)	Actual Key Size (Bits)	Plaintext / Ciphertext Block Size (Bits)
DES	64	8	56	64
3DES	192	24	168	64
AES	128 192 256	NA	128 192 256	128
ARC4	128	NA	128	8

The order expected by the NPE is in the network byte order (big endian). It is the responsibility of the client to ensure order.

#### 3DES

The order the keys are passed in should be Key 1, Key 2, and Key 3.

#### ARC4

The ARC4 algorithm can only be used in standalone mode or along with WEP-CRC algorithm. It cannot be combined with any other authentication algorithms, like HMAC-SHA-1 and HMAC-MD5. ARC4 keys used in WEP are generally 8 bytes (64-bit) or 16 bytes (128-bit). The ARC4 engine expects to be passed a key of 16 bytes in length, where it then copies the key to fill a 256-byte buffer. Therefore, if the key being used by the client is 8 bytes long, then the client should repeat it to fill the 16 bytes of key buffer.

SSL client applications can make use of the ARC4 processing features by registering an encryption-only or decryption-only crypto context and the `IxCryptoAccXScaleWepPerform()` or `IxCryptoAccNpeWepPerform()` functions. SSL clients should supply a full 128-bit key to the API.



## 7.8.2 Cipher Modes

There are four cipher modes supported by the NPE:

- Electronic code book (ECB)
- Cipher block chaining (CBC)
- Counter Mode (CTR)
- Counter-Mode / CBC-MAC Protocol (CCMP)

### 7.8.2.1 Electronic Code Book (ECB)

The ECB mode for encryption and decryption is supported for DES, Triple DES and AES. ECB is a direct application of the DES algorithm to encrypt and decrypt data.

When using the DES in ECB mode and any particular key, each input is mapped onto a unique output in encryption and this output is mapped back onto the input in decryption. The DES is an iterative, block, product-cipher system (that is, encryption algorithm). A product-cipher system mixes transposition and substitution operations in an alternating manner.

### 7.8.2.2 Cipher Block Chaining (CBC)

The CBC mode for encryption and decryption is supported for DES, Triple DES, and AES. It requires initialization vector (IV) of size 64-bit for DES and 128-bit for AES initialization vector (IV).

### 7.8.2.3 Counter Mode (CTR)

The counter mode (CTR) is only applicable for AES. The counter block consists of the SPI (the 32-bit value used to distinguish among different SAs terminating at the same destination and using the same IPSec protocol), IV, and a counter that is incremented per input block of plain text. The same AES key is used for the entire encryption process.

The counter block is always constructed by the client.

### 7.8.2.4 Counter-Mode Encryption with CBC-MAC Authentication (CCM) for CCMP in 802.11i

A protocol based on AES and Counter-Mode/CBC-MAC is being adopted for providing enhanced security in wireless LAN networks. This protocol is called Counter-Mode/CBC-MAC Protocol (CCMP). The standard defines the CCMP encapsulation/decapsulation processes, CCMP-MPDU formats, CCMP-states and CCMP-procedures. This section provides CCMP-procedure details for constructing CCM initial block (also called MIC-IV), MIC-Headers for performing CCMP MIC computation and CCM-CTR mode IV construction for performing CCM-CTR mode encryption/decryption.

The ixCryptoAcc API provides an interface for performing a single pass CCMP-MIC computation and verification with CTR mode encryption /decryption.

*Note:* The implementation of AES-CCM mode in IxCryptoAcc is designed to support 802.11i type applications specifically. As noted below, the API expects a 48-byte Initialization Vector and an 8-byte MIC value. These values correspond with an 802.11i AES-CCM implementation. IPSec implementations are expected to support 16- or 32-bit IV's and 8- or 16-bit MIC values, which are not supported by this component. Refer to [“Performing CCM \(AES CTR-Mode Encryption and AES CBC-MAC Authentication\) for IPSec” on page 115](#) for details on non-WEP AES-CCM operations.



The following should be noted regarding the support for CCMP:

- The NPE does not provide any support for:
  - constructing CCM initial block construction for MIC computation
  - constructing MIC-IV and MIC-Headers
  - constructing CTR-mode IV.
- The NPE expects that the initialization vector be 64 bytes of contiguous buffer consisting of 16 bytes of CTR-mode IV followed by 48 bytes of MIC-IV-HEADER. If the MIC-IV-HEADER constructed is less than 48 bytes, then it should be padded with zero to 48 bytes (3 AES blocks).
- Computed MIC is always 8 bytes and is not configurable to a different value.
- The coprocessor does the padding (with zeros, if required) of the data for the purposes of MIC computation. Once MIC is computed, and the data has been encrypted, the pad bytes are discarded and are not appended to the payload.
- CTR-mode IV, MIC-IV and MIC Headers are constructed by the client from RSN Header and other per-packet information.

### 7.8.3 Authentication Algorithms

Table 28 summarizes the authentication algorithms supported by IxCryptoAcc. The HMAC algorithms are offloaded to the hashing coprocessor on NPE C. The WEP-CRC algorithm may be performed using either NPE A or the Intel XScale® Processor WEP engine.

**Table 28. Supported Authentication Algorithms**

Authentication Algorithm Supported	Data Block Size (Bits)	Key Size (Bits)
HMAC-SHA-1	512	160-512
HMAC-MD5	512	128-512
WEP-CRC	8	-

## 7.9 Support for Large Number Arithmetic

ixCryptoAcc API supports large number arithmetic by offloading the computation to the Exponentiation Acceleration Unit (EAU) coprocessor in the PKE Crypto engine.

### 7.9.1 Large Number Arithmetic

The large number arithmetic supported by the EAU coprocessor in the PKE Crypto engine are:

- Modular exponential. Function is  $C = M^e \text{ mod } N$  where  $M$ ,  $e$ ,  $N$  are up to 2048 bits.
  - $M$ ,  $e$ ,  $N$  must be multiple of 32-bit.
  - $N$  must be greater than or equal to 96 bits with MSB and LSB = 1. Note: This is hardware limitation.
  - $M$  and  $e$  must be greater than or equal to 32-bit in size and less than  $N$  in value. Note that  $M$ ,  $e$ ,  $N$  can have same length.
  - $\text{sizeof}(C) \geq \text{sizeof}(N)$
- Modular reduction. Function is  $R = A \text{ mod } N$  where  $A$  is up to 4096 bits and  $N$  is up to 2048 bits.
  - $N$  must be multiple of 32-bit.



- N must be greater than or equal to 96 bits with MSB = 1.
- `sizeof(A)` must be greater than or equal to 32-bit and less than or equal to  $2 * \text{sizeof}(N)$ .
- `sizeof(R) >= sizeof(N)`
- Large number multiplication. Function is  $R=A*B$  where A & B are up to 2048 bits.
  - A, B must be multiple of 32-bit. Minimum size is 32-bit.
  - `sizeof(R) >= sizeof(A)+sizeof(B)`
- Large number addition. Function is  $R=A+B$  where A & B are up to 2048 bits.
  - A, B must be multiple of 32-bit. Minimum size is 32-bit.
  - `sizeof(R) >= max(sizeof(A), sizeof(B))`
  - Carry bit status is passed to the client in callback
- Large number subtraction. Function is  $R=A-B$  where A & B are up to 2048 bits.
  - A, B must be multiple of 32-bit. Minimum size is 32-bit.
  - `sizeof(R) >= max(sizeof(A), sizeof(B))`
  - Borrow bit status is passed to the client in callback

§ §



## 8.0 Access-Layer Components: DMA Access Driver (IxDmaAcc) API

---

This chapter describes the Intel® IXP400 Software v2.3's "DMA Access Driver" access-layer component.

### 8.1 What's New

The following new function has been added to the API:

***ixDmaAccUninit*** () This function will uninitialized the DMA Access component internals.

### 8.2 Overview

The IxDmaAcc provides DMA capability to offload large data transfers between peripherals in the Intel® IXP4XX Product Line of Network Processors memory map from the Intel XScale® Processor. The IxDmaAcc is designed to improve the Intel XScale® Processor system performance by allowing NPE to directly handle large transfers. The Direct Memory Access component (ixDmaAcc) provides the capability to do DMA transfer between peripherals that are attached to AHB buses (North AHB and South AHB buses). It also includes the APB bus, expansion bus, and PCI bus.

The ixDmaAcc component allows the client to access the NPEs' DMA services. The DMA service may run on one of the three NPEs. The appropriate NPE Microcode image with DMA services must be running on the NPE dedicated for DMA support.

**Note:** DMA transfers can be done only if the Intel XScale® Processor is operating in Big Endian mode.

### 8.3 Features

The IxDmaAcc component provides these features if the Intel XScale® Processor is operating in **Big Endian** mode:

- A DMA Access-layer API
- Clients' parameters validation
- Queues DMA requests (FIFO) to the Queue Manager

### 8.4 Assumptions

The DMA service is predicated on the following assumptions:

- IxDmaAcc has no knowledge about the Intel® IXP4XX product line processors memory map. The client must verify the validity of the source address and destination address of the DMA transfer.
- IxDmaAcc has no knowledge on the devices that involve in the DMA transfer. The client is responsible for ensuring the devices are initialized and configured correctly before request for DMA transfer.



- The Intel XScale® Processor is operating in Big Endian mode.

## 8.5 Fast Data Transfer for Little Endian Systems

The DMA Access Layer component can only operate in a Big Endian system because NPE and AHB Bus can only operate in big endian mode. Implementing Little Endian on the NPE for DMA requires implementing a very complex algorithm. Consider the source and destination address alignment, the memory alignment issue, number of bytes per transfer, and so forth. Implementing this complex algorithm on the NPE engine will trade-off the DMA performance significantly, and not deemed worth while in term of cycles consideration.

For users requiring data transfer that involve PCI and other IXP4XX product line of network processor peripherals, it is recommended that they use the hardware PCI-DMA engine (set up by using mmap reg).

For users requiring data transfers that DON'T involve PCI, use Intel XScale® Processor thread for memory copy since it is faster than NPE DMA-based DMA.

## 8.6 Builds

The current makefile does not have the DMA component specified in the LE component list, and the build would return **no such make target**.

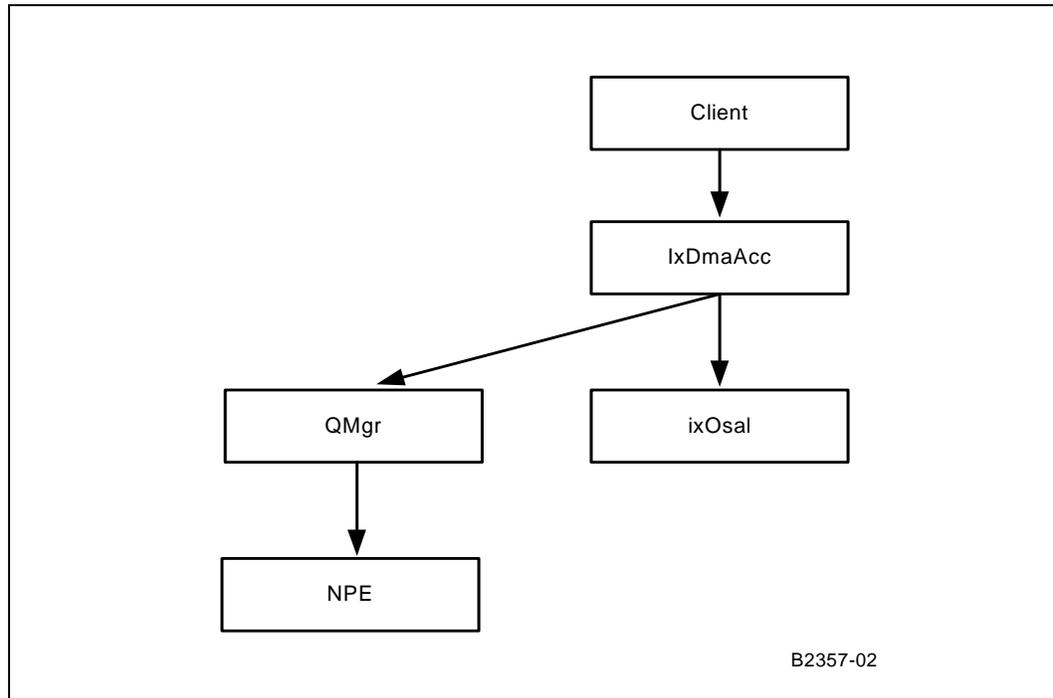
## 8.7 Dependencies

Figure 43 shows the functional dependencies of IxDmaAcc component. IxDmaAcc depends on:

- Client component using IxDmaAcc for DMA transfer access
- ixQMgr component to configure and use the Queue Manager hardware queues
- OSAL layer for error handling
- NPE to perform DMA transfer



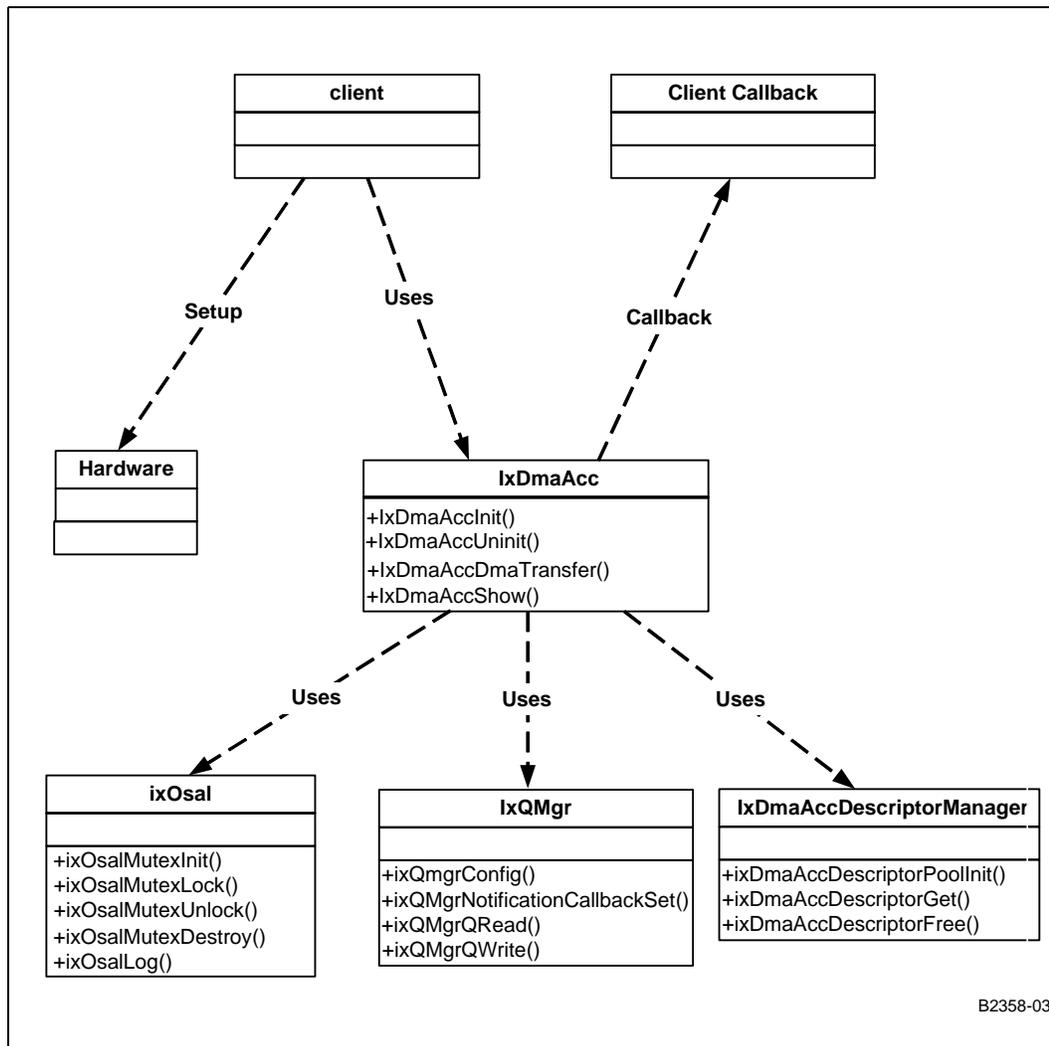
Figure 43. ixDmaAcc Dependencies



## 8.8 DMA Access-Layer API

One of the primary roles of the IxDmaAcc is to provide DMA services to different clients. These DMA services are offered through a set of functions that initialize, transfer, and display the data that needs direct memory access.

Figure 44. IxDmaAcc Component Overview



Note: IxDmaAcc components are in white.

Figure 44 shows the dependency between IxDmaAcc component and other external components (in grey). IxDmaAcc depends on:

- Client component using IxDmaAcc for DMA transfer access
- IxQMgr component for configuring and using the hardware queues to queue the DMA request and to get the 'DMA done' request status
- IxOSAL layer for mutual exclusion, error handling, and message log

The ixDmaAcc component consists of four APIs:

- **PUBLIC IX\_STATUS ixDmaAccInit (IxnpeDINpeld npeld)**  
This function initializes the DMA Access component internals.
- **PUBLIC IxDmaReturnStatus ixDmaAccDmaTransfer (IxDmaAccDmaCompleteCallback callback, UINT32 SourceAddr, UINT32 DestinationAddr, UINT16 TransferLength, IxDmaTransferMode**



### **TransferMode, IxDmaAddressingMode AddressingMode, IxDmaTransferWidth TransferWidth)**

This function performs DMA transfer between devices within the IXP4XX product line of network processor memory map.

- **PUBLIC IX\_STATUS IxDmaAccShow (void)**  
This function displays internal component information relating to the DMA service (for example, the number of the DMA requests currently pending in the queue)
- **PUBLIC IX\_STATUS IxDmaAccUninit (IxnPeDINpeId npeId)**  
This function will uninitialized the DMA Access component internals.

## **8.8.1 IxDmaAccDescriptorManager**

This component provides a private API that is used internally by the IxDmaAcc component. It provides a wrapper around the descriptor-pool-access to simplify management of the pool. This API allocates, initializes, gets, and frees the descriptor entry pool.

The descriptor memory pool is implemented using a circular buffer of descriptor data structures. These data structures hold references to the descriptor memory. The buffer is allocated during initialization. The buffer holds the maximum number of active DMA request the IxDmaAcc supports (16).

This data structure can be accessed by IxDmaAccDescriptorGet function to get an entry from the pool and IxDmaAccDescriptorFree to return the entry back to the pool.

These internal functions include:

- IxDmaAccDescriptorPoolInit(void) — Allocates and initializes the descriptor pool.
- IxDmaAccDescriptorPoolFree(void) — Frees the allocated the descriptor entry pool.
- IxDmaAccDescriptorGet(IxDmaDescriptorPoolEntry \*pDescriptor) — Returns pointer to descriptor entry.
- IxDmaAccDescriptorFree(void) — Frees the descriptor entry.

*Note:* The IxDmaAcc component addressing space for physical memory is limited to 28 bits. Therefore mBuf headers should be located in the first 256 Mbytes of physical memory.

## **8.9 Parameters Description**

The client must specify the source address, destination address, transfer mode, transfer width, addressing mode, and transfer length for each DMA transfers request. The following subsections describe the parameter details.

### **8.9.1 Source Address**

Source address is a valid Intel® IXP4XX product line processors memory map address that points to the first word of the data to be read. The client is responsible to check the validity of the source address because the access layer and NPE do not have information on the IXP4XX product line processors memory map.

### **8.9.2 Destination Address**

Destination address is a valid Intel® IXP4XX product line processors memory map address that points to the first word of the data to be written. The client is responsible to check the validity of the destination address because the access layer and NPE do not have information on the memory map.



### 8.9.3 Transfer Mode

Transfer mode describes the type of DMA transfers. There are four types of transfer modes supported:

- **Copy Only** — Moves the data from source to destination.
- **Copy and Clear Source** — Moves the data from source to destination and clears source to zero after the transfer is completed.
- **Copy and Bytes Swapping (endian)** — Moves the data from source to destination. The data written to the destination is byte swapped. The bytes are swapped within word boundary (for example, 0x 01 23 45 67 -> 0x 67 45 23 01 where the numbers indicate the source word and destination byte swapped word in the memory).
- **Copy and Bytes Reverse** — Moves the data from source to destination. The data written to the destination is byte reversed. The bytes are swapped across word boundary (for example, 0x 01 23 45 67 -> 0x 76 54 32 10 where the numbers indicate the source word and destination byte reversed word in the memory).

### 8.9.4 Transfer Width

The NPE can effect data transfer in **batches** to the AHB target, where the maximum number of bytes for each batch is 64 bytes. Only north AHB bus allows burst transfer. South AHB bus does not support burst transfer due to AHB-AHB Bridge limitation. So, transfer width indicates the **size** of the batches.

Transfer width describes how the data is transferred across the AHB buses. There are four transfer widths supported:

- **Burst** — Data may be accessed in a multiple of word per read or write transactions (normally used to access 32-bit devices).
- **8-bit** — Data must be accessed using an individual 8-bit *single* transaction (normally used to access 8-bit devices).
- **16-bit** — Data must be accessed using an individual 16-bit *single* transaction (normally used to access 16-bit devices).
- **32-bit** — Data must be accessed using an individual 32-bit *single* transaction (normally used to access 32-bit devices).

### 8.9.5 Addressing Modes

Addressing mode describes the types of source and destination addresses to be accessed. Two addressing modes are supported:

- **Incremental Address** — Address increments after each access, and is normally used to address a contiguous block of memory (for example, SDRAM).
- **Fixed Address** — Address remains the same for all access, and is normally used to operate on FIFO-like devices (for example, UART).

### 8.9.6 Transfer Length

For transfer length, it is actually the total size of data in *bytes* to be transferred from source location to destination location. However, there is some limitation for its usage. For example, a 16-bit device can only have the transfer length specified for multiple half-words. For example, a transfer length of **7** bytes on a 16-bit device would result in undetermined behavior of the NPE performing the data transfer. The values of the first 3 batches of 16 bits transferred are well-defined, but the value for the last byte is uncertain.



Transfer length restrictions are:

- Transfer length of 8-bit devices can be a multiple of byte, half-word, or word
- Transfer length of 16-bit devices can be a multiple of half-word or word
- Transfer length of 32-bit devices is a multiple of word

where a *word* is 32 bits long.

### 8.9.7 Supported Modes

This section summarizes the transfer modes supported by the IxDmaAcc. Some of the supported modes have restrictions. For details on restrictions, see “Restrictions of the DMA Transfer” on page 139.

**Table 29. DMA Modes Supported for Addressing Mode of Incremental Source Address and Incremental Destination Address**

Increment Source Address	Increment Destination Address	Transfer Mode			
		Copy Only	Copy and Clear	Copy and Bytes Swapping	Copy and Bytes Reverse
Transfer Width Source	Transfer Width Destination				
8-bit	8-bit	Supported	Supported	Supported	Supported
8-bit	16-bit	Supported	Supported	Supported	Supported
8-bit	32-bit	Supported	Supported	Supported	Supported
8-bit	Burst	Supported	Supported	Supported	Supported
16-bit	8-bit	Supported	Supported	Supported	Supported
16-bit	16-bit	Supported	Supported	Supported	Supported
16-bit	32-bit	Supported	Supported	Supported	Supported
16-bit	Burst	Supported	Supported	Supported	Supported
32-bit	8-bit	Supported	Supported	Supported	Supported
32-bit	16-bit	Supported	Supported	Supported	Supported
32-bit	32-bit	Supported	Supported	Supported	Supported
32-bit	Burst	Supported	Supported	Supported	Supported
Burst	8-bit	Supported	Supported	Supported	Supported
Burst	16-bit	Supported	Supported	Supported	Supported
Burst	32-bit	Supported	Supported	Supported	Supported
Burst	Burst	Supported	Supported	Supported	Supported



**Table 30. DMA Modes Supported for Addressing Mode of Incremental Source Address and Fixed Destination Address**

Increment Source Address	Increment Destination Address	Transfer Mode			
Transfer Width Source	Transfer Width Destination	Copy Only	Copy and Clear	Copy and Bytes Swapping	Copy and Bytes Reverse
8-bit	8-bit	Supported	Supported	Supported	Supported
8-bit	16-bit	Supported	Supported	Supported	Supported
8-bit	32-bit	Supported	Supported	Supported	Supported
8-bit	Burst	Not Supported	Not Supported	Not Supported	Not Supported
16-bit	8-bit	Supported	Supported	Supported	Supported
16-bit	16-bit	Supported	Supported	Supported	Supported
16-bit	32-bit	Supported	Supported	Supported	Supported
16-bit	Burst	Not Supported	Not Supported	Not Supported	Not Supported
32-bit	8-bit	Supported	Supported	Supported	Supported
32-bit	16-bit	Supported	Supported	Supported	Supported
32-bit	32-bit	Supported	Supported	Supported	Supported
32-bit	Burst	Not Supported	Not Supported	Not Supported	Not Supported
Burst	8-bit	Supported	Supported	Supported	Supported
Burst	16-bit	Supported	Supported	Supported	Supported
Burst	32-bit	Supported	Supported	Supported	Supported
Burst	Burst	Not Supported	Not Supported	Not Supported	Not Supported



**Table 31. DMA Modes Supported for Addressing Mode of Fixed Source Address and Incremental Destination Address**

Increment Source Address	Increment Destination Address	Transfer Mode			
		Copy Only	Copy and Clear	Copy and Bytes Swapping	Copy and Bytes Reverse
Transfer Width Source	Transfer Width Destination				
8-bit	8-bit	Supported	Supported	Supported	Supported
8-bit	16-bit	Supported	Supported	Supported	Supported
8-bit	32-bit	Supported	Supported	Supported	Supported
8-bit	Burst	Supported	Supported	Supported	Supported
16-bit	8-bit	Supported	Supported	Supported	Supported
16-bit	16-bit	Supported	Supported	Supported	Supported
16-bit	32-bit	Supported	Supported	Supported	Supported
16-bit	Burst	Supported	Supported	Supported	Supported
32-bit	8-bit	Supported	Supported	Supported	Supported
32-bit	16-bit	Supported	Supported	Supported	Supported
32-bit	32-bit	Supported	Supported	Supported	Supported
32-bit	Burst	Supported	Supported	Supported	Supported
Burst	8-bit	Not Supported	Not Supported	Not Supported	Not Supported
Burst	16-bit	Not Supported	Not Supported	Not Supported	Not Supported
Burst	32-bit	Not Supported	Not Supported	Not Supported	Not Supported
Burst	Burst	Not Supported	Not Supported	Not Supported	Not Supported

## 8.10 Data Flow

The purpose of the DMA access layer is to transfer DMA configuration information from its clients to the NPEs. It is a control component where the actual DMA data flow is transparent to the IxDmaAcc component.

## 8.11 Control Flow

For a DMA transaction to start, the client must initialize the DMA access layer, write to the queue manager, and receive a status of the transaction.

The IxDmaAcc component simultaneously supports multiple services. Consequently, a new request may be submitted before the confirmation of a previous DMA request is received from the NPE. The DMA Access layer API, however, assumes that all requests originate from the same Intel XScale® Processor task. The DMA request is queued in the AQM's request queue and waits to be serviced by the DMA NPE.

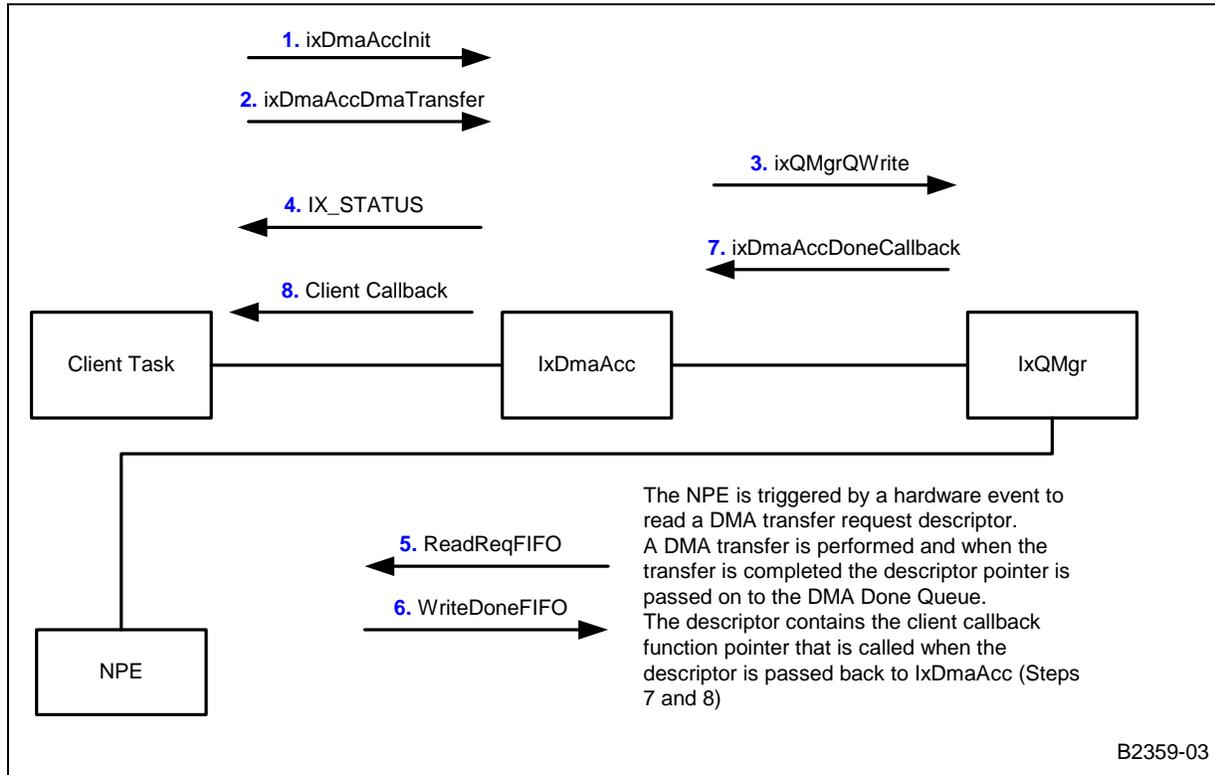
Upon completion of the DMA transfer, the NPE writes a message to the AQM-done queue. The AQM dispatcher then calls the ixDmaAcc callback and the access layer calls the client callback.

DMA require two queues from the QMgr. The queue entry size is 1-word and queue depth is 16. The depth of the queue restricts the support of multiple service requests simultaneously to 16. When the 17th entry is enqueued into the DMA Request queue, and if the queue is full, the DMA Access Layer would not enqueue this entry to the

queue. Instead, it will increment an overflow count and return DMA Request Queue FULL status to client, which has the responsibility to resend the same 17th request to the DMA Request queue later.

Figure 45 shows the overall flow of the DMA transfer operation between the client, the access layer, and the NPE.

**Figure 45. IxDmaAcc Control Flow**

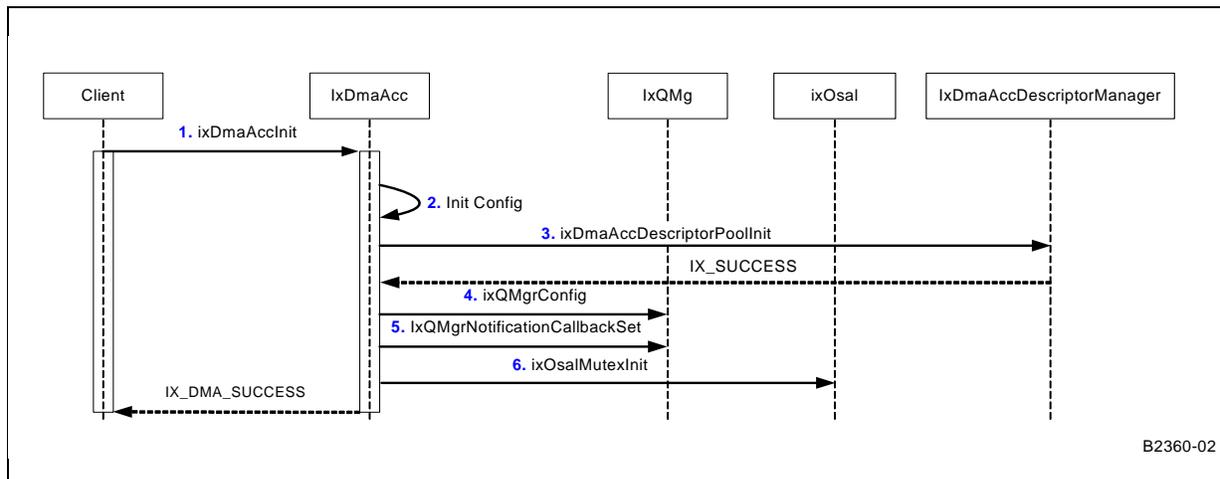


### 8.11.1 DMA Initialization

Figure 46 and the following steps describe the DMA access-layer initialization:



**Figure 46. IxDmaAcc Initialization**

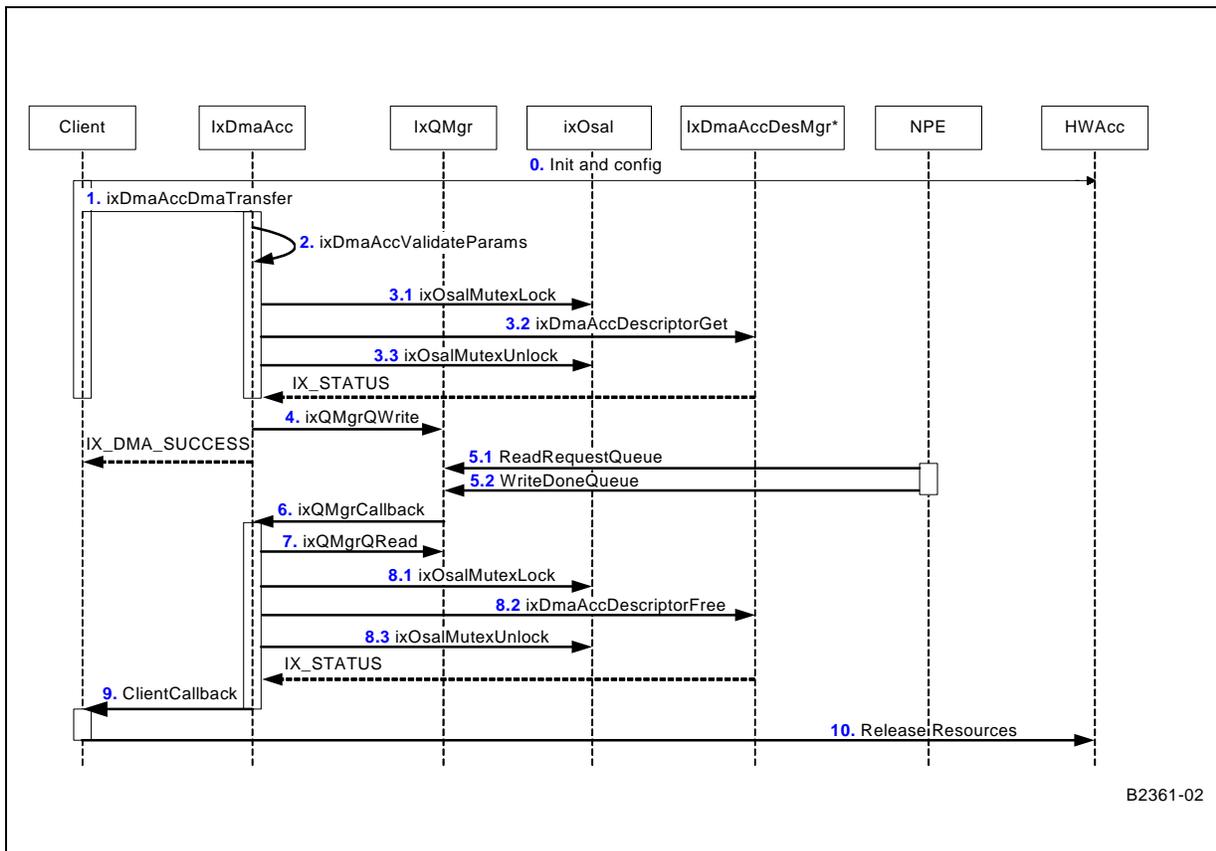


1. Client calls ixDmaAccInit to initialize the IxDmaAcc component with an NPE ID as a parameter. The NPE ID indicates which NPE is been used to provide the DMA functionality.
2. ixDmaAccInit checks if ixQmgr and the OSAL components have been initialized.
3. ixDmaAccInit calls ixDmaAccDescriptorPoolInit to allocate and initialize an array of descriptor data structures to store the DMA request and client’s callback function. (See the ixDmaAccDescriptorManager description.)
4. ixDmaAccInit calls ixQmgrConfig to configure the DMA request queue and the DMA done queue.  
The queue ID depends on which NPE the DMA component is loaded on. The selection of which NPE to run is made during run time by the client code. The client also need to initialize AQM (the Queue Manager).
5. ixDmaAccInit calls ixQMgrNotificationCallbackSet to register the callback function for the DMA-done queue.
6. ixDmaAccInit calls ixOsal to initialize mutex.  
The mutex ID is used to access queue descriptor entry pool.  
ixDmaAccInit returns IX\_DMA\_SUCCESS upon completion of the DMA initialization.

### 8.11.2 DMA Configuration and Data Transfer

Figure 47 describes the configuration and DMA data transfer between a client and an NPE.

Figure 47. DMA Transfer Operation



1. Client must initialize and configure the hardware for the DMA transfer to ensure that the devices are set up properly and ready for DMA transfer.
2. Client requests the DMA transfer by calling ixDmaAccDmaTransfer function.
3. Internally, ixDmaAccDmaTransfer function calls ixDmaAccValidateParams function to validate the client's input parameters.
4. If the client input parameters are valid, the ixDmaAccDmaTransfer function gets a descriptor entry from the descriptor manager. The descriptor pool must be guarded by mutual exclusion because there are two contexts that access the pool descriptor buffer. The ixDmaAcc component will get the pool entry and the AQM will free the entry pool (via callback).
5. The ixDmaAccDmaTransfer function composes the descriptor — based on the client's parameters — and calls ixQMgrQWrite to queue the descriptor to AQM.
6. ixDmaAccDmaTransfer returns and gets ready to process the new DMA transfer request.
7. The NPE reads the queue manager and does the DMA transfers. Upon completion of the DMA transfer, the NPE writes to AQM's done queue. The AQM dispatcher calls the IxDmaAcc's registered callback function.
8. IxDmaAccCallback calls ixQMgrQRead to read the result and that result is stored in the third descriptor. If the third word of the descriptor is zero, an AHB error is asserted by a peripheral having been accessed.



9. The descriptor pool must be guarded by mutual exclusion because there are two contexts that access the pool descriptor buffer (see Step 4.).
10. IxDmaAccCallback frees the descriptor.  
The descriptor pool must be guarded by mutual exclusion (see Step 4.).
11. IxDmaAccCallback calls client registered callback.
12. Client releases the resources allocated in Step 1.1.

## 8.12 Restrictions of the DMA Transfer

The client is responsible for ensuring that the following restrictions are followed when issuing a DMA request:

- The Intel XScale® Processor is operating in the big endian mode.
- The host devices are operating in big endian mode. This means that the valid bytes for 8-bit and 16-bit transfer width are in the most-significant bytes (MSB). For example, for the 16-bit transfer, the data is 0xAABBXXXX, where X is don't care value.
  - There is a slight difference in the access to the APB memory map region, specifically for UART accessed. A read from an APB target is a 32-bit read from a word-aligned address.
  - In the case of the UART Rx and Tx FIFOs, only the least significant byte (bits 7:0) of each word read/written contains valid data not in the MSB. Therefore, instead of using 0xC8000000 for UART1 and 0xC8001000 for UART2, any DMA request involving the UARTs must instead specify an address of 0xC8000003 for UART1 and 0xC8001003 for UART2 (in both cases the transfer width should be set to 8 bits). APB discards 1:0 bit address when decoding the AHB addresses. Therefore, valid data is read in MSB.
- Fixed address does not support burst mode. Fixed address associates with a single transaction. This means that the fixed address will either have a transfer width of 8-bit, 16-bit, or 32-bit single transaction. Fixed address (either fixed source address or fixed destination address) does not support burst transaction because burst transaction will always increment the address throughout the transaction. In addition, the AHB coprocessor does not have an instruction set to do burst transfer on fixed address mode.
- Fixed source address with copy and clear transfer mode, the source is clear only once after the transfer is completed.
- In the fixed source address mode, the client application is responsible to ensure that the data is available for transfer. For example, using FIFO with entry size 32-bit as a fixed address mode with the transfer length of 8 bytes, the client must ensure that the data is available before the DMA transfer is performed.
- Due to the asymmetric nature of the expansion bus, the incrementing source address and a "burst" transfer width will not support the "copy and clear" mode for expansion bus sources. The reason that this mode is not supported is that expansion bus targets can be read in burst mode, but they cannot be written in burst mode.
- If DMA transfer mode of "Byte-Swapped" or "Byte Reverse" is selected and if the Source DMA Addressing mode is "Incremental," the DMA Source address must be "word-aligned" and the DMA transfer length would be a multiple of words. The reason is that endianness swapping will always be done on the word boundary.
- Burst mode is not supported for DMA targets at AHB South Bus. This is due to hardware restriction. Therefore, all DMA transactions originated or designated the south AHB bus peripherals is carried out in *single* transaction mode.



- The DMA access component is fully tested on SDRAM and flash devices only. Even though the IxDmaAcc is designed to provide capability to offload large data transfers between peripherals in the IXP4XX product line processors' memory map.
- These DMA restrictions apply when a flash is a destination device:
  - Burst mode is not supported and only supports *single* mode.
  - Incremental source to fixed destination DMA addressing mode is not supported.
  - DMA transfer width for the destination must match the flash device data bus width.
  - Byte-reverse DMA mode with fixed source to incremental destination is not supported with the Flash write buffer mode.
- These DMA restrictions apply when a flash is a source device:
  - Copy and clear DMA mode is not supported
  - DMA transfer width for the source must match the Flash device data bus width.

### 8.13 Error Handling

IxDmaAcc returns an error type to the user when the client is expected to handle the error. Internal errors is reported using standard IXP4XX product line processors error-reporting techniques, such as the OSAL layer's error-reporting mechanism.

§ §



## 9.0 Access-Layer Components: Ethernet Access (IxEthAcc) API

---

This chapter describes the Intel® IXP400 Software v2.3's "Ethernet Access API" access-layer component.

### 9.1 What's New

The following changes and enhancements were made to this component in software release 2.3:

- Enhanced Ethernet functionality with Ethernet services and HSS Channelized services co-existence in NPE-A in Intel® IXP45X and Intel® IXP46X Product Line of Network Processors.
- Soft-error handling is introduced to restore/handle NPE from parity error detection. To enable this soft-error handling feature, new API function calls are introduced, refer to "New APIs".
- Changed of naming convention for existing APIs
  - **ixEthRxPriorityPoll()** changed to **ixEthAccRxPriorityPoll()**
  - **ixEthRxMultiBufferPriorityPoll()** changed to **ixEthAccRxMultiBufferPriorityPoll()**
- MAC transmit lock-up detection and recovery is introduced, additional information on MIB II extended statistics for MAC transmit lock-up detected and recovery is added.
- Changed in behavior of API **ixEthAccPortTxFrameSubmit()**.

### 9.2 New APIs

As mentioned above, the following new APIs have been added. More details regarding the input parameters, description, and return parameters can be found in the API reference document file, *APIReference.pdf*. This document is found in the doc directory of the software release.

- *PUBLIC IxEthAccStatus ixEthAccStartRequest(void)*
  - To request Eth traffic to start both Rx and Tx services.
- *PUBLIC IxEthAccStatus ixEthAccStopRequest(void)*
  - To request Eth traffic to stop both Rx and Tx services.
- *PUBLIC BOOL ixEthAccStopDoneCheck(void)*
  - To check whether Eth Traffic services are stopped.
- *PUBLIC IxEthAccStatus ixEthAccMacStateRestore(IxEthAccPortId portId)*
  - To re-update MAC register of a Ethernet ports to the state before the occurrence of soft-error.
- *PUBLIC IxEthAccStatus ixEthAccQMStatusUpdate(IxEthAccPortId portId)*



- To re-trigger the update of queue condition (interrupt and status flag) in order to restore the queue condition changes that are lost during soft-error handling. This ensures EthNPE to proceed it services accordingly.

### 9.3 IxEthAcc Overview

The IxEthAcc component (along with its related components, IxEthDB and IxEthMii) provides data plane, control plane, and management plane information for the Ethernet MAC devices residing on the Intel® IXP4XX Product Line of Network Processors. Depending on which processor variants are being used, the Intel® IXP4XX product line processors contain one, two, or three 10/100-Mbps Ethernet MAC devices.

The data path for each of these devices is accessible via dedicated NPEs. One Ethernet MAC is provided on each NPE. The NPEs are connected to the North AHB for access to the SDRAM where frames are stored. The control access to the MAC registers is via the APB Bridge, which is memory-mapped to the Intel XScale® Processor.

The IxEthAcc component is strictly limited to supporting the internal Ethernet MACs on the IXP4XX product line processors.

The services provided by the Ethernet Access component include:

- Ethernet Frame Transmission
- Ethernet Frame Reception
- Buffer management
- MAC control
- PHY control
- Statistics

PHY control is accomplished via the MII interface, which is accessible via the MAC control registers. This PHY control is not performed by the IxEthAcc component, but rather by the IxEthMii component. Although mechanisms to set the port operation state have been provided in the IxEthAcc module, true operating state-link indications should be obtained from IxEthMii.

Related components involved with EthAcc include QMgr and featureControl.

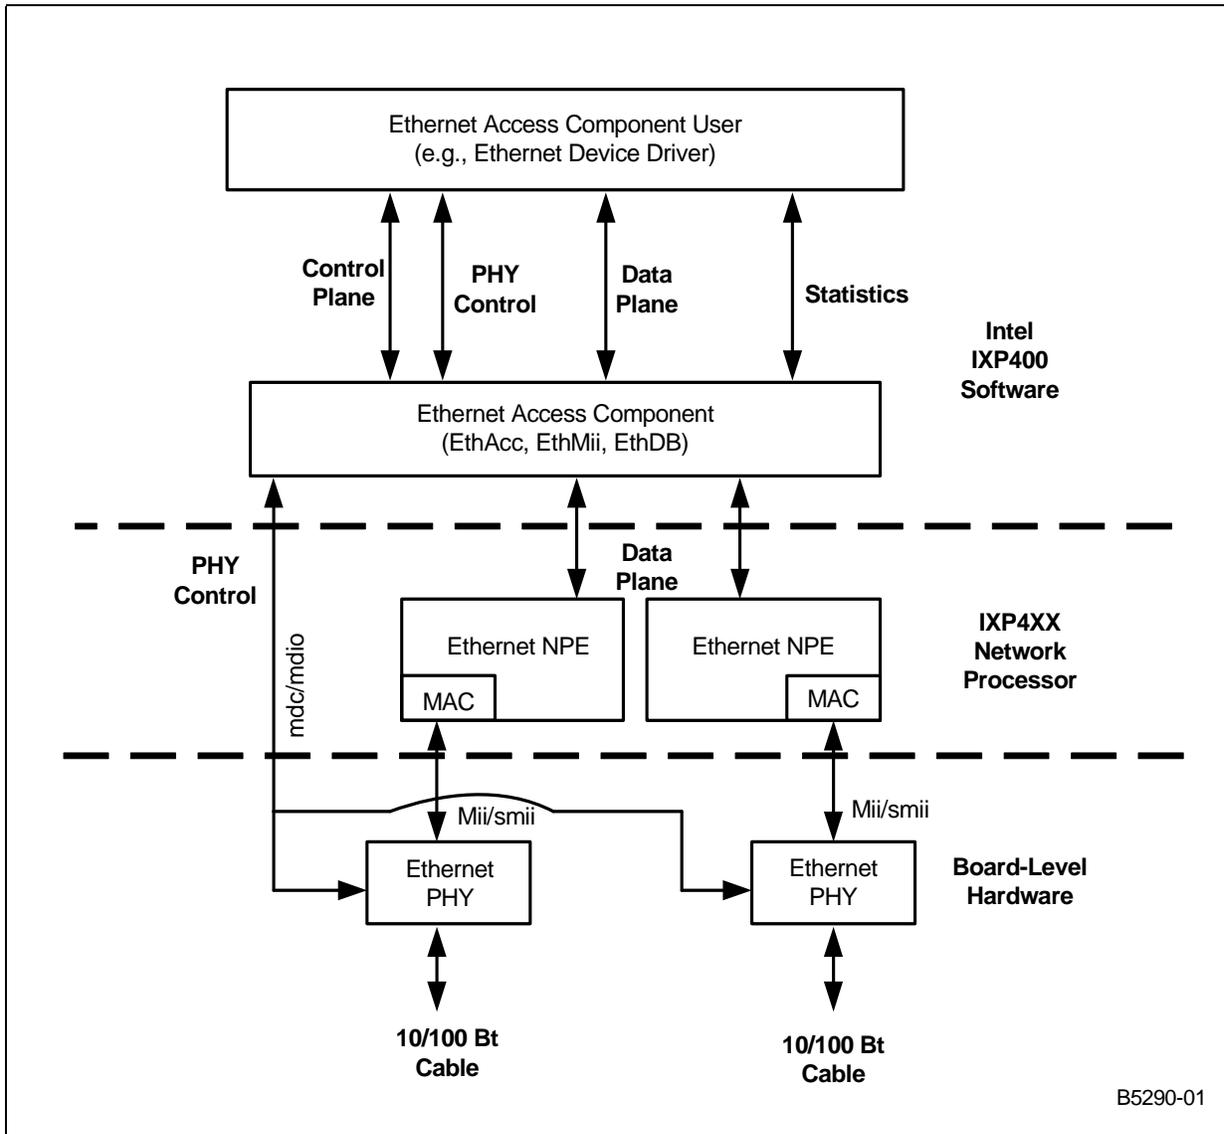
See [http://www.intel.com/intelpress/sum\\_ixp4.htm](http://www.intel.com/intelpress/sum_ixp4.htm), which describes an excellent book titled *Designing Embedded Networking Applications*. This book's "A Simple Application Using Ethernet Transmit/Receive" chapter focuses on the IXP400 software Ethernet access component, providing line-by-line code descriptions and other highly useful information.

### 9.4 Ethernet Access Layer and its Interface to Other Components: Architectural Overview

IxEthAcc is not a standalone API. It relies on services provided by number of other components such as IxEthDB, IxEthMii, NPE message handler, NPE downloader, Queue Manager, and Feature Control component. The primary role of IxEthAcc component is the transmission and reception of Ethernet traffic, managing the buffers, scheduling, and collecting statistics.



Figure 48. Ethernet Component Overview



### 9.4.1 Role of the Ethernet NPE

The Ethernet NPE microcode is responsible for moving data between an Ethernet MAC and external data memory where it can be made available to the Intel XScale® Processor. In addition, the Ethernet NPE microcode performs a number of data-processing operations.

There are many possible functions that can be performed by the NPE microcode, some examples of which are described here. On the Ethernet receive path, the Ethernet NPE microcode performs filtering, conversion of frame header to support VLAN/QoS, detects specific characteristics about a frame and notifies the client via IX\_OSAL\_MBUF header flags, and collects MAC statistics. On the Ethernet transmit path, the Ethernet NPE microcode can convert the frame header in support of VLAN/QoS, perform priority queuing of outgoing frames, and collect MAC statistics collection.



Note that filtering is not just destination-address-based; invalid source MAC address filtering is also performed. 802.3 -> 802.11 conversion on the receive side, and 802.11 -> 802.3 conversion on the transmit side are also handled by Ethernet NPE.

It is important to note that the Ethernet NPE microcode support for Ethernet data transport does not extend to support all Ethernet-related protocols and functions. For example, the NPE microcode does not automatically detect that a frame is part of an SMB protocol message and prioritize it automatically above incoming HTTP response data. However, the lack of NPE-level support for these features in no way inhibits the Intel XScale® Processor-based software from implementing them.

Communication between an Ethernet NPE and the Intel XScale® Processor is facilitated by two mechanisms. The IxQMgr component is used to handle the data path communications between the Intel XScale® Processor and NPEs. IxNpeMh is used to facilitate the communication of control messages between the Intel XScale® Processor and the NPEs.

### 9.4.2 Queue Manager

The AHB Queue Manager is a hardware block that communicates buffer pointers between the NPE cores and the Intel XScale® Processor. The IxQMgr API provides the queuing services to the access-layer and other upper level software executing on the Intel XScale® Processor. The primary use of these interfaces is to communicate the existence and location of network payload data and Ethernet service configuration information in external SDRAM.

The following is the data flow for Ethernet Ingress Frame. Ethernet frames are presented to an Ethernet-capable NPE via its Ethernet coprocessor, which serves as an interface between the Ethernet MAC and the NPE core block. Ethernet frame payloads are transferred from the Ethernet coprocessor to the host NPE in discrete blocks of data. The frames are buffered in NPE internal data memory, optionally filtered according to their destination MAC address, checked for errors, and then (assuming that no errors exist and that the frame is not filtered) transferred to external SDRAM. The Intel XScale® Processor client (via IxEthAcc) is notified of the arrival of new frames via the queue manager interface.

For more details about the data flow between the Queue Manager and Intel XScale® Processor as well as the Queue Manager and NPE, See the book referenced in [Section 9.3, "IxEthAcc Overview" on page 142](#).

### 9.4.3 Learning/Filtering

IxEthAcc relies on the IxEthDB component for the MAC learning and filtering required in a routing or bridging application.

The NPEs provide a function whereby MAC address-source learning is performed on received (ingress) Ethernet frames. Not all NPE microcode images provide the filtering capability. If source learning is enabled, the source MAC addresses are automatically populated in a learning database. For a frame to be filtered, there must be a filtering database entry whose MAC address matches the frame's destination MAC address and whose port ID matches that of the ingress MAC.

Each entry in the filtering database is composed of a MAC address and a logical port number. Whenever the bridge receives a frame, the frame is parsed to determine the destination MAC address, and the filtering database is consulted to determine the port to which the frame should be forwarded. If the destination MAC address of the frame being processed has been learned on the same interface from which it was received, it is dropped. Otherwise, the frame is forwarded from the NPE to the Intel XScale® Processor.



Learning and Filtering functionality is delivered via the EthAcc and EthDB components. For example, upon receiving NEW\_SRC\_MAC indication by NPE (a flag in IX\_BUF header), the EthAcc learns it and, with the help of the EthDB component, adds it to the database. For further details regarding EthDB Learning/Filtering, See [Section 10.4.3](#), “MAC Address Learning and Filtering” on page 174.

#### 9.4.4 MAC/PHY Configuration

IxEthMii is used primarily to manipulate a minimum number of necessary configuration registers on Ethernet PHYs supported on the Intel® IXDP425 / IXCDP1100 Development Platform, the Intel® IXDPG425 Network Gateway Development Platform, the Coyote\* Gateway Reference Design, and the Intel® IXDP465 Development Platform, without the support of a third-party operating system. Although this API is mainly intended for use with codelets and other software used for internal validation, it is provided as part of the software release 2.3 for public use.

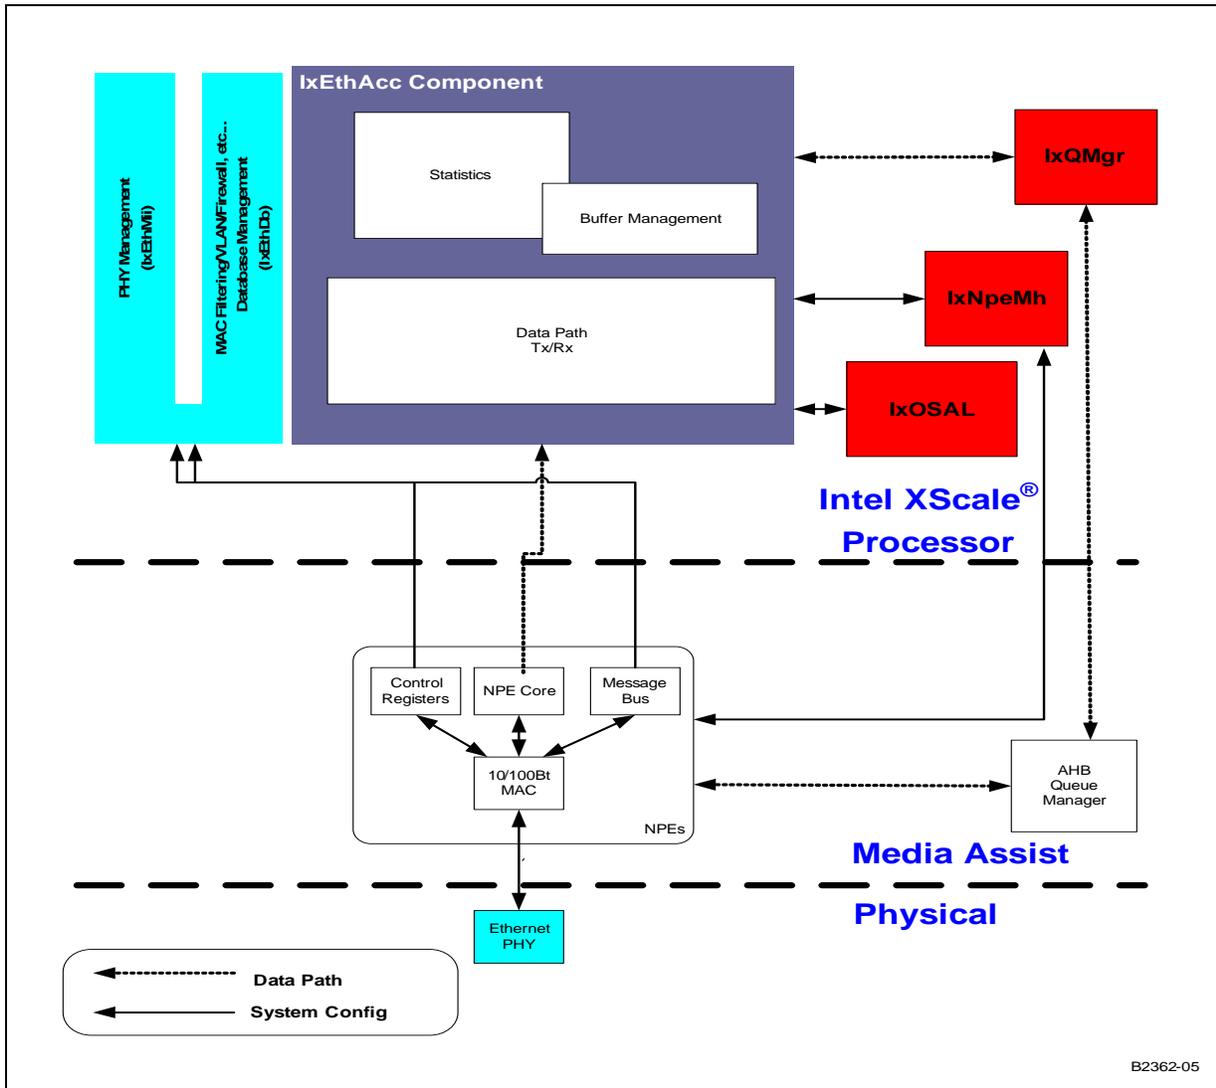
While the MAC configuration is performed within IxEthAcc, the PHY configuration requires both IxEthAcc and IxEthMii. Since the MAC also controls the MDIO interface that is used for configuring the PHY, IxEthMii must initialize the MAC in order for the PHY to be configured. IxEthAcc initializes the MAC and virtual memory mapping and executes all register reads/writes on the PHY. IxEthMii provides the register definitions for supported PHYs. Thus, IxEthMii and IxEthAcc are dependant upon each other.

### 9.5 Ethernet Access Layers: Component Features

The Ethernet access component features may be divided into three areas:

- **Data Path** — Responsible for the transmission and reception of IEEE 803.2 Ethernet frames. The Data Path is performed by IxEthAcc.
- **Control Path** — Responsible for the control of the MAC interface characteristics and some learning/filtering database functions. Control Plane functionality is included in both IxEthAcc and IxEthDB
- **Management Information** — Responsible for retrieving counter and statistical information associated with the interfaces. IxEthAcc provides this management support.

Figure 49. Ethernet Access Layers Block Diagram



## 9.6 Data Plane

The data plane is responsible for the transmission and reception of Ethernet frames. Note that the livelock prevention feature is discussed in [Section 18.10, "Livelock Prevention"](#) on page 296.

### 9.6.1 Port Initialization

Prior to any operation being performed on a port, the appropriate microcode must be downloaded to the NPE using the IxNpeDI component.

The IxEthAccPortInit() function initializes all internal data structures related to the port and checks that the port is present before initialization. The Port state remains disabled even after IxEthAccPortInit() has been called. The port is enabled using the IxEthAccPortEnable() function.

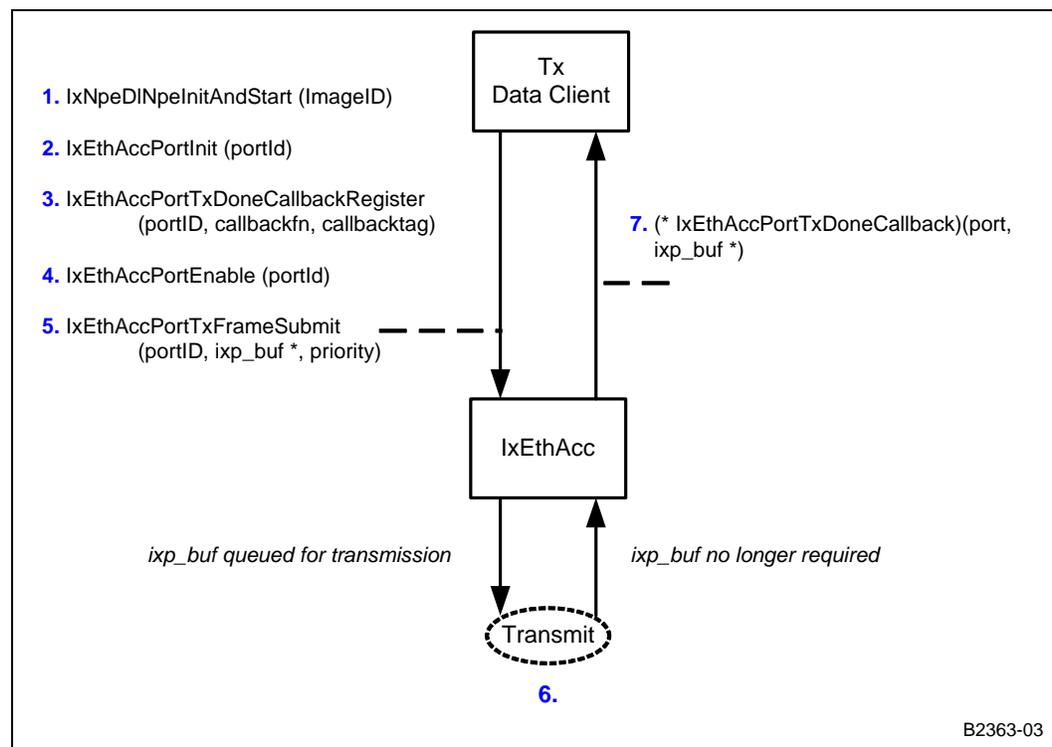


The number of Ethernet ports supported on the processor varies by processor model or variant. Note that the number of ports referred to in the EthAcc component and in EthDB component are different. Ports that are defined in IxEthAcc.h are the actual physical ports supported. See the IxEthAcc.h and IX\_ETH\_ACC\_NUMBER\_OF\_PORTS variable.

## 9.6.2 Ethernet Frame Transmission

The Ethernet access component provides a mechanism to submit frames with a relative priority to be transmitted on a specific Ethernet MAC. Once the IX\_OSAL\_MBUF is no longer required by the component, it is returned from the Ethernet access component via a free buffer callback mechanism. The flow of Ethernet frame transmission is shown in Figure 50.

Figure 50. Ethernet Transmit Frame API Overview



### 9.6.2.1 Transmission Flow

1. Proper NPE images must be downloaded to the NPEs and initialized.
2. The transmitting port must be initialized.
3. Register a callback function for the port. This function is called when the transmission buffer is placed in the TxDone queue.
4. After configuring the port, the transmitting port must be enabled in order for traffic to flow.
5. Submit the frame, setting the appropriate priority. This places the IX\_OSAL\_MBUF on the transmit queue for that port.
6. IxEthAcc interfaces to the Ethernet NPE, which uses the Ethernet Coprocessor and Ethernet MAC to transmit the frame on the wire. When transmission is complete, the IX\_OSAL\_MBUF is placed in the TxDone queue.



7. Frame transmission is complete when the TxDone callback function is invoked. The callback function is passed with a pointer to that IX\_OSAL\_MBUF.

The frame-transmission API is asynchronous in nature. Because the transmit frame request queues the frame for transmission at a later point, the call is non-blocking. There is no direct status indication as to whether the frame was successfully transmitted on the wire or not. Statistics, however, are maintained at the MAC level for failed transmit attempts.

### 9.6.2.2 Transmit Buffer Management and Priority

The overall queuing topology for the Ethernet transmission system is made up of the following queues:

- Software queues within IxEthAcc for buffering traffic when downstream queues are full, or for establishing priority queuing.
- IxQMgr queues for passing data to and from the NPEs. A maximum of 128 entries per port are supported for the TxEnet queues, and there is a single 128 entries queue for TxEnetDone.
- NPE microcode queues, used to hold IX\_OSAL\_MBUF header data for transmission. There are 64 entries in the NPE microcode queue(s).

Figure 51 provides a visual explanation of queue management for Ethernet transmission.

The IxQMgr queues are a maximum of 128 entries deep per port. The frame submit function must internally queue (in the IxEthAcc software) frames which are submitted in excess of a predefined limit. All internally queued buffers submitted for transmission but not queued to the hardware queues are stored in IxEthAcc software queues.

If priority FIFO queuing is being used, the frames is saved in individual per priority FIFOs.

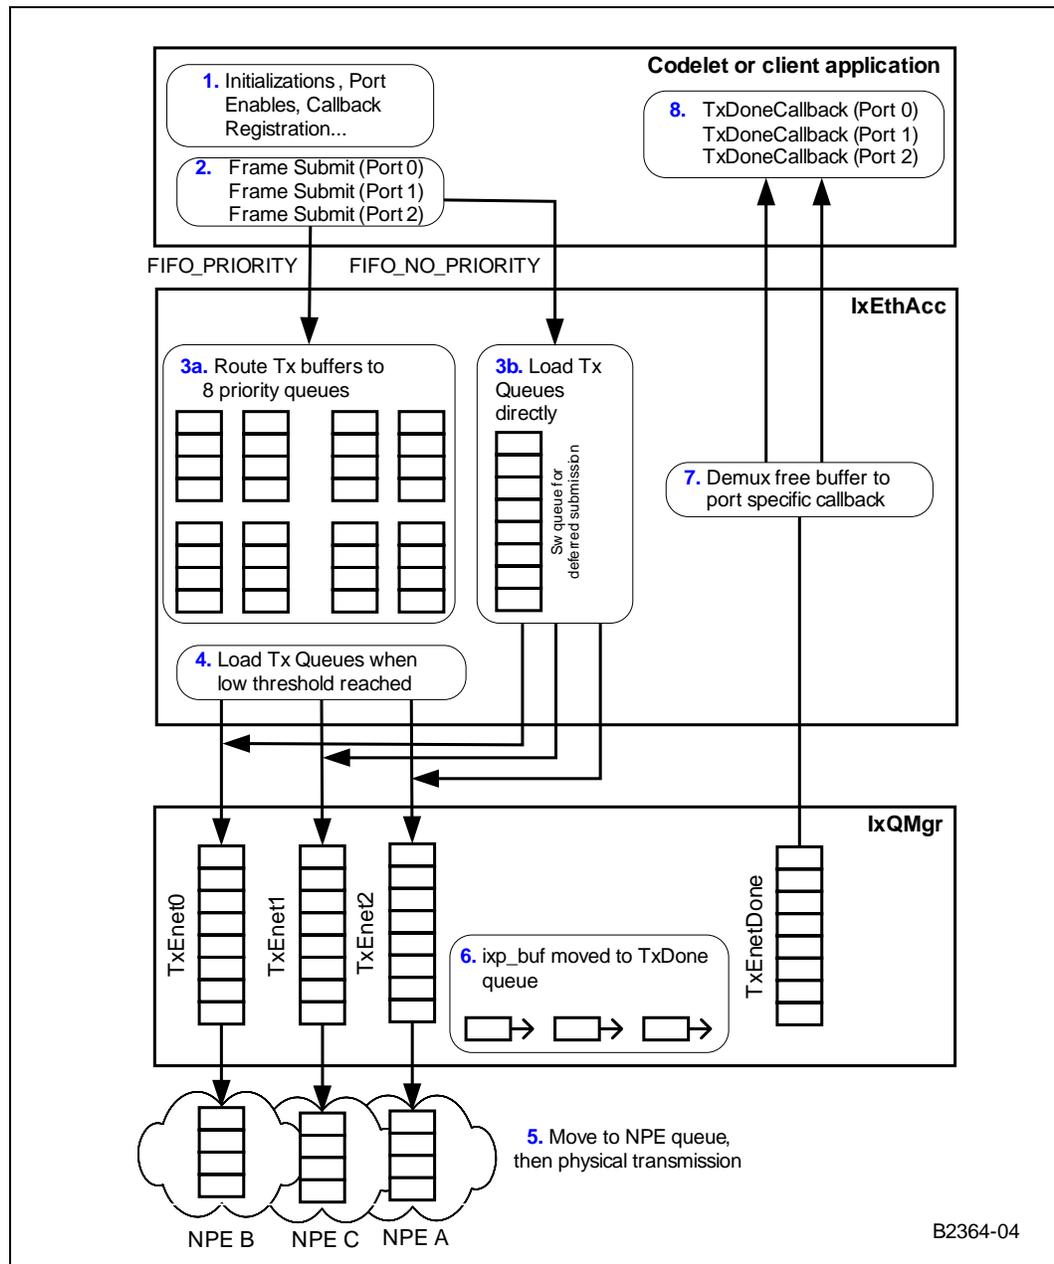
Frames is submitted to the port specific IxQMgr queue when a low/empty threshold is reached on the queue. From there, the buffer header is passed into the NPE queue that supports that respective port. If priority queueing is enabled, the NPE can re-order the frames internally to ensure that higher priority frames are transmitted before lower priority frames.

Once frame transmission has completed, the buffer is placed on the TxEnetDone IxQMgr queue. This queue contains multiplexed entries from both NPE ports. The IxEthAcc software consumes entries from this queue and returns the buffers to the client via the function previously registered by IxEthAccTxDoneCallbackRegister().

There is no specific port flush capability. To retrieve submitted buffers from the system, the port must be disabled, using the IxEthAccPortDisable() function. This has the result of returning all Tx buffers to the TxDone queue and then passed to the user via the registered TxDone callback.



Figure 51. Ethernet Transmit Frame Data Buffer Flow



There are two scheduling disciplines selectable via the `IxEthAccTxSchedulerDiscipline()`. The frame submit behavior is different for each case. Available scheduling disciplines are No Priority and Priority.

### Tx FIFO No Priority

If the selected discipline is `FIFO_NO_PRIORITY`, then all frames may be directly submitted to the `IxQMgr` queue for that port if there is room on the port. Frames that cannot be queued in the `IxQMgr` queue are stored in an `IxEthAcc` software queue for



deferred submission to the IxQMgr queue. The IxQMgr threshold in the configuration can be quite high. This allows the IxEthAcc software to burst frames into the IxQMgr queue and improve system performance due to the resultant higher cache hit rates.

### Tx FIFO Priority

If the selected discipline is FIFO\_PRIORITY, then frames are queued by IxEthAcc software in separate priority queues. The threshold in the IxQMgr must be kept quite low to improve fairness among packets submitted. Once the low threshold on the IxQMgr queue is reached, frames are selected from the priority queues in strict priority order (for example, all frames are consumed from the highest priority queue before frames are consumed from the next lowest priority).

The priority is controlled by the IxEthAccTxPriority value in the IxEthAccPortTxFrameSubmit() function as follows:

- IX\_ETH\_ACC\_TX\_PRIORITY\_0 is the lowest priority submission
- IX\_ETH\_ACC\_TX\_PRIORITY\_7 is the highest priority submission

*Note:* IxEthAccPortTxFrameSubmit() will reject frame submission if Ethernet services are stopped.

There are no fairness mechanisms applied across different priorities. Higher priority frames could starve lower-priority frames indefinitely.

### 9.6.2.3 Using Chained IX\_OSAL\_MBUFs for Transmission / Buffer Sizing

Submission of chained IX\_OSAL\_MBUF clusters for transmission is supported, but excessive chaining may have an adverse impact on performance. It is expected that chained buffers are used to add protocol headers and for large packet handling. The payload portion of large PDUs may also use chained IX\_OSAL\_MBUF clusters. The suggested minimum size for the buffers within the payload portion of a packet is 128 bytes. The “transmit done” callback function is called with the head of the cluster IX\_OSAL\_MBUF only when the entire chain has completed transmission.

The minimum size for the buffer payload is 64 bytes, including the Ethernet FCS. The ixEthAccPortTxFrameAppendPaddingEnable () function will append up to 60 bytes to an undersized frame, and will also enable FCS calculation and appending.

### 9.6.3 Ethernet Frame Reception

The Ethernet access component must be supplied with receive buffers prior to any receive activity on the Ethernet MAC. The flow of Ethernet frame reception is shown in Figure 52.

The Ethernet access component provides a mechanism to register a callback to receive Ethernet frames from a particular MAC. There are two types of user-level callbacks that can be registered:

1. A single-buffer callback is registered using ixEthAccPortRxCallbackRegister(port, callback). When the user level callback is registered using this function, it is called for each Ethernet frame received on the particular port. The prototype for the user level callback must match this format:  
void **IxEthAccPortRxCallback** (UINT32 callbackTag, IX\_OSAL\_MBUF \*buffer, UINT32 reserved)  
Note that the buffer pointer is a pointer to a single IX\_OSAL\_MBUF.
2. A multi-buffer callback is registered using ixEthAccPortMultiBufferRxCallbackRegister(port, callback). When the user-level callback is registered using this function, multiple buffers may be received in one



call as a separate array for each port. Some operating systems may perform better when the stack (Driver and OS stack) is not invoked for each frame (for example, trigger a context switch for each frame). The prototype for the user-level callback must match the following format:

**IxEthAccPortMultiBufferRxCallback** (UINT32 callbackTag, IX\_OSAL\_MBUF \*\*buffer)

Note that the buffer pointer is actually an array of IX\_OSAL\_MBUF pointers.

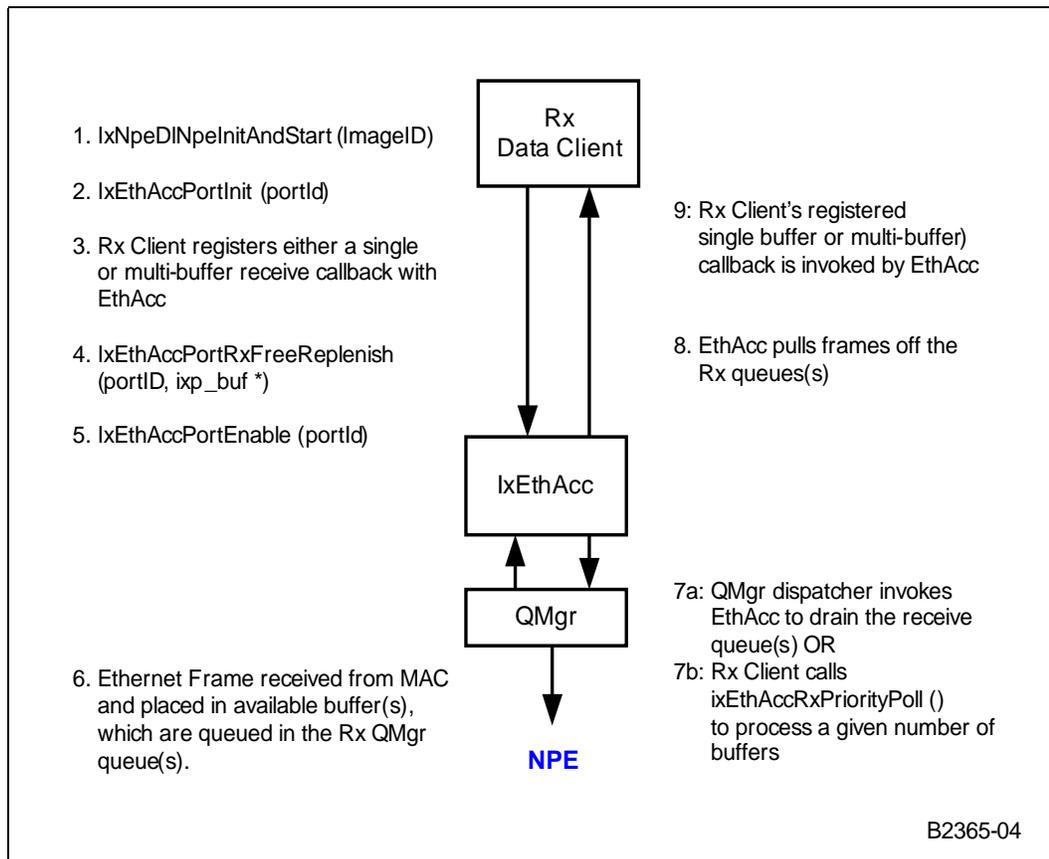
Only one type of callback can be registered at a given time so the same type of user-level callback must be registered for all ports.

As was previously mentioned, the Ethernet access component must be supplied with receive buffers via ixEthAccPortRxFreeReplenish() prior to any receive activity on the Ethernet MAC. The user must also ensure that there are sufficient buffers assigned to this component to maintain Ethernet receive performance. When a frame arrives and there is no buffer available, there is no callback indication and an rx\_buffer\_underrun counter is incremented by the NPE.

*Note:* The IX\_OSAL\_MBUF payload must not contain less than IX\_ETHACC\_RX\_MBUF\_MIN\_SIZE bytes in a single data cluster. The buffer format supported is IX\_OSAL\_MBUF provided by OSAL.

Receive frames may be pushed into a chained IX\_OSAL\_MBUF structure. Once this service calls the callback with the receive IX\_OSAL\_MBUF, 'ownership' of the buffer is transferred to the user of the access component, for example, the user has to free the buffer since the access component won't free the buffer.

Figure 52. Ethernet Receive Frame Overview



There are two methods for invoking EthAcc to process buffers on the receive queue(s):

1. The default method is that QMgr will invoke EthAcc to drain each queue individually. This process is initiated any time the QMgr dispatcher is run, typically in the client's interrupt service routine handling QMgr interrupts.
2. There is also an optional polling method in which the client can tell EthAcc directly to pull a certain number of entries off the QMgr receive queue(s) (always in priority order). The client can call either of two polling functions depending on which type of callback has been registered:
  - a. **ixEthAccRxPriorityPoll** (UINT32 reserved, UINT32 maxQEntries) should be called when a normal single buffer callback is registered by the client.
  - b. **ixEthAccRxMultiBufferPriorityPoll** (UINT32 reserved, UINT32 maxQEntries) should be called when a multi-buffer callback is registered by the client.

*Note:*

The user must ensure that the proper polling function is being called because EthAcc does not handle error checking in the datapath. Whether the user is directly calling the polling API or the normal queue dispatch interface is being used, a user callback (multi-buffer or single-buffer) must always be registered beforehand.

The default method (see 1, above) is always active thus it is suggested that any client wishing to use the polling method (see 2, above) should make sure the dispatcher will not be processing the receive queues simultaneously. This can be done by using the EthAcc Rx interrupt disable/enable APIs.



Buffers may also be returned to the user if the user flushes the receive path via the mechanism provided.

### 9.6.3.1 Receive Flow

1. Proper NPE images must be downloaded to the NPEs and initialized.
2. The receiving port must be initialized.
3. Register either a single or multi-buffer receive callback function for the port.
4. Preload free receive buffers for use by IxEthAcc.
5. After configuring the receiving port and pre-loading buffers, the receiving port is enabled, allowing traffic to be received.
6. An Ethernet frame is received on the wire and placed in the IxQMgr Rx queue.
7. (a) Queue Manager dispatcher invokes EthAcc to drain the receive queue(s)

**OR (b)** Rx client calls `ixEthAccPriorityPoll()` to process a given number of buffers.

8. EthAcc pulls frames off the Rx queue(s)
9. The Rx client's registered single buffer or multibuffer callback is invoked by EthAcc. The upper-level user or OS processes must recover the receive buffers once processing of the frame is completed, and replenish the RxFree queue using `IxEthAccPortRxFreeReplenish()` as needed.

*Note:* The process for multi-buffer receive callback is similar to single-buffer receive callback, with the exception that the multi-buffer callback should not be invoked for every frame. A polling dispatch mechanism should be used.

### 9.6.3.2 Receive Queue Interrupt Disable/Enable

EthAcc provides APIs to enable and disable queue manager interrupts for all receive queues. These functions are primarily targeted for applications making use of the receive polling interface and are highly streamlined for performance when used within the data path. The disable API clears the queue interrupt notification as well as the pending queue status for the receive queues. It is important that (when enabling and disabling interrupts using these APIs) the application confirms that it has completely drained the receive queues prior to reenabling them.

The APIs that are being referred to here are `VOID ixEthAccQMGrRxNotificationDisable()` for disabling interrupts and `VOID ixEthAccQMGrRxNotificationEnable()` for enabling interrupts.

### 9.6.3.3 Receive Buffer Management and Priority

The key interface from the NPEs to the receive data path (IxEthAcc) is a selection of queues residing in the queue manager hardware component. These queues are shown in [Figure 52](#).

#### Buffer Sizing

The receive data plane subcomponent must provide receive buffers to the NPEs. These `IX_OSAL_MBUFs` should be sized appropriately to ensure optimal performance of the Ethernet receive subsystem. The `IX_OSAL_MBUF` should contain `IX_ETHACC_RX_MBUF_MIN_SIZE` bytes in a single data cluster, though chained `IX_OSAL_MBUFs` are also supported. It is expected that chained `IX_OSAL_MBUFs` is used to handle large frames. Receive frames may be pushed into a chained `IX_OSAL_MBUF` structure, but excessive chaining will have an adverse impact upon performance. See [Figure 51](#) for a receive plane data buffer flow diagram.



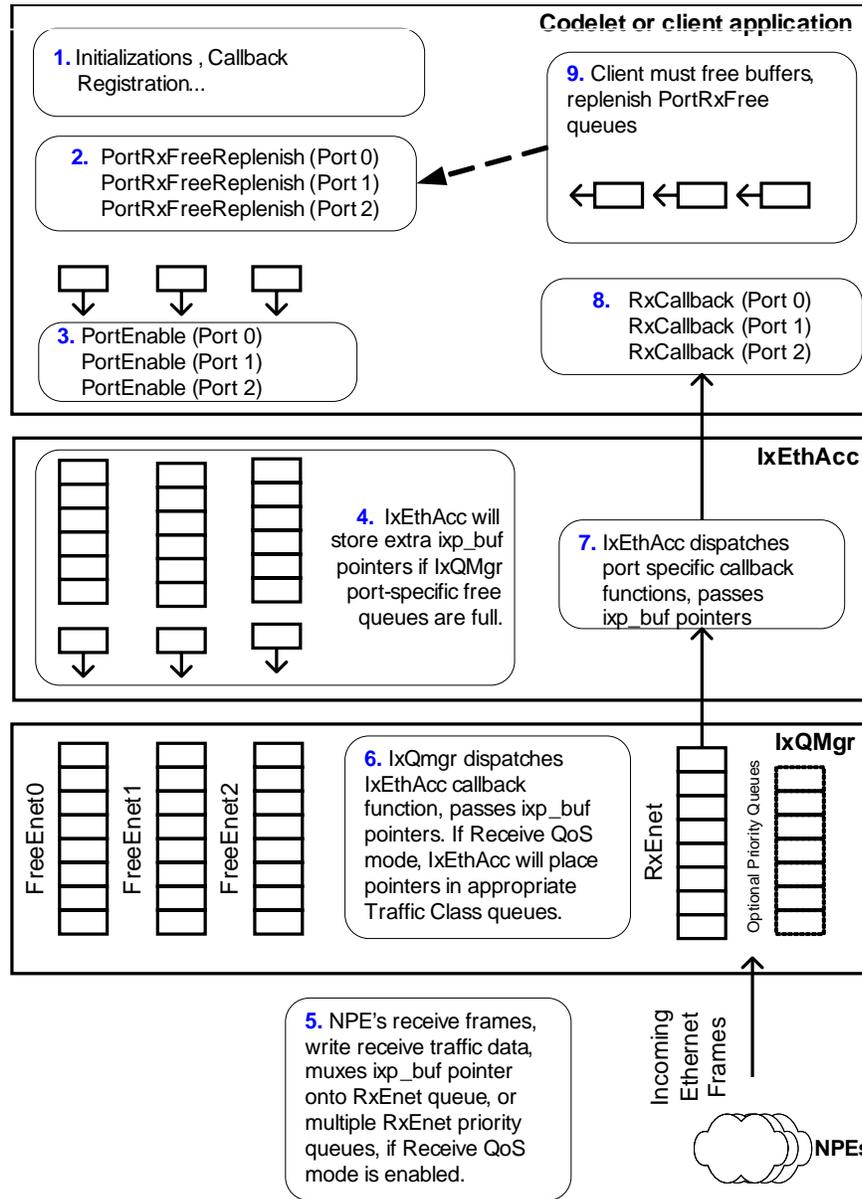
The NPEs write data in 64 byte words. For maximum performance, the IX\_OSAL\_MBUF size should be greater than the maximum frame size (Ethernet header, payload and FCS), rounded up to the next 64 byte multiple. Supplying smaller IX\_OSAL\_MBUFs to the service results in IX\_OSAL\_MBUF chaining and degraded performances.

The minimum buffer size for IEEE 802.3 Ethernet frame traffic without VLAN or IPSEC features should be 1536 bytes (1518 byte Ethernet frame + 18 bytes for 64 byte alignment). For IEEE 802.3 traffic, the recommended size is 2,048 bytes. This is adequate to support VLAN-tagged Ethernet 802.3 frames, IPsec encapsulated Ethernet frames, “baby jumbo” frames without chaining, and “jumbo” frames with chaining. The maximum 802.11 frame size is 2348 bytes. Therefore, if the 802.3 <-> 802.11 Frame Conversion feature is used, the IX\_OSAL\_MBUF should be sized at 2368 bytes (2348 + 20 for 64 byte alignment) or larger.

Buffers may not be filled up to their length. The NPE microcode will fill the IX\_OSAL\_MBUF fields up to the 64-byte boundary. The user should be aware that the length of the received IX\_OSAL\_MBUFs may be smaller than the length of the supplied IX\_OSAL\_MBUFs.



Figure 53. Ethernet Receive Plane Data Buffer Flow



B2366-04

### Supplying Buffers

There are three separate free buffer IxQMgr queues allocated to providing the NPEs with receive buffers (one per port). The buffers are supplied on a per port basis via the user level interface `ixEthAccPortRxFreeReplenish()` function. The replenish function loads the port specific free buffer IxQMgr queue with an `IX_OSAL_MBUF` pointer. The replenish function can provide checking to ensure that the `IX_OSAL_MBUF` is at least as large as `IX_ETHNPE_ACC_RXFREE_BUFFER_LENGTH_MIN`. If the port specific free buffer IxQMgr queue is full, the replenish function queues the buffer in a software queue.



Once a low threshold on the specific queue is reached the software reloads the port specific free buffer queue from its software queue if available. Frames greater in size than the size of the IX\_OSAL\_MBUF provided by the replenish function will trigger chaining.

*Note:* The `ixEthAccPortRxFreeReplenish()` function can receive chained IX\_OSAL\_MBUFs, which the NPEs are able to unchain as needed. This method may offer a performance improvement for some usage scenarios.

The user also must ensure that there are sufficient buffers assigned to this component to maintain wire-speed, Ethernet-receive performance. If the receive NPE does not have a receive buffer in advance of receiving an Ethernet frame, the frame is dropped. Should a frame arrive while there are no free buffers available, no callback indication is provided and a `rx_buffer_underrun` counter is incremented.

### Rx FIFO No Priority

Received frames from all NPEs are multiplexed onto one queue manager queue. The IxEthAcc component will de-multiplex the received frames and call the associated user level callback function registered via `IxEthAccRxCallbackRegister()`. The frames placed in the IxQMgr queue have already been validated to have a correct FCS. They are also free from all other types of MAC/PHY-related errors, including alignment errors and “frame too long” errors. Note that the receive callback is issued in frame-receive order. No receive priority mechanisms are provided. Errored frames (FCS errors, size overrun) are not passed to the user.

This is configured using the `ixEthAccRxSchedulingDisciplineSet()` function.

### Rx FIFO Priority (QoS Mode)

IxEthAcc can support the ability to prioritize frames based upon 802.1Q VLAN data on the receive path. This feature requires a compatible NPE microcode image with VLAN/QoS support. Enabling this support requires a two-part process: IxEthDB must be properly configured with support for this feature, and the Rx port in IxEthAcc must be configured using the `ixEthAccRxSchedulingDisciplineSet()` function.

In receive QoS mode, IxEthAcc will support up to four IxQMgr priority receive queues in configurations which involving only NPE-B and/or NPE-C. If NPE A is configured for Ethernet by selecting an Ethernet-enabled NPE microcode image for NPE A, then eight IxQMgr receive queues may be used. The NPE microcode will detect 802.1Q VLAN Priority data within an incoming frame or insert this data into a frame if configured to do so by IxEthDB. The NPE will then map the priority data to one of up to 8 traffic classes and places the IX\_OSAL\_MBUF header for each frame into its respective IxQMgr queue. IxEthAcc will service all frames in higher priority queues prior to servicing any entries in queues of a lower priority. Lower priority queues could be starved indefinitely.

The actual impact on system performance of the Rx FIFO priority mode is heavily influenced by the amount of traffic, priority level of the traffic, how often IxQMgr queues are serviced, and how many IxQMgr queues have entries during the time of servicing by the dispatcher loop.

If the `IxEthAccPortMultiBufferRxCallback()` function is used, it will return all currently available entries from all EthRx queues. If there are two entries in the Priority 3 EthRx queue and two entries in the Priority 1 EthRx queue, then four entries is returned with the multi-buffer callback.

Enabling the Rx QoS Mode generally involves the following process: initialize IxEthDB, enable VLAN/QoS on the desired ports, download the appropriate QoS->Traffic Class priority map (or use the default one, which is 802.1P compliant), initialize IxEthAcc and set the Rx discipline.



### Freeing Buffers

Once this service calls the callback with the receive IX\_OSAL\_MBUF, “ownership” of the buffer is transferred to the user of the access component (for example, the access component will not free the buffer). Once IxEthAcc calls the registered user-level receive callback, the receive IX\_OSAL\_MBUF “ownership” is transferred to the user of the access component. IxEthAcc will not free the buffer. Should a chain of IX\_OSAL\_MBUFs be received, the head of the buffer chain is passed to the Rx callback.

Buffers can also be freed by disabling the port, using the IxEthAccPortDisable() function. This has the result of returning all Rx buffers to the Rx registered callback, which may then de-allocate the IX\_OSAL\_MBUFs to free memory.

### Recycling Buffers

Buffers received (chained or unchained) on the Rx path can be used without modification in the Tx path. Rx and TxEnetDone buffers (chained or unchained) should have the length of each cluster reset to the cluster original size before re-using it in the ixEthAccPortRxFreeReplenish() function.

## 9.6.3.4 Additional Receive Path Information

### API function calls for Receive Polling

Intel® IXP400 Software v2.3 provides two new API function calls for receive priority polling:

1. UINT32 **ixEthAccRxPriorityPoll** (UINT32 reserved, UINT32 **maxQEntries**).
2. UINT32 **ixEthAccRxMultiBufferPriorityPoll**(UINT32 reserved, UINT32 **maxQEntries**)

There are APIs to enable and disable Queue Manager interrupts for all receive queues. These APIs are primarily targeted for applications making use of the receive polling interface.

### IPv6/IPv4 Payload Detection

Intel® IXP400 Software v2.3 supports identification of Ethernet Frames that carry IPv6 packets at NPE level.

For every received frame delivered to the Intel XScale® Processor, the NPE firmware reports whether the payload of the frame is an IPv4 or IPv6 packet by setting the **ixp\_ne\_flags.ip\_prot** flag (2-bit) in the buffer header. NPE Ethernet firmware examines the Length/Type field to determine whether the payload is IPv4 or IPv6. A value of 0x0800 indicates that the payload is IPv4 and 0x86DD indicates that the payload is IPv6.

**Table 32. IPv6/IPv4 Payload Detection**

ip_prot	IP type
00	Non IP payload
01	IPv4 payload
10	IPv6 payload
11	Reserved



The IPv4 and IPv6 payload detection service is always enabled in any NPE firmware load that supports it. However, the application software is in no way obligated to make use of the fields written by this service. The capability of this service is limited by the overall capabilities of the particular NPE firmware version. An NPE firmware version that is not VLAN-capable will not be intelligent enough to ignore the VLAN tag of a tagged frame and will always report such frames as non-IP.

#### 9.6.4 Data-Plane Endianness

All data structures provided to the IxEthAcc components, such as IX\_OSAL\_MBUF headers or statistic structures, are defined by the target system byte order. No changes to data structures are required in order to use the access component data path interfaces as IxEthAcc effects any conversion required to communicate to the NPEs. The data pointed to by the IX\_OSAL\_MBUF (the IX\_OSAL\_MBUF payload) is expected to be in network byte order (big endian). No byte swapping takes place on the data prior to transmission to the Ethernet MAC.

#### 9.6.5 Maximum Ethernet Frame Size

The maximum supported Ethernet frame size is 16,320 bytes. This value is set on a per-port basis using the IxEthDB API.

### 9.7 Control Path

The main control path functions are performed by two external components: IxEthMii and IxEthDB.

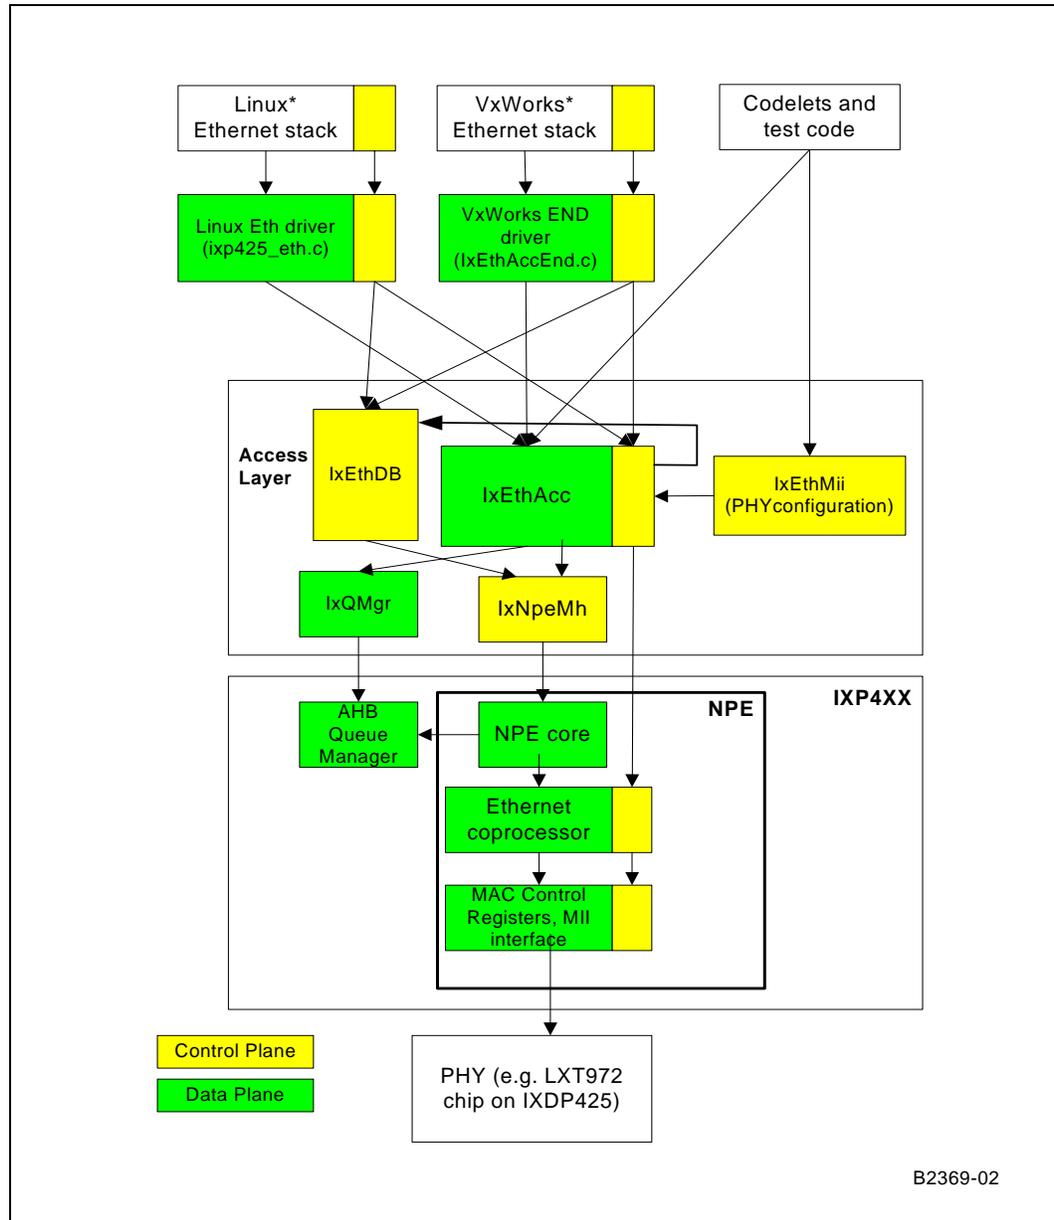
IxEthMii is used primarily to manipulate a minimum number of necessary configuration registers on Ethernet PHYs supported on the IXDP425 / IXCDP1100 platform, the IXDPG425 network gateway platform, and the IXDP465 platform without the support of a third-party operating system. IxEthMii exists as a separate function in order to make IxEthAcc independent of the specific PHY devices used in a system. However, IxEthAcc does retain control of configuring the Ethernet MAC devices on the NPEs and drives the MII and MDIO interfaces, which are used by IxEthMii to communicate physically with the PHYs.

IxEthDB is the learning and filtering database that runs within the context of the Intel XScale® Processor. The IxEthDB component handles the database structure, maintenance, searching, and aging, and has an API for the provisioning of dynamic and static addresses. This database populates filtering entries on the NPEs and also retrieves learning entries from the NPEs. An API is provided to the access layer.



The relationship between IxEthAcc, IxEthDB, and IxEthMii is shown in Figure 54.

**Figure 54. IxEthAcc and Secondary Components**



The control path component of IxEthAcc is responsible for the control of the MAC interface characteristics and some learning/filtering database functions.

### 9.7.1 Ethernet MAC Control

The role and responsibility of this module is to enable clients to configure the Ethernet coprocessor MACs for both NPES. This API permits the setting and retrieval of uni-cast and multi-cast addresses, duplex mode configuration, FCS appending, frame padding, promiscuous mode configuration, and reading or writing from the MII interface.



For more APIs related to the MAC control feature, refer to the API reference document file, APIReference.pdf, found in the doc directory of the software release.

### 9.7.1.1 MAC Duplex Settings

Functions are provided for setting the MACs at full or half duplex. This setting should match the setting of the connected PHYs.

### 9.7.1.2 MII I/O

IxEthAcc provides four functions that interact with the MII interfaces for the PHYs connected to the NPEs on the IXDP425 / IXCDP1100 platform, the IXDPG425 network gateway platform, and the IXDP465 platform. These functions do not support reading PHY registers of devices connected on the PCI interface. The MAC must be enabled with IxEthAccMacInit () first.

- IxEthAccMiiReadRtn () – Read a 16-bit value from a PHY
- IxEthAccMiiWriteRtn () – Write a 16-bit value from a PHY
- ixEthAccMiiAccessTimeoutSet() – Override the default timeout value (100 ms) and retry count when reading or writing MII registers using ixEthAccMiiWriteRtn() or ixEthAccMiiReadRtn(). This is useful for speeding up read/write operations to PHY registers.

*Note:* Note that the default values are set in accordance with the specifications of the LXT9XX, KS8995, and RTL8305 used in the IXDP425 / IXCDP1100 platform, the IXDPG425 network gateway platform, and the IXDP465 platform. The user can adjust the timeout value and number of retries according to the specific PHY being used in the application.

- IxEthAccMiiStatsShow () – Displays the values of the first eight PHY registers.

### 9.7.1.3 Frame Check Sequence

An API is provided to provision whether the MAC appends an IEEE-803.2 Frame Check Sequence (FCS) to the outgoing Ethernet frame or if the data passed to the IxEthAcc component is to be transmitted without modification.

An API is also provided to provision whether the receive buffer — sent to the Intel XScale® Processor's client — contains the frame FCS or not. The default behavior is to remove the FCS from Rx frames and to calculate and append the FCS on transmitted frames. Rx frames are still subject to FCS validity checks, and frames that fail the FCS check are dropped.

Both of these interfaces operate on a per-port basis and should be set before a port is enabled.

Special care should be taken when using the VLAN/QoS and 802.3/802.11 Frame Conversion features, as FCS behavior may be different with these features. See [Chapter 10.0](#) for clarification on these conditions.

### 9.7.1.4 Frame Padding

The IxEthAcc component by default will add up to 60-bytes to any Tx frames submitted that do not meet the Ethernet required minimum of 64-bytes. When padding is enabled, FCS appending will also be turned on.

Frame padding may not be desirable in all situations, such as when generating a "heartbeat" signal to other nodes on the network. To disable frame padding, the function IxEthAccPortTxFrameAppendPaddingDisable() is available.



This feature is available on a per-port basis and should be set before a port is enabled.

### 9.7.1.5 MAC Filtering

The MAC subcomponent within the Ethernet NPEs is capable of operation in either promiscuous or non-promiscuous mode. An API to control the operation of the MAC is provided.

**Warning:** Always use the ixEthAcc APIs to Set and Clear Promiscuous Mode. If the MAC Rx control register is modified directly, some flags in the IX\_OSAL\_MBUF header will not be populated properly.

#### Promiscuous Mode

All valid Ethernet frames are forwarded to the NPE for receive processing. NPE Learning/Filtering will not function in IxEthDB unless the MACs are configured in promiscuous mode.

#### Non-Promiscuous Mode

This allows the following frame types to be forwarded to the NPE for receive processing:

- Frame destination MAC address = Provisioned uni-cast MAC address
- Frame destination MAC address = Broadcast address
- Frame destination MAC address = Provisioned multi-cast MAC addresses. The MAC uses a mask and a multicast filter address. Packets where  $(dstMacAddress \& mask) = (mCastfilter \& mask)$  are forwarded to the NPE.

#### Address Filtering

The following functions are provided to manage the MAC address tables:

- IxEthAccPortMulticastAddressJoinAll() — all multicast frames are forwarded to the application.
- IxEthAccPortMulticastAddressLeaveAll() — Rollback the effects of IxEthAccPortMulticastAddressJoinAll().
- IxEthAccPortMulticastAddressLeave() — Unprovision a new filtering address.
- IxEthAccPortMulticastAddressJoin() — Provision a new filtering address.
- IxEthAccPortPromiscuousModeSet() — All frames are forwarded to the application regardless of the multicast address provisioned.
- IxEthAccPortPromiscuousModeClear() — Frames are forwarded to the application following the multicast address provisioned.

### 9.7.1.6 802.3x Flow Control

The Ethernet coprocessors adhere to the 802.3x flow control behavior requirements. Upon receiving a PAUSE frame, the Ethernet coprocessor will stop transmitting. PAUSE frames will not be forwarded to the NPE or Intel XScale® Processor. There is no software control for this feature.



### 9.7.1.7 NPE Loopback

Two functions are provided that enable or disable NPE-level Ethernet loopback for the NPE ports. This is useful for troubleshooting the data path.

**ixEthMiiPhyLoopbackEnable()** configures the PHY to operate in loopback mode, while **ixEthAccNpeLoopbackEnable()** can be used to test the capability of the Ethernet MAC coprocessor to loopback traffic.

### 9.7.1.8 Emergency Security Port Shutdown

Several functions are provided that may be used by an application to immediately shut down the Tx and/or Rx data path. The normal procedure is to gracefully shut down a port using the **ixEthAccPortDisable()** function, which will drain any traffic remaining in the Ethernet Tx or Rx queues prior to disabling the port. The **ixEthAccPortRxDisable()** and **ixEthAccPortTxDisable()** immediately disable the Ethernet MAC interface. These functions may be useful if a client application detects a security issue with some Ethernet traffic and must terminate any frames that may be in-process.

There are corresponding functions to re-enable the Ethernet MAC coprocessors and reset the NPE core, but recovery from an Emergency Security Port Shutdown is not guaranteed.

### 9.7.1.9 Soft-error Handling

The soft-error handling is introduced to restore/handle soft-error that is detected in Intel® IXP45X and Intel® IXP46X Product Line of Network Processors. The objective of this new handling is to allow soft-error handling module to stop or start Ethernet traffic service during error recovery, for more detail information, refer to [Chapter 17.0, "Access-Layer Components: Error Handler \(ixErrHdlAcc\) API"](#). In addition, there are functions available to error handling module to restore the MAC registers to the states in use before the occurrence of soft-error.

When soft-error being reported to Ethernet Access component, **ixEthAccStopRequest()** in the ixEthAcc component is called to sets a request to stop the Ethernet services within Ethernet access layer. API **ixEthAccStopDoneCheck()** is called to ensure that the ethernet traffic are stopped before the soft reset and recovery handling take place. After soft-error handling recovered, **ixEthAccStartRequest()** is called to resume the Ethernet services. API **EthAccMACStateRestore(portId)** is called to re-update MAC register of a Ethernet ports to the state before the occurrence of soft-error. In order to ensure Ethernet NPE to proceed it services accordingly, API **ixEthAccQMStatusUpdate(portId)** is called to re-trigger the update of queue condition (interrupt and status flag) in order to restore the queue condition changes that are lost during soft-error handling.

## 9.8 Initialization

IxEthAcc is dependent upon IxEthDB and provides for most of its initialization. The general initialization order for the Ethernet subsystem is as follows:

1. Initialize IxNpeMh, OSAL, IxQMgr.
2. Download the appropriate NPE microcode images, using IxNpeDI.
3. Configure IxEthDB.
  - a. define IxEthDBPortDefs, if necessary.
  - b. confirm capabilities and enable appropriate features using ixEthDBFeature\*() functions. It may be required to enable ports within IxEthDB using ixEthDBPortInit and ixEthDBPortEnable at this time. A specific example of this is that if the VLAN/QoS feature set is to be enabled, it must be done at this time.



4. Initialize IxEthAcc.
5. Initialize each port, and then configure port MAC addresses, PHY characteristics, and so forth, using IxEthAcc.
6. Enable traffic flow with ixEthAccPortEnable().
7. Manage Ethernet subsystem features (firewall, VLAN/QoS, Learning/Filtering, 802.11 header conversion) using IxEthDB functions.

## 9.9 Uninitialization

**ixEthAccUninit()** API is provided to uninitialized the IXP400 software Ethernet Access Service. The API **PUBLIC IxEthAccStatus ixEthAccUninit(void)** should be called once per module for uninitialization. All the resources being allocated or binded during initialization time is released by **ixEthAccUninit()**.

## 9.10 Shared Data Structures

The following section describes the data structures that are shared by the NPE Ethernet firmware and the Intel XScale® Processor client software (such as IxEthAcc, IxEthDB, and Ethernet device drivers). These data structures are used to pass information from the Intel XScale® Processor to the NPE or from the NPE to the Intel XScale® Processor. Some data structures serve to pass data in both directions.

### IX\_OSAL\_MBUFs

The buffer descriptor format supported is the IX\_OSAL\_MBUF, which is defined in [Chapter 3.0](#). The Ethernet NPE firmware expects that all such structures (for example, IX\_OSAL\_MBUF structures) are aligned to 32-byte boundaries.

The NPE is capable of handling chained IX\_OSAL\_MBUFs (for example, IX\_OSAL\_MBUFs making use of the `ixp_ne_next` field to link multiple buffers together to contain a single frame) on both the transmit and receive paths. However, for the sake of NPE performance, any use of IX\_OSAL\_MBUF chaining should be kept to a minimum. In particular, it is preferable that the IX\_OSAL\_MBUF data clusters (which are referenced by the `ixp_ne_data` structure members) to be used on the Ethernet receive path be sized so that they may contain the largest expected Ethernet frame.

It is important to note that the field definitions described within this section are valid only for the interface between the NPE Ethernet firmware and the interfacing Intel XScale® Processor client software. The Intel XScale® Processor client software is free to use these fields in any manner during the interval in which a frame is accessible only to Intel XScale® Processor software. If any IX\_OSAL\_MBUF fields are altered during Intel XScale® Processor-based processing, the Intel XScale® Processor client software must ensure that they are valid (according to the definitions in this section) before a frame is submitted to an EthTx queue.

The following tables list the specific IX\_OSAL\_MBUF fields used in the Ethernet subsystem. Note that IxEthAcc provides access to these fields via macros that are defined by the API. Those macros generally adhere to the terminology used in the following tables. Refer to the source code for specific syntax.

Many of the IX\_OSAL\_MBUF field features described are further explained in [Chapter 10.0](#).



**Table 33. IX\_OSAL\_MBUF Structure Format**

	Offset	+0	+1	+2	+3	
ixp_ne_header	0	ixp_ne_next				
	4	ixp_ne_len		ixp_ne_header		
	8	ixp_ne_data				
ixp_ne_if_eth	12		ixp_ne_scr_port	ixp_ne_flags		
	16	ixp_ne_qos_class	ixp_ne_pad_len	ixp_ne_vlan_tci		
	20	ixp_ne_dest_mac[0:5]			ixp_ne_src_mac[0:5]	
	24					
	28					

**Table 34. ixp\_ne\_flags Field Format**

Bit 15	Bit 14	Bit 13	Bit 12	Bit 11	Bit 10	Bit 9	Bit 8
new_src	vlan_en	vlan_prot		local_mac	tag_over	tag_mode	
Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
filter	st_prot	link_prot		ip_prot		multicast	broadcast

**Table 35. IX\_OSAL\_MBUF Header Definitions for the Ethernet Subsystem (Sheet 1 of 3)**

Field	Description	Queue			
		Eth Rx Free	Eth Rx	Eth Tx	Eth Tx Done
ixp_ne_next	Physical address of the next IX_OSAL_MBUF in a linked list (chain) of buffers. For the last IX_OSAL_MBUF in a chain (including the case of a single, unchained IX_OSAL_MBUF containing an entire frame), <i>ixp_ne_next</i> contains the value 0x00000000.	R	W	R	
ixp_ne_len	The interpretation of this field depends on how the IX_OSAL_MBUF is being used: <ul style="list-style-type: none"> <li>For IX_OSAL_MBUFs submitted to the EthTx or EthTxDone queues, <i>ixp_ne_len</i> represents the size (in bytes) of the valid frame data in the associated data cluster prior to any frame modifications that may occur on the NPE transmit data path. In this case, the value of <i>ixp_ne_len</i> must always be greater than 0, unless the frame length (as specified by the <i>ixp_ne_pkt_len</i> field in the first IX_OSAL_MBUF header of the current chain) is exhausted before the current IX_OSAL_MBUF is reached. In other words, it is acceptable for a number of zero-length IX_OSAL_MBUFs to be present at the end of a chain, provided that the frame ends before the first zero-length buffer is reached.</li> <li>For IX_OSAL_MBUFs submitted to the EthRx queues, <i>ixp_ne_len</i> represents the size (in bytes) of the valid frame data in the associated data cluster. In this case, the value of <i>ixp_ne_len</i> must always be greater than 0.</li> <li>For IX_OSAL_MBUFs submitted to the EthRxFree queue, <i>ixp_ne_len</i> represents the space in the associated data cluster (in bytes) available for buffering a received frame. In this case, its value must always be at least 128 bytes.</li> </ul>	R	W	R	



**Table 35. IX\_OSAL\_MBUF Header Definitions for the Ethernet Subsystem (Sheet 2 of 3)**

Field	Description	Queue			
		Eth Rx Free	Eth Rx	Eth Tx	Eth Tx Done
ixp_ne_pkt_len	<p>The value of this field depends on how the IX_OSAL_MBUF is being used:</p> <ul style="list-style-type: none"> <li>For IX_OSAL_MBUFs submitted to the EthTx, EthTxDone, and EthRx queues, <i>ixp_ne_pkt_len</i> represents the size (in bytes) of the frame contained within the IX_OSAL_MBUF. It is valid only in the first IX_OSAL_MBUF in a series of chained IX_OSAL_MBUFs. In the event that a frame is contained in a single, unchained IX_OSAL_MBUF, the value of this field is equal to the value of the <i>ixp_ne_len</i> field. For use with these queues, the value of <i>ixp_ne_pkt_len</i> must always be greater than 0. In the case of IX_OSAL_MBUFs submitted to the EthTx and EthTxDone queues, this field represents the length of the frame prior to any modifications that may occur on the NPE transmit data path.</li> <li>For IX_OSAL_MBUFs submitted to the EthRxFree queue, the value of <i>ixp_ne_pkt_len</i> must always be 0.</li> </ul>		W <sup>(5)</sup>	R	
ixp_ne_data	Physical address of the IX_OSAL_MBUF data cluster.	R		R	
ixp_ne_src_port	Either the physical MII port (see <a href="#">Table 38</a> and <a href="#">Table 36</a> ) through which an Ethernet frame was received or the port ID extracted from the VLAN TPID field of a VLAN-tagged frame (only if the port ID extraction service is enabled).		W <sup>(5)</sup>		
ixp_ne_flags.new_src	New source address flag. A value of 0 indicates that a matching entry for the frame's source MAC address exists in the filtering database; a value of 1 indicates that no matching entry could be found. For NPE Ethernet firmware versions not supporting an NPE Learning/Filtering Tree, this field is always set to 0.		W <sup>(5)</sup>		
ixp_ne_flags.filter	<p>Deferred filter flag. A value of 0 indicates a normal frame. A value of 1 indicates that the NPE would normally have dropped the frame due to a filtering operation, but that the frame was preserved and presented to the Intel XScale® Processor client because it contains a new source MAC address that must be learned. Furthermore, when this flag is set, the only IX_OSAL_MBUF fields that may be considered to be valid are <i>ixp_ne_next</i>, <i>ixp_ne_data</i>, <i>ixp_ne_dest_mac</i>, and <i>ixp_ne_src_mac</i>. For NPE firmware versions that do not support source MAC address learning, this flag is always set to 0.</p> <p><b>Note:</b> IxEthAcc will not forward these frames to the client application. After IxEthDB is notified of the new MAC address, the buffer is replenished to the EthRxFree queue.</p>		W <sup>(5)</sup>		
ixp_ne_flags.st_proto	Spanning tree protocol flag. A value of 0 indicates a normal frame; a value of 1 indicates a spanning tree protocol BPDU.		W <sup>(5)</sup>	R	
ixp_ne_flags.link_prot	Link layer protocol indicator. This field reflects the state of a frame as it exits an NPE on the receive path (and is placed into an EthRx queue) or enters an NPE on the transmit path (from the EthTx queue). It does not reflect the state of the frame when it is received or transmitted through an MII port. Its values are as listed in <a href="#">Table 38</a> .		W <sup>(5)</sup>	R	
ixp_ne_flags.vlan_prot	VLAN flag. A value of 0 indicates that the frame is not VLAN/priority-tagged when it is delivered to the host CPU; a value of 1 indicates that the frame is VLAN/priority-tagged. Note that this flag does not necessarily indicate the state of the frame when it was first received via the MII interface.		W <sup>(5)</sup>	R	
ixp_ne_flags.ip_prot	IP flag. 2 bits indicate Non IP payload, IPv4 payload, IPv6 payload as shown by <a href="#">Table 32</a> , "IPv6/IPv4 Payload Detection" on <a href="#">page 157</a> .		W <sup>(5)</sup>		
ixp_ne_flags.multicast	Multicast flag. A value of 0 indicates a non-multicast frame; a value of 1 indicates a multicast frame.		W <sup>(5)</sup>		
ixp_ne_flags.broadcast	Broadcast flag. A value of 0 indicates a non-broadcast frame; a value of 1 indicates a broadcast frame.		W <sup>(5)</sup>		



**Table 35. IX\_OSAL\_MBUF Header Definitions for the Ethernet Subsystem (Sheet 3 of 3)**

Field	Description	Queue			
		Eth Rx Free	Eth Rx	Eth Tx	Eth Tx Done
ixp_ne_flags.local_mac	A bit if set to 1 indicates that the frame is destined for the local device. For example, "Local MAC" flag and should be forwarded to Intel XScale® Processor with no conversion. If set to 0 then the frame is not be destined for Local Device.				
ixp_ne_flags.tag_over	Transmit VLAN tagging override flag. A value 0 indicates that the default tagging behavior for the port/VID should be followed; a value of 1 indicates that the default behavior should be overridden by the <i>ixp_ne_flags.tag_mode</i> flag.		W <sup>(4,5)</sup>	R	
ixp_ne_flags.tag_mode	VLAN tag behavior flag (ignored if the value of <i>ixp_ne_flags.tag_over</i> is 0). A value of 0 indicates that the output transmitted frame should be untagged; a value of 1 indicates that the output transmitted frame should be tagged.		W <sup>(4,5)</sup>	R <sup>(1)</sup>	
ixp_ne_flags.vlan_en	Transmit path VLAN functionality enable flag. A value of 0 indicates that all transmit path VLAN services, including VLAN ID-based filtering and VLAN ID-based tagging/untagging, should be disabled for the frame. A value of 1 indicates that these services should be enabled. This bit is unconditionally set by the NPE receive path firmware in VLAN-enabled builds and is unconditionally cleared by the NPE receive path firmware in non-VLAN-enabled builds.		W <sup>(4,5)</sup>	R	
ixp_ne_qos_class	The internal QoS class of the frame (set by the NPE Ethernet receive path firmware and used by the NPE transmit path firmware to queue the frame for transmission within the NPE-internal priority queue).		W <sup>(5)</sup>	(2)	
ixp_ne_pad_len	Length of the pad field for example, the number of bytes padded to 802.11 frames		W <sup>(5)</sup>	R	
ixp_ne_vlan_tci	The VLAN tag control information field of the frame (if any).		W <sup>(5)</sup>	R <sup>(3)</sup>	
ixp_ne_dest_mac	The destination MAC address of the frame.		W <sup>(5)</sup>		
ixp_ne_src_mac	The source MAC address of the frame.		W <sup>(5)</sup>		

(R) - A value of "R" in a particular column indicates that the *IX\_OSAL\_MBUF* header field is read by the Ethernet NPE firmware when it extracts the *IX\_OSAL\_MBUF* (more accurately, a pointer to the *IX\_OSAL\_MBUF*) from the AQM queue specified in the column header. The Intel XScale® Processor client software is responsible for ensuring that the field before inserting (a pointer to) the *IX\_OSAL\_MBUF* into the indicated AQM queue.

(W) - A value of "W" in a particular column indicates that the *IX\_OSAL\_MBUF* header field is written by the Ethernet NPE firmware before it inserts the *IX\_OSAL\_MBUF* (more accurately, a pointer to the *IX\_OSAL\_MBUF*) into the AQM queue specified in the column header. The Intel XScale® Processor client software may be certain that these fields are valid in *IX\_OSAL\_MBUFs* that it extracts from the indicated AQM queue.

(1) - The *ixp\_ne\_tag\_mode* field is read only if the *ixp\_ne\_flags.tag\_over* flag indicates that the behavior specified by the *VLAN Transmit Tagging Table* should be overridden.

(2) - The NPE Ethernet transmit path firmware ignores the *ixp\_ne\_qos\_class* field. Instead, it extracts the QoS class information from the *QoS* field of the *EthTx* queue entry, which must be set by the Intel XScale® Processor software before the entry is enqueued.

(3) - The *ixp\_ne\_vlan\_tci* field is read only if the output frame format is VLAN-tagged.

(4) - These fields are cleared by the NPE Ethernet receive path firmware, even though they have meaning only for the transmit path.

(5) - Although these fields may be considered to be valid only in the first *IX\_OSAL\_MBUF* in a chain of *IX\_OSAL\_MBUFs* containing a single received frame, the NPE Ethernet firmware may overwrite these fields in any and all *IX\_OSAL\_MBUFs* in the chain (regardless of their location within the chain).

**Table 36. IX\_OSAL\_MBUF "Port ID" Field Format**

7	6	5	4	3	2	1	0
NPE ID				PORT ID			



**Table 37. IX\_OSAL\_MBUF “Port ID” Field Values**

Field	Bit Position	Values
NPE ID	5.4	Ethernet-capable NPE identifier, defined as follows: <b>0x0</b> - NPE A (on IXP46X network processors only) <b>0x1</b> - NPE B <b>0x2</b> - NPE C <b>0x3</b> - Reserved
PORT ID	3.0	Sequential MII port number within the range of supported MII ports for the specified NPE. The valid ranges are as follows: Intel® IXP42X product line <ul style="list-style-type: none"> <li>• NPE A - none</li> <li>• NPE B - <b>0x0</b></li> <li>• NPE C - <b>0x0</b></li> </ul> IXP46X network processors <ul style="list-style-type: none"> <li>• NPE A - <b>0x0</b></li> <li>• NPE B - <b>0x0-0x3</b></li> <li>• NPE C - <b>0x0</b></li> </ul>

**Table 38. ixp\_ne\_flags.link\_prot Field Values**

Value	EthRx Frame Type	EthTx Frame Type
00	IEEE802.3 - 8802 (with LLC/SNAP)	IEEE802.3 - 8802 (with LLC/SNAP)
01	IEEE802.3 - Ethernet (w/o LLC/SNAP)	IEEE802.3 - Ethernet (w/o LLC/SNAP)
10	IEEE802.11 - AP -> STA	IEEE802.11 - STA -> AP
11	IEEE802.11 - AP -> AP	IEEE802.11 - AP -> AP

## 9.11 Management Information

The IxEthAcc component provides MIB II EtherObj statistics for each interface. The statistics are collected from Ethernet component counters and NPE collected statistics. Statistics are gathered for collisions, frame alignment errors, FCS errors, and so forth.

Note that each frame may be counted against a maximum of one statistic counter. In the case when more than one statistic may apply to a particular frame, it is the condition that causes the frame to be dropped at the earliest point in the data path that is recorded.

MII/RMII errors (for example, MII/RMII alignment errors, extra byte errors) take precedence over MAC errors (FCS errors, late collisions, and so forth). Next in precedence are buffer overrun errors, which take precedence over frame drops due to filtering operations. The filtering operations occur in the order of destination MAC address filtering, spanning tree, VLAN acceptable frame type filtering, VLAN ID-based filtering, firewall, and then internal queue under-run errors.

The statistics counters that are support by the Ethernet access component are shown in [Table 39](#) and [Table 40](#). For more details on these statistics objects, see RFC 2665.

These APIs are provided to retrieve these statistics:

IxEthAccMibIIStatsGet() – Returns the statistics maintained for a port

IxEthAccMibIIStatsGetClear() – Returns and clears the statistics maintained for a port

IxEthAccMibIIStatsClear() – Clears the statistics maintained for a port.



In software release 2.3, there is a handling of MAC TX lock-up, additional MIB II statistics are introduced as follows:

- Transmit
  - TxUnderrunDiscards
  - MacRecoveryTriggered

**Table 39. Managed Objects for Ethernet Receive (Sheet 1 of 2)**

Object	Increment Criteria
dot3StatsAlignmentErrors	RFC-2665 definition
dot3StatsFCSErrors	RFC-2665 definition
RxFrameTooLong	RFC-2665 definition
dot3StatsInternalMacReceiveErrors	RMII_FRM_ALN_ERROR    XTRA_BYTE    LEN_ERR    RX_LATE_COLL    (MII_FRM_ALN_ERR && !FCS_ERR)
RxOverrunDiscards	Received frames dropped because either the internal buffering capability of the NPE has been overrun (possibly because insufficient free IX_OSAL_MBUFs were available).
RxLearnedEntryDiscards	Received frame dropped due to MAC destination address filtering.
RxLargeFramesDiscards	Received frames dropped by the frame size filtering service.
RxSTPBlockedDiscards	Received frame dropped by the spanning tree port blocking service.
RxVLANTypeFilterDiscards	Received frame dropped by the VLAN ingress acceptable frame type filtering service.
RxVLANIdFilterDiscards	Received frame dropped by the VLAN ingress filtering service.
RxInvalidSourceDiscards	Received frames dropped by the invalid source MAC address filtering firewall service.
RxBLackListDiscards	Received frames dropped by the MAC address blocking firewall service.
RxWhiteListDiscards	Received frames dropped by the MAC address admission firewall service.
RxUnderflowEntryDiscards	<p>Received frame dropped due to replenishing starvation. An Underflow Discard occurs when the Ethernet Rx Free Queue becomes empty. When the NPE receives an Ethernet frame it looks to the "Rx Free" queue to find an empty buffer where it can place the incoming Ethernet packet. If no buffer is available (for example, the queue is empty) then the NPE drops the packet.</p> <p>To troubleshoot this problem, ensure the IxEthAccPortRxFreeReplenish is providing enough buffers to the Ethernet RxFree Queue. Possible root causes of replenish starvation can be that this function is either not getting the CPU time to execute with sufficient frequency, or buffers in the system are not being recycled in an efficient manner to allow the Rx Free queue replenishment to occur.</p>
RxValidFramesTotalOctets	Total bytes received from valid frames
RxUcastPkts	Number of received frames with unicast destination addressing
RxBcastPkts	Number of received frames with broadcast destination addressing
RxMcastPkts	Number of received frames with multicast destination addressing
RxPkts64Octets	Number of received frames with length <= 64 bytes
RxPkts65to127Octets	Number of received frames with 65bytes <= length <= 127 bytes
RxPkts128to255Octets	Number of received frames with 128bytes <=length<=255 bytes
RxPkts256to511Octets	Number of received frames with 256bytes<=length<=511 bytes

**Table 39. Managed Objects for Ethernet Receive (Sheet 2 of 2)**

Object	Increment Criteria
RxPkts512to1023Octets	Number of received frames with 512bytes<=length<=1023 bytes
RxPkts1024to1518Octets	Number of received frames with 1024bytes<=length<=1518 bytes
RxInternalNPEReceiveErrors	Received frame dropped due to a) replenishing starvation or b) NPE cycle starvation (for example, when crypto is enabled on the same NPE).

**Table 40. Managed Objects for Ethernet Transmit**

Object	Increment Criteria
dot3StatsSingleCollisionFrames	RFC-2665 definition
dot3StatsMultipleCollisionFrames	RFC-2665 definition
dot3StatsDeferredTransmissions	RFC-2665 definition Note that this statistic will erroneously increment when 64-byte (or smaller) frames are transmitted.
dot3StatsLateCollisions	RFC-2665 definition
dot3StatsExcessiveCollisions	RFC-2665 definition
dot3StatsInternalMacTransmitErrors	RFC-2665 definition
dot3StatsCarrierSenseErrors	RFC-2665 definition
TxLargeFrameDiscards	Transmit frames dropped by the frame size filtering service.
TxVLANIdFilterDiscards	Transmit frames dropped by the VLAN egress filtering service.
TxInternalNPETransmitErrors	Transmit frames dropped due to NPE cycle starvation.
TxUnderrunDiscards	Transmit frames dropped due to MAC Tx underrun.
MacRecoveryTriggered	Number of times MAC Transmit lock-up detected and recovered.

## 9.12 Ethernet and HSS Channelized services co-existence

NPE-A on IXP45X/IXP46X product line supports co-existence of Ethernet and HSS channelized services, the corresponding access layer components must support the co-existence capabilities. The details about the HSS access component exist in [Chapter 13.0, "Coexistence of HSS Channelized services and Ethernet services in NPE-A"](#)

**Notes:** The co-existence of these services does not support 3 ports IP routing.

In order to support this function, there are changes and consideration being take place in ixEthAcc component.

### 9.12.1 Queue Manager Queues

The following considerations need to be made in order for HSS channelized and Ethernet to co-exist

- HSS channelized and Ethernet co-existence feature should not cause queue manager queue overlaps between Ethernet and HSS services.
- Live lock prevention support to queue manager dispatcher.



- Prioritize the HSS channelized AQM queue servicing over Ethernet queues, in order to reduce the delay in voice related processing.
- Common mutex locking scheme to enable both HSS and Ethernet services. A common mutex locking scheme is introduced for message FIFO access between HSS & Ethernet access layers via NpeMh memory.

### 9.12.2 Live lock prevention scheme in HSS and Ethernet co-existence

This is a mechanism to ensure good QoS for the voice modules by reducing the AQM callback service delay on the HSS Channelized service when it coexist with other NPE-A services such as Ethernet service. This is done by making sure that the HSS Channelized operation has the priority of AQM queue dispatcher service over the Ethernet traffic. For more information about live lock mechanism, refer to [Chapter 18.0, "Livelock Prevention"](#).

§ §



## 10.0 Access-Layer Components: Ethernet Database (IxEthDB)

---

This chapter describes the Intel® IXP400 Software v2.3 “Ethernet Database API” access-layer component.

### 10.1 Overview

To minimize the unnecessary forwarding of frames, an IEEE 802.1d-compliant bridge maintains a filtering database. IxEthDB provides MAC address-learning and filtering database functionality for the Ethernet NPE interfaces. IxEthDB also provides the configuration and management of many of the Ethernet subsystem NPE-based capabilities, such as VLAN/QoS, MAC address firewall and (802.11) frame header conversion.

### 10.2 What’s New

The following change or enhancement was made to this component in software release 2.3:

- Soft-error handling is introduced to restore/handle NPE from parity error detection. To enable this soft-error handling feature, new API function calls are introduced, refer to “New APIs”

### 10.3 New APIs

As mentioned above, the following new APIs have been added. More details regarding the input parameters, description, and return parameter can be found in the API reference document file, APIReference.pdf. This document is found in the doc directory of the software release.

- *IX\_ETH\_DB\_PUBLIC IxEthDBStatus ixEthDBFeatureStatesRestore (IxEthDBPortId portId)*
  - Restores the state of EthDB based on latest settings, following the occurrence of an NPE soft-error. State is restored by re-downloading tables for the enabled features (for example Header Conversion and Firewall) or sending configuration messages to NPE.
- *IX\_ETH\_DB\_PUBLIC IxEthDBStatus ixEthDBPriorityMappingTableUpdate(IxEthDBPortId portId)*
  - Reloads the last port priority mapping table set by the user.
- *IX\_ETH\_DB\_PUBLIC IxEthDBStatus ixEthDBEventProcessorPauseModeSet (BOOL pauseEnable)*
  - Pauses/resumes Ethernet DB event processor.



## 10.4 IxEthDB Functional Behavior

There are two major elements involved in the IxEthDB subsystem: a software database that executes on the Intel XScale® Processor of the processor, and one or more Network Processing Engines (NPEs) that are capable of making decisions or performing manipulations on the Ethernet traffic that they encounter. While the capabilities of the NPEs are determined by the microcode that runs on them, the specifics related to how the NPE should drop, forward or manipulate the Ethernet traffic are managed by the IxEthDB component.

IxEthDB handles the configuration of several Ethernet subsystem features:

- MAC Address Learning and Filtering
- Frame Size Filtering
- Source MAC Address Firewall
- 802.1Q VLAN
- 802.1p QoS
- 802.3 / 802.11 frame conversion
- Spanning Tree Protocol port settings

IxEthDB also has several more generalized features that relate to the databases and the API itself:

- Database management
- Port Definitions
- Feature Control

### 10.4.1 Feature set

IxEthDB is structured in a feature set, which can be enabled, disabled and configured at run time. Since IxEthDB provides support for NPE features, the feature set presented to client code at any one time depends on the run-time configuration of the NPEs. IxEthDB can detect the capabilities of each NPE microcode image and expose only those features supported by that image.

Table 41. IxEthDB Feature Set

Feature	Description	Required NPE Capabilities	Relation to Other Features
Ethernet Learning (Source MAC Address Learning)	Implements a software database on the Intel XScale® Processor for storing and managing (searching, aging, and so forth) source MAC addresses detected on received packets.	NPE Learning Assistance feature is optional. Needed for automated population of database by IxEthAcc.	None
Ethernet Filtering (Destination MAC Address Filtering)	Provides Ethernet NPEs with MAC address data (learning/filtering trees) used to filter frames depending on frame destination MAC address.	NPE Learning Assistance and Filtering capabilities	Depends on Ethernet Learning. Mutually exclusive with 802.3/802.11 Frame Conversion.
VLAN / QoS	Configures VLAN and QoS support.	NPE VLAN and QoS support.	None

**Table 41. IxEthDB Feature Set (Continued)**

Feature	Description	Required NPE Capabilities	Relation to Other Features
Firewall (Source MAC Address Based)	Configures NPE firewall mode (white list/blacklist) and provides MAC address plus MAC address mask list for allowing/blocking.	NPE MAC-based Firewall	None
802.3 / 802.11 Frame Conversion	Configures NPE MAC address database, gateway access point database and frame conversion parameters (such as ToEth flag, LocalMac flag, ToSta flag, ToAP flag, VLANTag flag) and BSSIDs.  Configures pad length to add or remove bytes to/from incoming/outgoing 802.11 frames.  Configures logical port ID which is copied by NPE into IXP_NE	NPE 802.3 / 802.11 Frame Conversion.  In case the frame conversion requires VLAN tagging support for 802.11 frames, then NPE with VLAN/QoS capability must be used.	Only 802.3/802.11 frame conversion is mutually exclusive with Ethernet Filtering feature. However, VLAN feature can be enabled.
Spanning Tree Protocol	Sets Ethernet ports in blocked/unblocked STP state.	NPE STP support	None

The API can be used to enable or disable individual IxEthDB services on each NPE, assuming that an NPE has a given capability. For example, NPE A, NPE B and NPE C may all have microcode images with Ethernet Learning and Ethernet Filtering support. Using `ixEthDBFeatureEnable()`, the Ethernet Filtering capability could be disabled on NPE C.

Certain features are always functional and cannot be actually disabled. In these situations disabling the feature will cause its corresponding API to become inaccessible (returning `IX_ETH_DB_FEATURE_UNAVAILABLE`), and the feature is configured in such a way that the NPE behaves as if the feature is not implemented.

Ethernet Learning and Ethernet Filtering features are ENABLED by default when those capabilities are detected on NPE microcode. All remaining features are disabled by default.

## 10.4.2 Additional Database Features

In addition to the main features described in [Section 10.4.1, “Feature set” on page 172](#), the following subsections describe the additional database features available.

### 10.4.2.1 User-Defined Field

IxEthDB provides functions to associate a user-defined field to a database record, and later retrieve the value of that field. The user-defined field is passed as a (void \*) parameter, hence it can be used for any purpose (such as identifying a structure). Retrieving the user-defined field from a record is done using `ixEthDBUserFieldGet()`. Note that neither IxEthDB, nor the NPE microcode, ever uses the user-defined field for any internal operation and it is not aware of the significance of its contents. The field is only stored as a pointer.

The user-defined field may be added to any of the IxEthDB Intel XScale® Processor-based databases:

- XScale Learning/Filtering Database (including VLAN-related records)
- Ethernet Firewall Database
- Wi-Fi Header Conversion Database



#### 10.4.2.2 Database Clear

The IxEthDB component provides a function for removing port-specific records from each database listed above. It also provides the capability for removing one or more records from one, many, or all databases.

#### 10.4.3 MAC Address Learning and Filtering

There are two major elements involved in the IxEthDB MAC Address Learning and Filtering subsystem: a software database containing MAC address/port entries that resides on the Intel XScale® Processor of the processor, and a learning/filtering capability for each of the NPEs capable of Ethernet co-processing. Although it is possible to create static entries in the database via the IxEthDB API, most information is created dynamically via the MAC address learning process. The Intel XScale® Processor-based database aggregates all of the MAC address/port entries and can also push learning/filtering entries down to the NPEs.

The NPE-based data structure of MAC addresses learned or to be filtered is referred to throughout this document as the NPE Learning/Filtering Tree. Each NPE has its own NPE Learning/Filtering Tree. On a multiple-NPE processor, the trees for each port will usually have different data sets.

The Intel XScale® Processor-based database is referred to as the XScale Learning/Filtering Database. This database contains learning/filtering entries for all of the ports managed by the IxEthDB component. When IxEthDB is configured to enable learning, the IxEthDB component handles downloading data from the XScale Learning/Filtering Database to each NPE Learning/Filtering Tree automatically.

##### 10.4.3.1 Learning and Filtering

EthDB implements learning, filtering and Spanning Tree algorithm.

The NPEs provide a function whereby source MAC address learning is performed on received (ingress) Ethernet frames. If learning is enabled, the source MAC address of the received frame is compared against the entries in the *NPE Learning/Filtering Tree* and against the MAC address of the receiving port. If no matches are found, the MAC address of the receiving port is extracted from the frame, and the MAC address and receiving port ID are passed to the Intel XScale® Processor in the IX\_OSAL\_MBUF header, along with a notification flag. The EthDB component adds the new MAC address / port ID record into the *XScale Learning/Filtering Database*. The process of detecting new source MAC addresses and adding the new MAC address / port ID combination into the database is known as **learning**.

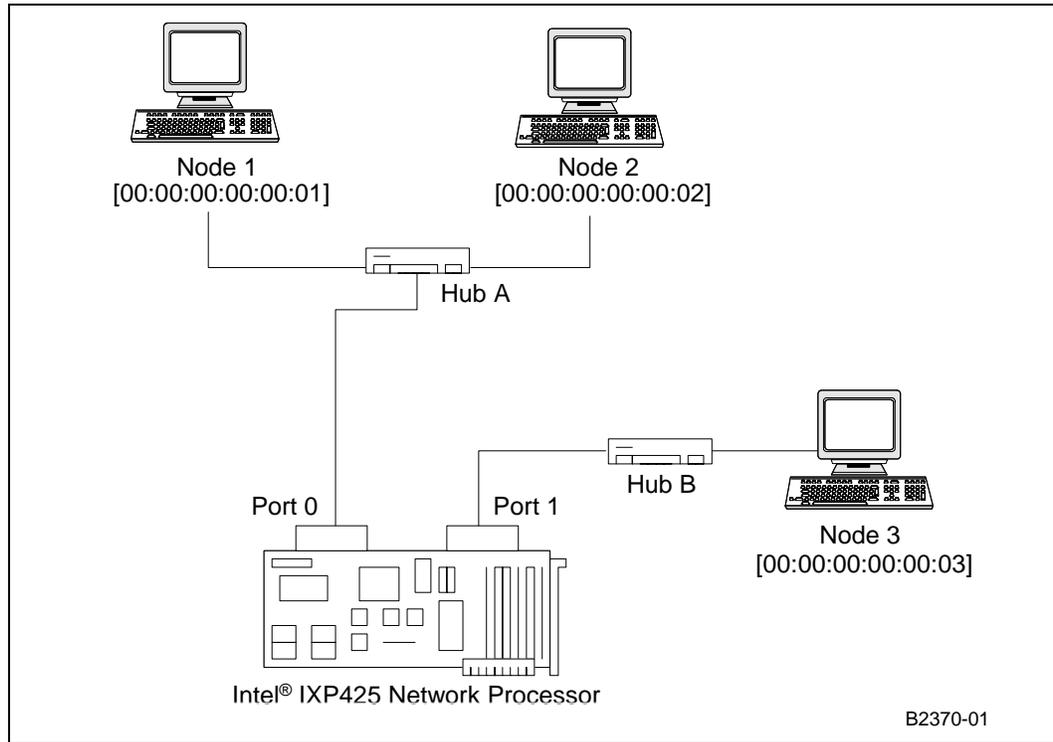
As per IEEE802.1D, an Ethernet bridge must filter frames that are received through a specific port but are destined for another station on the same LAN. To achieve this functionality, the NPE extracts the destination MAC address from every received frame and then attempts to find a match in the *NPE Learning/Filtering Tree*. If no match is found, the frame continues on to the next step of receive path processing. If a match is found, the NPE inspects the Port ID field of the matching *NPE Learning/Filtering Tree* entry. If the value of the Port ID field is equal to that of the port through which the frame was received, the frame is dropped and the RxLearnedEntryDiscards counter is updated; otherwise, the frame is not filtered and is allowed to continue on to the next step of receive path processing. This process of dropping a frame using the logic described here is called **filtering**.

Filtering can also be done according to some characteristics of a frame received on a port, such as frames exceeding a maximum frame size or frames that do not include VLAN tagging information. For example, EthDB provides a facility to set the maximum frame size that should be accepted for each NPE-based port. This means that if a port



receives a frame that is larger than the maximum frame size, that frame is filtered. An example of this type of filtering can be found in section titled “Filtering Example Based Upon Maximum Frame Size” on page 179.

**Figure 55. Example Network Diagram for MAC Address Learning and Filtering with Two Ports**



Assuming we start with blank (empty) learning trees, a possible scenario of filtering is the following:

- Node 1 sends a frame to Node 3 (source MAC 00:00:00:00:00:01, destination 00:00:00:00:00:03)
  - The frame is forwarded by Hub A to Node 2 (ignores the frame, as the destination does not match its own address) and Port 0
  - Port 0 adds the source address (00:00:00:00:00:00:01) to its learning tree
  - Port 0 searches for the destination address (00:00:00:00:00:03) in its learning tree, it is not found therefore the frame is forwarded to the other ports – in this case Port 1
  - Port 1 forwards the frame to Hub B
  - Hub B forwards the frame to Node 3, intended recipient of the frame
- Node 2 sends a frame to Node 1 (source MAC 00:00:00:00:00:02, destination 00:00:00:00:00:01)
  - The frame is sent to Hub A, which forwards it to Node 1 (intended recipient) and Port 0
  - Port 0 adds the source MAC address (00:00:00:00:00:02) to its learning tree
  - Port 0 searches for the destination address (00:00:00:00:00:01) in its learning tree, it is found therefore Port 0 knows that both Node 1 and Node 2 are connected on the same side of the network, and this network already has a



frame forwarder (in this case Hub A) – the frame is filtered (dropped) to prevent unnecessary propagation.

MAC address frame filtering based on learning trees is usable only when a port operates in promiscuous mode. Otherwise the frames is filtered at the MAC (not NPE) level according to normal MAC filtering rules — the frame is received only if the destination address matches the port address, if the destination is the broadcast address, or if the destination is a multicast address subscribed to by the port.

### 10.4.3.2 Other MAC Learning/Filtering Usage Models

If a terminal (source of Ethernet traffic on the network) is moved from one NPE port to another, IxEthDB is responsible for ensuring the consistency of the *XScale Learning/Filtering Database*. The Intel XScale® Processor database and *NPE Learning/Filtering Trees* are updated within one second of the terminal move being detected. The change is detected when traffic is first received from the terminal on the new NPE port. This behavior is described as “**migrating**”.

One of the advantages of the split NPE/XScale model is that the NPE can attempt to identify if an incoming frame is destined for another known port in the system. For example, the *NPE Learning/Filtering Tree* for port 1 may contain an entry that shows the frames destination MAC address as having been learned on port 2. The NPE will include the destination port id in the IX\_OSAL\_MBUF header fields as part of the receive callback.

There are some situations in which the *NPE Learning/Filtering Trees* may not have learned the proper destination port for a received packet. The NPEs will then pass the packet to the IxEthAcc component to allow it to search the *XScale Learning/Filtering Database* for the proper destination port. If the system is operating in a **bridging or switching** fashion, the *XScale Learning/Filtering Database* will know the appropriate port to send the packet out on. If the *XScale Learning/Filtering Database* does not know the appropriate destination port, the receive callback function will set the port ID field in the IX\_OSAL\_MBUF header to a value of IX\_ETH\_DB\_UNKNOWN\_PORT, indicating that the destination port of this packet is unknown. The client may then **broadcast** on all ports in the hopes that a node somewhere on the network will respond.

### 10.4.3.3 Learning/Filtering General Characteristics

#### Port Definitions

The port definition in the Ethernet Database component does not directly depend on the number of Ethernet ports available on the Intel® IXP4XX Product Line of Network Processors. The user can define up to 255 ports (including the Ethernet NPE ports), which is recognized by the component, although this definition is static and cannot be changed at run-time. The only requirement is that port ID 0-5 are re-served for Ethernet NPE ports and cannot be used for user ports (nor should they be removed). Port IDs therefore range between 0 and 0xFE.

Port definitions are placed in the public include file IxEthDBFeatures.c (located in ixp400\_xscale/src/ethDB). The main port definition table is an array having the following format:

```
IxEthDBPortDefinition ixEthDBPortDefinitions[IX_ETH_DB_NUMBER_OF_PORTS] =  
{  
    /* id    type    capabilities */  
    /* 0 */    IX_ETH_NPE, IX_ETH_NO_CAPABILITIES}, /*Ethernet NPE B*/
```



```

    /* 1 */ IX_ETH_NPE, IX_ETH_NO_CAPABILITIES}, /*Ethernet NPE C*/
    /* 2 */ IX_ETH_NPE, IX_ETH_NO_CAPABILITIES}/*Ethernet NPE A*/
};

```

The first six entries are reserved and the user can add additional ports starting with ID 6. The definitions listed above include the six Ethernet NPE ports and one example user-defined WAN port. Port numbers (IDs) are automatically determined from the definition location - they are written as comments above only for clarity reasons.

All user ports must be defined as ETH\_GENERIC with NO\_CAPABILITIES in order to accommodate future software features. Unlike the Ethernet NPEs, user-defined ports lack certain capabilities, namely ETH\_GENERIC describes a port with no automatic database update features, and NO\_CAPABILITIES is a descriptor indicating the port has no special capabilities. Unlike IXP400 software versions prior to 1.5, NPE software no longer ages MAC addresses automatically. Therefore the ENTRY\_AGING capability is no longer defined for NPE ports. This characterization will instruct the Ethernet Database component not to attempt to upload up-to-date learning trees in user ports and age the entries itself.

EthDB is not strictly limited to the NPE-based Ethernet ports available on the IXP4XX product line processor. The user can define up to 255 ports (including the Ethernet NPE ports), which is recognized by the component. Adding user-defined ports (such as one representing a PCI-based Ethernet adapter) allows the manual provision of MAC address/port records to the *XScale Learning/Filtering Database* and the *NPE Learning/Filtering Trees* via the IxEthDB API. The NPEs will then be able to detect that an incoming frame is destined for the user-defined port, and report the destination port ID in the IX\_OSAL\_MBUF header for the frame.

Do not change or remove the ports marked as **reserved** or **NPE**; the Ethernet Access component relies on this definition. Accordingly, IX\_ETH\_DB\_NUMBER\_OF\_PORTS should be set to at least two at any time. Other components may have also defined their own ports (see the related header file for up-to-date information).

**Warning:** The id value assigned to NPE ports in IxEthDbPortDefs.h may not be the same as the value used to identify ports in the IXP\_BUF fields written by the NPEs, as documented in [Table 37 on page 167](#). The Ethernet device driver for the supported operating systems may enumerate the NPE ports differently as well.

### Limits for Number of Supported Learning/Filtering Entries

Each NPE is capable of storing 511 MAC address entries in its *NPE Learning/Filtering Tree*. The *XScale Learning/Filtering Database* will handle all the addresses for all NPEs plus any number of addresses required for user-defined ports, up to 4096 records by default. This will suffice for the three NPEs and a considerable number of user-defined ports plus operating headroom. If the value is not large enough the user can tweak database pre-allocation structures by changing `ixp400_xscalesw/src/ethDB/include/IxEthDB_p.h`.

It is not recommended to add more than 511 addresses per NPE port. While IxEthDB itself can learn more than 511 entries per port, the NPEs cannot use more than 511. If more than 511 entries are defined for an NPE port, only the first 511 entries will be stored in NPE for filtering.

### Port Dependency Map

The IxEthDB API provides functions to set or retrieve Port Dependency Maps. The Port Dependency Maps are used to share filtering information between ports. By adding a port into another port's dependency map, the target port filtering data will import the



filtering data from the port it depends on. Any changes to filtering data for a port — such as adding, updating or removing records — will trigger updates in the filtering information for all the ports depending on the updated port.

For example, if ports 2 and 3 are set in the port 0 dependency map the filtering information for port 0 will also include the filtering information from ports 2 and 3. Adding a record to port 2 will also trigger an update not only on port 2 but also on port 0.

This feature is useful in conjunction with the NPE destination port lookup service, where the NPE searches for the destination MAC of a received frame in its tree and, if found, copies the port ID from the record into the buffer header. This saves the Intel XScale® Processor from having to perform this lookup in a switching application.

### Provisioning Static and Dynamic Entries

The IxEthDB API provides a function allowing the user to statically provision entries in the *XScale Learning/Filtering Database*. Dynamic entries may also be provisioned via the API. It is important to note that if a static MAC address is provisioned for port X, but later a frame having this source MAC address is detected arriving from port Y, the record in the database is updated from X to Y and the record will no longer be marked as static.

### Aging

Aging is the process through which inactive MAC addresses are removed from the filtering database. At periodic intervals, the *XScale Learning/Filtering Database* is examined to determine if any of the learned (or dynamically provisioned) MAC addresses have become inactive during the last period (for example, no traffic has originated from those MAC addresses/port pairs for a period of roughly 15 minutes). If so, they are removed from the *XScale Learning/Filtering Database*.

In the software release 2.3, if the NPE finds a match to a source MAC address in its *NPE Learning/Filtering Tree* as part of the learning process, the NPE will update the record to indicate that the transmitting station is still active. At defined intervals, the *NPE Learning/Filtering Tree* data is merged into the *XScale Learning/Filtering Database*, so that it reflects the current age of MAC address entries and can expire older entries as appropriate. This is tied into the database maintenance functionality, further documented in [“Database Maintenance” on page 179](#). When a record age exceeds the `IX_ETH_DB_LEARNING_ENTRY_AGE_TIME` definition, the record is removed at the next maintenance interval.

`IX_ETH_DB_LEARNING_ENTRY_AGE_TIME` is 15 minutes by default, but may be changed as appropriate.

The aging of entries is handled first in the XScale Learning/Filtering Database and propagated to the NPE Learning/Filtering Trees.

Static entries provisioned using the IxEthDB API are not subject to aging. Provisioned entries that are defined as dynamic (`ixEthDBFilteringDynamicEntryProvision()`) are subject to aging.

*Note:* Entries age only if their ingress port is explicitly configured to do so using the `ixEthDBPortAgingEnable()` function.

### Record Management

The IxEthDB component contains functions for managing records in its various databases. Capabilities specific to the MAC Address Learning/Filtering facility include:

- Add static or dynamic records.



- Remove records.
- Search for a given MAC address, with the option to reset the aging value in the record.
- Displaying the database contents, grouped by port.

### Database Maintenance

Maintenance is required to facilitate the aging of entries in the *XScale Learning/Filtering Database* and *NPE Learning/Filtering Trees*.

The IxEthDB component performs all database maintenance functions. To facilitate this, the `ixEthDBDatabaseMaintenance()` function must be called with a frequency of `IX_ETH_DB_MAINTENANCE_TIME`. It is the client's responsibility to ensure the `ixEthDBDatabaseMaintenance()` function is executed with the required frequency. The default value of `IX_ETH_DB_MAINTENANCE_TIME` is one minute.

If the maintenance function is not called, then the aging function will not run. An entry is aged at `IX_ETH_DB_LEARNING_ENTRY_AGE_TIME +/- IX_ETH_DB_MAINTENANCE_TIME` seconds.

## 10.4.4 Frame Size Filtering

The API provides the ability to set the maximum size of Ethernet frames supported per port, using the `ixEthDBFilteringPortMaximumFrameSizeSet()` function. When a maximum frame size value is set for a port, there are multiple effects:

- Any incoming (Rx) frames on the specified port larger than the set value is dropped. No further learning process will be done on this frame.
- In the Transmit data path, the NPE will check the size of an Ethernet frame during the final stage of processing the frame, just prior to transmission. If the NPE adds data (VLAN tag or FCS, for example) that causes the frame to exceed the maximum frame size, the frame will not be transmitted. The `TxLargeFramesDiscard` counter is incremented (see [Chapter 9.0](#)).

The maximum supported value is 16,320 bytes. For purposes of clarification, the number of bytes making up the Maximum Frame Size value is the Ethernet MSDU (Media Service Data Unit) and defined as the sum of the sizes of:

- the Ethernet header: dest MAC + src MAC + VLAN Tag and/or length/type field
- the Ethernet payload
- the Ethernet frame check sequence (FCS), if not stripped out by `IxEthAccPortRxFrameFcsDisable()`.

### 10.4.4.1 Filtering Example Based Upon Maximum Frame Size

On a system with three ports (0, 1, 2), execute:

```
ixEthDBFilteringPortMaximumFrameSizeSet(0, 9014);
ixEthDBFilteringPortMaximumFrameSizeSet(1, 9014);
ixEthDBFilteringPortMaximumFrameSizeSet(2, 1514).
```

The NPE on Ports 0 and 1 will filter all Rx frames over 9,014 bytes.

A frame of 1,000 bytes is received on Port 2. The NPE will determine the destination port based on learned MAC address, and:

- If the port is unknown, process the frame.



- If the destination port is 0 or 1, process the frame.
- If the port is 2, drop the frame according to the normal MAC filtering rules.

A frame of 3,000 bytes is received on Port 2, it is dropped according to the frame size setting.

#### 10.4.5 Source MAC Address Firewall

The Ethernet NPE firmware provides three firewall-related services, each of which is capable of filtering a frame based on the value of its source MAC address field:

- Invalid MAC address filtering
- MAC address block (**black list**)
- MAC address admission (**white list**)

This feature is dependent on the run-time NPE configuration and specific NPE image capabilities, described in “Dependencies” on page 199 and Chapter 14.0). Each NPE supporting this feature can be configured independently of the others.

Firewall MAC address filtering based on **MAC address mask** at the NPE level is supported.

##### MAC Address Block/Admission

IxEthDB supports per-NPE MAC address-based firewall lists and provides the API to add/remove these MAC addresses, as well as to configure the NPE firewall. There are two firewall operating modes:

- allow / white list state – only incoming packets with a source MAC addresses found in the firewall list are allowed
- deny / black list state – all incoming packets are allowed except for those whose source address is found in the firewall list.

As mentioned above, the per-NPE firewall lists have an address mask associated with each MAC address entry. With a mask, a segment of addresses can be allowed/denied using a single firewall entry. The standard API for adding entries will assume unmasked entries, but a new API, **ixEthDBFirewallMaskedEntryAdd()**, will also allow the addition of an address/mask pair. The new API **ixEthDBFirewallMaskedEntryRemove()** allows removal of address/mask pair.

The address/mask pair should be of the following format:

This address/mask pair would allow or deny all MAC addresses from 00:01:00:00:00:00 to 00:01:00:FF:FF:FF.

The firewall lists support a maximum of 31 addresses. This feature is disabled by default and there are no pre-defined firewall records. When enabled, it operates in black list mode until reconfigured. The firewall feature can be freely turned on or off and reconfigured at run time.

IxEthDB contains an *Ethernet Firewall Database* that contains MAC address and mask (and port ID records) for this firewall feature. Each MAC address/mask pair is also unique and the same port can have more than one entry with the same MAC address, providing the masks for each of these entries are unique.

Also, the firewall records are independent of the *XScale Learning/Filtering Database* and Wi-Fi header conversion records, for example, firewall records can co-exist with either Wi-Fi or filtering records but filtering and Wi-Fi records are mutual exclusive and they can not co-exist. Once configured, the API is used to download a firewall filtering table to the NPE.



A typical usage scenario of this feature would consist of the following steps:

1. Enable the IX\_ETH\_DB\_FIREWALL feature
2. Set the firewall operating mode (white list or black list)
3. Add addresses to be blocked (black list mode) or specifically allowed (white list mode). Or add address segments using addresses and masks.
4. Download the firewall configuration data using ixEthDBFirewallTableDownload(port)

### Invalid MAC Address Filtering

According to IEEE802, it is illegal for the source address of an Ethernet frame to be either a broadcast address or a multicast address. These broadcast/multicast addresses are distinguished by the value of their first bit (for example, the least significant bit of the first byte). If the first bit of the MAC address is 1, the MAC address is either a broadcast or multicast address.

IxEthDB can be used to enable invalid source MAC address filtering in the NPE. When this feature is enabled, the NPE will inspect the source MAC address of incoming packets and drop packets whose source MAC address is a multicast or broadcast address. IxEthDB disables this feature by default.

## 10.4.6 IPv4 and IPv6 Payload Detection

IXP400 software supports identification of Ethernet Frames that carry IPv6 packets at NPE level.

For every received frame delivered to the Intel XScale® Processor, the NPE firmware reports whether the payload of the frame is an IPv4 or IPv6 packet by setting the `ixp_ne_flags.ip_prot` flag (2-bit) in the buffer header according to [Table 42, “Possible IP Types” on page 181](#). NPE Ethernet firmware examines the Length/Type field to determine whether the payload is IPv4 or IPv6. A value of 0x0800 indicates that the payload is IPv4 and 0x86DD indicates that the payload is IPv6.

**Table 42. Possible IP Types**

<code>ip_prot</code> field value	IP Version
00	NonIP payload
01	IPv4 payload
10	IPv6 payload
11	Reserved

The IPv4 and IPv6 payload detection service is always enabled in any NPE firmware load that supports it. However, the application software is in no way obligated to make use of the fields written by this service. The capability of this service is limited by the overall capabilities of the particular NPE firmware version.

Only an NPE firmware version that is VLAN-capable is intelligent enough to understand the VLAN tag of a tagged frame and report such frames as IP frames. That is, an NPE firmware version that is not VLAN-capable will always report such frames as non-IP. This is because it has no intelligence to understand and ignore the VLAN field.



### 10.4.7 802.1Q VLAN

The IxEthDB component provides support for VLAN features when using NPE microcode images that include VLAN support. All the major VLAN features defined in IEEE 802.1Q are supported. These include:

- Acceptable frame type filtering for each ingress port
- VLAN tagging and tag removal for each ingress and egress port
- VLAN membership filtering for each ingress port
- VLAN tagging and tag removal control for individual egress packets
- Support for a maximum of 4095 VLAN groups
- Tagging or un-tagging control to 802.11 frames

This feature makes heavy use of the IX\_OSAL\_MBUF header flag fields to allow a client application to make VLAN-based processing decisions. Their NPE behavior for these header fields is documented in this section. However, refer to [Chapter 9.0](#) for a more comprehensive understanding of the data path.

#### 10.4.7.1 Background – VLAN Data in Ethernet Frames

According to IEEE802.3, an untagged or normal Ethernet frame has the fields listed in [Table 43](#).

**Table 43. Untagged MAC Frame Format**

0	1	2	3	4	5	6	7	8	9	10	11	12	13		
Destination address						Source address						Length/Type	MAC client data and pad (46–1500 bytes)	FCS	

The **Length/Type** field is differentiated by whether its numerical value is greater than or equal to 0x600. If it is greater than or equal to 0x600, the field is interpreted as *Type*, which additionally implies that there is no LLC/SNAP header in the frame. Otherwise, the field is interpreted as *Length*, for example the number of bytes in the **MAC client data** field. In this case, it is also implied that the first field in the **MAC client data** field is an LLC/SNAP header.

**Table 44. VLAN Tagged MAC Frame Format**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17		
Destination address						Source address						VLAN TPID	VLAN TCI	Length/Type	MAC client data and pad (46–1500 bytes)	FCS			

**Table 45. VLAN Tag Format**

12												13												14												15											
31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0																
VLAN TPID												VLAN TCI																																			
0x810												0x0/Port ID						Priority			CFI			VLAN ID																							

The VLAN tagged Ethernet frame format, as specified in IEEE802.3, is as listed in [Table 44](#). A received frame is considered to be VLAN-tagged if the two bytes at offset 12-13 are equal to 0x8100. Note that this definition of a “VLAN-tagged frame” is meant to include frames that are only priority-tagged (for example, frames whose VLAN ID is 0).



### 10.4.7.2 Database Records Associated With VLAN IDs

IxEthDB supports MAC-based VLAN classification for a bridging application by providing the API to associate a VLAN ID with a record (identified by a MAC address), and later retrieve the VLAN ID provided the MAC address is known. This data structure is essentially the *XScale Learning/Filtering Database* with an additional 802.1Q field for each record.

In a typical bridge scenario where MAC-based classification is used, the bridge would be provided with MAC address-VLAN ID association via a user-controlled configuration mechanism, which is stored in IxEthDB by using `ixEthDBVlanTagSet()`. Classification based on MAC addresses can be then achieved on the data path by searching the VLAN ID of each received buffer using `ixEthDBVlanTagGet()`.

Note that while theoretically it is possible to duplicate MAC addresses across VLANs, this is not supported by IxEthDB for the purpose of MAC-based classification support. Each record (hence each MAC address) can only be associated with one VLAN ID. It should also be noted that MAC duplication across a network is an error.

### 10.4.7.3 Acceptable Frame Type Filtering

IxEthDB defines an API for setting per-port acceptable frame type filtering policies. Frame identification and IEEE 802.1Q compliance are ensured by the NPE, which can detect and filter untagged, tagged and priority-tagged frame types. The filtering policies are defined as follows:

- Accept untagged (no 802.1Q tag).
- Accept tagged (802.1Q tag is detected, includes user priority and frame VLAN ID membership).
- Accept priority-tagged (802.1Q tag is detected, includes user priority and no VLAN membership – VLAN ID set to 0).

*Note:* Setting the acceptable frame type to `PRIORITY_TAGGED_FRAMES` is accomplished within the API by changing the frame filter to `VLAN_TAGGED_FRAMES` and setting the VLAN membership list of the port in question to include only VLAN ID 0. The membership list will need to be restored manually to an appropriate value if the acceptable frame type filter is changed back to `ACCEPT_ALL_FRAMES` or `VLAN_TAGGED_FRAMES`. Failure to do so will filter all VLAN traffic except those frames tagged with VLAN ID 0.

The acceptable frame type filter can be any of the values above. Additionally, filters can be combined (ORed) to achieve additional effects:

- Accept all frames – equivalent to accept tagged and accept untagged. Used to declare hybrid VLAN trunks.
- Accept only untagged and priority tagged frames – equivalent to discard frames pertaining to a VLAN. Used to declare trunks that are QoS aware but do not support VLAN.

By default all ports accept all the frame types. The frame type filter can be dynamically configured at run time.

### 10.4.7.4 Ingress Tagging and Tag Removal

Each port can be associated with a default 802.1Q tag control information field, which includes the **Priority**, **CFI**, and **VLAN ID** fields. Each port can be individually configured to tag all the incoming untagged frames, remove the tag from all the incoming tagged frames, or leave the frames unchanged.



Applying the default port 802.1Q tag to incoming untagged frames constitutes port-based VLAN classification. Untagged Ingress frames will automatically be associated with the default port VLAN of the port they were received on, if this feature is enabled.

Ports can be configured to remove the 802.1Q tag from the incoming frames, if the tag is present (type/len field set to 0x8100). This feature will guarantee that no packets received from the port is VLAN or priority tagged when the port is configured to remove tag for incoming frames as if it were an 802.1Q-unaware port.

By default each port is configured in pass-through mode. When using this mode no tags are applied or removed from the incoming frames. In this mode ports operate as hybrid VLAN trunks. Tagging and tag removal can be dynamically configured at run time.

The NPE microcode sets the *ixp\_ne\_flags.vlan\_en* field in the IX\_OSAL\_MBUF to 1 on all frames during ingress for all VLAN-enabled NPE images, or is set to 0 on all frames for all non-VLAN-enabled NPE images. This field value is useful on egress because the NPE microcode can use it to distinguish between regular untagged Ethernet frames and tagged frames that have Priority 0 + VLAN ID 0. The *ixp\_ne\_vlan\_tci* field value is 0 for both types of frames.

*Note:* The NPE cannot update the **FCS** field to reflect the changes made to frames modified by ingress tagging or tag removal. The client application should disable receive FCS appending (`ixEthAccPortRxFrameAppendFCSDisable()`), or ignore the FCS contents on received frames.

#### 10.4.7.5 Port-Based VLAN Membership Filtering

Ports can be individually configured to define their VLAN membership status and enable VLAN membership filtering of incoming and outgoing frames.

Port VLAN membership is a group of VLAN IDs which are allowed to be received and transmitted on the specified port. If the port is VLAN enabled (Port VLAN ID (PVID) — not set to 0), the minimum membership group for the port is its own PVID. Ports with no default VLAN membership (PVID set to 0) cannot have membership groups and cannot filter frames based on VLAN membership information. A VLAN membership group is a set of VLAN IDs to which the port adheres to.

For example, Port 1 is configured with a PVID set to 12 and VLAN membership group of {1, 2, 10, 12, 20 to 40, 100, 102, 3000 to 3010}. If VLAN membership filtering is enabled and acceptable frame type filtering is configured appropriately for the port, the following scenarios are possible:

- If tagging is not enabled, untagged frames is left untagged and passed through,.
- If tagging is enabled, untagged frames is tagged with a VLAN ID set from the port PVID (12) and passed through. Since the frame is tagged with the port VLAN ID, it will always be accepted by the same port's membership table.
- Tagged frames is checked against the port membership table, therefore:
  - frames with VLAN IDs of 2, 10, 25, 100 or 3009 is accepted,
  - frames with VLAN IDs of 0 (priority-tagged frame), 4, 15, 200 or 4072 is discarded.

The IxEthDB API allows the user to add and remove individual VLAN ID entries as well as entire VLAN ranges into each port's VLAN membership table. Also, membership checks can be enabled or disabled at run time.

Port membership filtering is disabled by default.



Note that a port will always have a non-empty membership table. By default the PVID, which is 0 at initialization time, is declared in the membership table. The PVID cannot be removed from the membership table at any time.

#### 10.4.7.6 Port and VLAN-Based Egress Tagging and Tag Removal

IxEthDB supports configuration of Egress frame tagging and tag removal, depending on the NPE image capabilities. Unlike Ingress tagging and tag removal, the egress tagging process adds a per-VLAN tagging configuration option. The port membership and egress tagging settings for each VLAN are stored in a structure called the Transmit Tagging Information (TTI) table.

Tagging and tag removal can also be individually overridden for each frame, using the following IX\_OSAL\_MBUF header flags:

- **ixp\_ne\_flags.vlan\_en** – This flag must be enabled if any tagging or untagging will occur.

If the `vlan_en` (VLAN ENABLE) flag is not set, the frame is treated as a non-VLAN frame, and no VLAN processing will take place for the frame. The frame is transmitted unmodified (no tagging or tag removal will take place), and membership filtering will not take place, irrespective of the port configuration and value of the `tci` field. VLAN-enabled NPE images always set this flag during frame Rx, and NPE images without VLAN capabilities always clear this flag during frame Rx. Manually changing this flag on the data path can be used to implement a hybrid VLAN bridge (for example, a bridge that can forward both untagged frames in their original untagged format, as well as VLAN frames).

- **ixp\_ne\_flags.tag\_over** – transmit VLAN override tag. A value of 0 indicates that the default tagging behavior for the port/VID should be used. A value of 1 indicates an override. The `ixp_ne_flags.tag_mode` flag can be set by the client application to override the Egress tagging behavior, and the `ixp_ne_vlan_tci` field can be populated with the proper TCI information for that frame.
- **ixp\_ne\_flags.tag\_mode** – VLAN tag behavior control. A value of 0 indicates that the frame is transmitted untagged. A value of 1 indicates that the frame is tagged. This flag can be set by the client application to override the default Egress tagging behavior.
- **ixp\_ne\_vlan\_tci** – tag control information. Frames are tagged using this tag, irrespective whether they already have a VLAN tag or not.

Use **ixp\_ne\_flags.vlan\_en** to override the special conditions listed below.

The `ixp_ne_vlan_tci` field is automatically populated on ingress with the 802.1Q tag present in the frame (if any), or with the ingress port VLAN ID tag (for untagged frames). This happens even if the frame is untagged during ingress, giving the client application a chance to inspect the original VLAN tag. If this field is not changed by the client code, the frame is re-tagged on transmission with the same tag.

Tagging frames on egress is determined in the following manner:

- The frame IX\_OSAL\_MBUF header can contain override information (flags – see above) explicitly stating whether the frame is to be tagged or not.
- Tagging information (802.1Q tag) is contained in the IX\_OSAL\_MBUF header.
- The frame VLAN ID, if any, is compared against the transmit port VLAN membership table and discarded if not found in the membership table.
- If the buffer header does not override the port tagging behavior, then the TTI table is consulted for the VLAN ID found in the `ixp_ne_vlan_tci` field of the frame header. If the bit corresponding to the VLAN ID is set, the frame is to be tagged by the NPE prior to transmission. Otherwise, the frame is transmitted without the tag



### Special Conditions

The NPE microcode uses the *ixp\_ne\_flags.vlan\_en* field to distinguish between regular untagged Ethernet frames and tagged frames that have Priority 0 + VLAN ID 0, since both will have an IX\_OSAL\_MBUF header *ixp\_ne\_vlan\_tci* value of 0.

If egress tagging is enabled on VLAN ID 0, then the *ixp\_ne\_flags.vlan\_en* field must be disabled for regular untagged Ethernet frames to prevent them from being tagged with Priority 0. Similarly, if Egress tagging is disabled on VLAN ID 0, then Priority 0 tagged frames must enable the *ixp\_ne\_flags.vlan\_en* field to override the default behavior of sending them as untagged frames.

*vlanTagFlag* in the header conversion table is valid only when *vlan\_en* (VLAN ENABLE) flag is set. Otherwise 802.11 header converted frames will also be treated as non-VLAN frames and will not be tagged.

*Note:* When using the egress VLAN-tagging feature, be sure to enable FCS appending ( *ixEthAccPortTxFrameAppendFCSEnable()* ) on the affected NPE ports so that a valid FCS is calculated and appended to the frame prior to transmission. Refer to [Section 10.5.4.2, “FCS Appending” on page 200](#)

An overview of the Egress tagging process is shown in [Figure 56](#). The figure shows the decision tree for an untagged frame. The process is identical for a tagged frame.

**Figure 56. Egress VLAN Control Path for Untagged Frames**

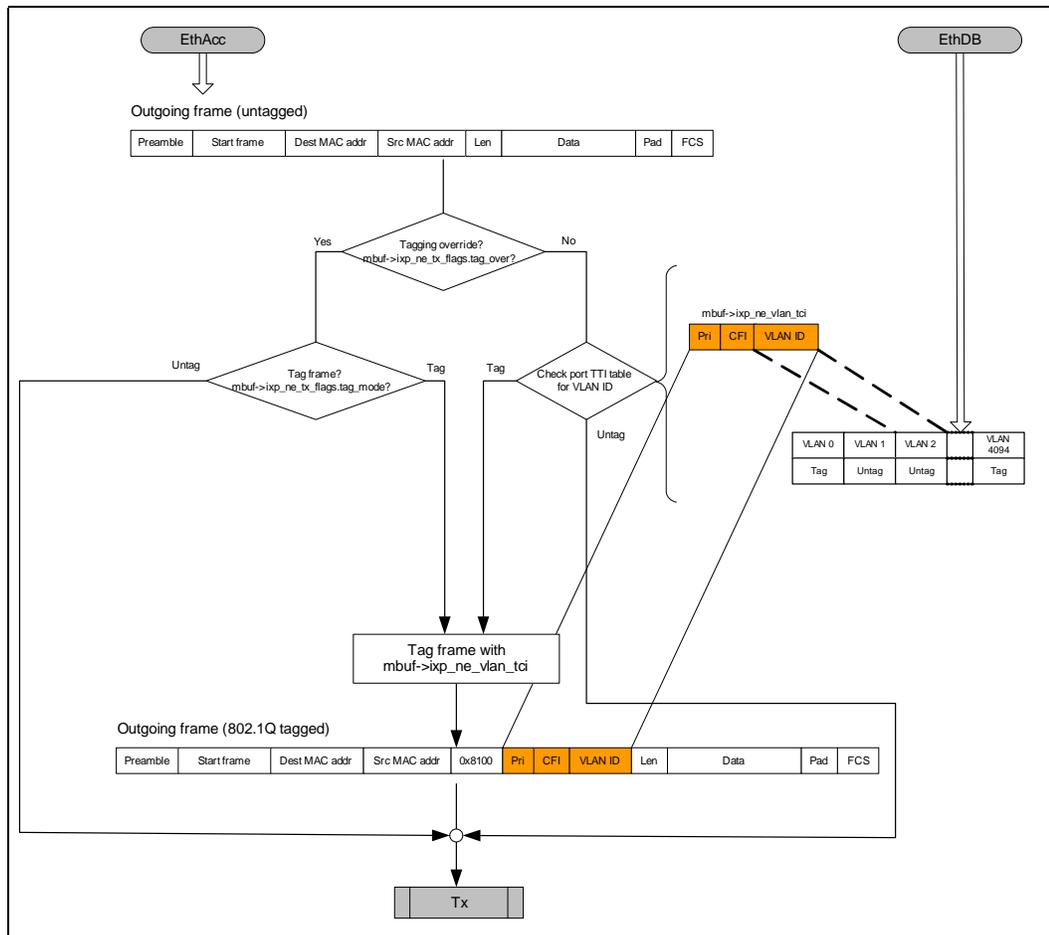




Table 46 presents an egress VLAN tagging/untagging behavior matrix.

**Table 46. Egress VLAN Tagging/Untagging Behavior Matrix**

Tag Mode (1)	Frame Status (2)	Action
Untag	Untagged	The NPE microcode does not modify the frame.
Untag	Tagged	The NPE microcode removes the VLAN tag from the frame.
Tag	Untagged	The NPE microcode inserts a VLAN tag into the frame. The VLAN tag to be inserted is created by concatenating a VLAN TPID field (always 0x8100) with the value of the <i>ixp_ne_vlan_tci</i> field from the IX_OSAL_MBUF header.
Tag	Tagged	The NPE microcode overwrites a VLAN TCI field of the frame with the value of the <i>ixp_ne_vlan_tci</i> field from the IX_OSAL_MBUF header.

(1) - The tag mode is the result obtained by consulting the Transmit Tagging Table, unless it is overridden by the *ixp\_ne\_tx\_flag* field of the frame's IX\_OSAL\_MBUF header, as described above.  
(2) - The (input) frame status is determined by examining the *ixp\_ne\_flags.vlan\_prot* flag from the frame's IX\_OSAL\_MBUF header.

#### 10.4.7.7 Port ID Extraction

A device connected to an MII interface can be a single one-port Ethernet PHY or a multi-port device (such as a switch). Some popular Ethernet switch chips use the **VLAN TPID** field (see Table 44) in VLAN-tagged frames to encode the port through which a frame is received. These devices encode the physical port from which a frame is received in the least significant 4 bits of this field.

IxEthDB provides the API for enabling the NPE to extract this port ID information. When enabled using the function `ixEthDBVlanPortExtractionEnable()`, the NPE will copy the port ID from the VLAN type field into the *ixp\_ne\_src\_port* field of the buffer header and restore the VLAN type field to 0x8100. This feature is disabled by default and can be switched on or off at run time.

When not enabled, the *ixp\_ne\_src\_port* value is the physical MII port ID (for example, always 0 or 1) (Note: This is not applicable for IXP46X/IXP45X Network Processors).

#### 10.4.8 802.1Q User Priority / QoS Support

The IxEthDB component provides support for QoS features when using NPE microcode images that include VLAN and QoS support. This support includes:

- Priority aware transmit and receive, using different priority queues for transmit and receive.
- QoS priority (for example, user priority, as per IEEE802.1Q) to traffic class mapping via priority mapping tables on received frames.
- Priority frame tagging and tag removal prior to transmission. This is discussed in "Port and VLAN-Based Egress Tagging and Tag Removal" on page 185.

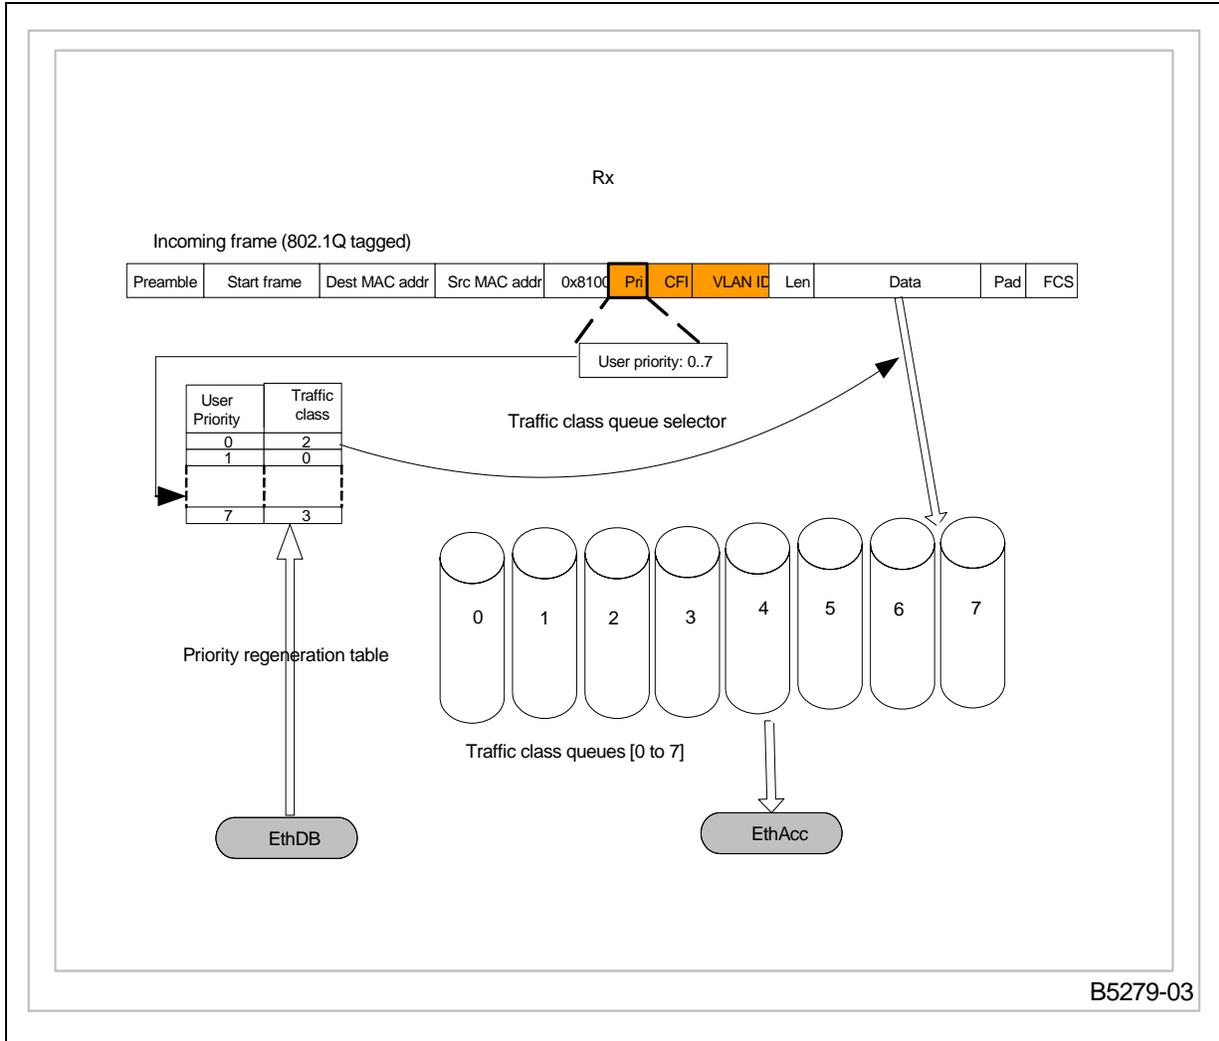
##### 10.4.8.1 Priority Aware Transmission

Submitting Ethernet frames for transmission is done by specifying a traffic class (priority) to be used for ordering frame transmission requests. This feature is covered in Section 9.6.2.2.

### 10.4.8.2 Receive Priority Queuing

Incoming frames is classified into an internal traffic class, either by mapping the 802.1Q priority field (if available) into an internal traffic class or by using the default traffic class associated with the incoming port. The incoming frame is placed on a receive queue depending on its traffic class. Up to eight traffic classes and associated queues are supported. Traffic classes are ordered in their priority order, with 0 being the lowest priority.

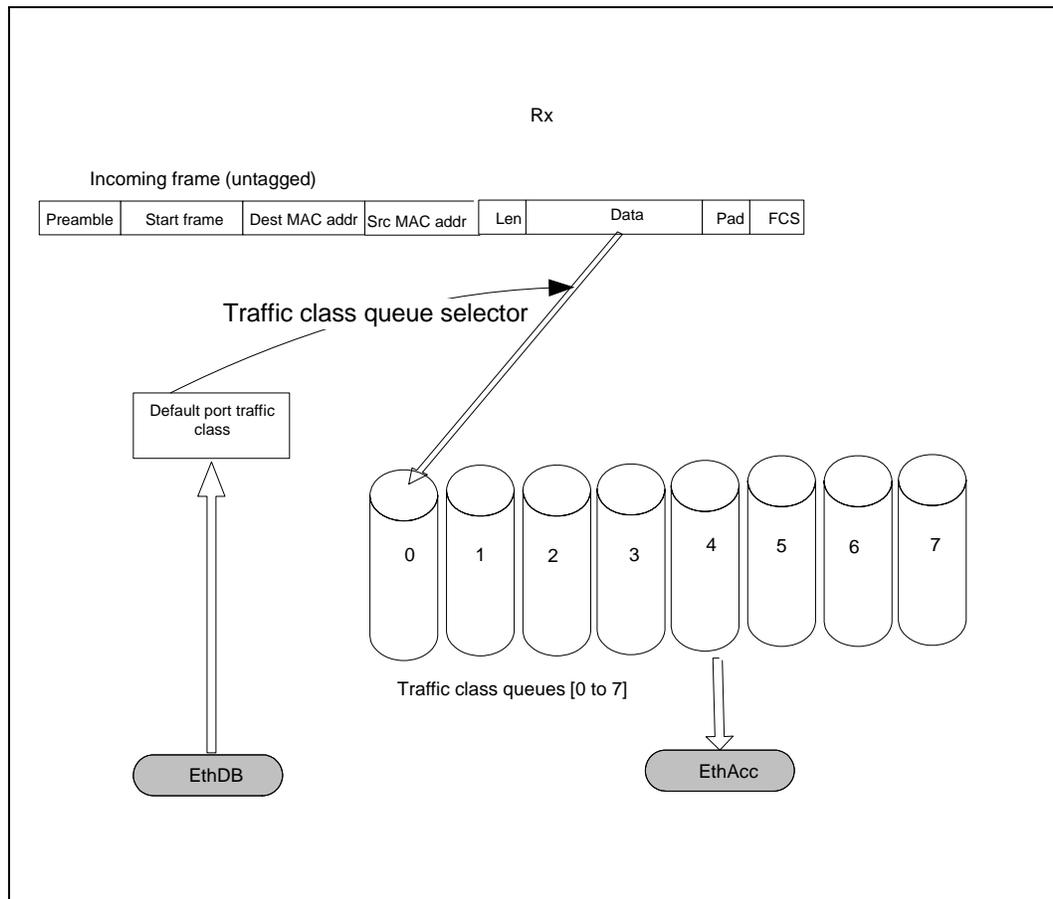
Figure 57. QoS on Receive for 802.1Q Tagged Frames



Traffic class for untagged frames (unexpedited traffic) is automatically selected from the default traffic class associated with the port. The default port traffic class is computed from the default port 802.1Q tagging information, configured as described in [“Ingress Tagging and Tag Removal” on page 183](#). The first three bits from the default 802.1Q tag constitute the default port user priority, which is mapped using the priority mapping table to obtain the default port traffic class.



**Figure 58. QoS on Receive for Untagged Frames**



*Note:* In order to use Receive QoS processing, IxEthAcc must be configured to operate in Receive FIFO Priority Mode. Refer to [Section 9.6.3.2](#).

### 10.4.8.3 Priority to Traffic Class Mapping

In order to associate the mapping of a frame's 802.1Q priority value to the receive traffic class, the IxEthDB API maintains a Priority Mapping Table. Functions are provided to modify individual priority mapping entries, or to define a completely new table definition.

At initialization, a default traffic class mapping is provided, as shown [Table 47](#). These values apply to NPE images that include eight default traffic classes. When using NPE images that provide a larger number of priority queues, the values may differ.

**Table 47. Default Priority to Traffic Class Mapping**

VLAN TCI Priority Field	Internal Traffic Class
0	3
1	1
2	2
3	4



**Table 47. Default Priority to Traffic Class Mapping**

VLAN TCI Priority Field	Internal Traffic Class
4	5
5	6
6	7
7	8

Some NPE images will not provide the eight IxQMgr queues that would allow the priority to traffic class mapping mentioned above. A header file is provided (`/src/include/IxEthDBQoS.h`) that defines the number of queues available for QoS processing in various NPE images, and provides the traffic class mapping default values. The above header file also provides static constant structures that are indexed to determine the queue assignments corresponding to the number of traffic classes.

### 10.4.9 802.3 / 802.11 Frame Conversion

The NPEs are capable of converting between IEEE 802.3 Ethernet and IEEE 802.11 wireless frame formats. IxEthDB provides support for configuring these NPE capabilities. Specific NPE microcode images are required to enable 802.3/802.11 conversion, and this feature is mutually exclusive with the MAC Address Filtering feature. Each NPE supporting this feature can have a unique 802.3 / 802.11 conversion configuration.

Specific to frames converted from 802.3 to 802.11 format, a destination MAC specific VLAN tagging service is provided.

Note that the Wi-Fi header conversion and learning / filtering are mutually exclusive. They cannot co-exist.

#### 10.4.9.1 Background — 802.3 and 802.11 Frame Formats

The 802.3 frame format is shown in [Figure 57](#) and [Figure 58](#). The 802.11 frame format is shown in [Table 48](#).

**Table 48. IEEE 802.11 Frame Format**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29		
FC	DID	Address1			Address2			Address3			SC	Address4			Frame Body (0–2312 bytes)				FCS												



**Table 49. IEEE802.11 Frame Control (FC) Field Format**

15	14	13	12	11	10	9	8	7	6	5	6	3	2	1	0
subtype				type		protocol version		order	WEP	more data	pwr mgr	retry	more flag	from DS	to DS

**Abbreviations:**

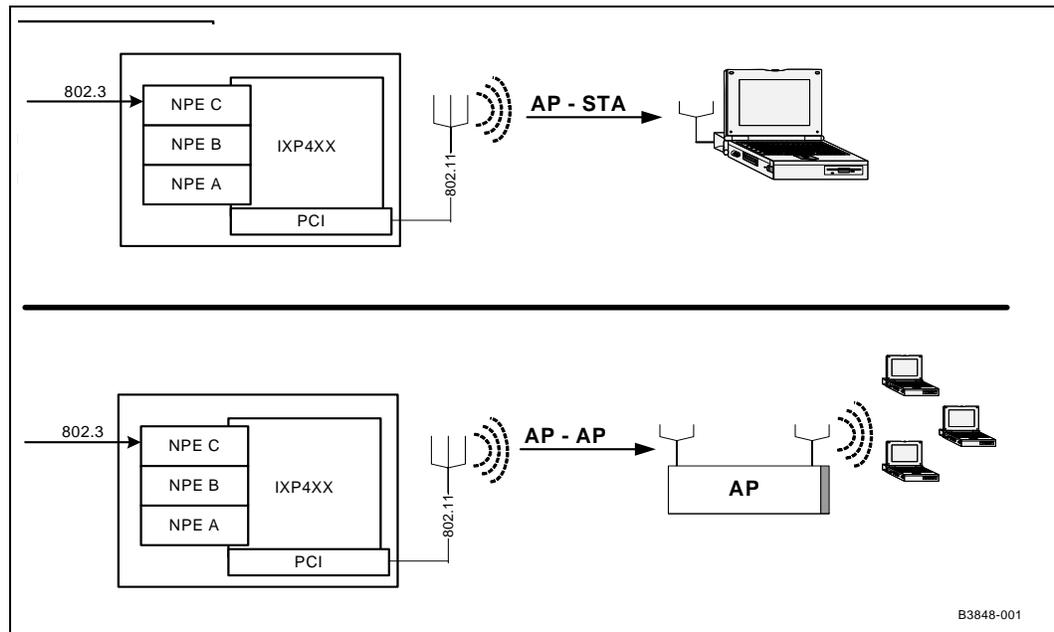
- **FC** - Frame Control
- **DID** - Duration / ID
- **SC** - Sequence Control

The usage of the 802.11 frame format depends heavily on the source and immediate destination for the frame. There are four distinct possibilities:

- From STA (station) to STA.
- From STA to AP (access point).
- From AP to STA.
- From AP to AP.

The APIs in the software release 2.3 focus on the two latter scenarios (AP → STA, and AP → AP).

**Figure 59. AP-STA and AP-AP Modes**



Conceptually, the idea of the platform running software release 2.3 to operate as a “Station” and also take advantage of the 802.3 / 802.11 Frame Conversion feature has limited applicability. This scenario would entail the platform sending or receiving 802.11 formatted frames via the Ethernet NPEs. Therefore the STA → STA and STA → AP modes are not discussed.



Table 50. STA Frame Format

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32
FC		DID		Address1				Address2				Address3				SC	Pad	LLC		OUI		Type										

The STA format has only three addresses whereas the AP format has four addresses, for example, address 4 is absent in the STA case. For more information on the STA format and protocols, refer to Table 51, “802.3 to 802.11 Header Conversion Rules” on page 194. And for specification details, refer to:

<http://standards.ieee.org/getieee802/download/802.11-1999.pdf>

For AP to STA mode, the following are the 802.11 field contents for Address1, Address2, Address3 (as defined by Table 51, “802.3 to 802.11 Header Conversion Rules” on page 194). The field Address 1 contains 802.3 destination MAC address. The field Address 2 contains value set by either *ixEthDBWiFiBSSIDSet()* or *ixEthDBWiFiRecordEntryAdd()*. The field Address 3 contains 802.3 source MAC address. Address 4 is absent in STA mode.

In 802.3 frames, there is a 2-byte Length/Type field, the interpretation of which depends on whether its value is smaller than 0x0600. When the value of this field is less than 0x0600, it is interpreted as Length, and the first 8 bytes of the MAC client data field is always the LLC/SNAP header, as defined in 802.2. Such frames are also known as “8802 frames”. When the value of the Length/Type field is greater than or equal to 0x600, it is interpreted as Type, and there is no LLC/SNAP header in the frame. Such frames are also known as “Ethernet frames”. Typically, IP packets are conveyed via Ethernet frames.

In 802.11 frames, there is always a LLC/SNAP header. This LLC/SNAP header always occupies the first 8 bytes of the Frame Body field (see Table 48). In addition to its dependence on the source and destination types, the process of converting from 802.3 frame headers to 802.11 frame headers also involves the complexity of LLC/SNAP sub-layer conversion. The appropriate conversion is handled by the NPE automatically.

### 10.4.9.2 Destination Port ID Indication for a Forwarding Frame

Destination Port ID support off loads XScale’s burden from table look up work to find the destination Port ID associated with the destination MAC address matched (and in the Wi-Fi Header Conversion Database). With Destination Port ID support, application s/w can indicate to which destination portID the received packet should be sent. It can do so by passing the Destination portID to TO\_STA and TO\_AP entries while calling *ixEthDBWiFiRecordEntryAdd()*.

Destination Port ID works with the ‘destination MAC address’ as a pair, similar to the pair of port ID and IP address in a router. NPE will copy the Destination Port ID value in the header conversion database into the ‘ix\_ne\_dest\_port’ in IXP\_NE header on the receive side. Thus the Destination PortID is reported back to the application upon receiving 802.11 frames.

In summary, the NPE will provide an indication of physical port ID (if the forwarded frame is 802.3 frame) or logical port ID (if the forwarded frame is 802.11 frame) in the IXP\_NE’s ‘ix\_ne\_dest\_port’ field. Note that this is only for *Wi-Fi Header conversion* support.

The permitted range of values for Destination PortID is 0-39.

Note that in TO\_ETH case, the physical port ID is copied to destination portID field of IX\_BUF but this port ID and logical port ID are different. In TO\_LOCAL case, no port ID is copied. Instead, the destination port ID of IX\_BUF is set to default value of 0xFF.



### 10.4.9.3 VLAN Support for 802.11 Frames

Specific to frames converted from 802.3 to 802.11 format, a destination MAC specific VLAN tagging service is provided. User can program this functionality using *vlanTagFlag* in **ixEthDBWiFiRecData** structure when calling *ixEthDBWiFiRecordEntryAdd()*. This flag indicates whether the frame should be VLAN Tagged if the frame is of type AP to STA or AP to AP. Thus the EthDB component provides an option to enable and disable the VLAN tagging 802.11 frames both at receive and transmit side based on this flag value.

VLAN tagging support for 802.11 frames is described next.

### 10.4.9.4 How the 802.3 / 802.11 Frame Conversion Feature Works

IXEthDb maintains two information structures for use in the 802.3/802.11 Frame Conversion feature:

- Wi-Fi Header Conversion Database. This database contains all of the per-MAC address information needed to perform the 802.3 to 802.11 header conversion that cannot be derived directly from the content of the IEEE802.3 header.
- In addition to the information contained in the Wi-Fi Header Conversion Database (*The Header Conversion Table*, *AP MAC Address Table* and *BSSID Table*), the NPE also needs several pieces of global information to perform header conversion. This information includes the global Frame Control and Duration ID for all frames that is converted. These two elements are referred to as the *802.11 Host Station Parameters*.

The above information is downloaded to each NPE performing 802.3/802.11 conversion via the *ixEthDBWiFiConversionTableDownload()* API, and is stored in an *NPE 802.3/802.11 Conversion Table*.

#### Receive Path

For every received 802.3 frame, once it passes all other checking, classification and validation, the NPE microcode will check the frame to see if the frame must be converted to the IEEE802.11 frame format. The NPE does this by comparing the destination MAC address against MAC addresses of the ultimate destination in the *NPE 802.3/802.11 Conversion Table*. If no match is found in the table, the frame is delivered to the client without conversion.

If a match is found, the NPE microcode inspects the matched table entry to determine whether the frame is “from AP to STA” or “from AP to AP” and then takes action accordingly. Note that the matched entry on the table is also checked to determine if VLAN tag is set. If this condition is true, then the existing 802.3 header is removed and a new 802.11 header is created using the rules and information listed in [Table 51 on page 194](#) (refer to [Table 48 on page 190](#) for locations of the 802.11 fields). The header will include a VLAN tag (described in [Table 44 on page 182](#) and [Table 45 on page 182](#)) located inside the “Frame Body” described in [Table 48 on page 190](#). NPE Ethernet firmware sets the *ixp\_ne\_flags.link\_prot* field in the buffer header to indicate the format of the converted frame header. If the VLAN flag is not set in the matched entry of the header conversion table then the 802.3 header is converted to 802.11 header in the same way but with the VLAN tag removed.

It is important to note that the IXP\_NEs extracted from the EthRxFree queue by the NPE may be used to deliver both IEEE802.3 and IEEE802.11 frames to the client software. The NPE microcode does not make any adjustment to the *ixp\_ne\_data* field from the IXP\_NE header before writing out the received frame, regardless of the header conversion operation performed.



Table 51. 802.3 to 802.11 Header Conversion Rules

802.11 Field	AP to STA mode	AP to AP mode
Frame Control	value set by ixEthDBWiFiFrameControlSet() (to DS=0)	value set by ixEthDBWiFiFrameControlSet() (to DS=1)
Duration / ID	value set by ixEthDBWiFiDurationIDSet()	value set by ixEthDBWiFiDurationIDSet()
Address 1	802.3 destination MAC address	gateway AP MAC address (from database)
Address 2	value set by ixEthDBWiFiBSSIDSet(). <b>NOTE:</b> Address 2 is also set by the new API ixEthDBWiFiRecordEntryAdd().	value set by ixEthDBWiFiBSSIDSet() (as transmitter MAC, TA)
Address 3	802.3 source MAC address	802.3 destination MAC address
Sequence Control	undefined <sup>1</sup>	undefined <sup>1</sup>
Address 4	absent <sup>2</sup>	802.3 source MAC address
LLC / SNAP	The conversion in this layer is dependant upon 802.3 – Ethernet, 8802, or 802.11 frame characteristics. The NPE handles this conversion appropriately.	

**Notes:**

1. Because the Sequence Control field is overwritten by the IEEE802.11 MAC/PHY, the NPE microcode does not attempt to set it to any particular value. Its value is undefined when returned to the client.
2. If the frame is of the type "from AP to STA", the Address4 field is not present, for example, the IEEE802.11 frame header is reduced to only 24 bytes total.

### Transmit Path

The NPE microcode converts input IEEE802.11 frames to IEEE802.3 frames prior to transmitting them to the PHY. Conversions are performed only if necessary (for example, input IEEE802.3 frames are not converted). Furthermore, conversions only apply to the data that is actually transmitted via the MII interface; the IXP\_NEs containing frames to be transmitted are never modified (for example, the content of an IXP\_NE is not altered between the time it is extracted from the EthTx queue and the time it is inserted into the EthTxDone queue). There is no table or global configuration variable associated with this service. All the information needed to perform 802.11 to 802.3 header conversion is contained within the submitted 802.11 frames and their associated IXP\_NE headers.

The NPE examines determines whether 802.11 header to 802.3 header conversion is required for each submitted frame by examining the *ixp\_ne\_flags.link\_prot* field of the IXP\_NE header associated with the frame.

If the NPE determines that no header conversion is required, it bypasses this service and continues with other transmit path processing. If the NPE determines that header conversion is requested, it performs the header conversion prior to performing additional transmit path processing (such as the VLAN-related processing). The NPE removes the 802.11 header, inserts an untagged 802.3 header, and conditionally removes the *LLC/SNAP* header as appropriate. The fields of the 802.3 header are filled according to the rules in [Table 52 on page 195](#).

The VLAN egress services (VLAN egress filtering and VLAN egress ID-based tagging/untagging) are available in any VLAN-enabled NPE firmware load. These services must be explicitly enabled on a per frame basis via the *ixp\_ne\_flags.vlan\_en* flag bit in the *ixp\_ne* header.

The NPE will set the VLAN tag in the Frame Body (see [Table 48 on page 190](#)) according to the format described in [Table 44 on page 182](#) and [Table 45 on page 182](#) if *ixp\_ne\_flags.vlan\_en* flag in the *ixp\_ne* header of the frame is set to 1.



Table 52. 802.11 to 802.3 Header Conversion Rules

Input 802.11 Frame Values				Output 802.3 Frame Field Values	
ixp_ne_flags.link_prot	From DS <sup>1</sup>	Frame Type	Header Size (bytes)	Destination Address	Source Address
10	0	From STA to AP	24	802.11 Address 3	802.11 Address 2
11	1	From AP to AP	39	802.11 Address 3	802.11 Address 4
<b>Note:</b> 1. The NPE does not actually inspect the <i>From DS</i> field to determine the 802.11 frame type. It relies exclusively on the value of the <i>ixp_ne_flags.link_prot</i> field.					

#### 10.4.9.5 Pad Field Addition/Removal for 802.11 Frames

Padding insertion and removal between 802.11 header and LLC/SNAP is supported. The `ixEthDBWiFiRecordEntryAdd()` API supports a user-configurable parameter, `padLength`, to specify the number of bytes to be padded. This is specified through `ixEthDBWiFiRecData` structure and it must be an even number. The default pad length is '0'. Its value must be minimum 0 bytes and maximum 16 bytes.

##### Padding insertion

Padding [2, 4, 6, ..., 16-byte, for example, only even numbers] may be added between 802.11 header and SNAP/LLC header during 802.3 to 802.11 header conversion. NPE will insert the specified number of zero padded bytes into the frame header. The number of bytes padded is passed to access layer via `ixp_ne_pad_len` field in `IXP_BUF`.

Odd-number pad length size is rejected.

##### Padding removal

During 802.11 to 802.3 header conversion, padding [2, 4, 6, ..., 16-byte, for example, only even numbers] may exist between 802.11 header and SNAP/LLC header. Based on the `ixp_ne_pad_len` value in `IXP_BUF` passed to NPE, the NPE will strip off these paddings during the header conversion.

If the removal pad length is an odd number, the frame is discarded.

#### 10.4.9.6 802.3 <-> 802.11 API Details

As mentioned previously, the `IxEthDB` component maintains a *Wi-Fi Header Conversion Database* to store MAC address/port entries and their respective 802.3/802.11 transformation mode. There are three functions used to add these entries:

- **`ixEthDBWiFiStationEntryAdd()`** – this function takes as parameters a port ID and the MAC address of a wireless station. This function should be used for AP-STA scenarios. Up to 511 station entries are supported per port.
- **`ixEthDBWiFiAccessPointEntryAdd()`** – this functions takes port ID, MAC address of a wireless station and MAC address of the gateway Access Point as parameters. Up to 40 entries of this type may be defined per port.
- **`ixEthDBWiFiBSSIDSet()`** – this function takes as parameters the port ID and BSSID (Basic Service Set ID).

IXP400 software defines a user Wi-Fi record structure called `IxEthDBWiFiRecData` and an API named `ixEthDBWiFiRecordEntryAdd()`. `IxEthDBWiFiRecData` is a versatile record structure with the following elements: a) record type to indicate AP, STA, ETHER or



Local type, b) Tag/untag flag for 802.11 frames, c) Destination Port ID (value can range of 0 to 39), d) size of pad fields in bytes (maximum value 16 and minimum 0), e) Peer AP Gateway address (40 entries), and f) BSSID (40 entries).

- **ixEthDBWiFiRecordEntryAdd** – This function takes as parameters a port ID, the MAC address to add and pointer to `IxEthDBWiFiRecData`. With the versatility of `IxEthDBWiFiRecData`, this function can be used to achieve many functionalities. For instance, this function can be used to a) set tagging parameters (`toEthFlag`, `localMacFlag`, `toStaFlag` or `toApFlag`), b) set `vlanTagFlag` to enable 802.1q tagging for the incoming 802.11 frames, c) set `padLength` parameter d) set `logicalPortID` (`logicalPortID` is also referred as Destination PortID).

In addition, `ixEthDBWiFiRecordEntryAdd` can be used to achieve the functionality of all the three functions (`ixEthDBWiFiStationEntryAdd()`, `ixEthDBWiFiAccessPointEntryAdd()`, `ixEthDBWiFiBSSIDSet()`). Note that to address a need to assign different BSSIDs to different end users, BSSID has been extended from 1 to 40 and user can specify BSSIDs ranging 0 to 39.

*Note:*

MAC addresses are unique database keys only within the configuration data of each port. Multiple ports can use the same MAC address entry if individually added to each port.

Additionally, two functions are provided that set the per port *802.11 Host Station Parameters*, namely Frame Control and Duration/ID fields in the 802.11 frame format.

The *NPE 802.3/802.11 Conversion Tables* are derived from the *WiFi Header Conversion Database* and must be downloaded to each NPE separately, using the `ixEthDBWiFiConversionTableDownload()` function.

The 802.3/802.11 Frame Conversion feature introduces specific requirements on when FCS Frame Appending should be enabled. Refer to “[FCS Appending](#)” on page 200.

A typical usage scenario of this feature would consist in the following steps:

1. Enable the `IX_ETH_DB_WIFI_HEADER_CONVERSION` feature.
2. Add BSSID (`ixEthDBWiFiBSSIDSet()` or the new API `ixEthDBWiFiRecordEntryAdd()`).
3. Add wireless station (`ixEthDBWiFiStationEntryAdd()` or the new API `ixEthDBWiFiRecordEntryAdd()`).
4. Add access point/gateway address (`ixEthDBWiFiAccessPointEntryAdd()` or the new API `ixEthDBWiFiRecordEntryAdd()`).

As mentioned previously, steps 2, 3, and 4 can be done by `ixEthDBWiFiRecordEntryAdd()`.

5. Set the 802.11 Host Station Parameters (Frame Control, Duration ID) using `ixEthDBWiFiFrameControlSet()` and `ixEthDBWiFiDurationIDSet()`.
6. Set tagging parameters to `EthFlag`, `localMacFlag`, `toStaFlag` or `toApFlag` to distinguish whether the incoming frame is destined to Ethernet, Local Device, Station point, or Access point interfaces.
7. Set `vlanTag` Flag to enable or disable 802.1q tagging for the incoming 802.11 frames.
8. Set `padLength` parameter in bytes. This must be an even number and its value must be minimum 0 and maximum 16 bytes.
9. Set `logicalPortID` for `TO_STA` and `TO_AP` entries. The range is between 0–39.
10. Download the Wi-Fi conversion configuration data using `ixEthDBWiFiConversionTableDownload(port)`.



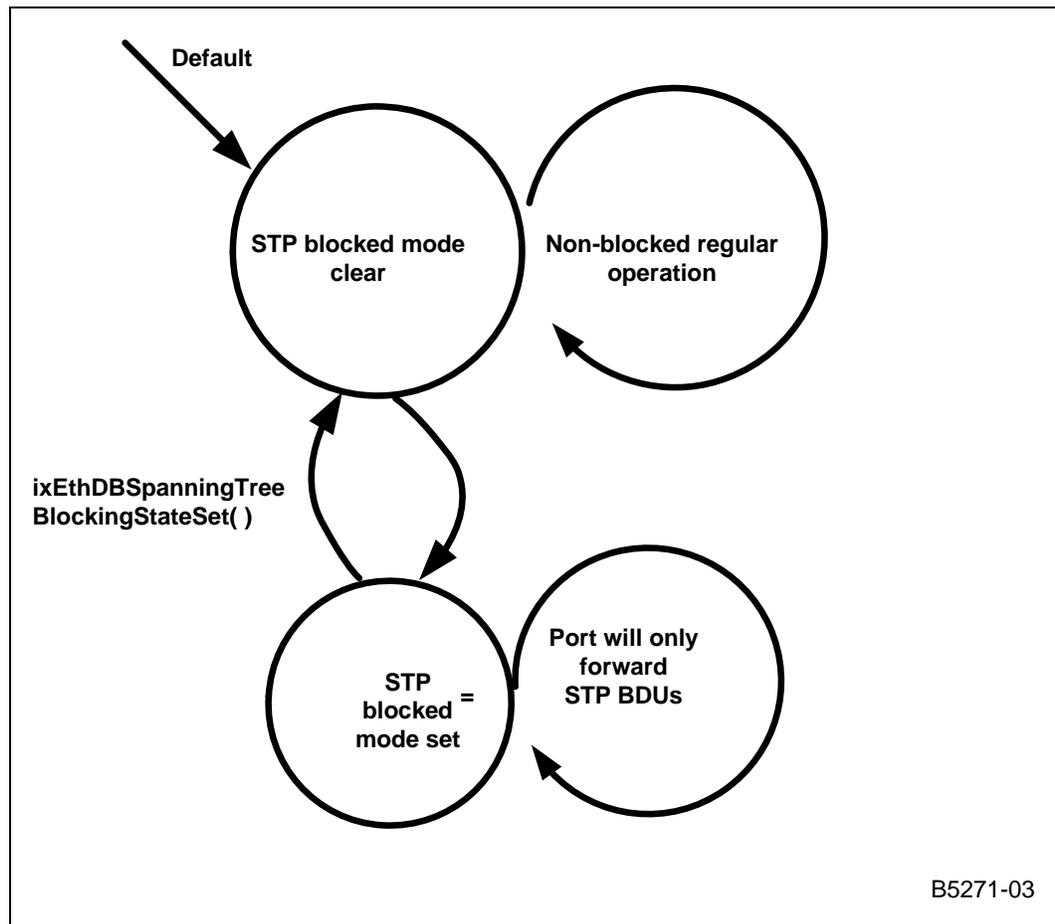
### 10.4.10 Spanning Tree Protocol Port Settings

The IxEthDB component provides an interface that can configure each NPE port to act in a “Spanning Tree Port Blocking State”. This behavior is available in certain NPE microcode images, and can be configured independently for each NPE.

Spanning-Tree Protocol (STP), defined in the IEEE 802.1D specification, is a link management protocol that provides path redundancy while preventing undesirable loops in the network. STP includes two special frame payload types that bridges use to help close loops in an Ethernet network. These frames are called a *configuration Bridge Protocol Data Unit (BPDU)* and a *topology change notification BPDU*.

The NPE tests every received frame to determine whether it is a configuration or topology change BPDU. Spanning tree BPDUs are delivered to the Intel XScale® Processor in the same manner as regular Ethernet frames, but the NPE firmware sets the *ixp\_ne\_flags.st\_prot* bit flag in the IX\_OSAL\_MBUF whenever the frame in the associated buffer is a spanning tree BPDU. Spanning tree BPDU frames are never subjected to any VLAN or 802.3 to 802.11 header conversion service.

Figure 60. Passing Only the Two Types of BPDUs



When IxEthDB configures a port to operate in an STP blocking state, using `ixEthDBSpanningTreeBlockingStateSet()`, the effect is that all frames EXCEPT STP configuration BPDUs and topology change BPDUs are dropped. A statistic counter is maintained to track the number of frames dropped while in this state.



## 10.4.11 Soft-error Handling

By introducing soft-error handling in Ethernet (refer to [Chapter 17.0, “Access-Layer Components: Error Handler \(ixErrHdlAcc\) API”](#)), new APIs calls are introduced. The objective of these new functions is to allow soft-error handling module to restore the EthDB functionality features (VLAN, Firewall, Header conversion and QoS) to the states before the occurrence of soft-error.

It is critical to pause Ethernet Database event processor during soft-error handling, the reason is to avoid the processor from flooding NPE message FIFOs. New API **ixEthDBEventProcessorPauseModeSet (BOOL pauseEnable)** provides the functionality to pause/resume Ethernet Database event processor at run-time without calling `ixOsalTreadKill()` and `ixOsalTreadCreate()`. Ethernet Database event processor can be resumed after the soft-error handling is recovered and database restored.

In the event of NPE soft-error, the state of traffic class and Rx queue assignments in NPE is lost when the impacted NPE is reset. In order to restore this state, **ixEthDBPriorityMappingTableUpdate(IxEthDBPortId portId)** is used during NPE soft-error recovery process, this is to restore the traffic class-priority queue mapping/configuration. It is also necessary to restore EthDB Basic settings and feature-specific database and configuration, example VLAN, Firewall, Header Conversion and STP, this can be done by calling function **ixEthDBFeatureStatesRestore(IxEthDBPortId portId)**. NPE learning/filtering database will not be restored after NPE soft error. The database will be re-generated via new MAC address learning process.

## 10.5 IxEthDB

### 10.5.1 Enabling/Disabling EthDB

IxEthDB performs an `ixFeatureCtrlSwConfigurationCheck()` to determine the value of `IX_FEATURECTRL_ETH_LEARNING`. IxEthDB is essentially disabled if this value is FALSE. Any component or codelet can modify the value prior to IxEthDB initialization using `ixFeatureCtrlSwConfigurationWrite(IX_FEATURECTRL_ETH_LEARNING, [TRUE or FALSE])`. Once IxEthDB has been initialized, the software configuration cannot be changed.

`IX_FEATURECTRL_ETH_LEARNING` is TRUE by default.

### 10.5.2 Initialization

IxEthAcc is dependent upon IxEthDB and provides for most of its initialization.

Feature capability scanning and initialization is done automatically during EthDB initialization (performed by `ixEthDBInit()`), which requires the NPEs to have been downloaded and started with suitable images for the application.

For backward compatibility purposes only learning and filtering are enabled by default when available. The remaining feature set (VLAN/QoS, Firewall, 802.3 <-> 802.11 frame header conversion) must be enabled manually, if needed, which is usually done at initialization time, using `ixEthDBFeatureEnable()`.

QoS is a case apart regarding initialization, as the Ethernet Access component (EthAcc) depends on the QoS initialization sequence in EthDB for the purpose of configuring the Rx queues. It is important to note that once EthAcc initialization is completed the Rx queues cannot be reconfigured at run-time. This has the effect that if VLAN/QoS is enabled before EthAcc initialization then EthAcc will configure and benefit from using all the available Rx queues associated with internal traffic classes (**currently 8 queues**



**are available**). If VLAN/QoS is enabled after EthAcc is initialized, only one Rx queue, associated with traffic class 0, is available. QoS cannot properly operate with only one traffic class, therefore only VLAN functionality is available.

The two initialization sequences are described below:

**a). to enable QoS (multiple traffic classes present)**

```
ixEthDBInit();
ixEthDBPortInit(portID);
ixEthDBFeatureEnable(portID, IX_ETH_DB_VLAN_QOS, TRUE);
ixEthAccInit(); /*Queue configuration uses all the available Rx queues */
```

Note that by default the priority mapping is the 802.1P standard 8-traffic class mapping. The mapping can be changed using ixEthDBPriorityMappingSet().

**b) to enable only VLAN functionality, without actual QoS functionality (one traffic class present)**

```
ixEthAccInit();
ixEthDBInit(); /*redundant, as this is also called by ixEthAccInit() */
ixEthDBPortInit(portID);
ixEthDBFeatureEnable(portID, IX_ETH_DB_VLAN_QOS, TRUE);
```

During the initialization sequence, and before enabling traffic on the port, EthDB should be specifically instructed to map all the QoS user priorities to traffic class 0, using a priority mapping table with all the values set to 0.

```
IxEthDBPriorityMap nullPriorityMap;
memset (nullPriorityMap, 0, sizeof (nullPriorityMap));
/* QoS priorities 0...7 are mapped to traffic class 0 */
ixEthDBPriorityMapSet(portID, nullPriorityMap);
```

This will ensure that the NPE will write all the incoming traffic into the only available traffic class queue (0).

### 10.5.3 Dependencies

The IxEthDB component relies on the following components:

- IxNpeMh component to send/receive control messages to/from the NPEs.
- IxNpeDI is used by IxEthDB to query the loaded NPE image IDs.
- IxOSAL to provide mutual exclusion mechanisms to the component.
- IxOSAL to provide multithreading.

### 10.5.4 Dependencies on IxEthAcc Configuration

One of the functions of IxEthAcc is to configure the MAC sub-component of each NPE. In order for many of the features provided in IxEthDB to work properly, the MAC must be configured appropriately.



#### 10.5.4.1 Promiscuous-Mode Requirement

Ethernet Filtering is operational only when a port is configured to operate in *promiscuous* mode. Otherwise the frames are filtered according to normal MAC filtering rules. Those filtering rules are that the frame is received only if one of the following is true:

- The destination address matches the port address
- The destination address is the broadcast address or if the destination is a multicast address subscribed to by the port
- The frame is a broadcast/multicast frame.

Configuration of promiscuous mode is described in the section for IxEthAcc, “MAC Filtering” on page 161.

#### 10.5.4.2 FCS Appending

Several NPE features controlled by IxEthDB cause changes to the frame data such that a previously calculated Frame Check Sequence is invalid. IxEthAcc provides a set of functions, (`ixEthAccPortRxFrameAppendFCSDisable()`, `ixEthAccPortTxFrameAppendFCSEnable()`) that can instruct the NPE to remove the FCS on received Ethernet frames, or calculate and append the FCS on frames prior to transmission. It is the responsibility of the client application to configure the FCS settings for each port properly.

##### Receive Traffic

FCS appending should be **disabled**, or the FCS data should be ignored when a port is configured for the following features:

- VLAN Ingress tagging/untagging
- 802.3 to 802.11 Frame Conversion

##### Transmit Traffic

For transmission services, the NPE calculates a valid FCS as its final step prior to transmitting the frame to the PHY. FCS appending should be **enabled** when a port is configured for the following features:

- VLAN Egress tagging/untagging
- 802.11 to 802.3 Frame Conversion





## 11.0 Access-Layer Components: Ethernet PHY (IxEthMii) API

---

This chapter describes the Intel® IXP400 Software v2.3's "Ethernet PHY API" access-layer component.

### 11.1 What's New

The following changes or enhancements were made to this component in software release 2.3.

- This component has been updated to support the Intel® LXT9785HC 10/100 Ethernet Octal PHY that is on the Intel® IXDP465 Development Platform
- The component has been updated to support the Realtek\* Ethernet PHY RTL8305, which is on the Intel® IXDPG425 Network Gateway Development Platform.

### 11.2 Overview

*IxEthMii* is used primarily to manipulate a minimum number of necessary configuration registers on Ethernet PHYs supported on the IXDP425 / IXCDP1100 platform, IXDPG425 network gateway platform, and IXDP465 platform without the support of a third-party operating system. Codelets and software used for Intel internal validation are the consumers of this API, although it is provided as part of the software release 2.3 for public use.

### 11.3 Features

The *IxEthMii* components provide the following features:

- Scan the MDIO bus for up to 32 available PHYs
- Configure a PHY link speed, duplex, and auto-negotiate settings
- Enable or disable loopback on the PHY
- Reset the PHY
- Gather and/or display PHY status and link state

### 11.4 Supported PHYs

The supported PHYs are listed in the table below. *IxEthMii* interacts with the MII interfaces for the PHYs connected to the NPEs on the IXDP425 / IXCDP1100 platform. These functions do not support reading PHY registers of devices connected on the PCI interface. Other Ethernet PHYs are also known to use the same register definitions but are unsupported by this software release (for example, Intel® 82559 10/100 Mbps Fast Ethernet Controller).

Register definitions are located in the following path:

```
ixp400_xscale_sw/src/ethMii/IxEthMii_p.h
```



**Table 53. PHYs Supported by IxEthMii**

Intel® LXT971 Fast Ethernet Transceiver
Intel® LXT972 Fast Ethernet Transceiver
Intel® LXT973 Low-Power 10/100 Ethernet Transceiver (LXT973 and LXT973A)
Micrel / Kendin* KS8995 5 Port 10/100 Switch with PHY
Realtek* Ethernet PHY RTL8305

## 11.5 Dependencies

*IxEthMii* is used by the *EthAcc* codelet and is dependant upon the *IxEthAcc* access-layer component and *IxOSAL*.

§ §



## 12.0 Access-Layer Components: Feature Control (IxFeatureCtrl) API

---

This chapter describes the Intel® IXP400 Software v2.3's "Feature Control API" access-layer component.

*IxFeatureCtrl* is a component that detects the capabilities of the Intel® IXP4XX Product Line of Network Processors. It provides a configurable software interface that can be used to simulate different processors variants in the IXP42X product line and IXP46X product line.

### 12.1 What's New

There are no changes or enhancements to this component in software release 2.3.

### 12.2 Overview

*IxFeatureCtrl* provides APIs for the following functions:

- Read product ID and available hardware components on the processor.
- Enable and disable hardware components on the processor
- Enable and disable miscellaneous software capabilities: The Ethernet Learning feature and Livelock Prevention feature in Queue Manager.

### 12.3 Hardware Feature Control

Detecting and controlling the hardware features of the processor is performed using several registers on the host processor. The registers include:

- CP15, Register 0, ID Register - This read-only register contains product identification data as shown in

**Table 54. Product ID Values (Sheet 1 of 2)**

Bits	Description
31:28	Reserved. Value: 0x6
27:24	Reserved. Value: 0x9
23:20	Reserved. Value: 0x0
19:16	Reserved. Value: 0x5
15:12	Reserved. Value: 0x4



**Table 54. Product ID Values (Sheet 2 of 2)**

Bits	Description
11:9	Device ID. IXP42X - 0x0 IXP46X - 0x1
8:4	Maximum Achievable Intel XScale® Processor Frequency for the IXP42X product line <b>only</b> . 533 MHz — 0x1C 400 MHz — 0x1D 266 Mhz — 0x1F <b>For the IXP46X product line, the value is 0x00.</b>
3:0	Si Stepping ID. A-step — 0x0 B-step — 0x1

- EXP\_UNIT\_FUSE\_RESET register in the Expansion Bus Controller - The value of this register is saved in a software register called Feature Control Register. The fields of this register are shown in Table 55.

**Table 55. Feature Control Register Values (Sheet 1 of 2)**

Bits	Description
31:24	(Reserved)
23:22	Processor frequency (IXP46X product line <b>only</b> ): 0x0 - 533 MHz 0x1 - 667 MHz 0x2 - 400 MHz 0x3 - 266 MHz
21 †	RSA Crypto Block coprocessor (IXP46X product line <b>only</b> )
20 †	NPE B Ethernet coprocessor 1-3 (IXP46X product line <b>only</b> )
19	IXP46X product line <b>only</b> 0 = NPE A Ethernet is enabled if Utopia bit is 1. 1 = NPE A Ethernet is disabled.
18 †	USB Host Coprocessor (IXP46X product line <b>only</b> )
17:16	UTOPIA PHY Limits. 32 PHYs: 0x0 16 PHYs: 0x1 8 PHYs: 0x2 4 PHYs: 0x3
15 †	ECC and 1588 Unit (IXP46X product line <b>only</b> )
14 †	PCI Controller
13 †	NPE C
12 †	NPE B
11 †	NPE A
10 †	Ethernet 1 Coprocessor (on NPE C)
9 †	Ethernet 0 Coprocessor (on NPE B)
8 †	UTOPIA Coprocessor
7 †	HSS Coprocessor

† For bit 0 through 15, 18, 20-21 the following values apply:

- 0x0 — The hardware component exists and is not software disabled.
- 0x1 — The hardware component does not exist, or has been software disabled.



Table 55. Feature Control Register Values (Sheet 2 of 2)

Bits	Description
6 †	AAL Coprocessor
5 †	HDLC Coprocessor
4 †	DES Coprocessor
3 †	AES Coprocessor
2 †	Hashing Coprocessor
1 †	USB Coprocessor
0 †	RComp Circuitry

† For bit 0 through 15, 18, 20-21 the following values apply:

- 0x0 — The hardware component exists and is not software disabled.
- 0x1 — The hardware component does not exist, or has been software disabled.

### 12.3.1 Using the Product ID-Related Functions

There are two functions to read the Product ID related information. They are:

- *ixFeatureCtrlProductIdRead()* returns the entire 32-bit value of the CP15, Register 0.
- *ixFeatureCtrlDeviceRead()* only returns an indication of the processor product line (the IXP46X product line or IXP42X product line).

*Note:* The only way to detect the core frequency on the IXP46X product line is to use the *ixFeatureCtrlHwCapabilityRead()*

### 12.3.2 Using the Feature Control Register Functions

There are four functions used to read, write the Software Feature Control Register and to check the availability of the hardware component on the system. See the following:

- *ixFeatureCtrlHwCapabilityRead()* function utilizes the EXP\_UNIT\_FUSE\_RESET register for detecting hardware components available on the system.
- *ixFeatureCtrlWrite()* function writes the contents of the Software Feature Control Register. This function can be used to disable a hardware component in software. The *IxFeatureCtrl* component does not actually write values to the EXP\_UNIT\_FUSE\_RESET register.
- *ixFeatureCtrlRead()* function returns the contents of the Software Feature Control Register.
- *ixFeatureCtrlComponentCheck()* checks for the availability of the specified hardware component. The other Access-Layer components in software release 2.3 use this function during their initialization routines to determine whether the required hardware components are available. Also, the *IxNpeDI* API uses the function to prevent the erroneous download of NPE microcode to disabled or unavailable NPEs.

## 12.4 Software Configuration

The provided software configuration structure and supporting functions can be modified at run-time. The software configuration structure is an array that stores the enable/disable state of particular global options. Other software components can be designed to read or write the software configuration array to enable or disable certain software features prior to initialization.



In software release 2.3, there are two entries in the software configuration array; IX\_FEATURECTRL\_ETH\_LEARNING and IX\_FEATURECTRL\_ORIGB0\_DISPATCHER.

### IX\_FEATURECTRL\_ETH\_LEARNING

*IxEthDb* performs an *ixFeatureCtrlSwConfigurationCheck()* to determine the value of IX\_FEATURECTRL\_ETH\_LEARNING. *IxEthDb* uses this value to decide whether or not to activate the NPE-based EthDB learning, and to spawn an Intel XScale® Processor thread to monitor it. Any component or codelet can modify the value prior to *IxEthDb* initialization using

*ixFeatureCtrlSwConfigurationWrite*(IX\_FEATURECTRL\_ETH\_LEARNING, [TRUE or FALSE]). Once *IxEthDB* has been initialized, the software configuration cannot be changed.

### IX\_FEATURECTRL\_ORIGB0\_DISPATCHER

*IxQMgr* performs a *ixFeatureCtrlSwConfigurationCheck* (IX\_FEATURECTRL\_ORIGB0\_DISPATCHER) to determine if the livelock prevention feature is required. Prior to start of the dispatcher, application users employ *ixQMgrDispatcherLoopGet()* to get the correct queue dispatcher. This feature is configured as TRUE by default, meaning that B0 versions of the IXP42X product line processors, and all versions of the IXP46X product line will use the standard *ixQMgrDispatcherLoopRunB0* dispatcher. To indicate that the *ixQMgrDispatcherLoopRunBOLLP* dispatcher with Livelock support is desired, use the *ixFeatureCtrlSwConfigurationWrite()* function to set this option to FALSE.

## 12.5 Dependencies

This component uses *IxOSAL* for memory mapping, reads, writes, and logging functions.





## 13.0 Access-Layer Components: HSS-Access (IxHssAcc) API

---

This chapter describes the Intel® IXP400 Software v2.3's "HSS-Access API" access-layer component.

### 13.1 What's New

The co-existence of HSS channelized services plus Ethernet services and the co-existence of HSS channelized services plus ATM services are available in IXP400 software.

### 13.2 Overview

The IxHssAcc component provides client applications with driver-level access to the High-Speed Serial (HSS) and High-Level Data Link Control (HDLC) coprocessors available on NPE A. This API and its supporting NPE-based hardware acceleration enable the Clarkspoint network processors to support packetized or channelized TDM data communications.

This chapter provides the details of how to use IxHssAcc to:

- Initialize and configure the HSS and HDLC coprocessors.
- Allocate buffers for transmitting and receiving data.
- Connect and enable packetized service and/or channelized service.
- Handle the transmitting and receiving process.
- Disconnect and disable the services.

#### Features

The HSS access component is used by a client application to configure both the HSS and HDLC coprocessors and to obtain services from the coprocessors. It provides:

- Access to the two HSS ports on the Intel® IXP4XX product line processors.
- Configuration of the HSS and HDLC coprocessors residing on NPE A.
- Support for TDM signals up to a rate of 8.192 Mbps (Quad E1/T1, byte interleave mode) on an HSS port.

#### Channelized Service

- Supports a single Channelized client for all logical T1/E1 trunks in one HSS port. A Channelized client:
  - Supports up to 32 channels, where each channel is composed of one timeslot. If there are more than one logical T1/E1 trunks then timeslots in excess of 32 channels should be unassigned or assigned to HDLC service.
  - Each channel is independently configurable for 56-Kbps or 64-Kbps mode  
For 56-Kbps mode:



- Configurable CAS bit position - least significant or most significant bit position. Configurable on a per-port basis only.
- Configurable CAS bit polarity for transmitted data
- Channels can be configured to Bypass mode specifying the source and destination timeslots to be switched.
- Support up to 2 pairs of bypassed channels on HSS port (port 0 only) dynamically.

### Packetized Service

- Support a single Packetized client (termination point) per T1/E1 trunk, up to maximum of four per HSS port. For each Packetized client:
  - One or more (max 32 timeslots for E1 and 24 timeslots for T1) contiguous or non-contiguous timeslots per T1/E1 trunk can be configured for the RAW or HDLC mode to carry the payload.
  - Configurable bit inversion - all data inverted immediately upon reception from and transmission to the trunk
  - Configurable for 56 Kbps or 64 Kbps mode. The maximum HDLC packet size for transmission is 64 Kbytes. The maximum recommended size of received HDLC packets is 16 Kbytes.
  - For 56-Kbps mode:
    - Configurable CAS bit position - least significant or most significant bit position
    - CAS bit always discarded for data received from trunk
    - CAS bit insertion for data transmitted to trunk
    - Configurable CAS bit polarity for transmitted data

## 13.2.1 Coexistence of HSS Services in NPE-A

Coexistence of HSS Packetized (4 HDLC packet pipes) and HSS Channelized services (32 voice channels and 2 pairs of bypass channels) is supported in NPE A. Refer to [Figure 60, “NPE-A Images” on page 243](#) for the supported different combination of NPE images on NPE-A. The processing of Packetized HDLC and Channelized voice services are handled by entirely different sets of NPE A contexts, and have individual software and hardware FIFOs and separate scheduling events and AQM queues. The segregation of the HDLC and Voice data are performed within the HSS Coprocessor. The timeslot that is marked as **voice** in the timeslot map during init time would have its data received by the HSS Coprocessor and put into the **hardware Voice FIFO** for channelized context processing. The timeslots that are marked as **“HDLC”** timeslot would have its data put into the HSS Coprocessor **“HDLC”** hardware FIFO. If the HDLC timeslot is in HDLC mode, it is sent to the HDLC coprocessor, and then passed on to Packetized services for processing. But if the HDLC timeslot is in raw mode, then it will bypass the HDLC coprocessor, and then passed on to Packetized services for processing. The HSS Coprocessor features a hardware register to assist the NPE core in processing these HSS services.

## 13.2.2 Coexistence of HSS services and ATM services in NPE-A

Coexistence of HSS services and ATM services are also supported in NPE-A. However, it should be noted that for some images which provide ATM services coexisting with HSS Packetized and Channelized services, only the following combination of HSS configuration can be supported. Refer to [Figure 60, “NPE-A Images” on page 243](#) for more details. This is due to internal NPE data memory constraints and NPE A processing bandwidth limitations.

- HSS Channelized service with Utopia Single PHY/1 port features enabled ATM services. The details are as below:



— **For HSS services**

On HSS Port 0 only

**Supports up to 32 Voice Channels**

**Supports up to 4 HDLC Packet Pipes**

**Note:** You can either choose combination of 16 Voice Channels and 4 HDLC Packet Pipes; or 32 Voice Channels and no HDLC Packet Pipe for the coexistence of HSS services and ATM services.

— **For ATM services**

**1 Port**

32 VCs

AAL5

AAL0

OAM

**UTOPIA SPHY**

HSS service with Utopia Multi PHY/1 port features enabled ATM services. The details are as below:

— **For HSS services**

On HSS Port 0 only

**Support up to 32 Voice Channels**

**Support up to 4 HDLC Packet Pipes**

**Note:** You can either choose combination of 16 Voice Channels and 4 HDLC Packet Pipes; or 32 Voice Channels and no HDLC Packet Pipe for the coexistence of HSS services and ATM services.

— **For ATM services**

**1 Port**

32 VCs

AAL5

AAL0

OAM

**UTOPIA MPHY**

- HSS services with Utopia Multi PHY/4 ports features enabled ATM services

— **For HSS services**

On HSS Port 0 only

**16 Voice Channels only**

**2 pairs HSS Bypass channels**

**NO Packetized service is provided**

— **For ATM services**

**4 Ports**

32 VCs

AAL5

AAL0

OAM

**UTOPIA MPHY**



### 13.2.3 Coexistence of HSS Channelized services and Ethernet services in NPE-A

Coexistence of HSS Channelized services and Ethernet services is supported in NPE-A. Due to NPE resources constraints and performance impact, it is not possible to support both HSS Packetized and Channelized services while allowing Ethernet running concurrently in NPE-A. There is no special API or procedure (no impact to customer experience) to invoke this coexistence. Currently there are 2 different combinations of configurations that are supported for HSS channelized and Ethernet coexistence.

- HSS Channelized service with MAC filtering/learning feature enabled Ethernet services. The details are as below:
  - **For HSS services**
    - On HSS Port 0 only
    - Support up to 32 Voice Channels
    - Support 2 pair of HSS bypass channels
  - **For Ethernet services**
    - Support MAC filter/Port ID**
    - MAC learning assist**
    - Spanning tree support
    - Frame size filtering
    - Mask-based Firewall support
    - VLAN/QoS support
    - Extended MIB-II
- HSS channelized services with 802.11 header conversion enabled Ethernet services. The details are as below:
  - **For HSS services**
    - HSS Port 0 only
    - 32 Voice Channels
    - Support up to 2 pairs of HSS Bypass channels
  - **For Ethernet Services**
    - 802.3 <-> 802.11 header conversion**
    - Spanning tree support
    - Frame size filtering
    - Mask-based Firewall support
    - VLAN/QoS support
    - Extended MIB-II

## 13.3 IxHssAcc API Overview

The IxHssAcc API is an access layer component that provides high-speed serial and packetized or channelized data services to a client application. This section describes the overall architecture of the API. Subsequent sections describe the component parts of the API in more detail and describe usage models for packetized and channelized data.

### 13.3.1 IxHssAcc Interfaces

The *client* application code executes on the Intel XScale® Processor and utilizes the services provided by IxHssAcc. In this software release, the IxHssAccCodelet is provided as an example of client software (HSS Bypass feature is not in



IxHssAccCodelet). As previously described, the *IxHssAcc API* is the interface between the client application code and the underlying hardware services and interfaces on the processor.

IxHssAcc presents two “services” to the client application. The *Channelized Service* presents the client with raw serial data streams retrieved from the HSS port, while the *Packetized Service* provides packet payload data that has been optionally processed according to the HDLC protocol.

*IxQMgr* is another access-layer component that interfaces to the hardware-based *AHB Queue Manager (AQM)*. The AQM is SRAM memory used to store pointers to data in SDRAM memory, which is accessible by both the Intel XScale® Processor and the NPEs. These items are the mechanism by which data is transferred between IxHssAcc and the NPE. Queues are handled in a different manner depending on whether packetized or channelized data services are being utilized. The queue behavior is described in subsequent sections of this chapter.

*IxNpeMh* is used to allow the IxHssAcc API to communicate to the NPE coprocessors described below. *IxNpeDI* is the mechanism used to download and initialize the NPE microcode.

The NPE provides hardware acceleration, protocol handling, and drives the physical interface to the High-Speed Serial ports. NPE A is the specific NPE that contains an HSS coprocessor and an HDLC coprocessor utilized by this API.

### 13.3.2 Basic API Flow

An overview of the data and control flow for IxHssAcc is shown in [Figure 61](#).

The client initializes and configures HSS using IxHssAcc to configure the HSS port signalling to match the connected hardware PHY's or framers. The HSS coprocessor on NPE A drives the HSS physical interfaces and handles the sending or receiving of the serial TDM data. Data received on ports configured for channelized data is sent up the stack from the HSS coprocessor. Received Packetized data — with the HDLC option turned on — is passed to HDLC coprocessor as appropriate. The IxHssAcc API uses callback functions and data buffers provided by the client to exchange NPE-to-Intel XScale® Processor data for transmitting or receiving with the help of the IxQMgr API.





The HSS coprocessor communicates with an external device using three signals per direction: a frame pulse, clock, and data bit. The data stream consists of frames — the number of frames per second depending on the protocol. Each frame is composed of timeslots. Each timeslot consists of 8 bits (1 byte) which contains the data and an indicator of the timeslot's location within the frame.

The maximum frame size is 1,024 bits and the maximum frame pulse offset is 1,023 bit. The line clock speed can be set using the API to one of the following values to support various E1, T1 or aggregated serial (MVIP) specifications:

- 512 KHz
- 1.536 MHz
- 1.544 MHz
- 2.048 MHz
- 4.096 MHz
- 8.192 MHz

The frame size and frame offsets are all programmable according to differing protocols. Other programmable options include signal polarities, signal levels, clock edge, endianness, and choice of input/output frame signal.

### HSS Output Clock Jitter and Error Characterization

The high-speed serial (HSS) port on the processors can be configured to generate an output clock on the HSS\_TXCLK pin. This output clock, however, is not as accurate as using an external oscillator. If the system is intended to clock a framer, DAA, or other device with a sensitive input PLL, an external clock should be used.

Clock signalling is defined in the file IxHssAccCommon.c. The following tables describe the error and jitter characteristics of signals based upon the values established in the software release 2.3.

**Table 56. HSS Tx Clock Output frequencies and PPM Error**

HSS Tx Freq.	Min. Freq. (MHz)	Avg. Freq. (MHz)	Max. Freq. (MHz)	Avg. Freq. Error (PPM)
512 KHz	0.508855	0.512031	0.512769	-60.0096
1.536 MHz	1.515	1.536	1.55023	-60.0096
1.544 MHz	1.515	1.5439	1.55023	+60.0024
2.048 MHz	2.01998	2.0481	2.08313	-60.0096
4.096 MHz	3.92118	4.0962	4.16625	-60.0096
8.192 MHz	7.4066667	8.1925	8.3325	-60.0096

*Note:* Characterization data of the HSS TX clock output frequency data was determined by silicon simulation. PPM parts per million error rate is calculated using average output frequency vs. ideal frequency.

**Table 57. HSS TX Clock Output Frequencies and Associated Jitter Characterization**

HSS Tx Freq.	Pj Max (ns)	Cj Max (ns)	Aj Max (ns)
512 KHz	12.189	15	18.283
1.536 MHz	9.063	15	86.102
1.544 MHz	12.359	15	210.099
2.048 MHz	-8.204	15	118.957
4.096 MHz	10.9	15	190.742
8.192 MHz	12.951	15	226.634



**Table 58. Jitter Definitions**

Jitter Type	Jitter Definition
Period Jitter (Pj)	$Pj_{(i)} = Period_{(i)} - Period_{average}$
Cycle to Cycle Jitter (Cj)	$Cj_{(i)} = Pj_{(i+1)} - Pj_{(i)}$
Wander or Accumulated Jitter (Aj)	$Aj_{(i)} = \sum_i Pj$

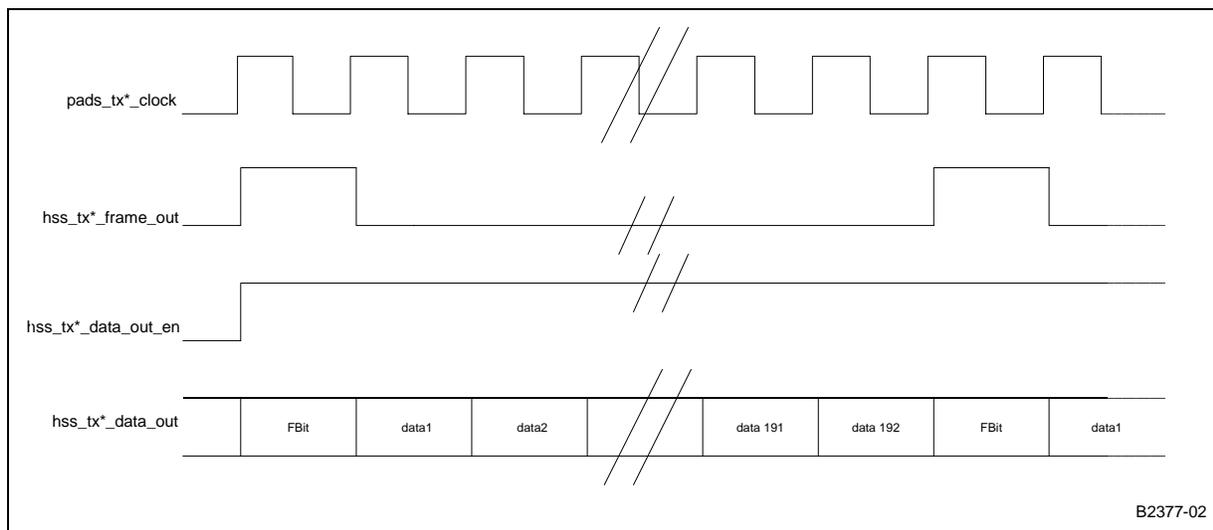
**Table 59. HSS Frame Output Characterization**

HSS Tx Freq.	Frame Size (Bits)	Actual Frame Length (µs)	Frame Length Error (PPM)
512 KHz	32	62.496249	-60.0096
1.536 MHz	96	62.496249	60.016
1.544 MHz	193	125.007499	60.0024
2.048 MHz	256	124.9925	-60.0096
4.096 MHz	512	62.496	-60.0096
8.192 MHz	1024	62.49624	-60.0096

*Note:* PPM frame length error is calculated from ideal frame frequency.

Figure 62 illustrates a typical T1 frame with active-high frame sync (level) and a posedge clock for generating data. If the frame pulse was generated on the negedge in the figure, it would be located one-half clock space to the right. The same location applies if the data was generated on the negedge of the clock.

**Figure 62. T1 Tx Signal Format**



The timeslots within a stream can be configured as packetized (raw or HDLC, 64 Kbps, or 56 Kbps), channelized voice64K, or channelized voice56K or left unassigned. "Voice" slots are those that is sent to the channelized services. For more details, see "HSS Port Initialization Details" on page 218.



For packetized timeslots, data is passed to the HDLC coprocessor for processing as packetized data. The HDLC coprocessor provides the bit-oriented HDLC processing for the HSS coprocessor and can also provide “raw” packets, those which do not require HDLC processing, to the client. The HDLC coprocessor can support up to four packetized services per HSS port.

The following HDLC parameters are programmable:

- The pattern to be transmitted when a HDLC port is idle.
- The HDLC data endianness.
- The CRC type (16-bit or 32-bit) to be used for this HDLC port.
- CAS bit polarity
- CAS bit position (LSB or MSB)
- Bit polarity (bits inverted or not)

For more details, see [“Packetized Connect and Enable”](#) on page 228.

### 13.3.4 Packetized Data Stream

The raw or HDLC timeslots in a logical T1/E1 trunk carry payload for the same data stream controlled by one packetized client. These timeslots carrying raw or HDLC data can be contiguous or non-contiguous. Maximum throughput is achieved by using all the timeslots in a T1/E1 trunk.

timeslots in more than one logical T1/E1 trunk cannot be used to carry one packetized data stream.

#### 13.3.4.1 56-Kbps, Packetized Raw Mode

When a packet service channel is configured for 56-Kbps, packetized Raw mode, byte alignment of the transmitted data is not preserved. All raw data that is transmitted by a device using IxHssAcc in this manner is received in proper order by the receiver (the external PHY device, for example). However, the first bit of the packet may be seen at any offset within a byte. All subsequent bytes has the same offset for the duration of the packet. The same offset also applies to all subsequent packets received on the service channel as well.

The receive data path is identical to the scenario described above.

While this behavior also occurs for 56-Kbps, packetized HDLC mode, the HDLC encoding/decoding preserves the original byte alignment at the receiver end.

### 13.3.5 High-Level API Call Flow

The steps describe the high-level API call-process flow for initializing, configuring, and using the IxHssAcc component.

1. The proper NPE microcode images must be downloaded to the NPEs and initialized accordingly. The IxNpeMh and IxQMgr components are then initialized.
2. Client calls `ixHssAccInit()`. This function is responsible for initializing resources for use by the packetised and channelised clients.
3. For HSS configuration, the client application calls function `ixHssAccPortInit()`. No channelized or packetized connections should exist in the HssAccess layer while this interface is being called. This configures each timeslot in a frame to provide either packetized or channelized service as well as other characteristics of the HSS port. Apart from that, it also help to register client callback to report last error.



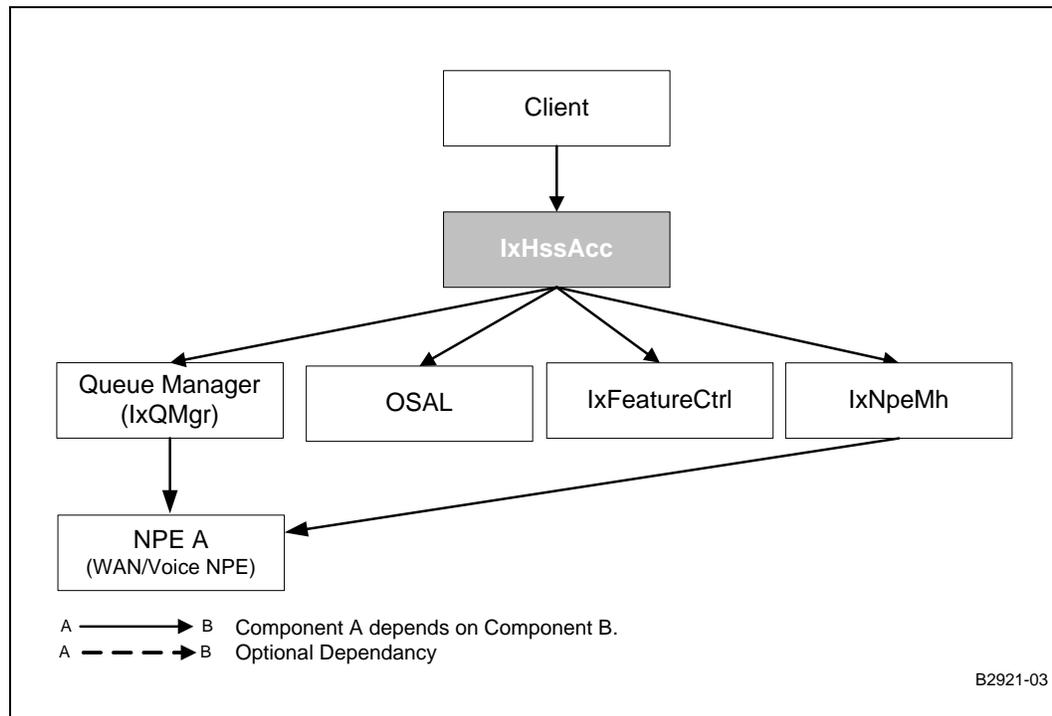
4. Next, the clients prepare data buffers to exchange data with the HSS component, for transmitting or receiving. Depending on whether it is channelized or packetized service, the data is exchanged differently, as described in [“HSS Port Initialization Details” on page 218](#).
5. The client then calls the `ixHssAccPktPortConnect()` (for packetized service) or `ixHssAccChanConnect()` (for channelized service) to connect the client to the IxHssAcc service. Additionally, the client may provide callback functions for the service to inform the client when data is received and ready to delivered to the client.
6. The client begins receiving data once a port is enabled. The functions to enable the packetized or channelized service ports are `ixHssAccPktPortEnable()` and `ixHssAccChanPortEnable()`.  
For channelized service, timeslot can be further configured to Bypass mode using `ixHssAccChanTslotSwitchEnable()` and disabling it using `ixHssAccChanTslotSwitchDisable()`. After enabling the Bypass mode, a Gain Control Table (GCT) must be downloaded into the NPE using `ixHssAccChanTslotSwitchGctDownload()` API. GCT has 1:1 mapping.  
As traffic is being transmitted and/or received on the HSS interfaces and passed to the client, via a channelized (Normal mode or Bypass mode) or packet service, a variety of tasks may be called by the client to check the status, replenish buffers, retrieve statistics, and so forth. Callback functions or a polling mechanism can be used in the transmission and reception process.  
The client processes the received data or provides new data for transmission. This is done by providing new buffer pointers or by adjusting the existing pointers. The data path and requisite buffer management are described in more detail in [“Buffer Allocation Data-Flow Overview” on page 235](#).
7. Finally, call `ixHssAccPktPortDisable()` and/or `ixHssAccChanPortDisable()` if the client just wanted to temporary halt the operation. If the connection is no longer needed, the client may call `ixHssAccPktPortDisconnect()` and/or `ixHssAccChanDisconnect()`. The disable functions instructs the NPE's to stop data handling, while the disconnect functions clears all port configuration parameters. By default the disconnect functions automatically disables the port. The port disable and enable function can be called repeatedly to disable and enable the port without loosing frame synchronization.

### 13.3.6 Dependencies

[Figure 63 on page 217](#) shows the component dependencies of the IxHssAcc component.



Figure 63. IxHssAcc Component Dependencies



The dependency diagram can be summarized as follows:

- Client component calls IxHssAcc for HSS and HDLC data services. NPE A performs the protocol conversion, signalling on the HSS interfaces, and data handling.
- IxHssAcc depends on the IxQMgr component to configure and use the hardware queues to pass data between the Intel XScale® Processor and the NPE.
- NpeMh is used by the component to configure the HSS and HDLC coprocessor operating characteristics.
- OSAL services are used for error handling, thread creation, memory buffer, semaphore and critical code protection.
- IxFeatureCtrl is used to detect the existence of the required hardware features on the host processor. Specifically, IxHssAcc detects the existence of NPE A.

### 13.3.7 Key Assumptions

The HSS service is predicated on the following assumptions:

- Packetized service (for HDLC Mode) is coupled with the HSS port. Packets transmitted using the packetized service access interface is sent through the HDLC coprocessor and on to the HSS coprocessor.
- Tx and Rx TDM slot assignments are identical.
- Packetized services use IX\_OSAL\_MBUF.
- Channelized services use raw circular buffers discussed further in [“Data Flow in Channelized Service” on page 237](#). (For additional details, see [Chapter 3.0](#)).
- All IX\_OSAL\_MBUFs provided by the client to the packetized receive service contains 2,048-byte data stores.



### 13.3.8 Error Handling

The IxHssAcc component use IxOsal to report internal errors and warnings. Parameters passed through the IxHssAcc API interfaces is error checked whenever possible.

HDLC CRC errors and byte alignment errors is reported to packetized clients on a per packet basis. Port disable and disconnect errors on a transmit or receive packetized service pipe is transmitted to the client as well.

HSS port errors such as over-run, under-run and frame synchronization is counted by NPE A, along with other NPE software errors. This count of the total number of errors since configuration is reported to packetized clients on a per packet basis and to channelized clients at the trigger rate.

IxHssAcc provides an interface to the client to read the last error from the NPE. There is no guarantee that the client is able to read every error. A second error may occur before the client has had the opportunity to read the first one. The client will, however, have an accurate total error count.

## 13.4 HSS Port Initialization Details

### *ixHssAccPortInit()*

The HSS ports must be configured to match the configuration of any connected PHY. No channelized or packetized connections should exist in the IxHssAcc layer while this interface is being called.

This includes configuring the timeslots within a frame in one of the following ways:

- Configuring as HDLC — For packetized service, including raw packet mode
- Configuring as Voice64K/Voice56K — For channelized service
- Configuring as unassigned — For unused timeslot
- Choosing the line speed, frame size, signal polarities, signal levels, clock edge, endianness, choice of input/output frame signal, and other parameters

This function takes the following arguments:

- IxHssAccHssPort hssPortId — The HSS port ID.
- IxHssAccConfigParams \*configParams — A pointer to the HSS configuration structure.
- IxHssAccTdmSlotUsage \*tdmMap — A pointer to an array defining the HSS time-slot assignment types.
- IxHssAccLastErrorCallback lastHssErrorCallback — Client callback to report the last error.

The parameter IxHssAccConfigParams has two structures of type IxHssAccPortConfig — one for HSS Tx and one for HSS Rx. These structures are used to choose:

- Frame-synchronize the pulse type (Tx/Rx)
- Determine how the frame sync pulse is to be used (Tx/Rx)
- Frame-synchronize the clock edge type (Tx/Rx)
- Determine the data clock edge type (Tx/Rx)
- Determine the clock direction (Tx/Rx)
- Determine whether or not to use the frame sync pulse (Tx/Rx)
- Determine the data rate in relation to the clock (Tx/Rx)



- Determine the data polarity type (Tx/Rx)
- Determine the data endianness (Tx/Rx)
- Determine the Tx pin open drain mode (Tx)
- Determine the start of frame types (Tx/Rx)
- Determine whether or not to drive the data pins (Tx)
- Determine the how to drive the data pins for voice56k type (Tx)
- Determine the how to drive the data pins for unassigned type (Tx)
- Determine the how to drive the Fbit (Tx)
- Set 56Kbps data endianness, when using the 56Kbps type (Tx)
- Set 56Kbps data transmission type, when using the 56Kbps type (Tx)
- Set the frame-pulse offset in bits w.r.t, for the first timeslot (0-1,023) (Tx/Rx)
- Determine the frame size in bits (1-1,024)

IxHssAccConfigParams also has the following parameters:

- The number of channelized timeslots (0 - 31)
- The number of packetized clients (0 - 3)
- The byte to be transmitted on channelized service, when there is no client data to Tx
- The HSS coprocessor loop-back state
- The data to be transmitted on packetized service, when there is no client data to Tx
- The HSS clock speed

IxHssAccTdmSlotUsage is an array that take the following values to assign service types to each timeslot in a HSS frame:

IX_HSSACC_TDMMAP_UNASSIGNED	Unassigned
IX_HSSACC_TDMMAP_HDLC	Packetized
IX_HSSACC_TDMMAP_VOICE56K	Channelized
IX_HSSACC_TDMMAP_VOICE64K	Channelized

IxHssAccTdmSlotUsage has a size equal to the number of timeslots in a frame.

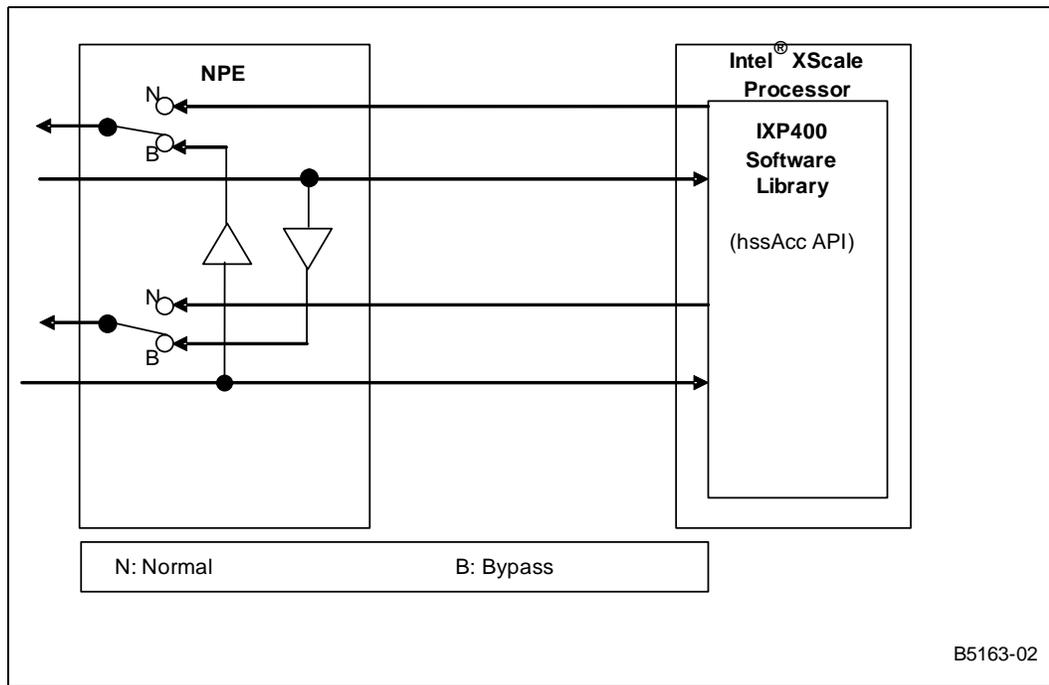
*IxHssAccLastErrorCallback()* is for error handling. The client initiates the last error retrieval. The HssAccess component then sends a message to the NPE through the NPE Message Handler. When a response to the error retrieval is received, the NPE Message Handler callbacks the HssAccess component, which executes *IxHssAccLastErrorCallback()* in the same IxNpeMh context. The client is passed the last error and the related service port.

When complete, the HSS coprocessor is running, although no access is given to the client until a connect occurs followed by an enable.

## 13.5 HSS Channelized Operation

In channelized operation, timeslots can be set to Normal mode or Bypass mode on port 0 only through HssAcc API. A graphical representation of both modes is shown in [Figure 64](#):

Figure 64. Normal and Bypass Mode Illustration



In Normal mode, data received on the HSS port for an enabled, channelized timeslot is always written to a predefined SDRAM area where the Intel XScale® Processor can read it and data timeslots that is to be transmitted to the HSS port is always taken from the Intel XScale® Processor. Whereas Bypass mode (which is also known as Voice Switching Mode) allows the received data from one TDM mapped voice timeslot (Rx) to transmit onto another voice timeslot (Tx) for the same HSS port at the NPE level, overwriting the data from the Intel XScale® Processor destined to the Tx bypassed voice timeslot. The data from the Rx bypassed voice timeslot is routed to the Intel XScale® Processor via the HSS Access Layer.

For example, Channel A is bypassed to Channel B bypass and Channel B is bypassed to Channel A bypass as shown above. A copy of the received data is sent to the Intel XScale® Processor in order to perform further optional voice related processing at client application layer. (eg. DSR functionality for tone detection, compression). Bypass mode can be enabled and disabled on a per timeslot basis and on the fly.

## 13.5.1 Channelized Connect and Enable

### 13.5.1.1 Normal Mode

#### **ixHssAccChanConnect()**

After the HSS component is configured, *ixHssAccChanConnect()* has to be called to connect the client application with the channelized service. This function is called once per HSS port, and there can only be one client per HSS port.

The client uses this function to:

- Register a Rx call-back function.
- Set up how often this callback function is called.



- Pass the pointer to the Rx data circular buffer pool.
- Set the size of the Rx circular buffers.
- Set the pointer to the Tx pointer lists pool.
- Set the size of the tx data buffers.

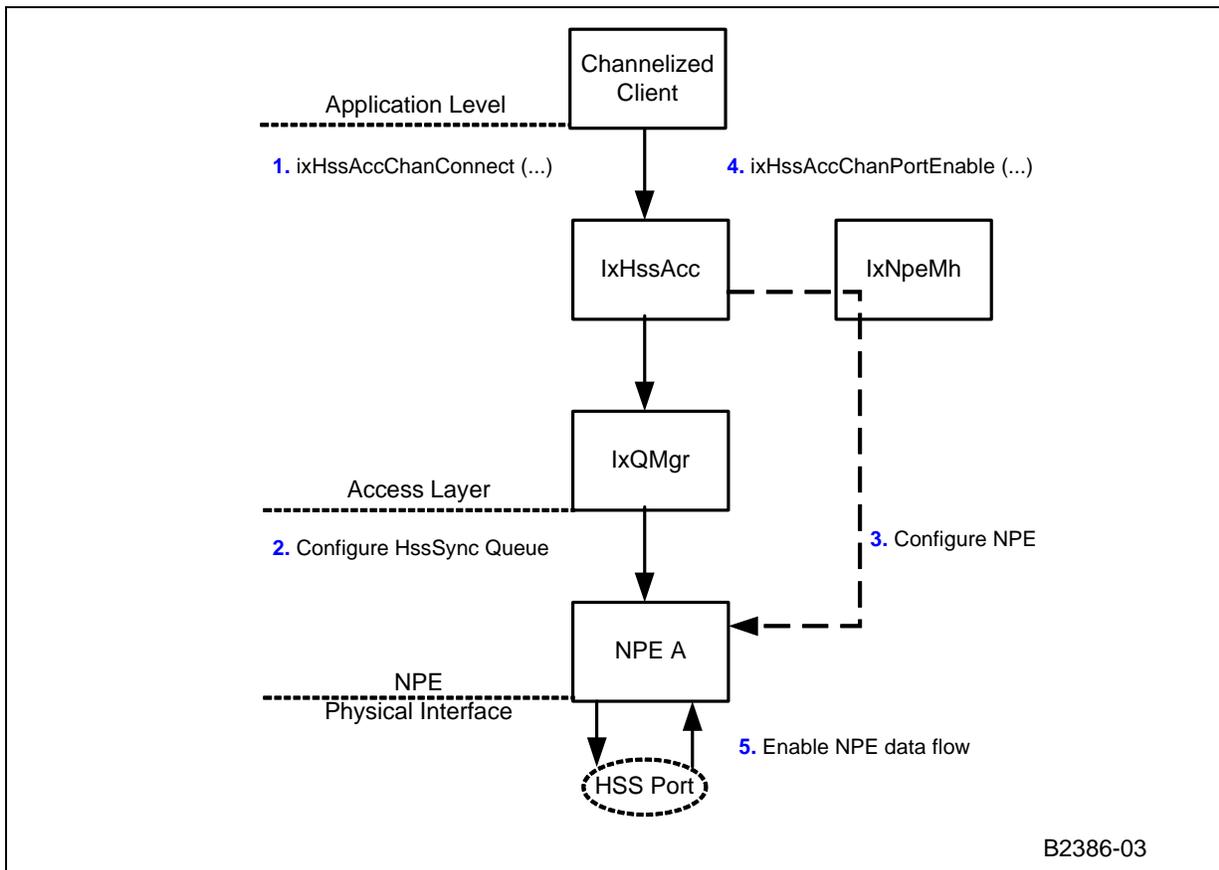
The parameters needed by `ixHssAccChanConnect()` include:

- `IxHssAccHssPort hssPortId` — The HSS port ID. There are two identical ports (0-1).
- `unsigned bytesPerTSTrigger` — The NPE triggers the access component to call the Rx call back function `rxCallback()` after `bytesPerTSTrigger` bytes have been received for all trunk timeslots. `bytesPerTSTrigger` is a multiple of eight. For example: 8 for 1-ms trigger, 16 for 2-ms trigger.
- `UINT8 *rxCircular` — A pointer to the Rx data pool allocated by the client as described in previous section. It points to a set of circular buffers to be filled by the received data. This address is written to by the NPE and must be a physical address.
- `unsigned numRxBytesPerTS` — The length of each Rx circular buffer in the Rx data pool. The buffers need to be deep enough for data to be read by the client before the NPE re-writes over that memory.
- `UINT32 *txPtrList` — The address of an area of contiguous memory allocated by the client to be populated with pointers to data for transmission. Each pointer list contains a pointer per active channel. The `txPtrs` points to data that has to be transmitted by the NPE. Therefore, they must point to physical addresses.
- `unsigned numTxPtrLists` — The number of pointer lists in `txPtrList`. This number is dependent on jitter.
- `unsigned numTxBytesPerBlk` — The size of the Tx data, in bytes, that each pointer within the `PtrList` points to.
- `IxHssAccChanRxCallback rxCallback` — A client function pointer to be called back to handle the actual tx/rx of channelized data after `bytesPerTSTrigger` bytes have been received for all trunk timeslots. If this pointer is NULL, it implies that the client uses a polling mechanism to detect when the Tx and Rx of channelized data is to occur.

After the client application is connected with the channelized service, the HSS component then can be enabled by calling `ixHssAccChanPortEnable()` with the port ID provided to enable the channelized service from that particular HSS port.

The following figure shows what is done in `IxHssAcc` when the `ixHssAccChanPortConnect()` and `ixHssAccChanPortEnable()` functions are called.

Figure 65. Channelized Connect For Normal Mode



1. The client issues a channelized connect request to IxHssAcc.
2. If an rxCallback is configured, the client expects to be triggered by events to drive the Tx and Rx block transfers. IxHssAcc registers the function pointer with IxQMgr to be called back in the context of an ISR when the HssSync queue is not empty.
3. IxHssAcc configures the NPE appropriately.
4. The client enables the channelized service through IxHssAcc.
5. IxHssAcc enables the NPE flow.

If the service was configured to operate in polling mode (for example, the rxCallback pointer is NULL), the client must poll the IxHssAcc component using the `ixHssAccChanStatusQuery()` function. IxHssAcc checks the HssSync queue status and returns a pointer to the client indicating an offset to the data in the Rx buffers, if any receive data exists at that time.

### 13.5.1.2 Bypass Mode

The procedures for enabling the Bypass mode still be the same as the Normal mode. The difference is, the Bypass mode needs an extra setup after the `ixHssAccChanPortEnable()` is called. `ixHssAccChanTslotSwitchEnable()` is called and is responsible for enabling timeslot switching (bypass) channel between two voice channels for the specified HSS port. The voice channels must have already been configured as channelised timeslot for the specified HSS port. In current release, only up to 2 pairs of timeslot switching channels can be enabled on port 0 only at any one



time. In order to minimize bypass delay and ensure better voice quality, this function requires at least 8 TDM timeslots on the specified HSS port to be setup as channelised timeslots. Client can choose to ignore the unwanted channelized timeslots. In Bypass mode, data received on srcTimeslot is transmitted onto a partner timeslot (for example destTimeslot) at NPE level. A copy of the received data on srcTimeslot is also sent to client via HssAccess component. No other HssAccChannelised interface should be called while this interface is being processed

The following are the parameters that needed by *ixHssAccChanTslotSwitchEnable()*:

- IxHssAccHssPort hssPortId (in) - The HSS port Id. There are two identical ports (0-1). Only port 0 is supported
- UNIT32 srcTimeslot (in) - The voice channel Id whose its receive side is used in the bypass (0-127)
- UNIT32 destTimeslot (in) - The voice channel Id whose its transmit side is used in the bypass (0-127)
- UNIT32 \*tsSwitchHandle (out) - A pointer where an handle ID value is returned to client through it during the enabling. It hooks to the bypass channel established between srcTimeslot and destTimeslot. This handle is the mean by which client disables or downloads gain control table to NPE for the bypass channel that associates with this handle. Client must ignore the value returned through this handle if bypass channel fails to setup.

*ixHssAccChanTslotSwitchEnable()* returns

- ..... IX\_SUCCESS if the function executed successfully or
- ..... IX\_FAIL if the function did not execute successfully or
- ..... IX\_HSSACC\_PARAM\_ERR if the function did not execute successfully due to a parameter error.

*IxHssAccChanTslotSwitchGctDownload ()* is then called to download a gain control table (256 bytes) to NPE for the specified timeslot switching (bypass) channel, associated with its tsSwitchHandle, on the specified HSS port. The bypass voice channel must have already been enabled for the specified HSS port before this API can be called to download the gain control table to NPE. No other HssAccChannelised interface should be called while this interface is being processed. Gain control table is a look up table (an array) where client defines the tune values (there is no restriction on the range of tune values and it is up to customer application) for every timeslots respectively. Timeslots in the Gain control table has a 1:1 mapping on its entry values.

- IxHssAccHssPort hssPortId (in) - The HSS port Id. There are two identical ports (0-1). Only port 0 is supported.
- UINT8 \*gainCtrlTable (in) - A pointer to an array of size IX\_HSSACC\_ENTRIES\_PER\_GAIN\_CTRL\_TABLE, defining each entry for a gain control table for the specified bypass voice channel
- UINT32 tsSwitchHandle (in) - The handle that hooks to the bypass channel. This handle is the parameter returned to client by *ixHssAccChanTslotSwitchEnable* during timeslot switching channel enabling operation

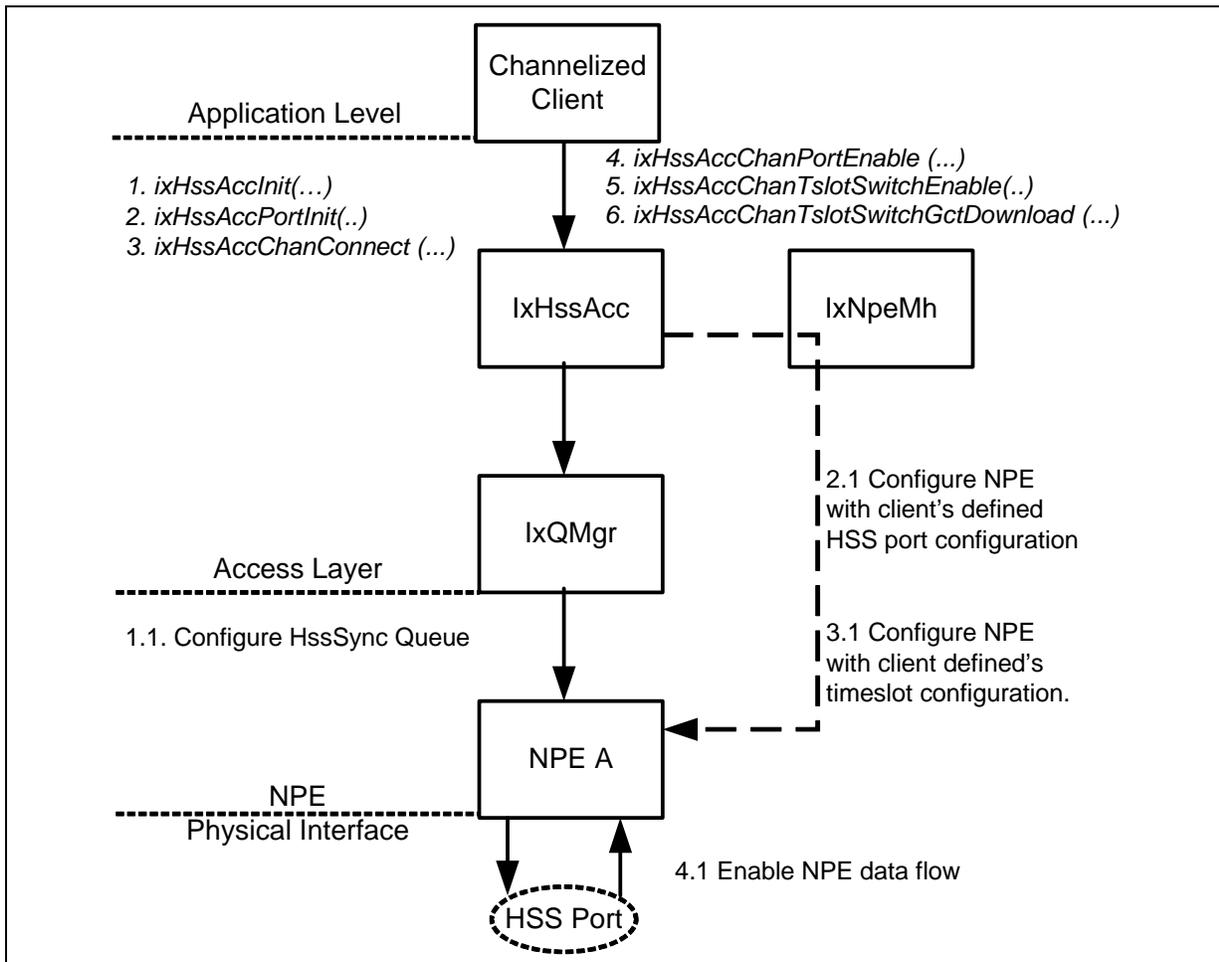
*ixHssAccChanTslotSwitchGctDownload ()* returns

- IX\_SUCCESS if the function executed successfully or
- IX\_FAIL if the function did not execute successfully or
- IX\_HSSACC\_PARAM\_ERR if the function did not execute successfully due to a parameter error

In Bypass mode, IxHssAccChanTsSwitchConf (it is associated to tsSwitchHandle) is used as a structure that tracks HSS bypass configuration. It is initialized in ixHssAccChanConnect() and updated in ixHssAccChanTslotSwitchEnable(), ixHssAccChanTslotSwitchDisable() and ixHssAccChanTslotSwitchGctDownload(). It tracks:

- whether the bypass channel has been enabled
- whether the gain control table has been downloaded
- the voice channel Id which its receive side is used in the timeslot switching channel
- the voice channel Id which its transmit side is used in the timeslot switching channel

**Figure 66. Channelized Connect For Bypass Mode**



**Pre-Condition for Enabling a HSS Bypass Channel**

The Intel XScale® Processor should not send the NPE any bypass channel enabling message and shall respond with error message to the client when the following scenarios happen:

- The client is enabling the same bypass channel the second time before disabling it first
- The client is trying to enable more than two bypass channels at the same time



- The client is trying to enable the bypass channels (on the Intel XScale® Processor side) with timeslot number other than the valid range (0-127)
- The client is trying to enable a bypass channel with a timeslot value that has not been enable as “voice timeslot” in the HSS Voice Channel Map; although this timeslot value fall within the valid range (0-127)
- The client has not configured at least 8 channels in the HSS Voice Channel Map as “voice timeslot”. This is to ensure latency less than or equal to 2ms for the bypass channels.
- The client uses NPE A images other than the functionality ID specified for HSS Voice Channel Switching (bypass) service

### HSS Bypass Channel Enabling Sequence

1. The Intel XScale® Processor downloads the Gain Control Table (GCT) to NPE for the corresponding bypass channel A, for example.
2. The Intel XScale® Processor waits for the acknowledgement from NPE for the success of current gain control word downloaded before sending in another 4 bytes. This sequence is repeated until the 256-bytes GCT has been fully downloaded.
3. The Intel XScale® Processor shall then send the HSS message to NPE to enable a bypass channel A.
4. NPE shall send a response to the Intel XScale® Processor when the current bypass channel has been configured.
5. The Voice Channel Switching service shall be enabled internally in the NPE at the beginning of the next voice frame.

Additional bypass channel can be configured following steps 1-5 above.

### HSS bypass Channel Disabling Sequence

The Intel XScale® Processor shall send a HSS BypassChanDisable message to NPE when:

- The client is trying to disable a bypass channel that has already been enabled previously
- When the client is trying to “disconnect” the HSS Channelised service.

## 13.5.2 Channelized Tx/Rx Methods

After being initialized, configured, connected, and enabled, the HSS component is up and running. There are two methods to handle channelized service Tx/Rx process: interrupt mode via callback and polling mode.

### 13.5.2.1 Interrupt Mode via Callbacks

If the pointer to the *rxCallback()* is not *NULL* when *ixHssAccChanConnect()* is called, an ISR calls *rxCallback()* to handle Tx/Rx data. It is called when each of *N* channels receives *bytesPerTStrigger* bytes.

Usually, a Rx thread is created to handle the HSS channelized service. The thread is waiting for a semaphore. When *rxCallback()* is called by IxHssAcc, *rxCallback()* puts the information from IxHssAcc into a structure, and send a semaphore to the thread. Then *rxCallback()* returns so that IxHssAcc can continue its own tasks. The Rx thread — after receiving the semaphore — wakes up, takes the parameters passed by *rxCallback()*, and perform Rx data processing, Tx data preparation, and error handling.



For Rx data processing, *rxCallback()* provides the offset value *rxOffset* to indicate where data is just written into each circular buffer. *rxOffset* is shared for all the circular buffers in the pool. The client has to make sure the Rx data are processed or moved to somewhere else before overwritten by the NPE since the buffers are circular.

For TX data preparation, *rxCallback()* provides the offset value *txOffset* to indicate which pointer list in the pointer lists pool is pointing to the data buffers currently being or is transmitted. As a result, the client can use *txOffset* to determine where new data must be put into the data buffer pool for transmission. For example, data can be prepared and moved into buffers pointed by the *(txOffset-2)*th pointer list.

*rxCallback()* also provides the number of errors NPE receives. The client can call function *ixHssAccLastErrorRetrievalInitiate()* to initiate the retrieval of the last HSS error.

### 13.5.2.2 Polling Mode

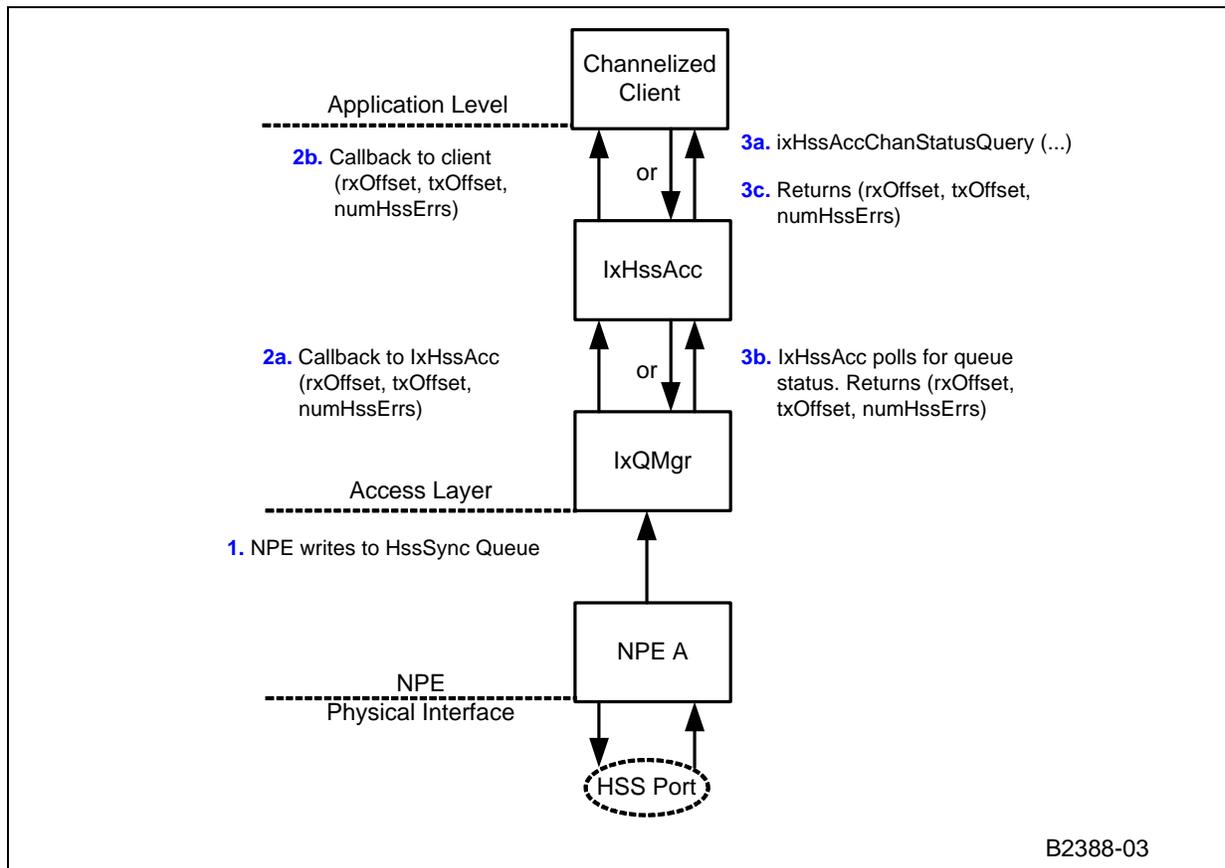
If the pointer to the *rxCallback()* is NULL when *ixHssAccChanConnect()* is called, it implies that the client uses a polling mechanism to detect when the Tx and Rx of channelized data is to occur. The client uses *ixHssAccChanStatusQuery()* to query whether channelized data has been received. If data has been received, IxHssAcc returns the details in the output parameters of *ixHssAccChanStatusQuery*.

*ixHssAccChanStatusQuery()* returns a flag *dataRecvd* that indicates whether the access component has any data for the client. If *FALSE*, the other output parameters will not have been written to. If it is *TRUE*, then *rxOffset*, *txOffset*, and *numHssErrs* — returned by *ixHssAccChanStatusQuery()* — are valid and can be used in the same way as in the callback function case above.

Figure 67 shows the Tx/Rx process.



Figure 67. Channelized Transmit and Receive



1. After reading a configurable amount of data from the HSS port and writing the same amount of data to the HSSport, the NPE writes to the hssSync queue. There are two possible paths after that depending on how the client is connected:
  2. Interrupt Mode via callback
    - a. Through an interrupt, the IxQMgr component calls back IxHssAcc with details of the hssSync queue entry.
    - b. IxHssAcc initiates the registered callback of the client.
- OR**
3. Polling Mode
    - a. The client polls IxHssAcc using `ixHssAccChanStatusQuery()`.
    - b. IxHssAcc will, in turn, poll IxQMgr's hssSync queue for status.
    - c. If IxHssAcc reads an entry from the hssSync queue, it returns the details to the client.



### 13.5.3 Channelized Disconnect

When the channelized service is not needed any more on a particular HSS port, `ixHssAccChanPortDisable()` is called to stop the channelized service on that port, then `ixHssAccChanDisconnect()` is called to disconnect the service. The calling of `ixHssAccChanPortDisable()` is optional because `ixHssAccChanDisconnect()` automatically checks and disable the port accordingly.

## 13.6 HSS Packetized Operation

### 13.6.1 Packetized Connect and Enable

#### **`ixHssAccPktPortConnect()`**

After the HSS component is configured, `ixHssAccPktPortConnect()` has to be called to connect the client application with the packetized services. This function is responsible for connecting a client to one of the four available packetized ports on a configured HSS port.

There are four packetized services per HSS port, so this function has to be called once per packetized service.

The client uses this function to:

- Pass data structures to configure the HDLC coprocessor
- Register a Rx call back function for Rx data processing
- Register a callback function to request more Rx buffers
- Register a callback function to indicate Tx done
- Pass a flag to turn HDLC processing on or off

The HDLC configuration structure sets up:

- What to transmit when an HDLC port is idle
- HDLC data endianness
- CRC type to be used for this HDLC port.

The parameters for `ixHssAccPktPortConnect()` include:

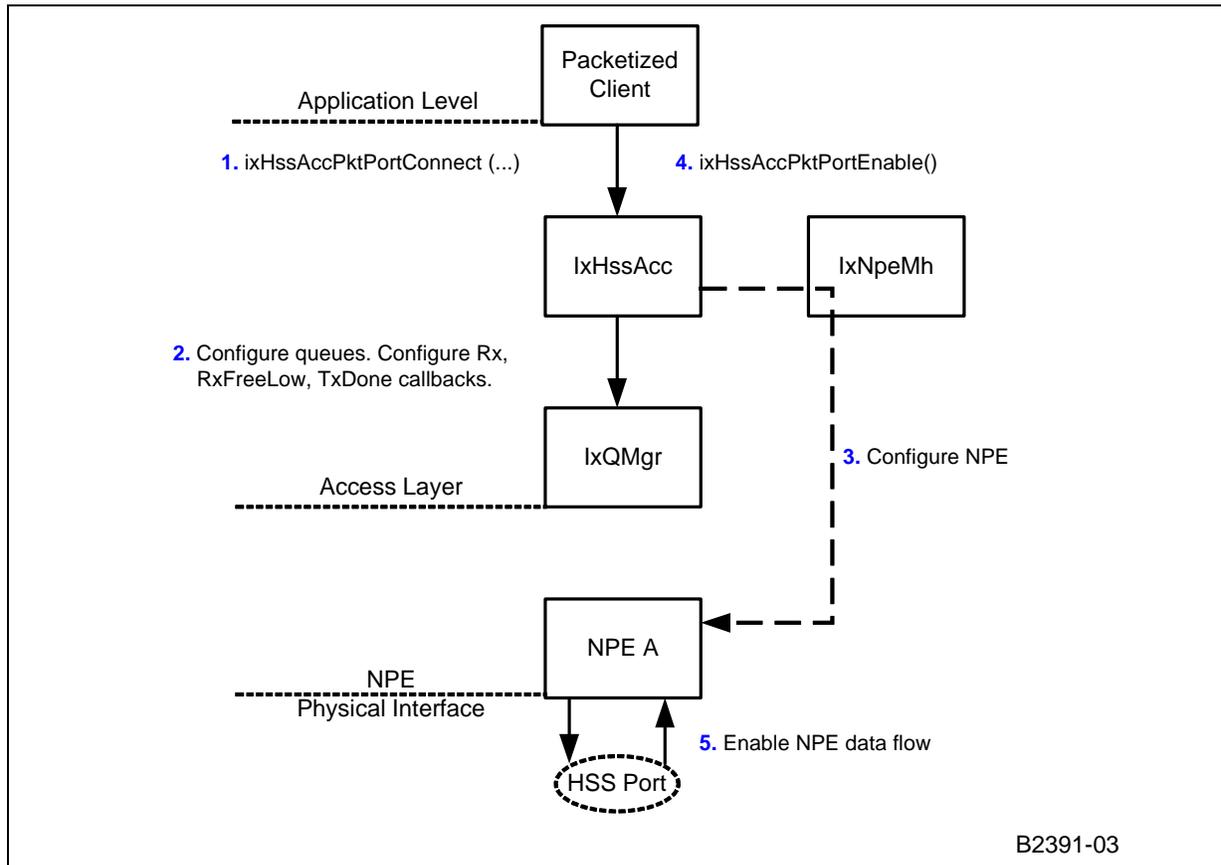
- `IxHssAccHssPort hssPortId` — The HSS port ID. There are two identical ports (0-1).
- `IxHssAccHdlcPort hdlcPortId` — This is the number of the HDLC port and it corresponds to the physical E1/T1 trunk (for example, 0, 1, 2, 3).
- `BOOL hdlcFraming` — This value determines whether the service uses HDLC data or the raw data type (for example, no HDLC processing).
- `IxHssAccHdlcMode hdlcMode` — This structure contains 56Kbps, HDLC-mode configuration parameters.
- `BOOL hdlcBitInvert` — This value determines whether bit inversion occurs between HDLC and HSS coprocessors (for example, post-HDLC processing for transmit and pre-HDLC processing for receive, for the specified HDLC termination Point).
- `unsigned blockSizeInWords` — The max tx/rx block size.
- `UINT32 rawIdleBlockPattern` — Tx idle pattern in raw mode.
- `IxHssAccHdlcFraming hdlcTxFraming/ hdlcRxFraming` — This structure contains the information required by the NPE to configure the HDLC coprocessor for Tx/ Rx.
- `unsigned frmFlagStart` — Number of flags to precede to transmitted flags (0-2).



- IxHssAccPktRxCallback rxCallback — Pointer to the clients packet receive function.
- IxHssAccPktUserId rxUserId — The client supplied Rx value to be passed back as an argument to the supplied rxCallback.
- IxHssAccPktRxFreeLowCallback rxFreeLowCallback — Pointer to the clients Rx-free-buffer request function. If NULL, it is assumed that the client frees Rx buffers independently.
- IxHssAccPktUserId rxFreeLowUserId — The client supplied RxFreeLow value to be passed back as an argument to the supplied rxFreeLowCallback.
- IxHssAccPktTxDoneCallback txDoneCallback — Pointer to the clients Tx done callback function.
- IxHssAccPktUserId txDoneUserId — The client supplied txDone value to be passed back as an argument to the supplied txDoneCallback.

Now the HSS component can be enabled by calling *ixHssAccPktPortEnable()* with the port ID provided. Figure 68 shows what is done in IxHssAcc when the packetized service connect function is called.

**Figure 68. Packetized Connect**



1. The client issues a packet service connect request to IxHssAcc.
2. IxHssAcc instructs IxQMgr to configure the necessary queues and register callbacks.
3. IxHssAcc configures the NPE with the HDLC parameters passed by the client.
4. The client enables the packet service.



5. IxHssAcc enables the NPE flow.

### 13.6.2 Packetized Tx

When the client has nothing to transmit, the HSS transmits the idle pattern provided in the function *ixHssAccPktPortConnect()*.

When the client has data for transmission, the client calls *IX\_OSAL\_MBUF\_POOL\_GET()* to get a *IX\_OSAL\_MBUF*, put the data into the *IX\_OSAL\_MBUF* using *IX\_OSAL\_MBUF\_MDATA()*. If the client data is too large to fit into one buffer, multiple buffers can be obtained from the pool, and put into a chained buffers by using *IX\_OSAL\_MBUF\_PKT\_LEN()* and *IX\_OSAL\_MBUF\_NEXT\_BUFFER\_IN\_PKT\_PTR()*. The whole chained buffer is passed to IxHssAcc for transmission by calling *ixHssAccPktPortTx()*.

When the transmission is done, the TxDone call back function, registered with *ixHssAccPktPortConnect()*, is called, and the buffer can be returned to *IX\_OSAL\_MBUF* pool using *IX\_OSAL\_MBUF\_POOL\_PUT\_CHAIN()*.

The following is example Tx code showing how to send an *IX\_OSAL\_MBUF*:

```
IX_OSAL_MBUF *txBuffer;

IX_OSAL_MBUF *txBufferChain = NULL;

// get a IX_OSAL_MBUF
IX_OSAL_MBUF_POOL_GET(poolId, &txBuffer);

// set the data length in the buffer
IX_OSAL_MBUF_MLEN(txBuffer) = NumberOfBytesToSend;

/* set the values to transmit */
for (byteIndex = 0; byteIndex < IX_OSAL_MBUF_MLEN(txBuffer); byteIndex++)
    ((UINT8 *)IX_OSAL_MBUF_MDATA(txBuffer))[byteIndex] =userData[byteIndex];

//send the buffer out
ixHssAccPktPortTx (hssPortId, hdlcPortId, txBuffer);
```

The following is example Tx code showing how to chain *IX\_OSAL\_MBUFs* together:



```
IX_OSAL_MBUF *txBufferChain = NULL;

IX_OSAL_MBUF *lastBuffer = NULL;

if (txBufferChain == NULL)
{
    // the first buffer

    txBufferChain = txBuffer;

    /* set packet header for buffer */

    IX_OSAL_MBUF_PKT_LEN(txBufferChain) = 0;
}

else
{
    // following buffers

    IX_OSAL_MBUF_NEXT_BUFFER_IN_PKT_PTR(lastBuffer) = txBuffer;
}

// update the chain length

IX_OSAL_MBUF_PKT_LEN(txBufferChain) += IX_OSAL_MBUF_MLEN(txBuffer);

IX_OSAL_MBUF_NEXT_BUFFER_IN_PKT_PTR(txBuffer) = NULL;

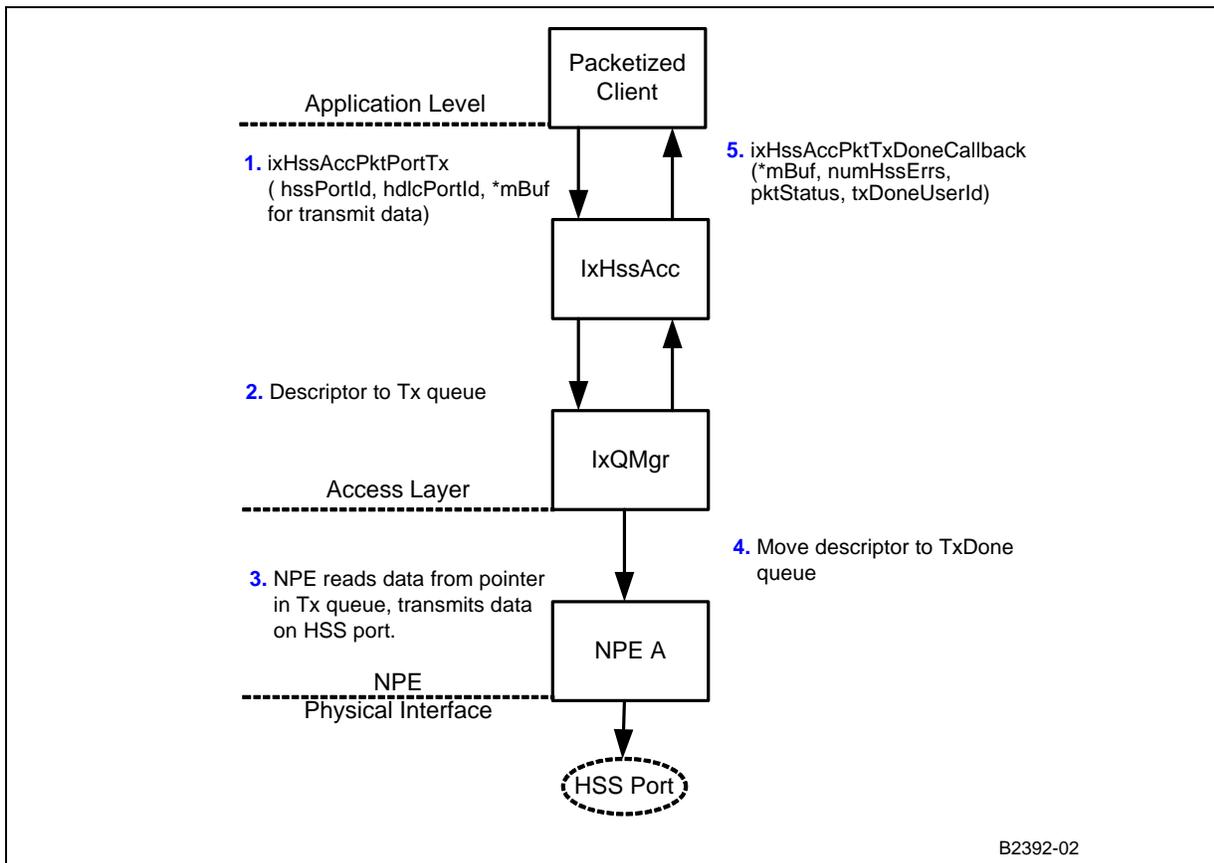
lastBuffer = txBuffer;

// send the bubble out

ixHssAccPktPortTx (hssPortId, hdlcPortId, txBufferChain);
```

The process is shown in [Figure 69](#).

Figure 69. Packetized Transmit



1. The client presents an IX\_OSAL\_MBUF to IxHssAcc for transmission.
2. IxHssAcc gets a transmit descriptor from its transmit descriptor pool, fills in the descriptor, and writes the address of the descriptor to the Tx queue.
3. The NPE reads a transmit descriptor from the Tx queue and transmits the data on the HSS port.
4. On completion of transmission, the NPE writes the descriptor to the TxDone queue.
5. IxHssAcc is triggered by this action, and the registered callback is executed. The descriptor is freed internally.
6. IxHssAcc initiates the TxDoneCallback on the client, passing it back its IX\_OSAL\_MBUF pointer.

### 13.6.3 Packetized Rx

1. Before packetized service is enabled, the Rx queue in the IxHssAcc component has to be replenished. This can be done by calling `IX_OSAL_MBUF_POOL_GET()` to get an IX\_OSAL\_MBUF and calling `ixHssAccPktPortRxFreeReplenish()` to put the buffer into the queue. This is repeated until the queue is full.

Here is an example:



```

// get a buffer

IX_OSAL_MBUF *rxBuffer;

rxBuffer = IX_OSAL_MBUF_POOL_GET(poolId);

//IxHssAcc component needs to know the capacity of the IX_OSAL_MBUF
IX_OSAL_MBUF_MLEN(rxBuffer) = IX_HSSACC_CODELET_PKT_BUFSIZE;

// give the Rx buffer to the HssAcc component

status = ixHssAccPktPortRxFreeReplenish (hssPortId,
hdlcPortId, rxBuffer);

```

Usually, an Rx thread is created to handle the HSS packetized service, namely, to handle all the callback functions registered with *ixHssAccPktPortConnect()*. The thread is waiting for a semaphore. When any one of the call back functions is executed by the HSS component, it puts the information from IxHssAcc into a structure, and send a semaphore to the thread. Then the callback function returns so that IxHssAcc can continue its own tasks. The Rx thread, after receiving the semaphore, wakes up, takes the parameters from the structure passed by the callback function, and performs Rx data processing and error handling.

When data is received, *rxCallback()* is called. It passes the received data in the form of a IX\_OSAL\_MBUF to the client. The IX\_OSAL\_MBUF passed back to the client could contain a chain of IX\_OSAL\_MBUF, depending on the packet size received. *IX\_OSAL\_MBUF\_NEXT\_BUFFER\_IN\_PKT\_PTR()* can be used to get access to each of the IX\_OSAL\_MBUF in the chained buffer, and *IX\_OSAL\_MBUF\_MDATA()* can be used to get access to each data value. The IX\_OSAL\_MBUF is returned to the buffer pool by using *IX\_OSAL\_MBUF\_POOL\_PUT\_CHAIN()*.

Here is an example:

```

IX_OSAL_MBUF *buffer,
IX_OSAL_MBUF *rxBuffer;

// go through each buffer in the chained buffer
for (rxBuffer = buffer;
    (rxBuffer != NULL) && (pktStatus == IX_HSSACC_PKT_OK);
    rxBuffer = IX_OSAL_MBUF_NEXT_BUFFER_IN_PKT_PTR(rxBuffer))
for (wordIndex =0;wordIndex<(IX_OSAL_MBUF_MLEN(rxBuffer) / 4);
    wordIndex++)
{
    // get the values in the buffer IX_OSAL_MBUF
    value = ((UINT32 *)IX_OSAL_MBUF_MDATA(rxBuffer))[wordIndex];
}

// free the chained buffer
IX_OSAL_MBUF_POOL_PUT_CHAIN(buffer);

```

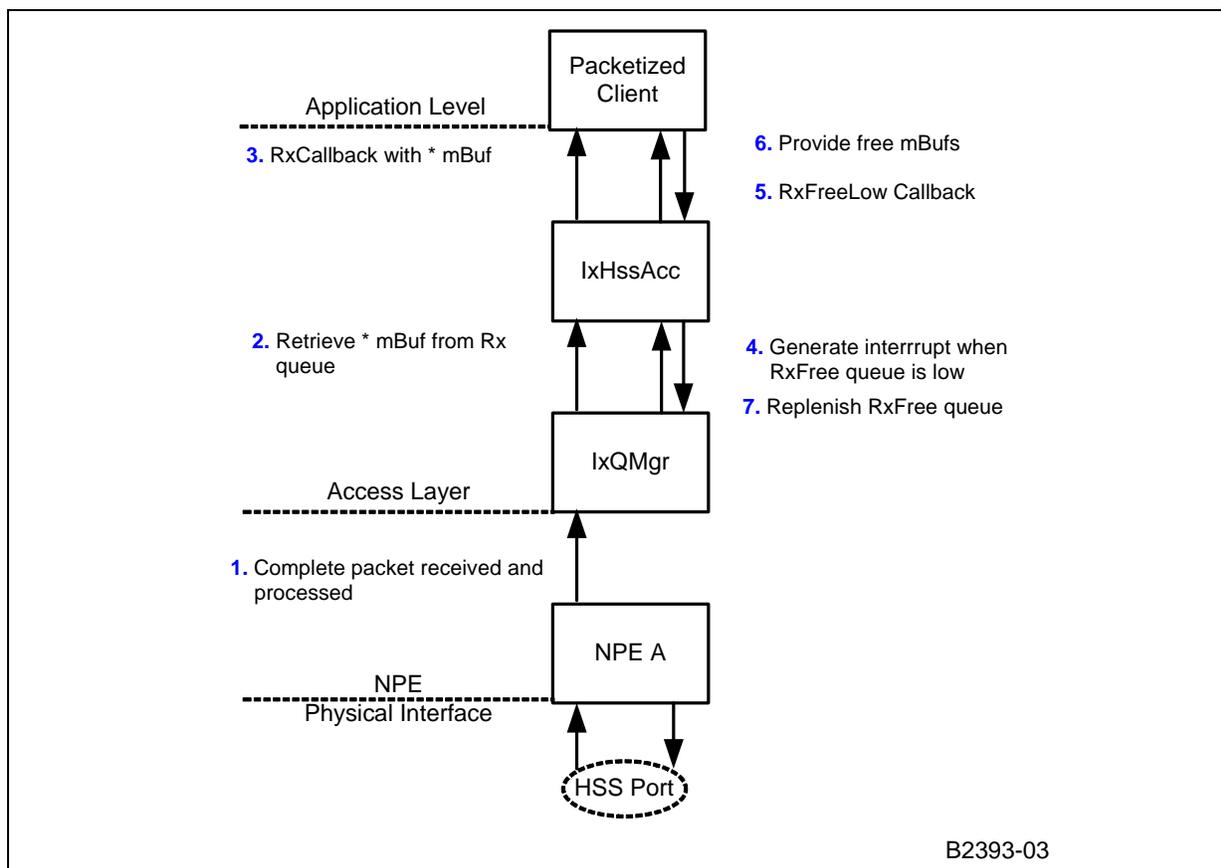
*rxCallback()* also passes the packet status and the number of errors that NPE receives. The packet status is used to determine if the packet received is good or bad, and the client can call function *ixHssAccLastErrorRetrievalInitiate()* to initiate the retrieval of the last HSS error.

When the Rx buffer queue is running low, *rxFreeLowCallback()* is called. Then, the client can call *IX\_OSAL\_MBUF\_POOL\_GET()* and *ixHssAccPktPortRxFreeReplenish()* to fill up the Rx queue again.

Alternatively, the client can use its own timer for supplying *IX\_OSAL\_MBUFs* to the queue. This is the case if the pointer for *rxFreeLowCallback()* passed to *ixHssAccPktPortConnect()* is *NULL*.

The process is show in Figure 70.

**Figure 70. Packetized Receive**



1. When a complete packet is received, the Rx queue call-back function is invoked in an interrupt.
2. The descriptor is pulled from the Rx queue and the callback for this channel is invoked with the descriptor. The descriptor gets recycled.
3. The buffer is transmitted to the client.
4. When the RxFree queue is low, IxQMgr triggers an interrupt to IxHssAcc.
5. IxHssAcc triggers the client *rxFreeLowCallback* function, which was registered during the client connection process.
6. Provide free mBufs
7. RxFreeLow Callback



6. The client provides free IX\_OSAL\_MBUFs for specific packetized channels.
7. Free IX\_OSAL\_MBUFs are stored in the RxFree queue, and listed within the IxHssAcc Rx descriptors.

#### 13.6.4 Packetized Disconnect

When packetized service channel is not needed any more, the function `ixHssAccPktPortDisable()` is called to stop the packetized service on that channel, and `ixHssAccPktPortDisconnect()` is called to disconnect the service.

This has to be done for each packet service channel. The client is responsible for ensuring all transmit activity ceases prior to disconnecting, and ensuring that the replenishment of the rxFree queue ceases before trying to disconnect.

### 13.7 Buffer Allocation Data-Flow Overview

Prior to connecting and enabling ports in IxHssAcc, a client must allocate buffers to the IxHssAcc component. IxHssAcc provides two services, packetized and channelized, and the clients exchange data with IxHssAcc differently for transmitting and receiving, depending on which service is chosen.

#### 13.7.1 Data Flow in Packetized Service

Data in the timeslots configured for HDLC or raw services forms packets for packetized service. IxHssAcc supports up to four packetized services per HSS port. The packetized service uses IX\_OSAL\_MBUFs to store data, or chains IX\_OSAL\_MBUFs together into chained IX\_OSAL\_MBUFs for large packets.

The client is responsible for allocating these buffers and passing the buffers to IxHssAcc.

An IX\_OSAL\_MBUF pool should be created for packetized service by calling function `IX_OSAL_MBUF_POOL_INIT()` of the `IxOsBuffMgt` API with the IX\_OSAL\_MBUF size and number of IX\_OSAL\_MBUF needed. For example:

```
IxHssAccCodeletMbufPool **poolIdPtr;

UINT32 numPoolMbufs;

UINT32 poolMbufSize;

*poolIdPtr = IX_OSAL_MBUF_POOL_INIT(numPoolMbufs, poolMbufSize,
"HssAcc Codelet Pool");
```

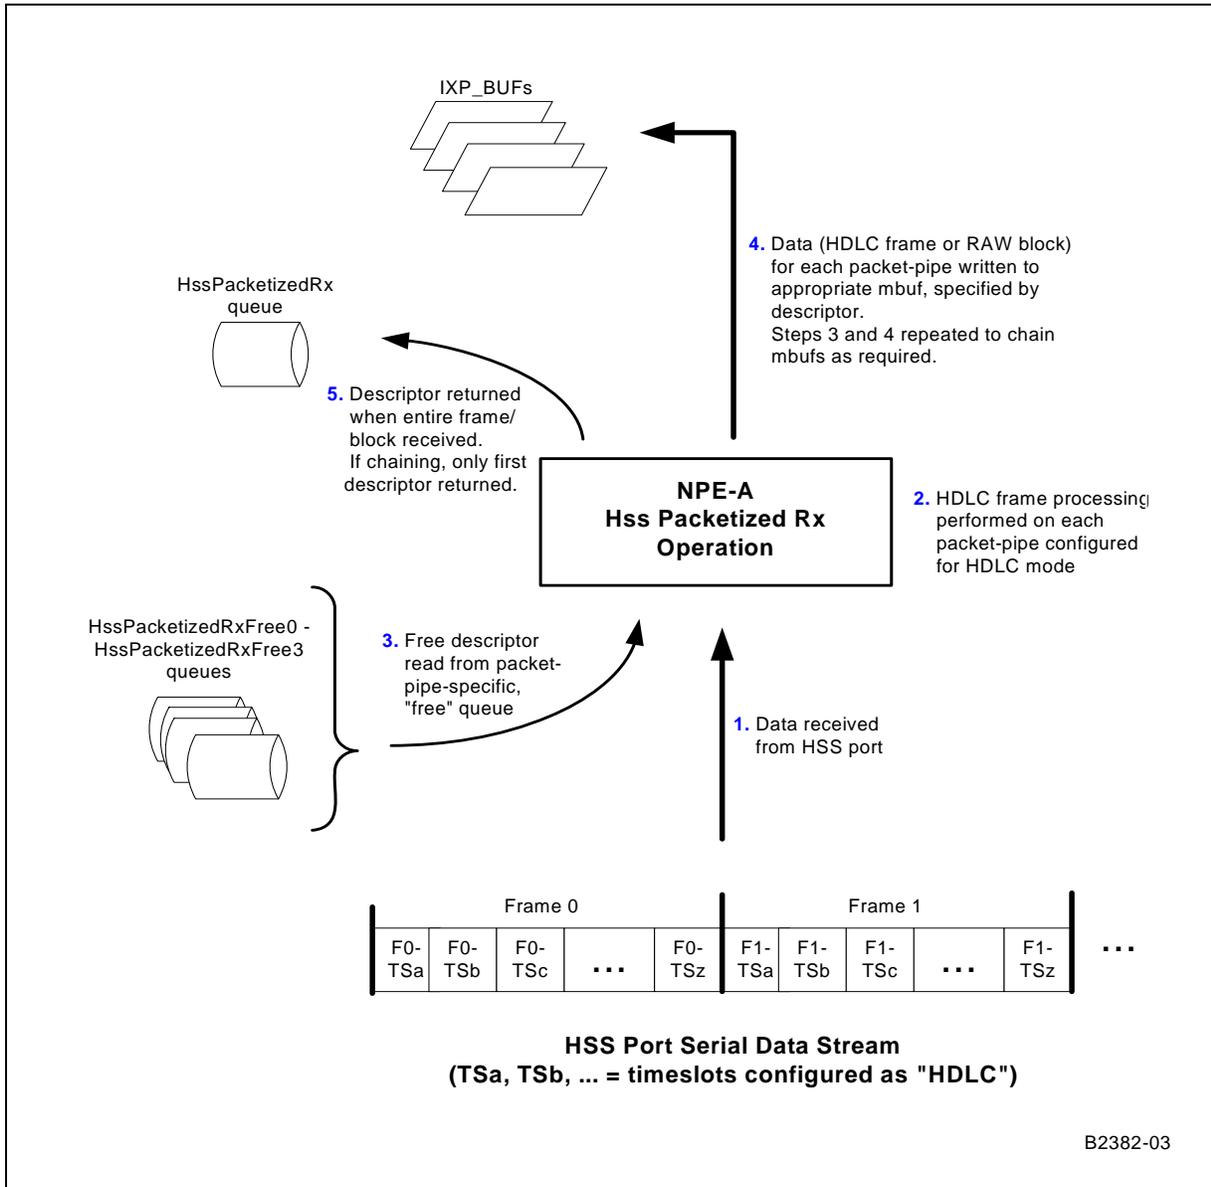
A IX\_OSAL\_MBUF can be obtained from the pool by calling `IX_OSAL_MBUF_POOL_GET()`. This Buffer pool is shared by the Tx and Rx processes.

For Rx, before the packetized service is enabled, the Rx buffer queue in IxHssAcc has to be replenished. This can be done by calling `ixHssAccPktPortRxFreeReplenish()`.

When packetized service starts, it is the client's responsibility to ensure there is always an adequate supply of IX\_OSAL\_MBUFs for the receive direction. This can be achieved in two ways. A call-back function can be registered with IxHssAcc to be called back when the free IX\_OSAL\_MBUFs queue is running low. This call back function is registered with the IxHssAcc packetized service when `ixHssAccPktPortConnect()` is called. Alternatively, the client can use its own timer to regularly supply buffers to the queue.

The client also provides a receive call-back function to accept packets received through the HSS. After the data in the IX\_OSAL\_MBUF is processed, IX\_OSAL\_MBUF\_POOL\_PUT\_CHAIN() can be called to put the Rx buffer back into the IX\_OSAL\_MBUF pool. The Rx packetized data flow is shown in Figure 71 on page 236.

Figure 71. HSS Packetized Receive Buffering

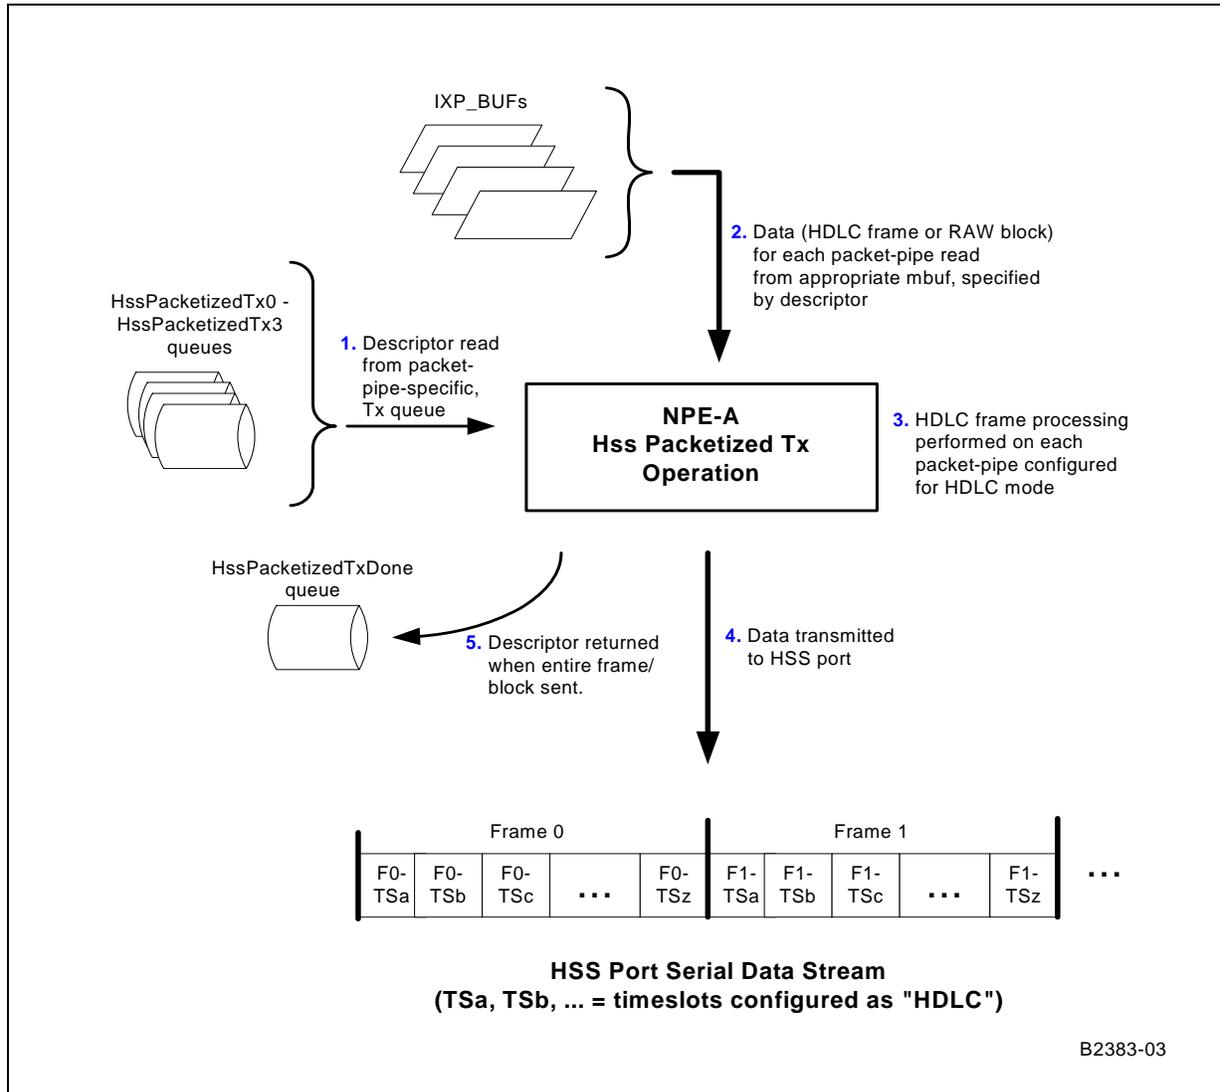


For Tx, buffers are allocated from the IX\_OSAL\_MBUF pool by calling IX\_OSAL\_MBUF\_POOL\_GET(). Data for transmitting can be put into the IX\_OSAL\_MBUF by using IX\_OSAL\_MBUF\_MDATA(). If the client data is too large to fit into one buffer, multiple buffers can be obtained from the pool and made into a chained IX\_OSAL\_MBUFs by using IX\_OSAL\_MBUF\_PKT\_LEN() and IX\_OSAL\_MBUF\_NEXT\_BUFFER\_IN\_PKT\_PTR(). The whole chained IX\_OSAL\_MBUF can be passed to IxHssAcc for transmission by calling ixHssAccPktPortTx().



A Tx callback function is also registered when `ixHssAccPktPortConnect()` is called before the service is enabled. When a chained IX\_OSAL\_MBUF is done with transmitting, the callback function is called and the buffers can be returned to the IX\_OSAL\_MBUF pool. The packetized Transmit data flow is described in Figure 72.

Figure 72. HSS Packetized Transmit Buffering



### 13.7.2 Data Flow in Channelized Service

Data in the timeslots configured as Voice64K/Voice56K types is provided to the client via the IxHssAcc channelized service. There are up to 32 such channels per HSS port. The channelized service uses memory that is shared between the Intel XScale® Processor and the NPEs. The client is responsible for allocating the memory for IxHssAcc to transmit and receive data through the HSS port.

For receive, `ixOsaiCacheDmaMalloc()` of the IxOSCacheMMU component can be used to create a pointer to a pool of contiguous memory from the shared memory of the Intel XScale® Processor and the NPEs. The pointer to this Rx data pool must be a physical



address because NPE directly writes data into this memory area. The memory pool is divided into  $N$  circular buffers, one buffer per channel.  $N$  is the total number of channels in service.

All the buffers have the same length. When the channelized service is initialized by `ixHssAccChanConnect()`, the pointer to the pool, the length of the circular buffers, and a parameter `bytesPerTStrigger` are passed to `IxHssAcc`, as well as a pointer to the `ixHssAccChanRxCallback()` Rx callback function.

Figure 73 shows how the circular buffers are filled with data received through the HSS ports. When each of the  $N$  channels receive `bytesPerTStrigger` bytes, the Rx callback function is called, and an offset value `rxOffset` is returned to indicate where data is written into the circular buffer. Note that `rxOffset` is shared for all the circular buffers in the pool. `rxOffset` is adjusted internally in the HSS component so that it is wrapped back to the beginning of the circular buffer when it reaches the end of the circular buffer.

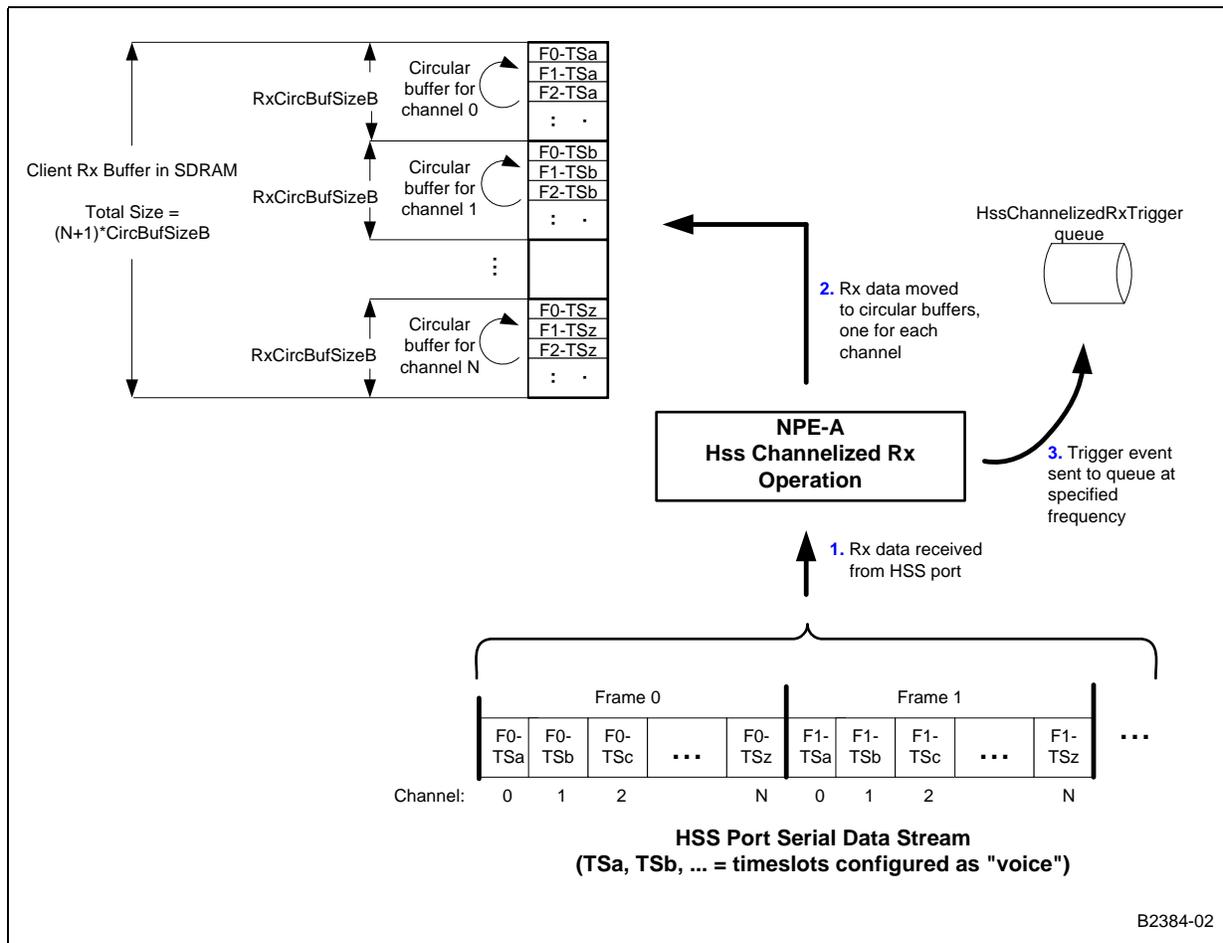
The client has to make sure the Rx data is processed or moved elsewhere before being overwritten by the HSS component. Hence the length of the circular buffers has to be chosen properly. The buffer need to be large enough for data to be read by the client and complete any possible in-place processing that would need to occur before the NPE rewrites over that memory. Understanding the client application's read and processing latency, the size of the data unit needed by the client application for processing, and the rate at which the NPE writes data to a buffer at a given channel rate, are useful in making this calculation.

In the case of Bypass mode, data flow for channelized Receive service is still the same but the data is transmitted to its partnering channel at NPE level.

Figure 73 on page 239 shows the data flow of the channelized Receive service.



Figure 73. HSS Channelized Receive Operation



For transmission, *ixOsaiCacheDmaMalloc()* is used to allocated two pools: a data buffer pool and a pointer list pool. The data buffer pool has  $N$  buffers — one for each channel. Each buffer is divided into  $K$  sections and each section has  $L$  bytes. The pointer list pool has  $K$  pointer lists. Each list has  $N$  pointers, each pointing to a section in a data buffer.

Before channelized service is enabled, the pointers must be initialized to point to the first section of each data buffer in the data buffer pool, and data for transmission is prepared and moved to the data buffer. The pointers to the data buffer pool and pointer list pool are passed to *IxHssAcc* when *ixHssAccChanConnect()* is called.

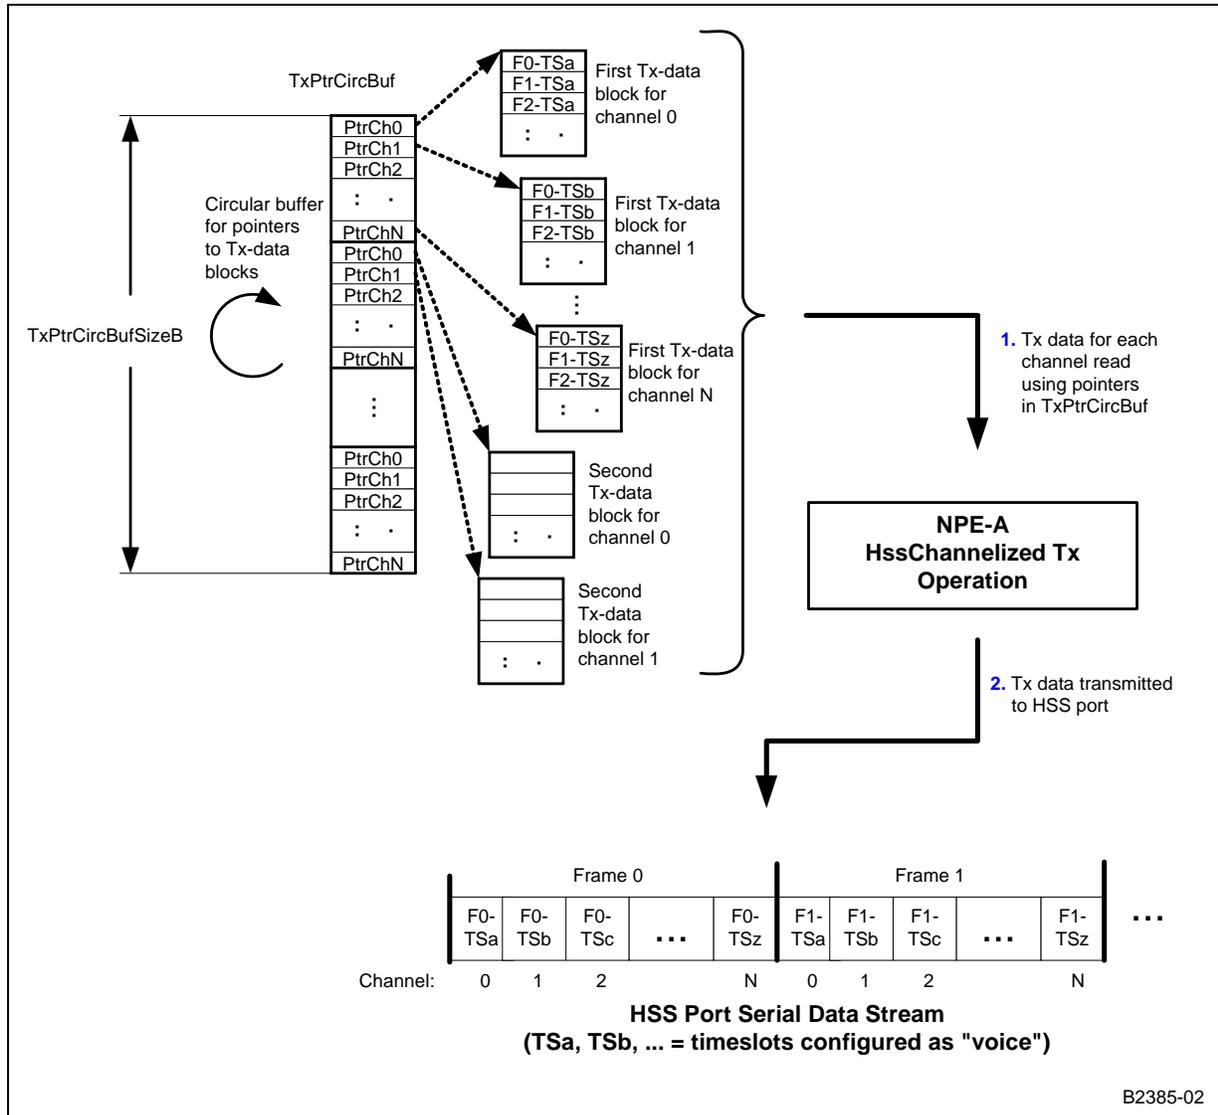
The client can check the current location of data being transmitted by using the registered *ixHssAccChanRxCallback()* function. When the Rx callback function is called, an offset value *txOffset* is returned.

*txOffset* indicates which pointer list in the pointer list pool is pointing to the sections of the data buffers currently being transmitted. Thus the client can use *txOffset* to determine where new data must be put into the data buffer pool for transmission. For example, data can be prepared and moved into sections pointed by the  $(txOffset-2)$ th pointer list. The length of the buffer,  $K * L$ , must be large enough so that the client has enough time to prepare data for transmission.



In the case of Bypass mode, there is no data path from the Intel XScale® Processor to NPE. The bypassed channel sources its data from its partnering channel to transmit at the NPE level to HSS port.

Figure 74. HSS Channelized Transmit Operation



§ §



## 14.0 Access-Layer Components: NPE-Downloader (IxNpeDI) API

---

This chapter describes the Intel® IXP400 Software v2.3's **NPE-Downloader API** access-layer component.

### 14.1 What's New

The *ixNpeDI* component provides a new API

***ixNpeDIDataMemRead()***: Reads one WORD from the DMEM of the NPE.

The following API was deprecated but it is reintroduced in *ixNpeDI* component.

***ixNpeDILoadedImageGet()***: Get the Id of the image currently loaded on a particular NPE.

Both of the above APIs are recommended to use only in the event of NPE soft- reset.

### 14.2 Overview

The NPE downloader (*ixNpeDI*) component is a stand-alone component providing a facility to download a microcode image to NPE A, NPE B, or NPE C in the system. The *ixNpeDI* component defines the default library of NPE microcode images. An NPE microcode image is provided to support each software release 2.3 release, and will contain up-to-date microcode for each NPE.

The *ixNpeDI* component also enables a client to supply a custom microcode image to use in place of the default images for each NPE. This **custom image** facility provides increased testability and flexibility, but is not intended for general use.

### 14.3 Microcode Images

All microcode images available for download to the NPEs are contained in a microcode image library. Each image contains a number of blocks of instruction, data, and state-information microcode that is downloaded into the NPE memory and registers. Each image also contains a download map that specifies how to extract the individual blocks of that image's microcode.

Given a microcode image library in memory, the NPE Downloader can locate images from that image library in memory, extract and interpret the contained download map, and download the code accordingly.

#### Loading NPE Microcode from a File Versus Loading from Memory

The NPE microcode library contains a series of NPE images. This microcode library can be compiled into the software release 2.3 object code at build time, or it can take the form of a single binary file. The method of operation is depending on type of operating system used.



The **Microcode from File** feature is only available for Linux\*. All other supported operating systems use obtain the NPE microcode from the compiled object code.

The purpose of providing the **Microcode from File** feature is to allow distribution of software release 2.3 and the NPE microcode under different licensing conditions for Linux\*. Refer to the *Intel® IXP400 Software Release 1.5 Software Release Notes* for further instructions on using this feature.

### NPE Microcode Library Customization

The NPE microcode library contains a series of NPE images. By default, all of these are included in the build. However, some of these images may not be required, and as such may be taking up excess memory. To omit one or more specific images, the user must edit `IxNpeMicrocode.h` and follow the instructions within. Essentially, by **undefining** an NPE image identifier, the corresponding NPE image is omitted from the overall build.

*Note:* If multiple image identifiers are provided for the same image, **all** of those identifiers need to be undefined to omit that image from the build.

### NPE Image Compatibility

The software releases do **not** include tools to develop NPE software. The supplied NPE functionality is accessible through the APIs provided by the software release 2.3 library. The NPE images are provided in the form of a single.C file. Corresponding NPE image identifier definitions are provided in an associated header file. Both are incorporated as part of the software release package.

NPE microcode images are assumed compatible for only the specific release they accompany.

## 14.4 Standard Usage Example

The initialization of an NPE has been made relatively easy. Only one function call is required.

Users call the `ixNpeDlNpeInitAndStart` function, which loads a specified image and begins execution on the NPE. Here is a sample function call, which starts NPE C with Ethernet and Crypto functionality:

```
ixNpeDlNpeInitAndStart(IX_NPEDL_NPEIMAGE_NPEC_CRYPTO_ETH);
```

The parameter is a UINT32 that is defined in the NPE image ID definition. [Table 60](#) to [Table 62](#) lists the parameters for the standard images.



**Table 60. NPE-A Images (Sheet 1 of 3)**

Image Name	Description
IX_NPEDL_NPEIMAGE_NPEA_HSSO	NPE Image ID for NPE-A with HSS-0 Only feature. It supports 32 channelized and 4 packetized.
IX_NPEDL_NPEIMAGE_NPEA_HSSO_ATM_SPHY_1_PORT	NPE Image ID for NPE-A with HSS-0 and ATM feature. For HSS, it supports <ul style="list-style-type: none"> <li>• 16/32 channelized</li> <li>• 4/0 packetized.</li> </ul> For ATM, it supports <ul style="list-style-type: none"> <li>• AAL 5,</li> <li>• AAL 0</li> <li>• OAM for UTOPIA SPHY,</li> <li>• 1 logical port</li> <li>• 32 VCs.</li> </ul>
IX_NPEDL_NPEIMAGE_NPEA_HSSO_ATM_MPHY_1_PORT	NPE Image ID for NPE-A with HSS-0 and ATM feature. For HSS, it supports <ul style="list-style-type: none"> <li>• 16/32 channelized</li> <li>• 4/0 packetized.</li> </ul> For ATM, it supports <ul style="list-style-type: none"> <li>• AAL 5,</li> <li>• AAL 0</li> <li>• OAM for UTOPIA MPHY,</li> <li>• 1 logical port</li> <li>• 32 VCs.</li> </ul>
IX_NPEDL_NPEIMAGE_NPEA_HSSO_ATM_MPHY_4_PORT.	NPE Image ID for NPE-A with HSS-0 and ATM feature. For HSS, it supports <ul style="list-style-type: none"> <li>• 16 channelized</li> <li>• 2 pairs of HSS bypass channels.</li> </ul> For ATM, it supports <ul style="list-style-type: none"> <li>• AAL 5,</li> <li>• AAL 0</li> <li>• OAM for UTOPIA MPHY</li> <li>• 4 logical port</li> <li>• 32 VCs</li> </ul>
IX_NPEDL_NPEIMAGE_NPEA_HSS_TSLOT_SWITCH	NPE Image ID for NPE-A with HSS-0 and timeslot switching feature. It supports 32 channelized and 4 packetized on HSS-0.
IX_NPEDL_NPEIMAGE_NPEA_ATM_MPHY_12_PORT	NPE Image ID for NPE-A with ATM-Only feature. It supports AAL 5, AAL 0 and OAM for UTOPIA MPHY, 12 logical ports, 32 VCs.
IX_NPEDL_NPEIMAGE_NPEA_HSS_2_PORT	NPE Image ID for NPE-A with HSS-0 and HSS-1 feature. Each HSS port supports 32 channelized and 4 packetized.
IX_NPEDL_NPEIMAGE_NPEA_HSSCHAN_ETH_MACFILTERLEARN_COEXIST	NPE Image ID for NPE-A with HSS-0 and Ethernet feature. For HSS, it supports <ul style="list-style-type: none"> <li>• 32 channelized and</li> <li>• 2 pairs of HSS bypass channels</li> </ul> For Ethernet, it supports <ul style="list-style-type: none"> <li>• MAC filtering / Port ID</li> <li>• MAC learning assist</li> <li>• Spanning tree support</li> <li>• Frame size filtering</li> <li>• Mask-based Firewall support</li> <li>• VLAN/QoS support and</li> <li>• Extended MIB-II</li> </ul>



Table 60. NPE-A Images (Sheet 2 of 3)

Image Name	Description
IX_NPEDL_NPEIMAGE_NPEA_HSS_ETH_HDRCONV_COEXIST	NPE Image ID for NPE-A with HSS-0 and Ethernet feature. For HSS, it supports <ul style="list-style-type: none"> <li>• 32 channelized and</li> <li>• 2 pairs of HSS bypass channels.</li> </ul> For Ethernet, it supports <ul style="list-style-type: none"> <li>• 802.3 &lt;=&gt; 802.11 header conversion</li> <li>• Spanning tree support</li> <li>• Frame size filtering</li> <li>• Mask-based Firewall support</li> <li>• VLAN/QoS support and</li> <li>• Extended MIB-II</li> </ul>
IX_NPEDL_NPEIMAGE_NPEA_DMA	NPE Image ID for NPE-A with DMA-Only feature.
IX_NPEDL_NPEIMAGE_NPEA_WEP	NPE Image ID for NPE-A with ARC4 and WEP CRC engines
IX_NPEDL_NPEIMAGE_NPEA_ETH	NPE Image ID for NPE-A with Ethernet-Only feature. This image definition is identical to the image below: IX_NPEDL_NPEIMAGE_NPEA_ETH_LEARN_FILTER_SPAN_FIREWALL.
IX_NPEDL_NPEIMAGE_NPEA_ETH_LEARN_FILTER_SPAN_FIREWALL	NPE Image ID for NPE-A with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• MAC_FILTERING</li> <li>• MAC_LEARNING</li> <li>• SPANNING_TREE</li> <li>• FIREWALL</li> </ul>
IX_NPEDL_NPEIMAGE_NPEA_ETH_LEARN_FILTER_SPAN_FIREWALL_VLAN_QOS	NPE Image ID for NPE-A with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• MAC_FILTERING</li> <li>• MAC_LEARNING</li> <li>• SPANNING_TREE</li> <li>• FIREWALL</li> <li>• VLAN/QoS</li> </ul>
IX_NPEDL_NPEIMAGE_NPEA_ETH_SPAN_FIREWALL_VLAN_QOS_HDR_CONV	NPE Image ID for NPE-A with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• SPANNING_TREE</li> <li>• FIREWALL</li> <li>• VLAN/QoS</li> <li>• 802.3/802.11 Frame Header Conversion</li> </ul>



**Table 60. NPE-A Images (Sheet 3 of 3)**

Image Name	Description
IX_NPEDL_NPEIMAGE_NPEA_ETH_LEARN_FILTER_SPAN_MASK_FIREWALL_VLAN_QOS_EXTMIB	NPE Image ID for NPE-A with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• MAC filtering / Port ID</li> <li>• MAC learning assist</li> <li>• Spanning tree support</li> <li>• Frame size filtering</li> <li>• Mask-based Firewall support</li> <li>• VLAN/QoS support</li> <li>• Extended MIB-II</li> <li>• Ethernet MAC Recovery Support</li> <li>• Ethernet MAC Address Filter Enhanced</li> </ul>
IX_NPEDL_NPEIMAGE_NPEA_ETH_SPAN_VLAN_QOS_HDR_CONV_EXTMIB	NPE Image ID for NPE-A with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• Spanning tree support</li> <li>• Frame size filtering</li> <li>• 802.3/802.11 header conversion</li> <li>• VLAN/QoS support</li> <li>• Extended MIB-II</li> </ul>
IX_NPEDL_NPEIMAGE_NPEA_ETH_SPAN_MASK_FIREWALL_VLAN_QOS_HDR_CONV_EXTMIB	NPE Image ID for NPE-A with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• Spanning tree support</li> <li>• Mask-based Firewall support</li> <li>• Frame size filtering</li> <li>• 802.3/802.11 header conversion</li> <li>• VLAN/QoS support</li> <li>• Extended MIB-II</li> <li>• Ethernet MAC recovery support</li> <li>• AHB Transfer Optimizer enabled</li> <li>• Ethernet MAC Address Filter enhanced</li> </ul>

**Table 61. NPE-B Images (Sheet 1 of 2)**

Image Name	Description
IX_NPEDL_NPEIMAGE_NPEB_DMA	NPE Image ID for NPE-B with DMA-Only feature.
IX_NPEDL_NPEIMAGE_NPEB_ETH	NPE Image ID for NPE-B with Ethernet-Only feature. This image definition is identical to the image below: IX_NPEDL_NPEIMAGE_NPEB_ETH_LEARN_FILTER_SPAN_FIREWALL.
IX_NPEDL_NPEIMAGE_NPEB_ETH_LEARN_FILTER_SPAN_FIREWALL	NPE Image ID for NPE-B with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• MAC_FILTERING</li> <li>• MAC_LEARNING</li> <li>• SPANNING_TREE</li> <li>• FIREWALL</li> </ul>
IX_NPEDL_NPEIMAGE_NPEB_ETH_LEARN_FILTER_SPAN_FIREWALL_VLAN_QOS	NPE Image ID for NPE-B with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• MAC_FILTERING</li> <li>• MAC_LEARNING</li> <li>• SPANNING_TREE</li> <li>• FIREWALL</li> <li>• VLAN/QoS</li> </ul>



**Table 61. NPE-B Images (Sheet 2 of 2)**

Image Name	Description
IX_NPEDL_NPEIMAGE_NPEB_ETH_SPAN_FIREWALL_VLAN_QOS_HDR_CONV	NPE Image ID for NPE-B with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• SPANNING_TREE</li> <li>• FIREWALL</li> <li>• VLAN/QoS</li> <li>• 802.3/802.11 Frame Header conversion</li> </ul>
IX_NPEDL_NPEIMAGE_NPEB_ETH_LEARN_FILTER_SPAN_MASK_FIREWALL_VLAN_QOS_EXTMIB	NPE Image ID for NPE-B with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• MAC filtering / Port ID</li> <li>• MAC learning assist</li> <li>• Spanning tree support</li> <li>• Frame size filtering</li> <li>• Mask-based firewall support</li> <li>• VLAN/QoS support</li> <li>• Extended MIB-II</li> <li>• Ethernet MAC Recovery support</li> <li>• Ethernet MAC Address Filter enhanced</li> </ul>
IX_NPEDL_NPEIMAGE_NPEB_ETH_SPAN_VLAN_QOS_HDR_CONV_EXTMIB	NPE Image ID for NPE-B with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• Spanning tree support</li> <li>• Frame size filtering</li> <li>• 802.3/802.11 header conversion</li> <li>• VLAN/QoS support</li> <li>• Extended MIB-II</li> </ul>
IX_NPEDL_NPEIMAGE_NPEB_ETH_SPAN_MASK_FIREWALL_VLAN_QOS_HDR_CONV_EXTMIB	NPE Image ID for NPE-B with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• Spanning tree support</li> <li>• Mask-based Firewall support</li> <li>• Frame size filtering</li> <li>• 802.3/802.11 header conversion</li> <li>• VLAN/QoS support</li> <li>• Extended MIB-II</li> <li>• Ethernet MAC recovery support</li> <li>• AHB Transfer Optimizer enabled</li> <li>• Ethernet MAC Address Filter enhanced</li> </ul>

**Table 62. NPE-C Images (Sheet 1 of 3)**

Image Name	Description
IX_NPEDL_NPEIMAGE_NPEC_DMA	NPE Image ID for NPE-C with DMA-Only feature.
IX_NPEDL_NPEIMAGE_NPEC_ETH	NPE Image ID for NPE-C with Eth-Only feature. This image definition is identical to the image below: IX_NPEDL_NPEIMAGE_NPEC_CRYPT0_ETH_LEARN_FILTER_SPAN_FIREWALL.
IX_NPEDL_NPEIMAGE_NPEC_CRYPT0_ETH_LEARN_FILTER_SPAN_FIREWALL	NPE Image ID for NPE-C with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• MAC_FILTERING</li> <li>• MAC_LEARNING</li> <li>• SPANNING_TREE</li> <li>• FIREWALL</li> </ul>



**Table 62. NPE-C Images (Sheet 2 of 3)**

Image Name	Description
IX_NPEDL_NPEIMAGE_NPEC_ETH_LEARN_FILTER_SPAN_FIREWALL_VLAN_QOS	NPE Image ID for NPE-C with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• MAC_FILTERING</li> <li>• MAC_LEARNING</li> <li>• SPANNING_TREE</li> <li>• FIREWALL</li> <li>• VLAN/QoS</li> </ul>
IX_NPEDL_NPEIMAGE_NPEC_ETH_SPAN_FIREWALL_VLAN_QOS_HDR_CONV	NPE Image ID for NPE-C with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• SPANNING_TREE</li> <li>• FIREWALL</li> <li>• VLAN/QoS</li> <li>• 802.3/802.11 Frame Header conversion</li> </ul>
IX_NPEDL_NPEIMAGE_NPEC_ETH_LEARN_FILTER_SPAN_MASK_FIREWALL_VLAN_QOS_EXTMIB	NPE Image ID for NPE-C with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• MAC filtering / Port ID</li> <li>• MAC learning assist</li> <li>• Spanning tree support</li> <li>• Frame size filtering</li> <li>• Mask-based Firewall support</li> <li>• VLAN/QoS support</li> <li>• Extended MIB-II</li> <li>• Ethernet MAC Recovery support</li> <li>• Ethernet MAC Address Filter enhanced</li> </ul>
IX_NPEDL_NPEIMAGE_NPEC_ETH_SPAN_VLAN_QOS_HDR_CONV_EXTMIB	NPE Image ID for NPE-C with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• Spanning tree support</li> <li>• Frame size filtering</li> <li>• 802.3/802.11 header conversion</li> <li>• VLAN/QoS support</li> <li>• Extended MIB-II</li> </ul>
IX_NPEDL_NPEIMAGE_NPEC_ETH_SPAN_MASK_FIREWALL_VLAN_QOS_HDR_CONV_EXTMIB	NPE Image ID for NPE-C with Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• Spanning tree support</li> <li>• Mask-based Firewall support</li> <li>• Frame size filtering</li> <li>• 802.3/802.11 header conversion</li> <li>• VLAN/QoS support</li> <li>• Extended MIB-II</li> <li>• Ethernet MAC recovery support</li> <li>• AHB Transfer Optimizer enabled</li> <li>• Ethernet MAC Address Filter enhanced</li> </ul>
IX_NPEDL_NPEIMAGE_NPEC_CRYPTO_ETH_LEARN_FILTER_SPAN_FIREWALL	NPE Image ID for NPE-C with Basic Crypto and Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>• MAC_FILTERING</li> <li>• MAC_LEARNING</li> <li>• SPANNING_TREE</li> <li>• FIREWALL</li> </ul> For Crypto, it supports DES, SHA-1, SHA-256, SHA384, SHA-512 and MD5.



Table 62. NPE-C Images (Sheet 3 of 3)

Image Name	Description
IX_NPEDL_NPEIMAGE_NPEC_CRYPT0_AES_ETH_LEARN_FILTER_SPAN	NPE Image ID for NPE-C with AES Crypto and Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>MAC_FILTERING</li> <li>MAC_LEARNING</li> <li>SPANNING_TREE</li> </ul> For Crypto, it supports AES, DES, SHA-1, SHA-256, SHA384, SHA-512 and MD5. AES-CCM mode is not supported.
IX_NPEDL_NPEIMAGE_NPEC_CRYPT0_AES_ETH_LEARN_FILTER_FIREWALL	NPE Image ID for NPE-C with AES Crypto and Basic Ethernet Rx/Tx, which includes: <ul style="list-style-type: none"> <li>MAC_FILTERING</li> <li>MAC_LEARNING</li> <li>FIREWALL</li> </ul> For Crypto, it supports AES, DES, SHA-1, SHA-256, SHA384, SHA-512 and MD5. AES-CCM mode is not supported.
IX_NPEDL_NPEIMAGE_NPEC_CRYPT0_AES_CCM_ETH	NPE Image ID for NPE-C with AES & AES-CCM Crypto and Basic Ethernet Rx/Tx. For Crypto, it supports AES, CCM, DES, SHA-1, SHA-256, SHA384, SHA-512 and MD5.

### 14.5 Custom Usage Example

Using a custom image is the second option for starting an NPE. This feature is only useful to those parties that have NPE microcode development capabilities, and thus does not apply to most users. The majority of users will use the Intel-provided NPE library.

This allows the use of an external library of images, if needed. External libraries come in the form of a header file. The header file defines the image library as a single array of type UINT32, and it is that array symbol to be used as a **imageLibrary** parameter for the function.

The function used for this procedure is as follows:

```
ixNpeDIcustomImageNpeInitAndStart(UINT32 *imageLibrary, UINT32 npeImageId);
```

### 14.6 IxNpeDI Uninitialization

After the first NPE has been started using one of the above methods, *IxNpeDI* is initialized and the specified NPEs will begin execution.

The *IxNpeDI* should be uninitialized prior to unloading an application module or driver. (This will unmap all memory that has been mapped by *IxNpeDI*.) If possible, *IxNpeDI* should be uninitialized before a soft reboot.

Here is a sample function call to uninitialized *IxNpeDI*:

```
ixNpeDIUnload();
```

*Note:* Calling *ixNpeDIUnload* twice or more in succession will cause all subsequent calls after the first one to exit harmlessly but returns a FAIL status.



## 14.7 Deprecated APIs

The functions listed below have been deprecated and may be removed from a future software release of this component. Additionally, the functions listed below will not work with the new microcode image format provided in software release 2.3. As of software release 1.3, the functions *ixNpeDINpeInitAndStart* and *ixNpeDICustomImageNpeInitAndStart* have replaced the functions listed below:

- *ixNpeDIImageDownload*
- *ixNpeDIAvailableImagesCountGet*
- *ixNpeDIAvailableImagesListGet*
- *ixNpeDILatestImageGet*
- *ixNpeDIMicrocodeImageLibraryOverride*

§ §





## 15.0 Access-Layer Components: NPE Message Handler (IxNpeMh) API

---

This chapter describes the Intel® IXP400 Software v2.3's "NPE Message Handler API" access-layer component.

### 15.1 What's New

IxNpeMh component provides a new API to re-update or restores the NPE core's OutFifo/InFifo interrupt enabling configuration after an event of an NPE soft reset. This API is called by the soft error recovery task during the re-initialization of the NPE after a NPE soft reset.

**ixNpeMhConfigStateRestore: Re-initialize the NPE after a NPE soft reset.**

### 15.2 Overview

This chapter contains the necessary steps to start the NPE message-handler component. Additionally, information has been included about how the Message Handler functions from a high-level view.

This component acts a pseudo service layer to other access components such as IxEthAcc. In the sections that describe how the messaging works, the "client" is an access component such as IxEthAcc. An application programmer will not need to do any coding to directly control message handling, just the initialization and uninitialization of the component.

The IxNpeMh component is responsible for sending messages from software components on the Intel XScale® Processor to the three NPEs (NPE A, NPE B, and NPE C). The component also receives messages from the NPEs and passes them up to software components on the Intel XScale® Processor. This encapsulates the details of NPE messaging in one place and provides a consistent approach to NPE messaging across all components. Message handling is a collaboration of Intel XScale® Processor software (IxNpeMh) and the NPE software.

When sending a message that solicits a response from the NPE, the client must provide a callback to the IxNpeMh component to hand the response back. The client should also register appropriate callbacks with the IxNpeMh component to handle unsolicited messages from the NPE.

The IxNpeMh component relies on the IDs of solicited and unsolicited messages to avoid "over-lapping" and determine if a received message is solicited or unsolicited.

Each NPE has two associated data structures — one for unsolicited message callbacks and another for solicited message callbacks.

Messages are sent to the NPEs inFIFOs, while messages are received from the NPEs outFIFOs. Both the inFIFO and outFIFO have a depth of two messages, and each messages is two words in length.



When sending a message that solicits a response, the solicited callback is added to the end of the list of solicited callbacks. For solicited messages, the first ID-matching callback in the solicited callback list is removed and called. For unsolicited messages, the corresponding callback is retrieved from the list of unsolicited callbacks.

The solicited callback list contains the list of callbacks corresponding to solicited messages not yet received from the NPE. The solicited messages for a given ID are received in the same order that those soliciting messages are sent, and the first ID-matching callback in the list always corresponds to the next solicited message that is received.

## 15.3 Initializing the IxNpeMh

The IxNpeMh has two modes of operation, interrupted or polled. This refers to whether the IxNpeMh will attach a message handler to service the 'outFIFO not empty' interrupts from the NPEs. If a message handler is not attached then the client must use **ixNpeMhMessagesReceive()** to control message receiving and processing.

### 15.3.1 Interrupt-Driven Operation

This is the preferred method of operation for the message handler. Here is a sample function call to initialize the IxNpeMh component for interrupt driven operation:

```
ixNpeMhInitialize (IX_NPEMH_NPEINTERRUPTS_YES);
```

The function takes a yes/no value from an enum, and all messages from all the NPEs is serviced by IxNpeMH. The IxNpeMh handles messages from all NPEs and should only be initialized once.

### 15.3.2 Polled Operation

Here is a sample function call to initialize the IxNpeMh component for polling operation:

```
ixNpeMhInitialize (IX_NPEMH_NPEINTERRUPTS_NO);
```

The function takes a yes/no value from an enum, and all messages from the NPEs must be manually checked. The IxNpeMh handles messages from all NPEs, and should only be initialized once.

After setting up polled operation the client **MUST** check for messages coming out of the NPEs in a loop fashion. Here is a sample function call that will check to see if NPE A has a message to send.

```
ixNpeMhMessagesReceive (IX_NPEMH_NPEID_NPEA);
```

Three separate function calls are required to check all three of the NPEs.

*Note:*

This function call cannot be made from inside an interrupt service routine as it will use resource protection mechanisms.

## 15.4 Uninitializing IxNpeMh

The IxNpeMh should be uninitialized prior to unloading a kernel module (this will unmap all memory that has been mapped by IxNpeMh). If possible, IxNpeMh should also be uninitialized before a soft reboot.



Here is a sample function call to uninitialized IxNpeMh:

```
ixNpeMhUnload();
```

*Note:* IxNpeMh can only be initialized from an uninitialized state and can only be uninitialized from an initialized state. If this order is not followed, for example by uninitialized an uninitialized IxNpeMh, then unpredictable behavior will result. Calling any other IxNpeMh API functions after unloading will also cause unpredictable results.

## 15.5 NPE Parity Error Handling

The IxNpeMh plays a part in the event of NPE soft-reset. It restores the interrupt enabling configuration for the outFIFO/inFIFO of the NPE used if the IxNpeMh is set to an interrupt mode during initialization. This API must be called once the NPE is reset which may occur during an NPE parity error.

```
ixNpeConfigStateRestore(IxNpeMhNpeId npeId);
```

It is not recommended to use this function other than in the event of NPE soft-reset.

## 15.6 Sending Messages from an Intel XScale® Processor Software Client to an NPE

Access-layer components — such as ixEthAcc and ixHssAcc — do all of their own message handling. This section describes the process of how messages are sent and processed so someone who is using IxNpeMh can understand what is going on in the background and gain insight into some performance issues.

There are two types of messages to send to an NPE: unsolicited and solicited. The first is just a simple message — that is, all it does is send a block of data. The second type sends data, but also registers a function to handle a response from the NPE.

The following sections give an overview of the process.

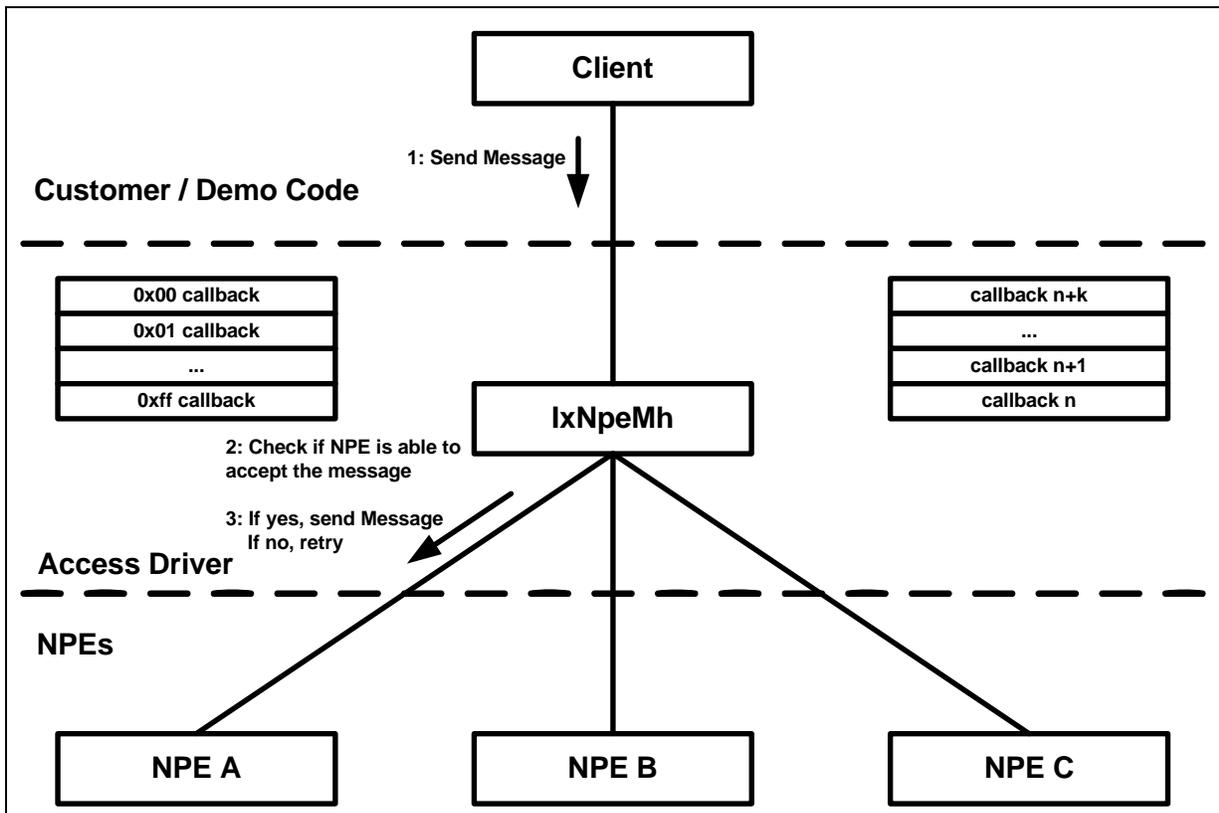
### 15.6.1 Sending an NPE Message

The scenario of sending a messages from an Intel XScale® Processor software client to an NPE (as shown in [Figure 75](#)) is:

1. The client sends a message to the IxNpeMh component, specifying the destination NPE.
2. The IxNpeMh component checks if the NPE is able to accept a message.
3. If yes, the IxNpeMh component sends the message to the NPE. If not, the specified number of retries is attempted by the IxNpeMh to repeat step 2. If the maximum number of retries has achieved, IxNpeMh will return fail to the client.

*Note:* ixNpeMh will recheck the NPE inFIFO status if NPE inFIFO is full during its first attempt. ixNpeMh will continue its attempts until the specified number of retries has been achieved. The action of rapidly messaging the NPE will consume the AHB bandwidth hence a specified task delay is used in between the retries. The number of retries is passed as a parameter to the send function; the default value is 3 retries.

Figure 75. Message from Intel XScale® Processor Software Client to an NPE



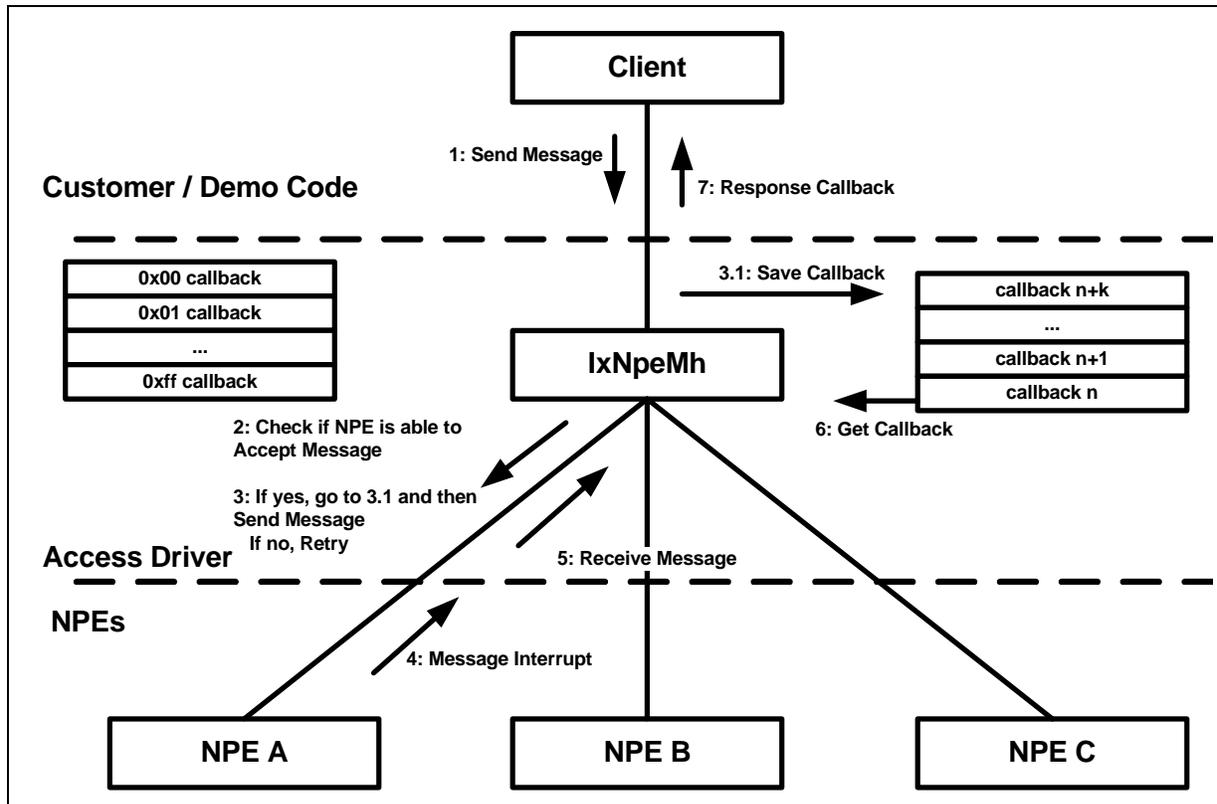
### 15.6.2 Sending an NPE Message with Response

In this case, the client's message requires a response from the NPE. The scenario (as shown in Figure 76) is:

1. The client sends a message to the IxNpeMh component, specifying the destination NPE and a response callback.
2. The IxNpeMh component checks if the NPE is able to accept a message. If the NPE cannot accept a message after the maximum number of retries, the send fails.
3. The IxNpeMh component adds the response callback to the end of the solicited callback list and sends the message to the NPE.
4. After some time, the NPEs "outFIFO not empty" interrupt invokes the IxNpeMh component's ISR.
5. Within the ISR, the IxNpeMh component receives a message from the specific NPE.
6. The IxNpeMh component checks if this message ID has an unsolicited callback registered for it. If the message has an unsolicited callback registered, the message is unsolicited. (See Section 15.7, "Receiving Unsolicited Messages from an NPE to Client Software" on page 255.)
7. Because this is a solicited message, the first ID-matching callback is removed from the solicited callback list and invoked to pass the message back to the client. If no ID-matching callback is found, the message is discarded and an error reported.



Figure 76. Message with Response from Intel XScale® Processor Software Client to an NPE

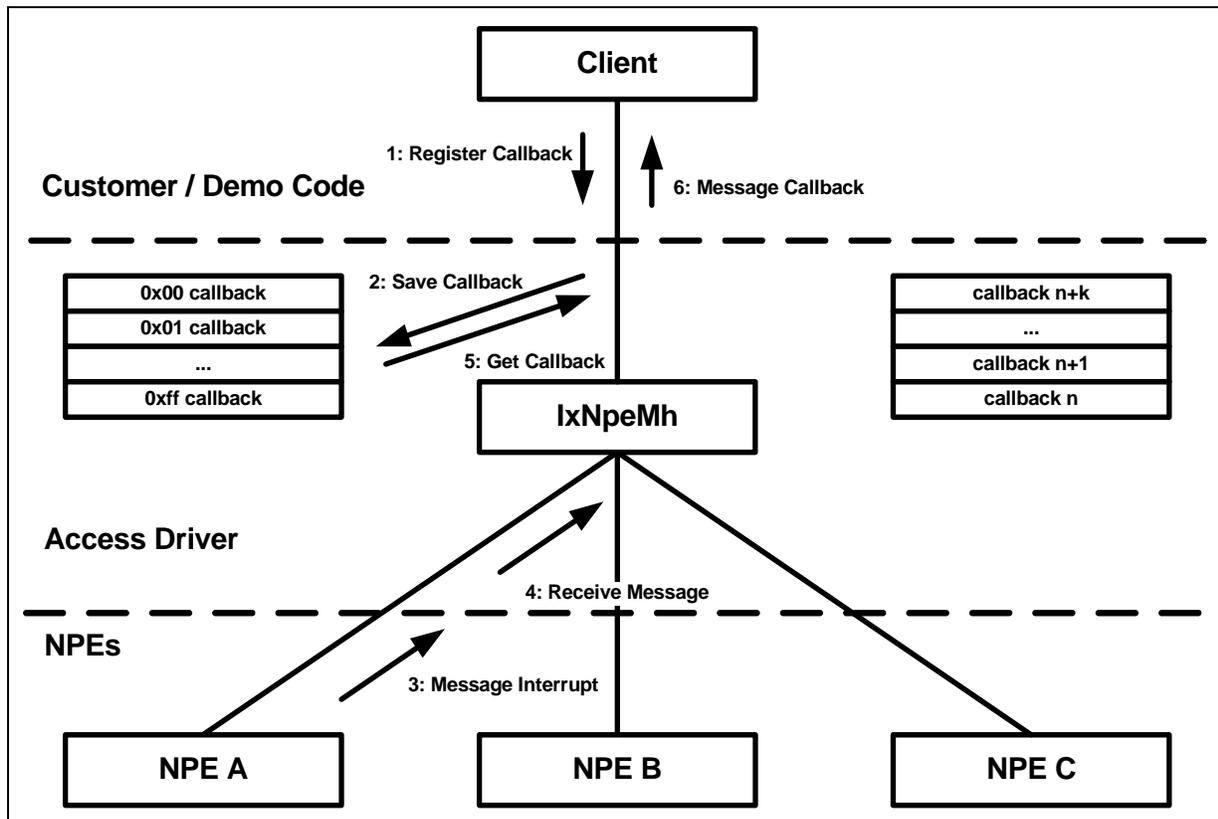


## 15.7 Receiving Unsolicited Messages from an NPE to Client Software

The scenario of receiving unsolicited messages from an NPE to client software (as shown in Figure 77) is:

1. At initialization, the client registers an unsolicited callback for a particular NPE and a message ID
2. This call back is then saved in the unsolicited callback list.
3. After some time, the NPEs “outFIFO not empty” interrupt invokes the IxNpeMh component’s ISR.
4. Within the ISR, the IxNpeMh component receives a message from the specific NPE.
5. The IxNpeMh component determines if this message ID has an unsolicited callback registered for it.  
If the message ID does not have a registered unsolicited callback, the message is solicited. (See “Sending an NPE Message with Response” on page 254.)
6. Since this is an unsolicited message, the IxNpeMh component invokes the corresponding unsolicited callback to pass the message back to the client.

Figure 77. Receiving Unsolicited Messages from NPE to Software Client



When ixNpeMh component receives a message from NPE, it will determine if this message ID has an unsolicited callback registered for it. If the ixNpeMh component can not find the corresponding callback, it will then search in the solicited call back list. If the ixNpeMh component still can not find the corresponding callback, it will return a warning message and the message from the NPE is discarded. This may either indicate that the message has not registered with appropriate callback or client is not interested in the response received.

The IxNpeMh component does not interpret message IDs. It only uses message IDs for comparative purposes, and for passing a received message to the correct callback function. This makes the IxNpeMh component immune to changes in message IDs.

The IxNpeMh component relies on the message ID being stored in the most-significant byte of the first word of the two-word message (IxNpeMhMessage).

*Note:* It is the responsibility of the client to create messages in the format expected by the NPEs.

Multiple clients may use the IxNpeMh component. Each client should take responsibility for handling its own range of unsolicited message IDs. (See the ixNpeMhUnsolicitedCallbackRegister.)

The IxNpeMh component handles messaging for the three NPEs independently. A problem or delay in interacting with one NPE will not impact interaction with the other NPEs.

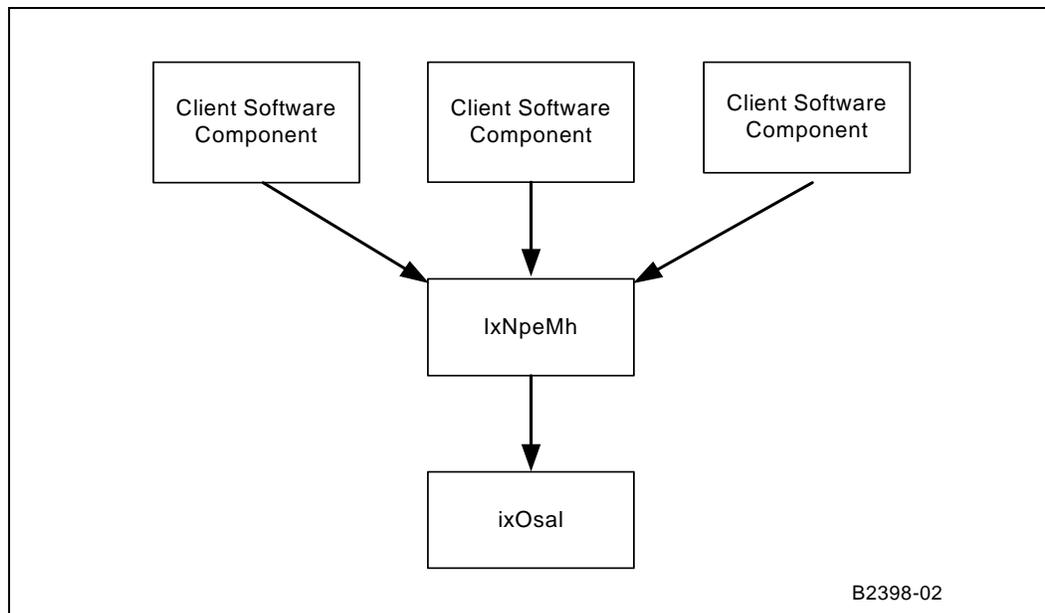


## 15.8 Dependencies

The IxNpeMh component’s dependencies (as shown in Figure 78) are:

- Client software components must use the IxNpeMh component for messages to and from the NPEs.
- The IxNpeMh component must use IxOSAL for error-handling, resource protection, and registration of ISRs.

Figure 78. ixNpeMh Component Dependencies



## 15.9 Error Handling

The IxNpeMh component uses IxOSAL to report errors and warnings. Parameters passed to the IxNpeMh component are error-checked whenever possible. Interface functions of the IxNpeMh component return a status to the client, indicating success or failure.

The most important error scenarios — when using the IxNpeMh — are:

- Failure to send a message if the NPE is unable to accept one. This failure implies there is a problem with the NPE. Failure to receive a message means the message is discarded.
- Failure to receive a message if no suitable callback can be found. To avoid message loss, clients should ensure that unsolicited callbacks are registered for all unsolicited message types.

The IxNpeMh component will return critical error status (timeout status) if it fails to perform the requested operation due to NPE hang. Upon receiving the timeout status, client application will need to ensure all the components and NPE return to normal / initial state before calling any of IxNpeMh APIs for that particular NPE.







## 16.0 Access-Layer Components: Parity Error Notifier (IxParityENAcc) API

---

This chapter describes Intel® IXP400 Software v2.3's "Parity Error Notifier (IxParityENAcc) API" access-layer component.

### 16.1 What's New

2 new APIs have been added to ixParityENAcc to support the NPE soft-reset feature.

***ixParityENAccParityNPEConfigReUpdate ()***: Allows the client to restore and update the NPE Error enabling configuration.

***ixParityENAccNPEParityErrorCheck ()***: Allows the client to check for any NPE specified error status.

### 16.2 Introduction

Many components in the IXP46X network processors provide parity error detection capabilities. These include:

- Instruction and Data Memory of the Network Processing Engines (NPEs)
- Switching Coprocessor in NPE B (SWCP)
- AHB Queue Manager SRAM (AQM)
- PCI Controller
- Expansion Bus Controller
- DDR SDRAM Memory Controller Unit (MCU). Additionally, the MCU on the IXP46X network processors provides Error Correction Code capabilities.

The IxParityENAcc access-layer component allows a client application to configure and enable/disable the parity error detection the blocks listed above on the Intel® IXP46X Product Line of Network Processors. It enables a client application to receive notification when a parity error is detected, along with information on the type and source of the error.

#### 16.2.1 Background

The processor or its external memory could be operating in an environment where bits in memory may be corrupted by electromagnetic radiation. All the above-mentioned blocks can be affected by unexpected corruptions. Errors that are not the result of a permanent hardware error, but are encountered as random errors in the state of individual memory cells, are called "soft errors". Parity and ECC are mechanisms to detect and provide corrective or restorative action from these soft errors.



For the purposes of this document, the following terms is used as defined below.

#### **Error Correction Logic/Error Correction Code**

The Error Correction Logic in the Memory Controller Unit (MCU) generates the ECC code (which requires additional bits for the code word) for DDR SDRAM reads and writes. For reads, this logic compares the ECC code read with the locally generated ECC code. If the codes do not match, then the Error Correction Logic determines the error type. For a single-bit error, this block determines which bit resulted in the error and corrects the error before the data is presented onto the bus. However, the error still remains in the memory location and must be fixed by writing the corrected data to the memory location. For writes, ECC logic in the MCU generates the ECC and sends it with the data to the memory.

#### **Scrubbing/Memory Scrub**

Scrubbing is the process of correcting an error in a memory location. When the MCU detects an error during a read, the MCU logs the address where the error occurred and interrupts the Intel XScale® Processor. The Intel XScale® Processor must then write back to the memory location to fix the error through a software handler. Note that the scrub rectifies only single-bit parity errors detected by the DDR MCU.

#### **Parity Error Context**

This refers to the type of the parity error, the source of the parity error (for example, the block which has the parity error) and the address of the failed word where applicable. The IxParityENAcc API provides a Parity Error Context to the client application when a parity or ECC error is detected.

#### **Parity and Error Correction Code**

Parity error detection is a simple and reliable mechanism to detect a single-bit error in a memory location. In general this mechanism is implemented by using an additional single bit along with the data bits in a memory location so that the bits set are of even/odd number and there is an even/odd number of '1's in the memory location. The MCU hardware will also detect multiple bit errors, but cannot detect whether the MCU is configured for odd or even parity.

## **16.2.2 Parity and ECC Capabilities in the Intel® IXP45X and Intel® IXP46X Product Line of Network Processors**

The IXP46X network processors can detect a variety of parity or ECC errors. The individual hardware blocks raise an interrupt to notify the Intel XScale® Processor about these failures. The interrupt controller on IXP46X network processors has a set of interrupts classified as 'error' class. These interrupts take unconditional high-priority from the normal positional priority interrupts. This section summarizes the interrupt behavior as it applies when a parity or ECC error is detected.

*Note:* For detailed information regarding the specific parity and ECC capabilities and interrupt mechanisms of IXP46X network processors, refer to the *Intel® IXP45X and Intel® IXP46X Product Line of Network Processors Developer's Manual*.

### **16.2.2.1 Network Processing Engines**

The NPE will lock and cease to operate immediately when affected by a parity error in its internal memories or due to external errors in coprocessors (AHB Coprocessor, or Switching Coprocessor). External NPE ports is disabled. An interrupt is sent to the Intel



XScale® Processor through the interrupt controller, and the parity context will provide information on whether the interrupt is related to internal memory parity errors or an external coprocessor error.

#### 16.2.2.2 Switching Coprocessor in NPE B (SWCP)

The Switching Coprocessor generates 8-bit parity – 1 bit per each byte of the 64 bit (8-byte) entries in its SRAM. These parity bits is generated and captured along with the 64 bits of data during a write operation. The subsequent read operation will again generate parity bits from the 64 bits of data and compare against the ones stored. If there is a mismatch, an interrupt is issued to the Intel XScale® Processor through the interrupt controller.

A parity error in this component would also generate an NPE-B interrupt as an external coprocessor error.

#### 16.2.2.3 AHB Queue Manager (AQM)

The AQM, on identifying a parity error from its internal memory, will return an ‘AHB Error’ response on the AHB bus to the requesting master. The interrupt context then refers to the address of either a queue entry or queue configuration entry, whose access resulted in failure. For queue entry address cases, the client application should treat the queue entry as invalid. The client should respond to a queue configuration parity error by rendering the entire queue invalid.

#### 16.2.2.4 DDR SDRAM Memory Controller Unit (MCU)

When the MCU detects a single-bit error, the word is corrected before it is delivered so that the Intel XScale® Processor gets a correct copy of the defective memory location contents (which still contains the uncorrected value). For multiple-bit errors, no correction is possible and an error response is placed on the bus visible to the Intel XScale® Processor. In either case the interrupt context refers to the address of the access that failed. The MCU keeps track of two such parity errors at any point in time and notifies of an overflow if more than two parity errors occurred at the same time, in which case the address will not be logged.

#### 16.2.2.5 Expansion Bus Controller

The Expansion Bus Controller, upon receiving a parity error on the Expansion Bus, terminates the transaction and responds on the South AHB bus with an “AHB Error” response for an outbound read initiated by an internal master. It will respond similarly in situations where an inbound write is initiated by an external master. It then provides an interrupt to the Intel XScale® Processor with a context containing a reference to the address of the access that contained the invalid data.

#### 16.2.2.6 PCI Controller

The PCI Controller will send an interrupt to the Intel XScale® Processor upon detecting a parity error in the following scenarios:

- read and write data transfers from AHB devices to PCI
- write data transfer from PCI to AHB devices.

For a read transaction initiated from PCI onto AHB, the MCU would detect any parity errors and send an interrupt to the Intel XScale® Processor.



### 16.2.2.7 Secondary Effects of Parity Interrupts

If the Intel XScale® Processor detects an error on the AHB bus or on its private DDR memory interface (MPI), an exception is generated that is serviced by its fault handler (such as data abort, or prefetch abort exception handler). The MCU will also generate a parity interrupt in this case.

**Caution:** There is no guarantee as to the arrival order at the Intel XScale® Processor of the data abort notification versus the parity interrupt. The client application should respond accordingly. For guidance in resolving the race condition between the data abort and the interrupt, refer to the scenarios described in “Parity Error Notification Detailed Scenarios” on page 267.

The AHB-AHB bridge unit, upon receiving an “AHB Error” from the South AHB caused by a parity error from a South AHB device, will respond with an AHB error to the originating master (an NPE) on the North AHB. The AHB Coprocessor in the NPE will abort the transaction and assert an error condition to the NPE, which will cause the NPE to lock up. This will result in an NPE external coprocessor interrupt event to the Intel XScale® Processor, as described in “Network Processing Engines” on page 260.

An NPE will report an ‘external’ error in the situation described above even though the chain of events started with a parity error on a South AHB device (such as an AQM, Expansion Bus Controller).

### 16.2.3 Interrupt Prioritization

Table 63 shows the list of interrupts that the Intel XScale® Processor would receive in the event of a parity error. IxParityENAcc applies only the software defined priority as indicated; the top priority being the priority 0 of the MCU.

**Table 63. Parity Error Interrupts**

Interrupt Bit <sup>1</sup>	Default Priority <sup>2</sup>	Software Defined Priority <sup>3</sup>	Source	Description
Int0	0	1	NPE-A	IMEM, DMEM or External Errors
Int1	1	2	NPE-B	IMEM, DMEM or External Errors <sup>4</sup>
Int2	2	3	NPE-C	IMEM, DMEM or External Errors
Int8	8	6	PCI	PCI Interrupt <sup>5</sup>
Int58	58	4	SWCP	Switching Coprocessor Interrupt <sup>4</sup>
Int60	60	5	AQM	AHB Queue Manager Interrupt
Int61	61	0	MCU	Single or Multi-Bit ECC Error. Multi-bit is serviced first in IxParityENAcc.
Int62	62	7	EXP	Expansion Bus Parity Error

- Notes:**
1. Interrupts 32-61 are higher-priority (error class) interrupts than 0-31. For example, MCU interrupt will take priority over NPE...even though the “Default Priority” table suggests otherwise.
  2. The interrupt controller applies the default priorities and accordingly asserts the parity error interrupts to the Intel XScale® Processor.
  3. The software defined priority is implemented by the access layer and is predefined.
  4. A SWCP interrupt is also seen as an NPE-B external interrupt.
  5. PCI Interrupts are those generated by the PCI Interrupt controller, and not the PCI Interrupt lines A,B,C and D.



## 16.3 IxParityENAcc API Details

### 16.3.1 Features

The parity error access component provides the following features:

- Interface to the client application to register a call back handler for application-specific processing with respect to the source of failure in the notification.
- Interface to the client application to individually enable and disable parity detection in the following hardware blocks, which are capable of generating parity errors. This interface can be invoked multiple times either to enable/disable or query parity error detection capabilities.
  - Instruction and Data Memory of the Network Processing Engines (NPEs)
  - Switching coprocessor in NPE B (SWCP)
  - AHB Queue Manager’s SRAM (AQM)
  - PCI Controller
  - Expansion Bus Controller
  - DDR SDRAM Memory Controller Unit (MCU).
- Interface to query the parity error detection status (whether enabled or not) on each of the above components.
- Interface to get the parity error detection statistics for each of the above-mentioned components.
- Interface to restore and update the NPE Error enabling configuration in the event of NPE soft-reset
- Interface to check for any NPE specified error status.
- Interface exchanges the data structures defined in the host byte order with the client application. This module operates in both big endian and little endian mode.

**Table 64. Parity Capabilities Supported by IxParityENAcc**

Feature	Hardware Component	Software Support	Recoverable
Error Correction Code	Memory Controller Unit - SDRAM	Single Bit Parity Error Notification	Yes
		Multi-Bit Parity Error Notification	No
Parity Error Detection	AHB Queue Manager SRAM	Parity Error Notification	No
Parity Error Detection	NPE IMEM, DMEM, AHB Coprocessor, Switch Coprocessor	Parity Error Notification	No
Parity Error Detection	Switch Coprocessor	Parity Error Notification	No
Parity Error Detection	PCI Controller	Parity Error Notification	No
Parity Error Detection	Expansion Bus Controller	Parity Error Notification	No

### 16.3.2 Dependencies

The client application at the time of initialization registers the parity error handler callback with IxParityENAcc. The client application also makes use of the parity error detection API to enable the underlying hardware blocks for parity error detection.

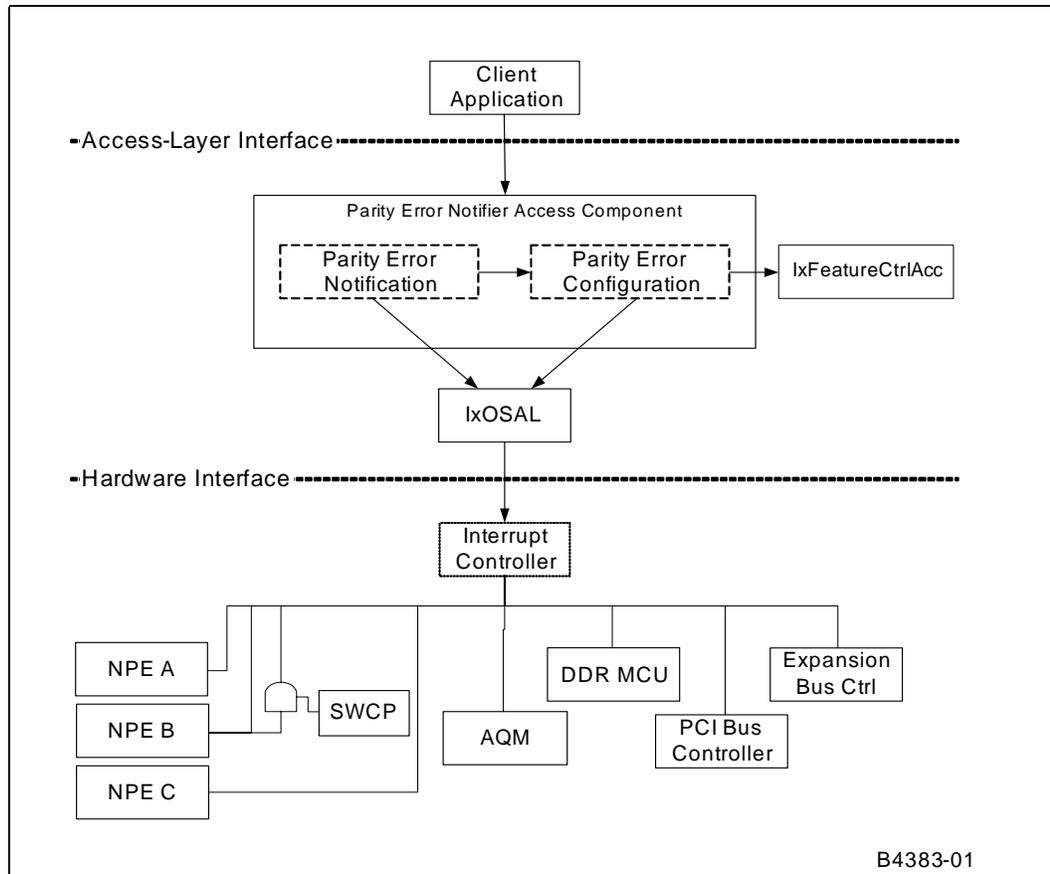
IxParityENAcc depends on various hardware registers to fetch the parity error information upon receiving an interrupt due to parity error. It then notifies the client application through the means of the callback handler with parity error context information.

IxParityENAcc also makes use of IxOSAL to access the underlying Operating System features such as IRQ registration, locks, and register access. IxOSAL is an abstracted interface, which is portable across different underlying OS.

Note that the client application may have dependencies on other access components when attempting to resolve the parity error issues. Indirect dependencies are not captured here.

Figure 79 presents a IxParityENAcc Dependency diagram.

**Figure 79. IxParityENAcc Dependency Diagram**



## 16.4 IxParityENAcc API Usage Scenarios

The following scenarios present usage examples of the interface by a client application.

There are three general tasks that would normally be provided by a client application with respect to parity events:

- Parity Error Notification
- Parity Error Recovery
- Parity Error Prevention

This section summarizes the high-level activities involved with these high-level tasks, and then presents specific usage scenarios.

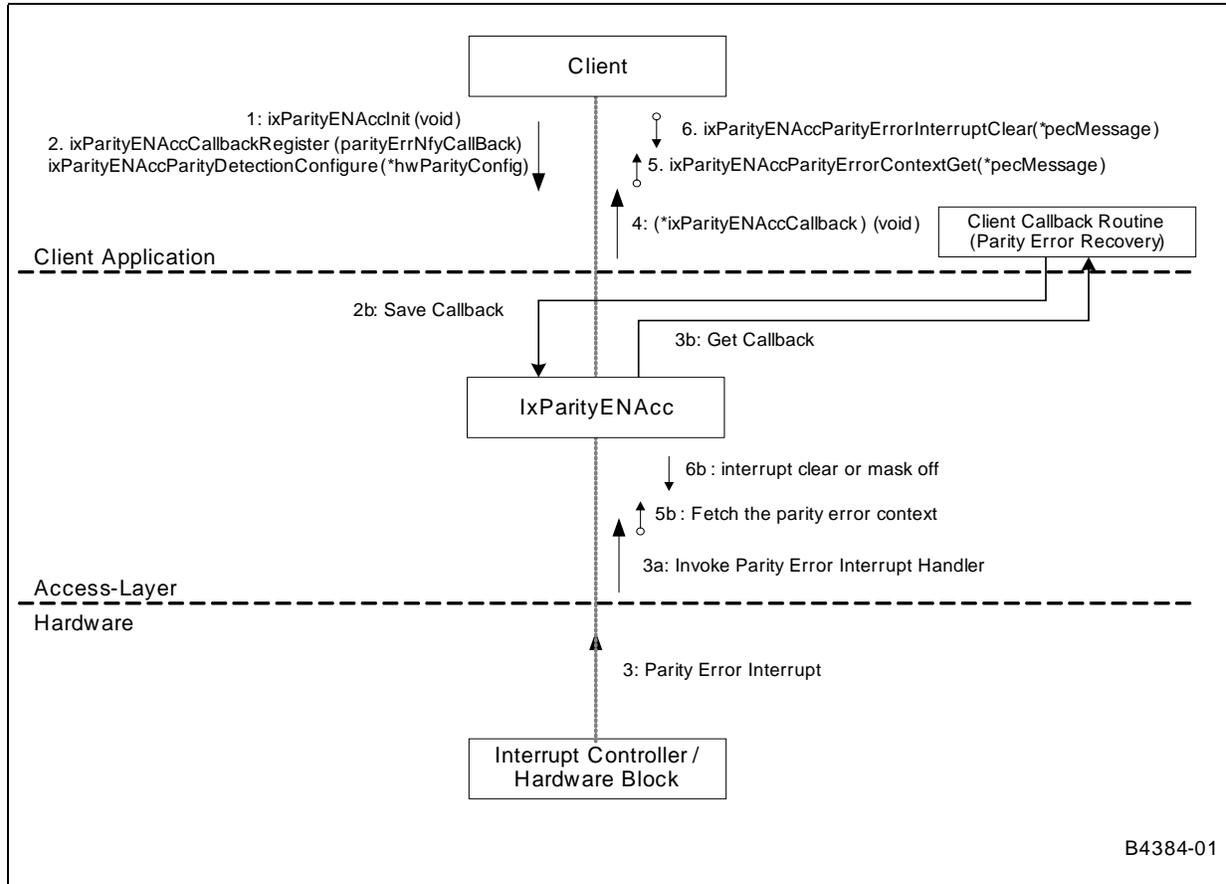


### 16.4.1 Summary Parity Error Notification Scenario

The interface between the client application and IxParityENAcc is explained in detail in the API source-code documentation. However, the following important scenario (shown in Figure 80) captures the usage of interface(s) by the client application.

The parity error context is represented with the data flow direction arrow with an open bubble at the end. The numbers at the beginning of each of the APIs and internal steps define their execution sequence in that order.

Figure 80. Parity Error Notification Sequence



1. The client application will initialize the component.
2. After initialization the client application will register callback and configure the parity error detection for the specified hardware blocks.
3. When a parity error occurs, the interrupt will fire and invoke the ISR of the IxParityENAcc component.
4. IxParityENAcc, in turn, invokes the client callback.
5. The client or data abort handler callback routine will then fetch the parity error context details and take appropriate action.
6. The client will then request to clear the interrupt condition.

The Parity Error Context will provide the following details:

- Source where the parity error detected



- Access type – Read/Write
- Faulty memory address
- Data from the faulty location if available
- Interface on which the request is made (AHB Bus or MPI)
- Master and Slave of the last erroneous AHB transaction

Table 65 describes the actions that should be taken when the client callback or data abort handler invokes the API to clear the parity interrupt conditions for the specified parity error context.

**Table 65. Parity Error Interrupt Deassertion Conditions**

Interrupt Bit	Source	API Invoked by...	Action Taken During Interrupt Clear
Int0 Int1 Int2	NPE-A NPE-B NPE-C	Client callback	Interrupt is masked off at the interrupt controller so that it will not trigger continuously. Client application has to take appropriate action and must reconfigure the parity error detection subsequently so that it is notified of the interrupts.
Int8	PCI	Client Callback	Interrupt condition is cleared at the PCI bus controller for the following: - PCI Initiator - PCI Target
Int58	SWCP	Client Callback	Interrupt is masked off at the interrupt controller so that it will not trigger continuously. Client application has to take appropriate action and must reconfigure the parity error detection subsequently so that it is notified of the interrupts.
Int60	AQM	Client Callback	Interrupt is masked off at the interrupt controller so that it will not trigger continuously. Client application has to take appropriate action and must reconfigure the parity error detection subsequently so that it is notified of the interrupts.
Int61	MCU	Client Callback of Data Abort Handler	Parity interrupt condition is cleared at the SDRAM MCU for the following: <ul style="list-style-type: none"> <li>• Single-bit</li> <li>• Multi-bit</li> <li>• Overflow condition, for example, more than two parity conditions occurred.</li> </ul> Note that single-parity errors do not result in data abort and not all data aborts are caused by multi-bit parity error. Refer to <a href="#">“Parity Error Notification Detailed Scenarios” on page 267.</a>
Int62	EXP	Client Callback	Parity interrupt condition is cleared at the Expansion Bus Controller for the following: <ul style="list-style-type: none"> <li>• External master initiated Inbound write</li> <li>• Internal master (IXP46X network processors) initiated Outbound read.</li> </ul>



## 16.4.2 Summary Parity Error Recovery Scenario

IxParityENAcc does not perform parity error recovery tasks. This can be done either by the ixErrHdlAcc or client application. refer to [Chapter 17.0, “Access-Layer Components: Error Handler \(ixErrHdlAcc\) API”](#) for further details on the hardware blocks that ixErrHdlAcc is supporting and the method used to recover parity errors.

When notified of any failure, the client application should identify the affected components by calling a function to fetch the Parity Error Context and decide on the appropriate course of action considering the impact on its functionality. For example:

- Reset the whole system immediately.
- Graceful shutdown of the system after taking the necessary actions to minimize the impact (informing the peers that it is about to shut down, tear down communication channels, and so forth)
- Other means, depending on the application and data integrity requirements.

The internal memories of NPEs, the Switching Coprocessor and AHB Queue Manager do not provide for an error correction facility. The DDR SDRAM controller implements a single-bit error correction mechanism that requires the Intel XScale® Processor to read and write the faulty memory location.

When the DDR controller notifies the Intel XScale® Processor about an error, error handling may vary slightly, depending on the operating system and Intel XScale® Processor MMU configurations. The user application should provide a scrub routine for single-bit parity errors. This routine is responsible for disabling interrupts, memory mapping, flushing of cache lines before reading the faulty word and after writing back the correct word onto it and finally enabling the interrupts.

For multi-bit parity errors, no error correction is possible and the Intel XScale® Processor is notified. The client application should handle such notifications.

## 16.4.3 Summary Parity Error Prevention Scenario

IxParityENAcc does not perform parity error prevention tasks. This should be done by the client application.

Since the DDR SDRAM controller provides the facility to correct single-bit parity errors, it is possible to run a background process/task to read the SDRAM locations at regular intervals and to fix the single-bit parity errors when encountered. This may be beneficial by reducing the chance of parity problems affecting the application code.

*Note:* In order to scrub single-bit parity error notification due to a read transaction, the scrub routine should first disable single-bit parity error detection and then perform a read and write access onto the faulty memory location. Otherwise the read memory access will result in another single-bit parity error notification and will result in an infinite number of iterations.

The scrub routine should ignore single-bit parity errors notified due to write transactions since the MCU will have scrubbed the data during the write transaction itself.

## 16.4.4 Parity Error Notification Detailed Scenarios

This section describes recommended usage of the IxParityENAcc component in several interrupt scenarios involving data aborts and parity error interrupts. The scenarios and possible implementations provided here are from the client application perspective



only, and could be resolved in an alternate manner. It is the client application's responsibility to implement an enhanced/modified data abort exception handler and the callback routine.

Note that the treatment of prefetch aborts may be very similar to that of data aborts, and is not described separately.

An Intel XScale® Processor access will result in data abort after experiencing problems in address translation, memory access protection, and so forth. These data aborts may not be specifically related to a parity error. In some situations, however, a parity error will also cause a data abort. Intel XScale® Processor accesses of South AHB bus targets that receive an AHB error response will result in a data abort. For example, an attempt to read from the AQM or Expansion Bus results in an AHB error response due to parity error at the AQM/Expansion Bus Controller.

Any non-Intel XScale® Processor access to faulty SDRAM memory will result in the Parity Error notification reaching the Intel XScale® Processor, but will not cause a data abort. However, an Intel XScale® Processor access to an SDRAM memory location that has a multi-bit parity problem will always result in the MCU triggering a Data Abort and may also result in a multi-bit Parity Error notification if the MCU is configured to detect the parity error.

The parity error context information also include details of the last error observed on the AHB bus. The information provided may be of help for the client application to decide which course of action to take. This information is retrieved from a Performance Monitoring Unit register, which might have been overwritten by another error by the time it is retrieved. The PMU may or may not include the information related to the parity event. This is because it may include data from previous errors. For example, an AHB transaction error has been locked into the PMU register, or there may be a parity event and the register data was retrieved or cleared by another process.

It is important to note that if an interrupt condition is not cleared then it will result in the parity interrupt being triggered again.

Figure 81–Figure 87 show the process flow that occurs in several data abort and parity error scenarios.

Figure 81. Data Abort with No Parity Error

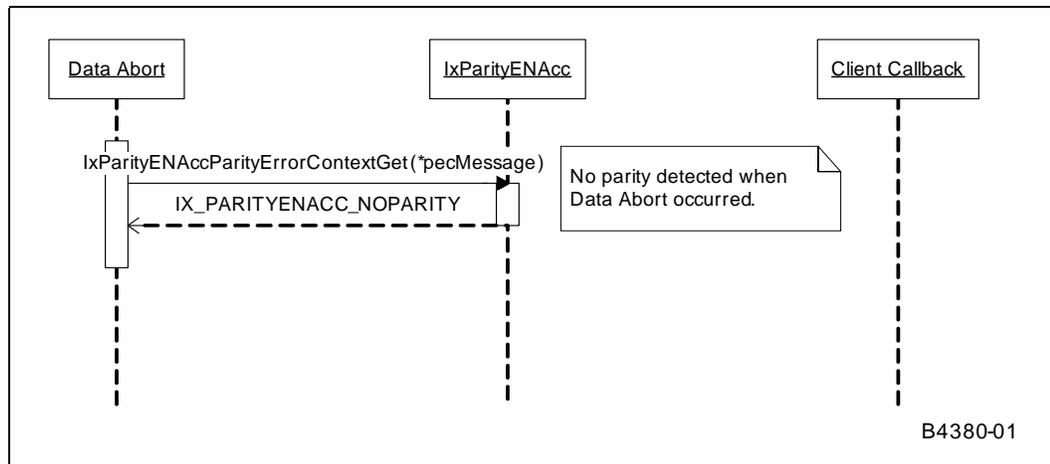




Figure 82. Parity Error with No Data Abort

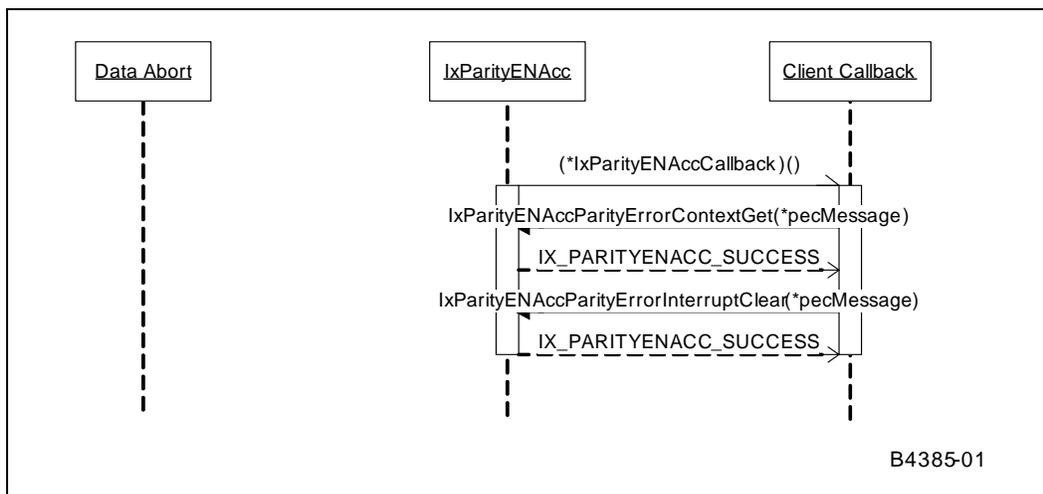
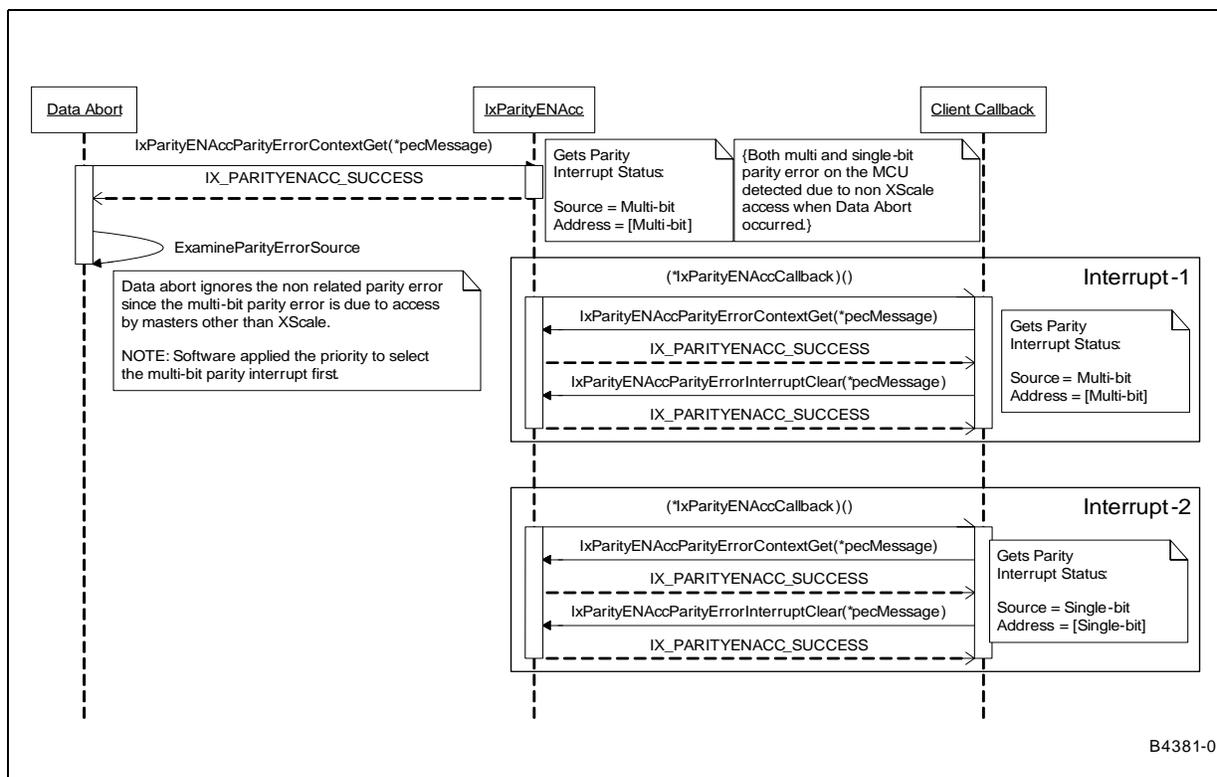
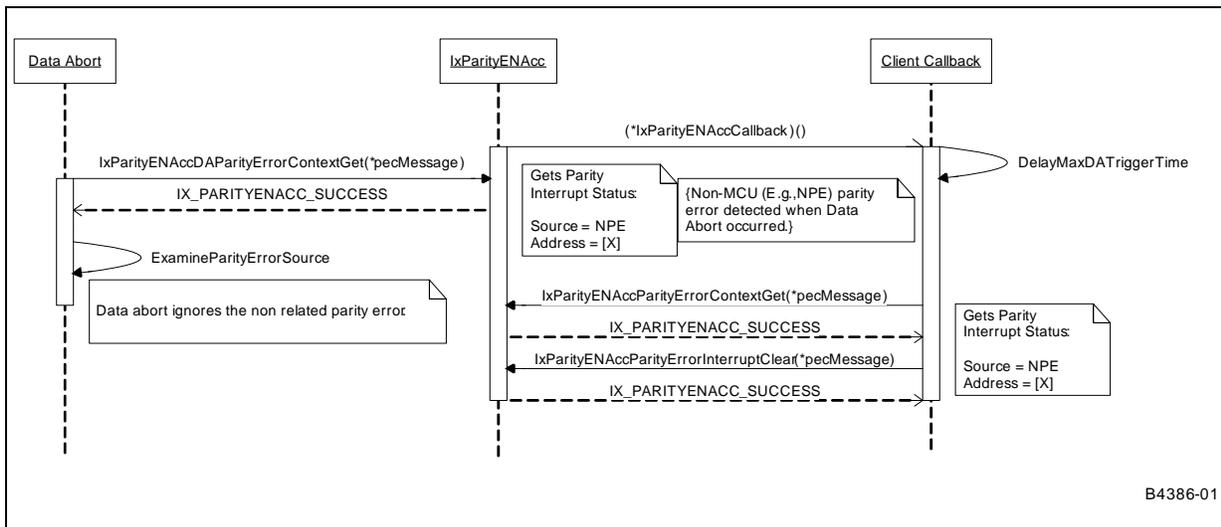


Figure 83. Data Abort followed by Unrelated Parity Error Notification

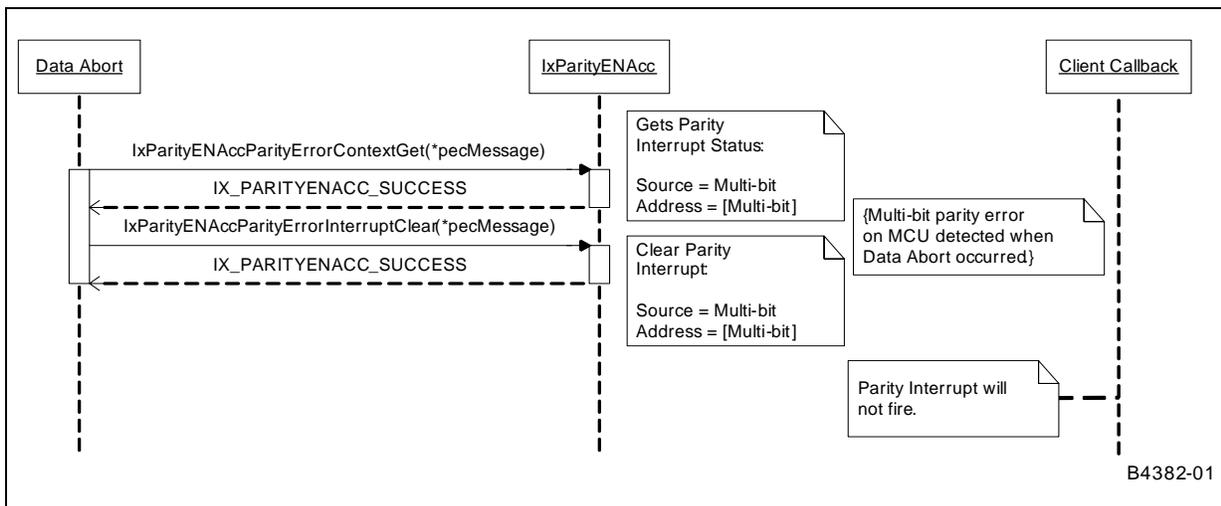


**Figure 84. Unrelated Parity Error Followed by Data Abort**



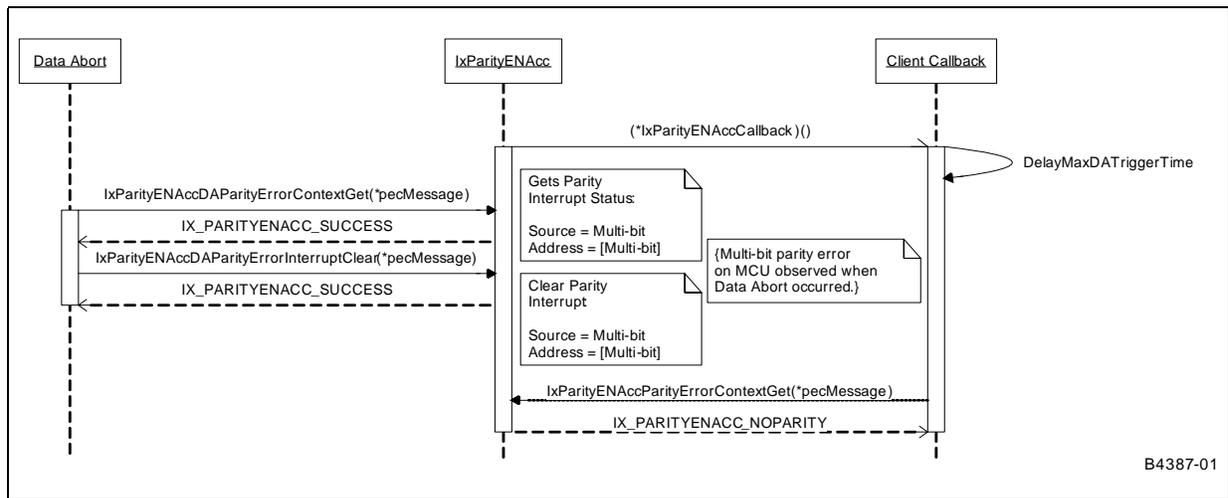
In order to avoid a race condition between the data abort handler and the parity error callback, delay has been introduced in the MCU parity event interrupt service routine of the access-layer component. This allows the data abort handler to complete prior to the interrupt service routine returning the parity context information.

**Figure 85. Data Abort Caused by Parity Error**



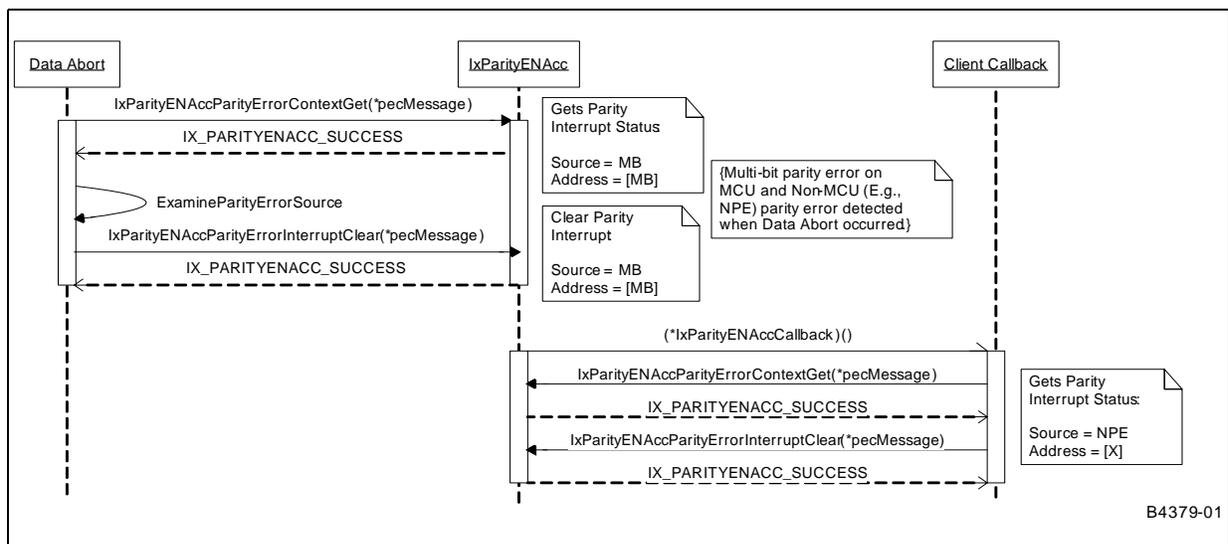


**Figure 86. Parity Error Notification Followed by Related Data Abort**



This scenario shown in [Figure 87](#) can occur because the order in which the interrupts are triggered for a parity error and a related data abort are not guaranteed.

**Figure 87. Data Abort with both Related and Unrelated Parity Errors**



§ §





## 17.0 Access-Layer Components: Error Handler (ixErrHdlAcc) API

---

This chapter describes the soft error handling and recovery in Intel® IXP400 Software on Intel® IXP4XX Product Line of Network Processors.

### 17.1 What's New

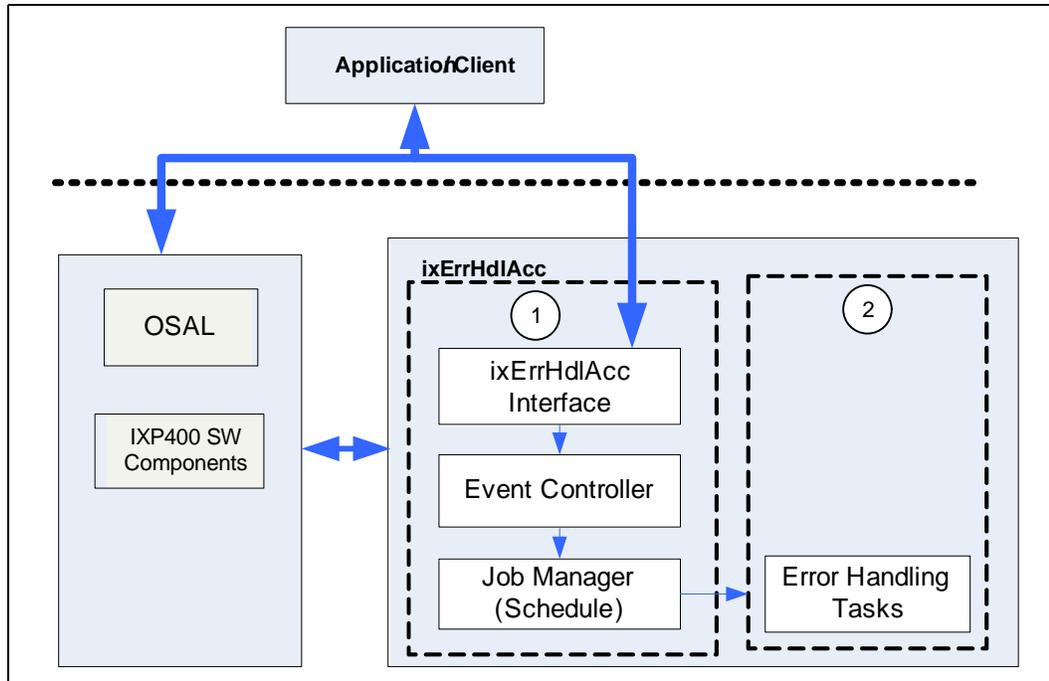
This is a new feature available in software release 2.3.

### 17.2 ixErrhdlAcc Overview

This is a new access library component for running and performing soft error recovery of a detected soft error on the IXP hardware platform. It provides a framework to enqueue error handling requests made by the device drivers or applications via a set of entry points. This is defined by the set of APIs provided by the call to the function API of *ixErrHdlAccErrorHandlerGet()* and shall be described further. The list of API to perform configuration, setup, initialization, and other operations shall be described in subsequent section.

## 17.3 Architectural Overview

Figure 88. Architectural View of ixErrHdlAcc Component



*ixErrHdlAcc* component consists of two sub-components

1. Job Queue Manager
2. Error handling module

Job Queue Manager identifies and detects soft error in queue, and error handling module consists of error handling routines for the following soft errors:

- Parity error in NPE IMEM and DMEM.
- NPE AHB error

## 17.4 Functional Details for NPE Error and AQM Parity Error

The error handling & recovery component provides the following functional features:

1. Functional interface feature:
  - a. Functional interface feature to the client application to enable/disable the soft error recovery of the soft errors on the AQM SRAM Parity error, NPE IMEM and DMEM parity error and NPE AHB error individually.
  - b. Functional interface feature to the client application to get the soft error interrupt service routines for AQM Interrupt error and NPE soft error interrupt.
  - c. Functional Interface feature to the client application to register a call-back notification after the completion of the error handling and recovery.
  - d. Functional Interface feature to the IXP400 software components to indicate a soft error condition state (details on the next sections). The components are:
    - *ixQMgr*
    - *ixNpeMh*



- ixNpeDI
- ixEthDB
- ixEthAcc
- ixEthMii

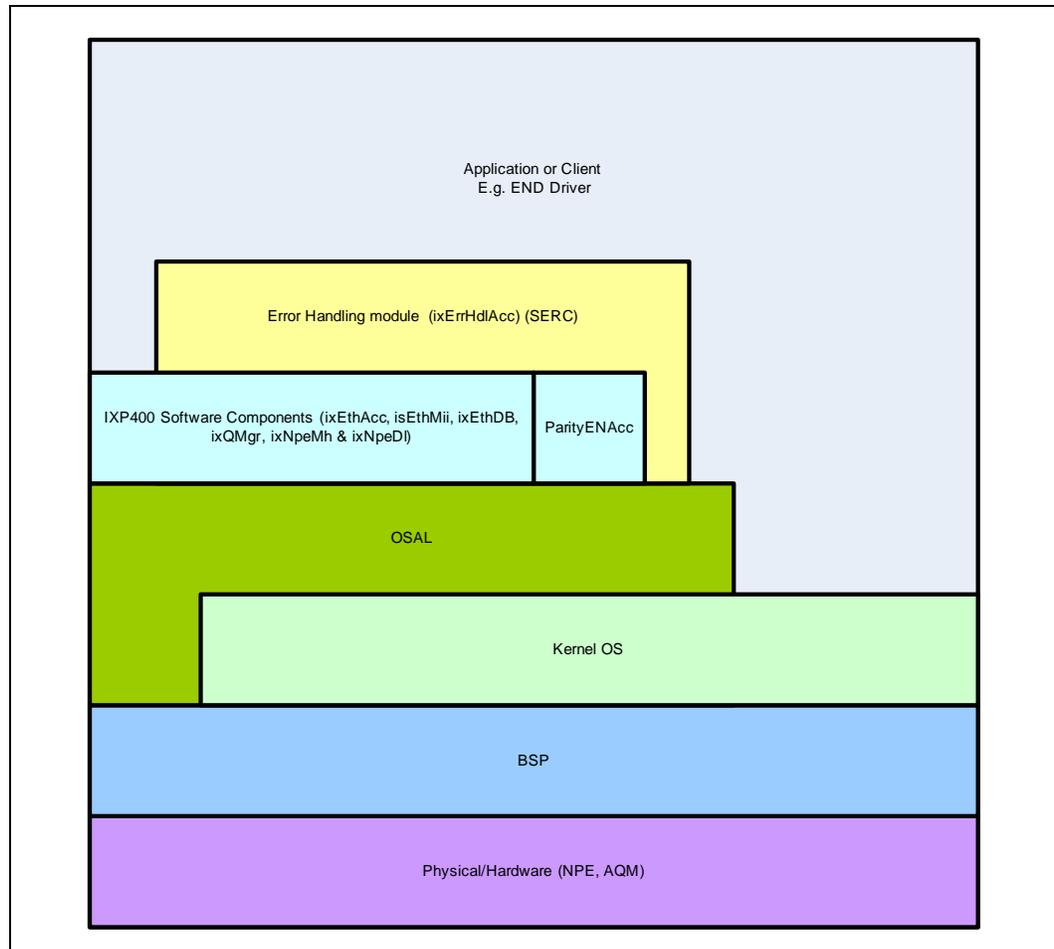
2. Functional soft error recovery feature in the detectable errors from NPE DMEM parity error, NPE IMEM parity error and NPE AHB Error. The feature recovery mechanism alleviates and fully recovers the IXP400 software driver components on the Intel XScale® Processor and NPE A, NPE B and NPE C micro-processors.

*Note:* The current design only supports Ethernet applications. In future, there shall be plans to add support for HSS, Crypto and ATM.

The details of the functional features is discussed in the following sections.

## 17.5 Dependencies

**Figure 89. Layered Block Diagram Depicting the Dependencies of Intel® IXP400 Software ParityENAcc Access Component with Notification and Soft Error Recovery**



As depicted in Figure 89, software in a layer is allowed to use any software layer it touches from above.

## 17.6 Software Interfaces

There are 3 types of interactions between the soft error recovery component or *IxErrHdlAcc* access component (SERC) with other elements, for example, the application, the IXP400 software Ethernet access components and the IXP400 software *ParityENAcc* component.

The 3 interfaces are as depicted in Figure 90. The SERC shall provide resources to elements on the interface 1 and 2, and uses resources provided by *parityENAcc* components as shown in interface 3.

Figure 90. Interface Architecture

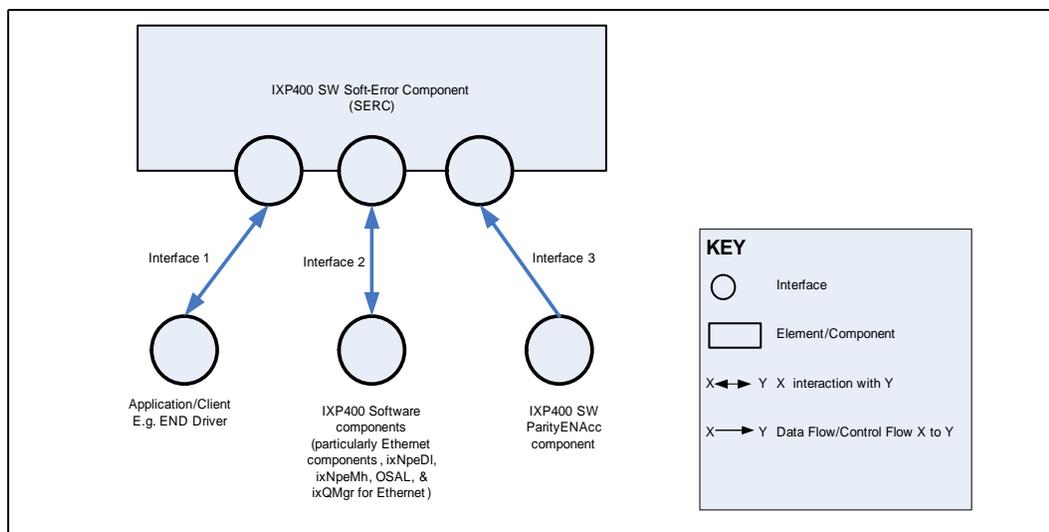




Table 66. Interface Identity and Types

Interface	Elements	Description
1	Application/client	Control interface that does <ul style="list-style-type: none"> <li>• configuration and initialization that initializes the <i>ixErrHdlAcc</i> component and setup that configures that enabling of the soft error handling.</li> <li>• soft error recovery completion call-backs for the soft error events.</li> </ul>
2	Ethernet Access Component	Data interface for <ul style="list-style-type: none"> <li>• Ethernet Access Component used to probe the error status</li> <li>• SERC to use the Ethernet Access component's resource to re-initialize the Ethernet application</li> </ul>
3	<i>parityENAcc</i>	Data interface used by the <i>ixErrHdlAcc</i> component to access the physical hardware layer for error status and error type <ul style="list-style-type: none"> <li>• Read the internal NPE memory mapped registers (NPESTAT), to check the NPE error type whether it is from DMEM or IMEM or whether it is from an external bus (AHB) error.</li> <li>• Read the Queue manager Internal mapped registers to get the last location of the memory of the last known parity error in it's SRAM.</li> </ul>

## 17.7 Intel® IXP400 Software Enabling of Soft Error Detection and Handling on the Intel® IXP4XX Development Platforms

This section shall provide a detailed description of how to implement soft error detection and handling using the IXP400 software access component library.

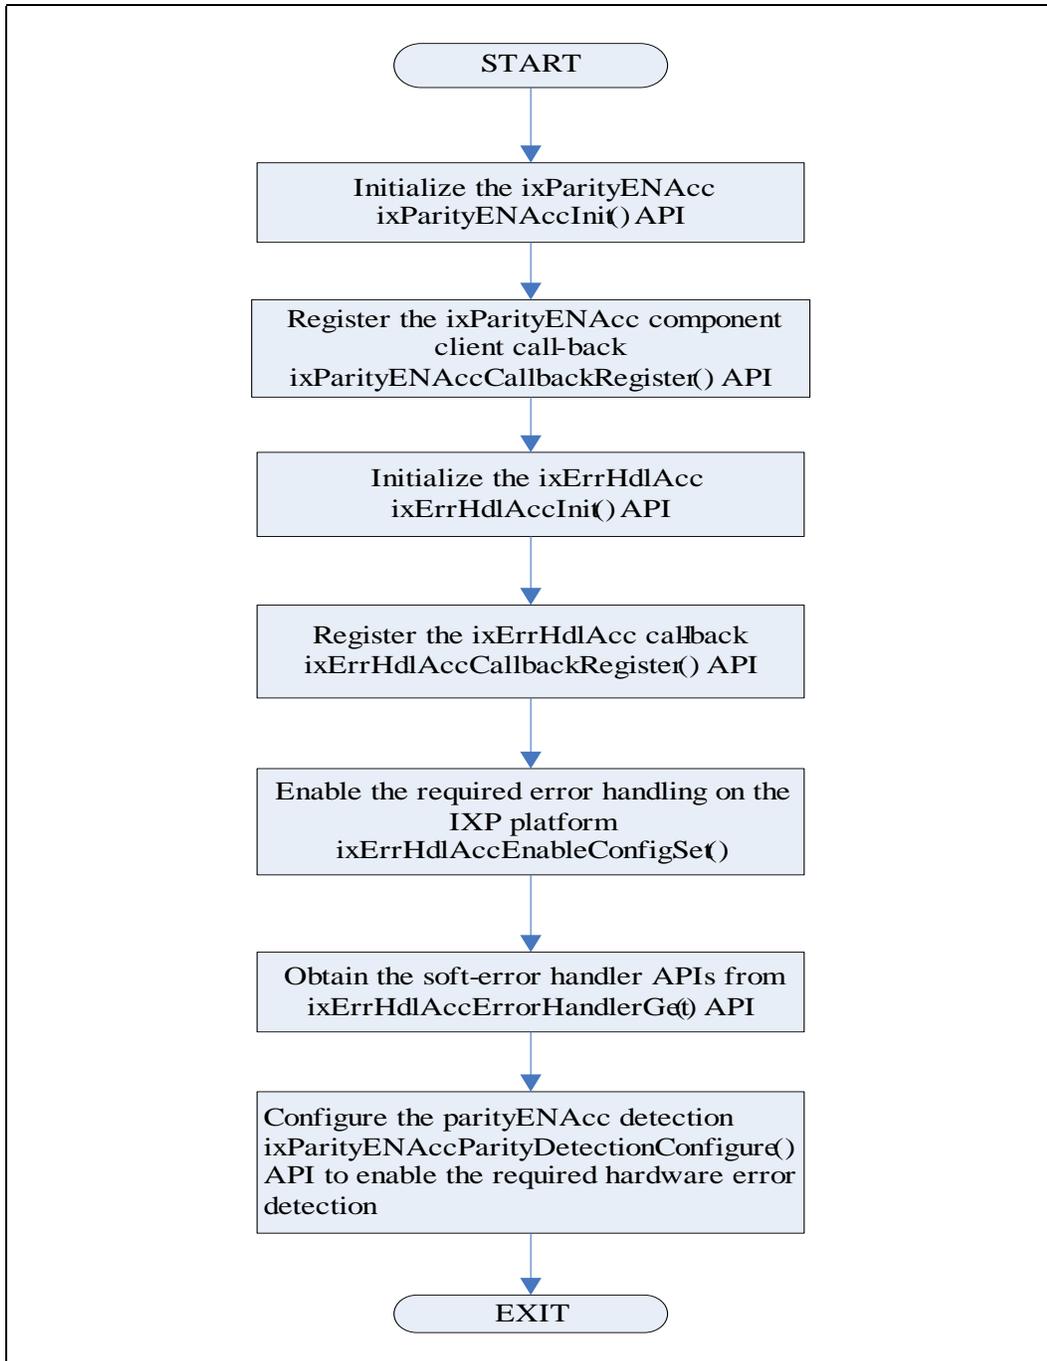
### 17.7.1 Soft Error Detection, Handling Configuration and Initialization Pseudo Code

Basic pseudo-code for initializing the SERC on the IXP400 software.

1. Initialize the IXP400 software *ParityENAcc* by calling *ixParityENAccInit()* API.
2. Register a call-back with the *parityENAcc* by calling *ixParityENAccCallbackRegister()* API.
3. Initialize the IXP400 software *ixErrHdlAcc* by calling *ixErrHdlAccInit()* API.
4. Register a call-back with the *ixErrHdlAcc* by calling *ixErrHdlAccCallbackRegister()* API
5. Enable the required error handling on the Intel® IXP4XX Development Platforms by calling the *ixErrHdlAccEnableConfigSet()* API. (*ixErrHdlAccEnableConfigGet()* API to query the previous configuration if there is one)
6. In the call-back registered with the *ParityENAcc* as in step 2, the various *ixErrHdlAcc* error handlers API can be obtained by the call to the *ixErrHdlAccErrorHandlerGet()* API and called only upon the detection of the same error type as the handler.
7. Configure the *ParityENAcc* detection by a call to the *ixParityENAccParityDetectionConfigure()* API to enable the required hardware error detection on the Intel® IXP4XX Development Platforms.

A flow chart depicting the above is shown in the following [Figure 91](#).

Figure 91. Flow Chart Depicting the Initialization and Configuration of the Soft Error Detection and Handling





## 17.7.2 Initializing the Intel® IXP400 Software Soft Error Modules

The IXP400 software *ixParityENAcc* and *ixErrHdlAcc* are components used in the Intel® IXP4XX Development Platforms to detect soft errors to alleviate the need for hard reboot of the system. All required modules or components must be initialized at least once in the system. The initialization sequence should be executed only after a full hardware initialization by the kernel OS or the boot loaders and recommended during the start-up of the application layer or device layer that utilizes the IXP400 software libraries. The mandatory pre-cautions above ensure that, all hardware peripherals and components of the IXP silicon is reset and initialized before the error detection is enabled. The enabling of the error detection on the Intel® IXP4XX Development Platforms is done by the *ixParityENAcc* access library. The *ixParityENAcc* access library performs read and write to the memory mapped configuration registers in the IO device memory space for setup. This setup however can be reset or cleared if any of the devices or peripherals is reset. One such example of a hardware device is the NPE.

In the application layer, an important note is that the *ixParityENAcc* component must only be initialized after the NPE firmware download and startup. This is to make sure that the entire used IMEM and DMEM is fully initialized before enabling the parity Error detection on the NPE. Another reason is that the NPE firmware download and start-up stops and resets the NPE and thus, all previous setup of the NPE device are cleared.

The *ixParityENAcc* component provides a common call-back that is to be registered by the call to the *ixParityENAccCallbackRegister ()* API for every supported error types, for example, ECC Error, NPE Error, AQM Error, Expansion bus error and others. Because of that, this restricts the usage of multiple applications registering their individual call-backs for an error event. Thus, the soft error detection and error handling must be driven only by a single driver/application that acts as a core controller for enabling error type detections and supervising the recovery process. The software architecture of the soft error feature on an IXP platform using the IXP400 software access libraries is shown in the [Figure 92](#). The IXP400 software Driver 0 does the call to *ixParityENAccInit* and *ixErrHdlAccInit* respectively to initialize the components.

The error handling and recovery must be enabled prior to calling the *ixErrHdlAcc* error handler function that is return by the *ixErrHdlAccErrorHandlerGet()* API (More on this in the next section). The enabling or disabling of the handler is done by the *ixErrHdlAccEnableConfigSet()* API. The input to the first argument is a configuration word mask. The bit 0, 1, 2, and 3 of the configuration mask sets the enable or disable configuration for NPE A, NPE B, NPE C and AQM error respectively. The IXP400 software does not support error handling for all error types. As such, Refer to the *ixErrHdlAcc* Module API list reference located in the Intel® IXP400 Software API reference manual to determine the supported error handling types.

An example of the *ixErrHdlAccEnableConfigSet()* API usage for enabling the NPE A and NPE C error handling is as shown:

```

/* Initialize ixErrHdlAcc Component*/
ixErrHdlAccInit();
ixErrHdlAccCallbackRegister(callbackfunc);

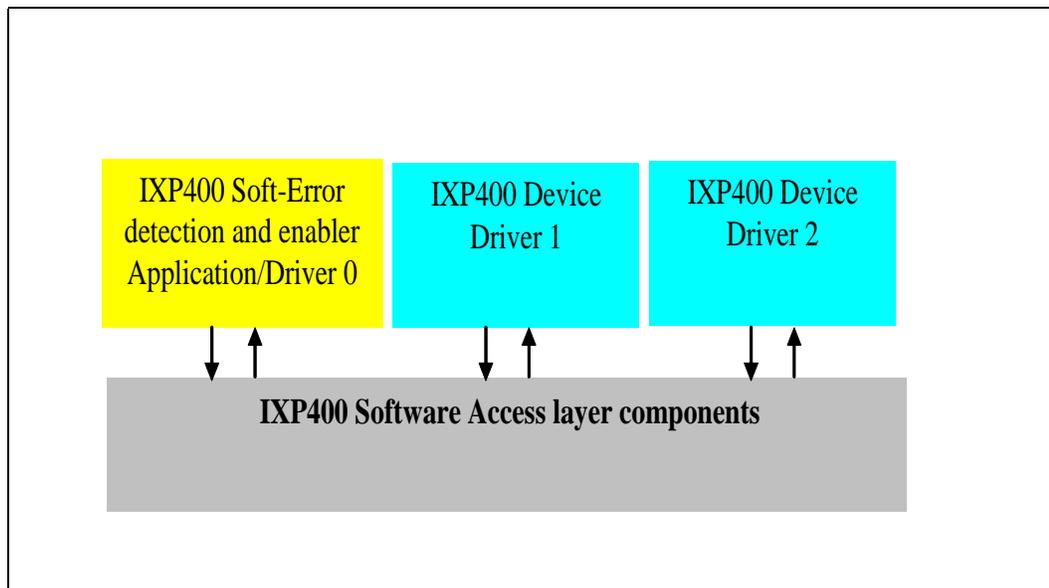
/* Enable all error handling by default*/
configEnable =(IX_ERRHDLACC_NPEA_ERROR_MASK_BIT|
               IX_ERRHDLACC_NPEB_ERROR_MASK_BIT|
               IX_ERRHDLACC_NPEC_ERROR_MASK_BIT|
               IX_ERRHDLACC_AQM_SRAM_PARITY_ERROR_MASK_BIT);
ixErrHdlAccEnableConfigSet(configEnable);

```

The current configuration setup in the *ixErrHdlAcc* component can be obtained using the *ixErrHdlAccEnableConfigGet()* API.

The enabling of the error detections must be performed lastly by using the *ixParityENAccParityDetectionConfigure()* API of the *ixParityENAcc* component. This enables or disables the detection of any error detectable from the IXP4XX development platforms that comes in as an interrupt error. The hardware error detection is disabled by default during the initialization of the IXP400 software *ixParityENAcc* access component. By enabling the hardware error detection, any occurrence of an error on the IXP4XX development platforms shall assert a high priority level interrupt error to the Intel XScale® Processor. (Refer to “Access-Layer Components: Parity Error Notifier (*ixParityENAcc*) API” on page 259 for detail of the *ixParityENAcc* access component) Subsequently, notifying the client or application via the *ixParityENAcc*'s registered call-back.

**Figure 92. A Basic Software Architecture Showing the Soft Error Detection and Handling Implementation using the Intel® IXP400 Software Access Libraries**



### 17.7.3 Detection and handling of Soft Error

The registered client call-back to the *parityENAcc* component shall be called by the ISR of the interrupt error. Within this call-back, the application/driver must call the *ixParityENAccParityErrorContextGet()* API to obtain the triggered error event type and appropriately summons the error handler provided by the *ixErrHdlAcc* component. The error handlers that are supported by the *ixErrHdlAcc* component are provided by the *ixErrHdlAccErrorHandlerGet()* API. Calling this pre-defined error handlers shall wake-up an internal recovery task/thread that enqueues the request into an internal job queue that shall be scheduled on a round-robin turn basis. The sample code of how to use this API and as shown in the sample source codes. The 1<sup>st</sup> argument parameter to this function is defined as macros, an enumerator of *ixErrHdlAccErrorEventType* type. For example, to obtain the NPE A soft error handler, one should use `IX_ERRHDLACC_NPEA_ERROR` as the 1<sup>st</sup> argument to the function *ixErrHdlAccErrorHandlerGet()*. The returned function pointer (2<sup>nd</sup> argument) is recommended to be called only from an interrupt context or a thread or task with a high priority running in the kernel mode. This is to prevent pre-emption of the internal operations of the function handler after a detected error.



The *ErrHdlAcc* component client registered call-back shall be asserted when the handling and recovery task ends. This call-back serves as a notification to the application layer of an impending error and the results or status of the recovery task. The error recovery result or status can be obtained by the call to the *ixErrHdlAccStatusGet()* API. The API returns the status of the error handling as in pass or fail. As such, any failure reported by the *ixErrHdlAcc* component can be dealt with by the device driver. For instance, the system can be hard-reboot or the faulted device can be unloaded.

A basic *parityENAcc* component client call-back implementation would be as in the sample source codes shown below. The codes, does an interrupt (IRQ) lock to protect the critical sections of the codes. The critical section of the codes performs a check on the error type by calling *ixParityENAccParityErrorContextGet ()* API. This may result in the NPE (A, B & C) soft error handler being called if an NPE source error is detected. The OSAL API of *ixOsallrqLock* and *ixOsallrqUnlock* does an OS Kernel system call that incidentally masks the Intel XScale® Processor maskable interrupts. Thus, this prevents any interrupt context preemption while the critical sections are being executed.

```

void parityENAccCallbackExample(void)
{
    IxParityENAccParityErrorContextMessage pENContext;

    ixParityENAccStatus                pENStatus;

    UINT32 irqLock;

    irqLock = ixOsallrqLock();

    /* initialize IxParityENAccParityErrorContextMessage buffer */
    memset (&pENContext, 0xFF, sizeof (IxParityENAccParityErrorContextMessage));
    /* get parity error context */
    pENStatus = ixParityENAccParityErrorContextGet (&pENContext);

    if (IX_PARITYENACC_SUCCESS == pENStatus)
    {
        /* Error detected*/
        IxErrHdlAccFuncHandler func;

        /* Attempt to recover from the error*/

        switch (pENContext->pecParitySource)
        {
            case IX_PARITYENACC_MCU_MBIT:
                break;

            case IX_PARITYENACC_AQM:
                break;

            case IX_PARITYENACC_EBC_CS:

                break;

            case IX_PARITYENACC_NPE_A_IMEM:

            case IX_PARITYENACC_NPE_A_DMED:

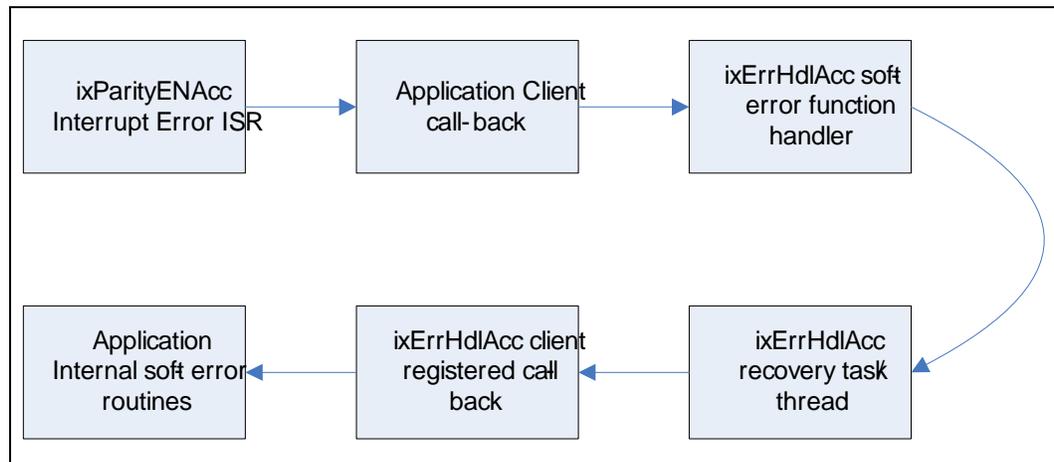
            case IX_PARITYENACC_NPE_A_EXT:

```



```
/* Get the soft-error handler function for NPE A Error*/
    ixErrHdlAccErrorHandlerGet(IX_ERRHDLACC_NPEA_ERROR, &func);
    (*func)();
    break;
case IX_PARITYENACC_NPE_B_IMEM:
case IX_PARITYENACC_NPE_B_DMEM:
case IX_PARITYENACC_NPE_B_EXT:
/* Get the soft-error handler function for NPE B Error*/
    ixErrHdlAccErrorHandlerGet(IX_ERRHDLACC_NPEB_ERROR, &func);
    (*func)();
    break;
case IX_PARITYENACC_NPE_C_IMEM:
case IX_PARITYENACC_NPE_C_DMEM:
case IX_PARITYENACC_NPE_C_EXT:
/* Get the soft-error handler function for NPE C Error*/
    ixErrHdlAccErrorHandlerGet(IX_ERRHDLACC_NPEC_ERROR, &func);
    (*func)();
    break;
case IX_PARITYENACC_SWCP:
    break;
default:
    break;
}
ixOsalIrqUnlock(irqLock);
return;
}
```

Figure 93 shows the data-flow from the time of assertion of the interrupt error till it is notified and recovered.


**Figure 93. Data Flow Diagram Depicting the Start of an Interrupt Error as it Recovers**


The callback is defined as in the sample code below.

```

void callback(IxErrHdlAccErrorEventType eventType)
{
}
  
```

The input parameter notifies the application of the error event type. In cases whereby, multiple error handling requests were made to the *ixErrHdlAcc* component; there shall be multiple notification callbacks. The input parameter of *eventType* is then used to differentiate callbacks for every error event type supported.

At any one time and after the callback notification by the *ixErrHdlAcc* component, the error recovery result can be obtained by the call to the API of *ixErrHdlAccStatusGet()*. The API outputs a 32 bit word mask (1<sup>st</sup> argument is of type 32 bit word data pointer) whereby a bit clear and set indicates a failure and a pass respectively. The currently supported bit mask is defined as shown below.

```

#define IX_ERRHDLACC_NPEA_ERROR_MASK_BIT
#define IX_ERRHDLACC_NPEB_ERROR_MASK_BIT
#define IX_ERRHDLACC_NPEC_ERROR_MASK_BIT
#define IX_ERRHDLACC_AQM_SRAM_PARITY_ERROR_MASK_BIT
  
```

A typical sample code is as shown below. The code obtains the current *ixErrHdlAcc* error status and checks if the NPE A device is at fault.



```
ixErrHdlAccStatusGet(&status);
```

```
if(status & IX_ERRHDLACC_NPEA_ERROR_MASK_BIT)
{
    /* NPE A Fail to recovery*/
}
```

On the contrary, the function *ixErrHdlAccStatusSet()* API can be used to set a failure in the callback or in the application for some unique cases only. One such scenario is that when there is a serious memory leakage that could significantly impede the normal operations of the device at the application layer after the *ixErrHdlAcc* error recovery. Another such case is, where an attempt is made to recover a custom made NPE firmware written by the user that failed. A note of caution is that, it is strongly recommended to not use the *ixErrHdlAccStatusSet()* API if none of the above scenarios occurred.

The API *ixErrHdlAccStatisticsShow* function can be used to display debug information regarding the operation of the *ixErrHdlAcc* component real-time. A sample display print output is as defined below. The data recorded is an accumulative history record of the target. However, if desired, the recorded debug data can be cleared by the call to the *ixErrHdlAccStatisticsClear* API.

Events Statistics

```
-----
Total requested NPEA Error event = 0
Total requested NPEB Error event = 8
Total requested NPEC Error event = 0
Total requested AQM Error event = 0
Total passed NPEA Error event = 0
Total passed NPEB Error event = 8
Total passed NPEC Error event = 0
```

### 17.7.4 Unloading the *ixErrHdlAcc*

The *ixErrHdlAcc* is unloaded by using the API *ixErrHdlAccUnload*. The unloading process does a thread kill or termination, de-allocation of any allocated memory, and the destruction of any created synchronization thread objects like semaphores and mutexes.



## 17.7.5 Ethernet Device Driver Modifications for Soft Error Detection and Handling

### 17.7.5.1 Resolving NPE IRQ Sharing Conflicts

As of Intel® IXP400 Software v2.3 and above, the IXP400 software access libraries shall be supporting the detection and handling of all NPE errors. For NPE error detection and handling, with respect to IXP4XX development platforms, there is an IRQ sharing conflict between the *ixNpeMh* Access library component and the *ixParityENAcc* Access library component. Both access libraries binds an ISR to the same IRQ of NPE A, B and C after initialization (call to *ixNpeMh ixNpeMhInitialize()* or a call to *ixParityENAcc ixParityENAccInit()*) thus, overwriting each other's ISR binding. For Intel® IXP400 Software v2.3 (not inclusive), the *ixNpeMH* component's *ixNpeMhInitialize* binds an ISR to the NPE IRQs irrespective of the mode that was configured. The *ixNpeMh* inspection of the out FIFO status can be either be configured as an interrupt binding to the NPE IRQ or in polling mode. However, this was subsequently fixed in the Intel® IXP400 Software v2.3. To resolve the IRQ conflicts between the 2 components, the method used today is to initialize the *ixNpeMH* component in a poll mode by a call of *ixNpeMhInitialize(IX\_NPEMH\_NPEINTERRUPTS\_NO)* during the IXP Ethernet hardware initialization sequence in the device driver. By doing so, only the *parityENAcc* access library component does the ISR binding to the NPE IRQs. To check for receive NPE messages from the out FIFO, the device driver then has to perform a polling task/thread to check the NPE out FIFO at intervals no less than 200 times per second or every 5ms intervals for every NPEs that was enabled. This can be fine tune accordingly to lower values and are highly dependent on the operating system, CPU load, and applications.

The application or device driver has to call the *ixNpeMH* API of *ixNpeMhMessagesReceive* providing the NPE ID as the only argument or parameter input to be check or polled. A pre-requisite is that *ixNpeMhMessagesReceive* API cannot be summoned from an interrupt context due to the usage of mutex within its internal implementation. The detail description of the API is as defined in the IXP400 Software API reference manual.

### 17.7.5.2 QM Queue Dispatcher Binding Setup

The queue dispatcher function is obtained by the call to the *ixQMgr* access library component API as shown below.

```
PUBLIC void ixQMgrDispatcherLoopGet(IxQMgrDispatcherFuncPtr
*qDispatcherFuncPtr)
```

Get QMgr *DispatcherLoopRun* for respective silicon device.

According to specification, the *qDispatcherFuncPtr* function can be either called from an interrupt service routine bind to the QM IRQs or called by a polling task or at an interval by a timer interrupt. The API of *ixQMgrDispatcherInterruptModeSet* must be called before the QM Queue dispatcher ISR IRQ binding and after *ixQMgr* component initialization. In short, this should be done before the dispatcher is bind to the ISR of the QM IRQ. The first argument to the API must be set to TRUE indicating an interrupt service routine binding of the QM IRQ. The default setup shall be FALSE indicating non-binding to QM IRQ at start-up or after initialization of the *ixQMgr* component.



This new API was conjured up to provide the QM Queue dispatcher setup data to be provided to the ixQMGr in order to prevent continuous pre-emption by the ISR on the system during NPE soft-reset whereby, the QM queue status (Nearly Empty, Empty, Nearly Full and Full) remains static till the NPE completely resets and starts.

## 17.8 Functional Soft Error Recovery and Handling of AQM SRAM Parity Error and NPE on an Ethernet Application

Generally, both AQM and NPE has hardware parity checker detection mechanism that provides notification to Intel XScale® Processor via interrupts, as well as the type of error, the error value of the memory, and address of the memory.

*Note:* AQM SRAM Parity error handling and recovery shall not be supported.

### 17.8.1 NPE Soft Error Recovery

#### 17.8.1.1 Background

As a start, a little back-ground, the Intel® IXP46X Network Processor introduces a soft error protection and detection mechanism that detects soft errors on the memory data structures which is the IMEM and DMEM and propagation soft errors that may affect AHB transactions. The detection uses the parity bits to validate the correctness or validity of the IMEM and DMEM. This detects soft errors on memory data structures. These are DUE type errors.

In the aftermath of a DUE on the NPE, the NPE enters into a halt or lock-up mode that is recoverable with a NPE reset. An interrupt error signal is sent to the system interrupt controller that triggers a fast IRQ or normal IRQ to the system CPU (Intel XScale® Processor). The Intel XScale® Processor CPU shall then appropriately handle this DUE. The implementation of the software shall be in the IXP400 software access layer components. There shall be error handling provision that then must be added to the IXP400 software access layer components.

The implementation is to avoid resorting to any hard re-boot of the system and to alleviate the impact of the soft error to the system, for example, Ethernet gateway, Ethernet router or setup box and so on.

#### 17.8.1.2 Intel® IXP400 Software Access Layer Components Error Handling

The error handling implementation shall be on an Ethernet application running on the IXDP465 platform, for example, routers, bridges and gateways. The Ethernet device driver, for example, END Driver (VxWorks\*) uses the IXP400 software Ethernet access libraries to enable and utilize the Intel® IXP46X Network Processor hardware functionality. By assuming that, the error handling control has been enabled prior to the execution of the system

During the soft error recovery, the read and write access to the AQM hardware queues has to be avoided or prevented. Also, if the data plane is in the midst of transmitting or receiving Ethernet data frames, it is not interrupted by the soft error recovery tasks till it's completed.



The preventive steps above are conjured up for the purpose to prevent the application or client from accessing any of the physical hardware resource of the IXP4XX product line of network processor till the error condition is cleared or calling any of the Ethernet Access libraries APIs. Also, by not interrupting critical sections of the Ethernet data path, this avoids us from placing too many soft error handling codes in the Ethernet data plane codes that could lower Ethernet performance.

### 17.8.1.3 Scalability

For scalar ability to other application, each soft error handler bind towards a particular NPE must check the features available base on the ImageID (for example, Ethernet NPE, HSS NPE or ATM NPE) that was downloaded to the NPE. If it's an Ethernet ImageID that was downloaded, the appropriate recovery shall be done. We used `ixNpeDILoadedImageFunctionalityGet` to obtain this information.

### 17.8.1.4 Multiple Event Errors

In the advent of a scenario whereby, multiple soft error events or more than one NPE has error at the same time or before the completion of the soft error recovery, the re-enabling of the Ethernet data plane shall be stopped till all the other faulted Ethernet NPE errors are cleared. Dependable upon the detail design, there can be more than 1 instance of NPE Error recovery task/threads running simultaneously and there shall be more than one call-back notification to the application/client. Alternatively, a single thread can be used to perform the recovery base upon an event queue that queues the recovery request from the application.

### 17.8.1.5 Memory Leakage Prevention Methods

Memory leakage or lost are attributed to the missing information of network buffer pointers or Ethernet frame descriptors. At any one time, the network buffers can be owned by the Intel XScale® Processor and NPEs processors for transmitting or receiving packets. Once the Ethernet frame descriptors are enqueued into the AQM hardware queues, they are considered disowned from Intel XScale® Processor and owned by the NPEs.

Based on the Ethernet NPEs functional specification, the NPE Tx logic empties the Ethernet Tx queue into its internal software queue while the NPE Rx logic always dequeues one Ethernet frame descriptor to be used to store the incoming Ethernet frames.

Also, freed network buffers must be recycle back to the RxFREE hardware FIFO Queue and to the transmit network buffer pool that may reside in the Ethernet device driver or it's TCP/IP network stack or any other protocol stack that binds to the Ethernet device driver.

If the Ethernet NPE is to be reset, those information stored in the DMEM shall be permanently lost upon DMEM re-initialization. The Intel XScale® Processor CPU has no means of retrieving the missing Ethernet frame descriptors held by the NPEs. Therefore, the only logical method is to do a probe or scan the internal software queue of the NPEs via the debug control registers after a NPE halt or lock-up and before a NPE reset. The implementation shall require the function specification modification of the Ethernet NPE to provide real-time dynamic information on its internal software queue. Finally, the retrieved Tx and RxFREE Ethernet frame descriptors shall be re-written or enqueued back to its AQM hardware queue before the Ethernet NPE is soft-reset and re-initialized.



However, the above memory leakage preventive methods above can only be exploited and guaranteed if the NPE error is not of a DMEM or IMEM parity error. The integrity of the data is in doubt if we are to read the DMEM's data that stores the internal software queue during DMEM parity error or IMEM parity error. Though this risk can be accepted given the severity of memory leakage and the chances or likelihood of the soft errors in the DMEM's vital memory section.

During DUE, Intel XScale® Processor shall also not re-initialize the AQM hardware queues, and shall only force a re-triggering of the event conditional bus to the NPEs during the recovery process. This updates the queue status of nearly full, nearly empty, full and empty to all the affected Ethernet NPEs. This prevents any loss of vital Ethernet frame descriptors in the AQM which causes memory leakage and also Ethernet network buffer starvation in the application or client. Starvation happens if the network buffers given to the AQM hardware queue was not re-cycle or returned to the application or client resulting in a growing depletion of freed network buffers in the network buffer pool and is finite.

§ §



## 18.0 Access-Layer Components: Queue Manager (IxQMgr) API

---

This chapter describes the Intel® IXP400 Software v2.3's "Queue Manager API" access-layer component.

### 18.1 What's New

Four new APIs have been added to the IxQMgr to support the new NPE soft-reset feature (ixErrHdlAcc)

**ixQMgrDispatcherLoopDisable()**: Disables the Queue Manager Dispatcher to prevent QM Read/Write access

**ixQMgrDispatcherLoopEnable()**: Re-enable the Queue Manager Dispatcher.

**ixQMgrDispatcherInterruptModeSet()**: Records whether the Queue Dispatcher is bind to the Queue Manager Interrupt notification IRQ or polling mode.

**ixQMgrDispatcherLoopStatusGet()**: Gets status of the Queue Manager Dispatcher.

### 18.2 Overview

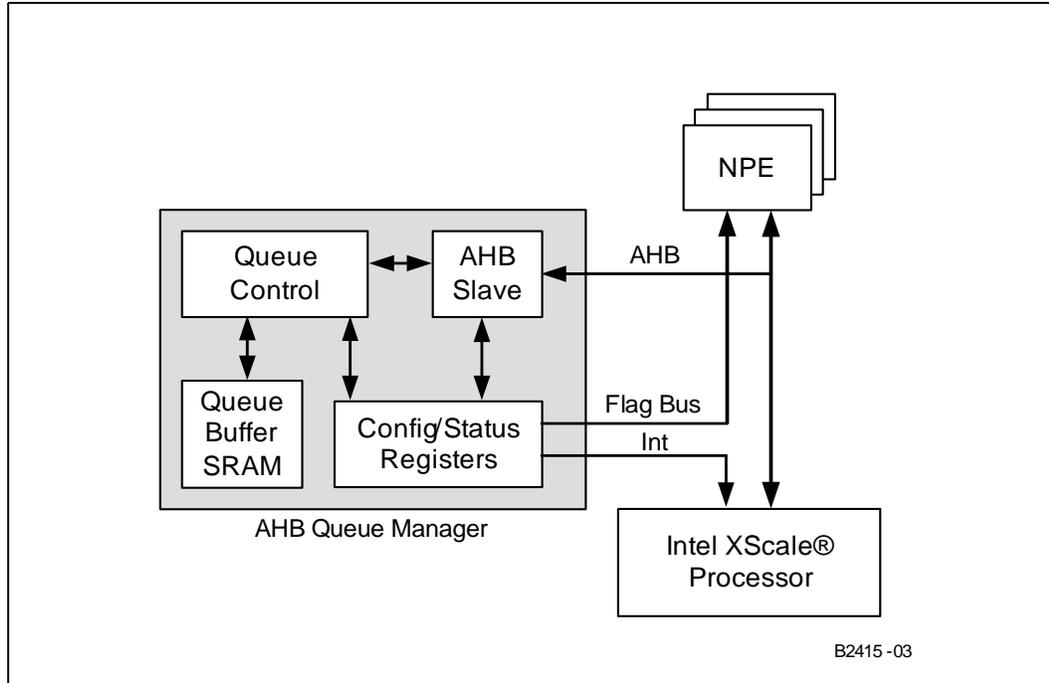
The IxQMgr (Queue Manager) access-layer component is a collection of software services responsible for configuring the Advanced High-Performance Bus (AHB) Queue Manager (also referred to by the combined acronym AQM). IxQMgr is also responsible for managing the flow of IX\_OSAL\_MBUF buffer pointers to and from the NPEs and client Intel XScale® Processor software. To do this, the IxQMgr API provides a low-level interface to the AQM hardware block of the Intel® IXP4XX Product Line of Network Processors. Other access-layer components can then use IxQMgr to pass and receive data to and from the NPEs through the AQM.

This chapter presents the necessary steps to start the IxQMgr component. A high-level overview of AQM functions is also provided. The IxQMgr component acts as a pseudo service layer to other access-layer components such as IxEthAcc.

In the sections that describe how the Queue Manager works, the "client" is an access component such as IxEthAcc. Application programmers will only need to write code for the initialization and uninitialization of the IxQMgr component, not for directly controlling the queues and the AQM.

### 18.3 Features and Hardware Interface

Figure 94. AQM Hardware Block



The IxQMGr provides a low-level interface for configuring the AQM, which contains the physical block of static RAM where all the data structures (queues) for the IxQMGr reside. The AQM provides 64 independent queues in which messages, pointers, and data are contained. Each queue is configurable for buffer and entry size and is allocated a status register for indicating relative fullness.

The AQM maintains these queues as circular buffers in an internal, 8-Kbyte SRAM. Status flags are implemented for each queue. The status flags for the lower 32 queues are transmitted to the NPEs via the flag data bus because the notification sources are programmable. There are no status flags for the upper 32 queues because the notification source cannot be changed. Two interrupts — (QM1) one for the lower 32 queues and (QM2) one for the upper 32 queues — are used as queue status interrupts.

The AHB interface provides for complete queue configuration, queue access, queue status access, interrupt configuration, and SRAM access.

IxQMGr provides the following services:

- Configures AQM hardware queues.  
Configuration of a queue includes queue size, entry size, watermark levels, and interrupt-source-select flag. IxQMGr checks the validity of the configuration parameters and rejects any configuration request that presents invalid parameters.
- Allows callbacks to be registered for each queue. This is also referred as notification callback.
- Enables and disables notifications for each queue.
- Sets the priority of a callback.
- Provides queue-notification source-flag select.



- For queues 0-31, the notification source is programmable as the assertion or de-assertion of one of four status flags: Empty, Nearly Empty, Nearly Full, and Full.
- For queues 32-63, the notification source is fixed — the assertion of the Nearly Empty flag.
- Performs queue-status query.
  - For queues 0-31, the status consists of the flags Nearly Empty, Empty, Nearly Full, Full, Underflow and Overflow.
  - For queues 32-63, the status consists of the flags Nearly Empty and Full.
- Determines the number of full entries in a queue.
- Determines the size of a queue in entries.
- Reads and writes entries from/to AQM.
- Dispatches queue notification callbacks registered by clients. These are called in a defined order, based on a set of conditions.
- Enable/Disable AQM sticky interrupt register for IXP42X product line B0 and IXP46X product line silicon.

## 18.4 IxQMgr Initialization and Uninitialization

The initialization of IxQMgr first requires a call to `ixQMgrInit()`, which takes no parameters and returns success or failure. No other `ixQMgr` functions may be called before this. Following initialization, the queues must be configured, and the dispatcher function should be called. Only one dispatcher can be invoked per each set of upper and lower 32 queues.

To uninitialize the IxQMgr component, call the `ixQMgrUnload()` function, which also takes no parameters and returns success or failure. Uninitialization should be done prior to unloading components that are dependant on IxQMgr. Uninitialization will unmap kernel memory mapped by the component.

Uninitialize the QMgr will basically clear the Queues' configuration. For example, ATM's Queues configuration like `TxDoneQ`, `TxQ`, `RxFreeQ`, and so forth, is cleared. Once this is cleared, the ATM component will not be able to use the QMgr. This uninitialization should be done only when the client is ready to un-init other components that use this component.

To avoid unpredictable results, the `ixQMgrUnload` function should not be called twice in sequence before a call to `ixQMgrInit`. No other `ixQMgr` functions may be called after `ixQMgrUnload` except for `ixQMgrInit`.

## 18.5 Queue Configuration

The queue base address in AQM SRAM is calculated at run time. The IxQMgr access-layer component must be initialized by calling `ixQMgrInit()` before any queue is configured. Queue configurations include queue size, queue entry size, queue watermarks, interrupt enable/disable and callback registration. A check is performed on the queue configuration to ensure that the amount of SRAM required by the configuration does not exceed the amount available. The Queue configuration function `ixQMgrQConfig()` provides a configuration interface to the AQM queues. With the exception of `ixQMgrQWatermarkSet()`, the queue-configuration information to which this interface provides access can only be set once.



## 18.6 Queue Identifiers

An AQM hardware queue is identified by one of the 64 unique identifiers. Each IxQMGr interface function that operates on a queue takes one of the 64 identifiers (defined in `ixp400_xscale_sw/src/include/IxQMGr.h`) as a parameter and it is the clients responsibility to provide the correct identifier.

## 18.7 Configuration Values

Table 67 details the attributes of a queue that can be configured and the possible values that these attributes can have (word = 32 bits).

Table 67. AQM Configuration Attributes

Attribute	Description	Values (word)
Queue Entry Size	The size of a queue entry in words.	1, 2 or 4
Queue Size	The size of the buffer used to store the queue entries in words. max entries= (Queue size) / (Queue entry size).	16, 32, 64, or 128
NE Watermark	When the number of entries is less than or equal to this value, the queue is considered NE - nearly empty.	0, 1, 2, 4, 8, 16, 32, or 64
NF Watermark	When the number of entries is greater than or equal to this value, the queue is considered NF - nearly full.	0, 1, 2, 4, 8, 16, 32, or 64

## 18.8 Dispatcher

The IxQMGr access-layer component provides a dispatcher to enable clients to register notification callbacks to be called when a queue is in a specified state. A queue's state is defined by the queue status flags E, NE, NF, F, NOTE, NOTNE, NOTNF, and NOTF. Each queue will have its own watermark level defined, which triggers a change in its status flag and generates an interrupt to the Intel XScale® Processor. The QM1 Queue Manager interrupt to the Intel XScale® Processor represents a change in the queue status for lower queues 0-31, and the QM2 interrupt represents a change in the queue status for upper queues 32-63.

In case of the upper queues 32-63, the notification occurs on change of the Nearly Empty flag and the watermark levels cannot be changed. The watermark level triggers the change of the status flag for a particular queue, and the lower queues 0-31 provide additional control when the interrupt gets triggered.

Prior to start of the dispatcher, `ixQMGrDispatcherLoopGet()` is used to get a pointer to the correct queue dispatcher. The function pointer being returned in response to `ixQMGrDispatcherLoopGet()` is — in the remainder of this section — referred to as the "dispatcher". There are three dispatchers in the IxQMGr component that may be returned to `ixQMGrDispatcherLoopGet()`.

- **`ixQMGrDispatcherLoopRunB0`** - This is the default dispatcher for IXP42X product line B0 stepping and all IXP46X product line processors are detected.
- **`ixQMGrDispatcherLoopRunB0LLP`** - This dispatcher is a variation of the `ixQMGrDispatcherLoopRunB0` dispatcher that adds LiveLock Prevention support (refer to "LiveLock Prevention" on page 296). The `IxFeatureCtrl` component is used to select whether this dispatcher is to be selected or not, as described in Section 12.4.

There is no assumption made about how the dispatcher is called. For example, `ixQMGrDispatcherLoopRunB0()` or `ixQMGrDispatcherLoopRunB0LLP()` may be registered as an ISR for the AQM interrupts, or it may be called from a client polling mechanism,



which calls the dispatcher to read the queues status at regular intervals. In the first example, the dispatcher is called in the context of an interrupt and the dispatcher gets invoked when the queue status change.

A parameter passed to the *ixQMgrDispatcherLoopRun()* function determines whether the lower set of 32 queues, queues 0-31 or the upper set of 32 queues, queue 32-63 are serviced by the dispatcher each time *ixQMgrDispatcherLoopRun()* is called. The order in which queues are serviced depends on the priority specified by calling *ixQMgrQDispatchPrioritySet()*.

**Note:** Application software does not need to access the queues directly. The underlying access-layer component software (for example, EthAcc, HssAcc, and so forth) handles this. However, the application software does need to initialize the queue manager software using *ixQMgrInit* and set up the dispatcher operation.

## 18.9 Dispatcher Modes

The Codelet/Customer code must first initialize the IxQMgr by making a call to *ixQMgrInit()*, which takes no parameters and returns success or failure. No other IxQMgr functions may be called by other access-layer components before this. After initialization, the queues must be configured before they can be used.

**Note:** The *ixQMgrInit()* function should only be called once for a system. Once the IxQMgr has been started all other access-layer components can register to use the services it provides without calling *ixQMgrInit()*.

The access-layer provides the following services for the application by performing the following functions:

- Perform Queue configuration
- Set the watermark levels. (The NF watermark is ignored for upper 32 queues.)
- Reads and writes entries to and from AQM
- Provides register-notification callbacks for a queue
- Set the priority of a dispatcher callback

Once the IxQMgr is initialized, the access component configures the queues. Queue configuration is done by setting up the attributes for respective queues. These attributes are typically set in the access components by using *ixQMgrQConfig()* and *ixQMgrWatermarkSet()* functions. Depending upon whether the queue is half full, nearly full, and so forth, the watermark level triggers the change of the status flag for a particular queue. The queue configuration and setting of the watermark levels and queue priority should be performed prior to enabling of the queue notification status flag. Once the queues are configured, the notification callback must be set or else it will go to a dummy callback. The Queue dispatcher loop can be started at any time following a *ixQMgrInit()*. However, the dispatcher function will service the callback only once the queue notification is enabled.

The IxQMgr governs the flow of traffic in software release 2.3. Depending upon the OS, the application, and the performance required, there are three different ways the dispatcher can be called: Busy loop, event-, or timer-based interrupt. The dispatcher can be called either in context of an interrupt or through a busy loop (which might run as a low-priority task). In case of an interrupt-driven mechanism, the interrupt can be triggered either by a timer or upon generation of QM hardware interrupts (which are event-driven). There is no single way to determine the best mechanism, although the choice of implementation would depend upon the OS, the application, and the nature of the traffic. The following includes factors to be considered in selecting the appropriate mechanism:

- Event-based interrupt – Interrupt driven through QM1 or QM2 interrupt:



- system is interrupted only when there is traffic to service
- suitable for low traffic rates
- provides lowest latency
- Timer-based interrupt (using the Performance Monitoring Unit Hardware timer)—  
polled from timer-based interrupt:
  - suitable for high traffic rates
  - minimizes the ISR overhead
  - most efficient use of the Intel XScale® Processor
- Polling mode – Busy loop to poll the queues:
  - suitable for higher traffic rates
  - throttles traffic automatically when additional cycles are not available on the Intel XScale® Processor

The status flag gets cleared within the dispatcher loop prior to servicing of the callback function. The QM1 and QM2 interrupt gets cleared when all the status flags for all the interrupt enabled queues are cleared. There can only be one dispatcher loop that can be defined for each set of queues.

Once the IxQMGr is initialized and the queues are configured, the Codelet/Customer code must determine how to invoke the dispatcher. Prior to invoking the dispatcher function, as stated before, the `ixQMGrDispatcherLoopGet(&dispatcher)`, returns a function pointer for the appropriate dispatcher. The dispatcher is invoked with an argument that points to the upper or lower 32 queues to determine if any queues in that group require servicing.

*Note:*

Only one dispatcher can be invoked per each set of upper and lower 32 queues. The client can register multiple callbacks as long as each of the callbacks are for different queues. When interrupted, the dispatcher will read the status flag register from the AQM and service only one of the callbacks that was registered for a given queue. In the event that multiple callbacks are registered for the same queue, the dispatcher will service the last registered callback.

Figure 95 shows the following sequence of events that occur when a dispatcher is run in the context of an interrupt.

At the start of the dispatcher, the interrupt register is read and written back immediately except in case of a livelock dispatcher. Since livelock prevention uses sticky interrupt, the interrupt gets cleared only when the queue threshold falls below the set watermark level. See “Livelock Prevention” on page 296.

1. The user registers a callback function with the access-layer component (for example, EthAcc). The dispatcher invokes callback in the access-layer component, and the access-layer component then invokes the user callback.
2. When the NPE receives a packet it updates the Rx Queue with location of the buffer.
3. Provided the Interrupt bit is set, when the water mark is crossed the status flag gets updated corresponding to that queue and it triggers an interrupt to the Intel XScale® Processor.
4. The Intel XScale® Processor vectors the interrupt to the corresponding ISR.
5. The ISR invokes the dispatcher.

*Note:*

In the context of an interrupt, the dispatcher can also be invoked through a timer-based mechanism.

6. The IxQMGr reads the status flag.



7. The IxQMGr access-layer component calls the registered notification.
8. The client gets the buffer pointer on the Rx queue from the access-layer through the callback. The access-layer, in turn, accesses the Rx queue through the IxQMGr access-layer component. The IxQMGr accesses the AQM hardware.

Following this, the Intel XScale® Processor may allocate a free buffer from the memory pool to the RxFree queue for the next incoming packet from the NPE.

**Figure 95. Dispatcher in Context of an Interrupt**

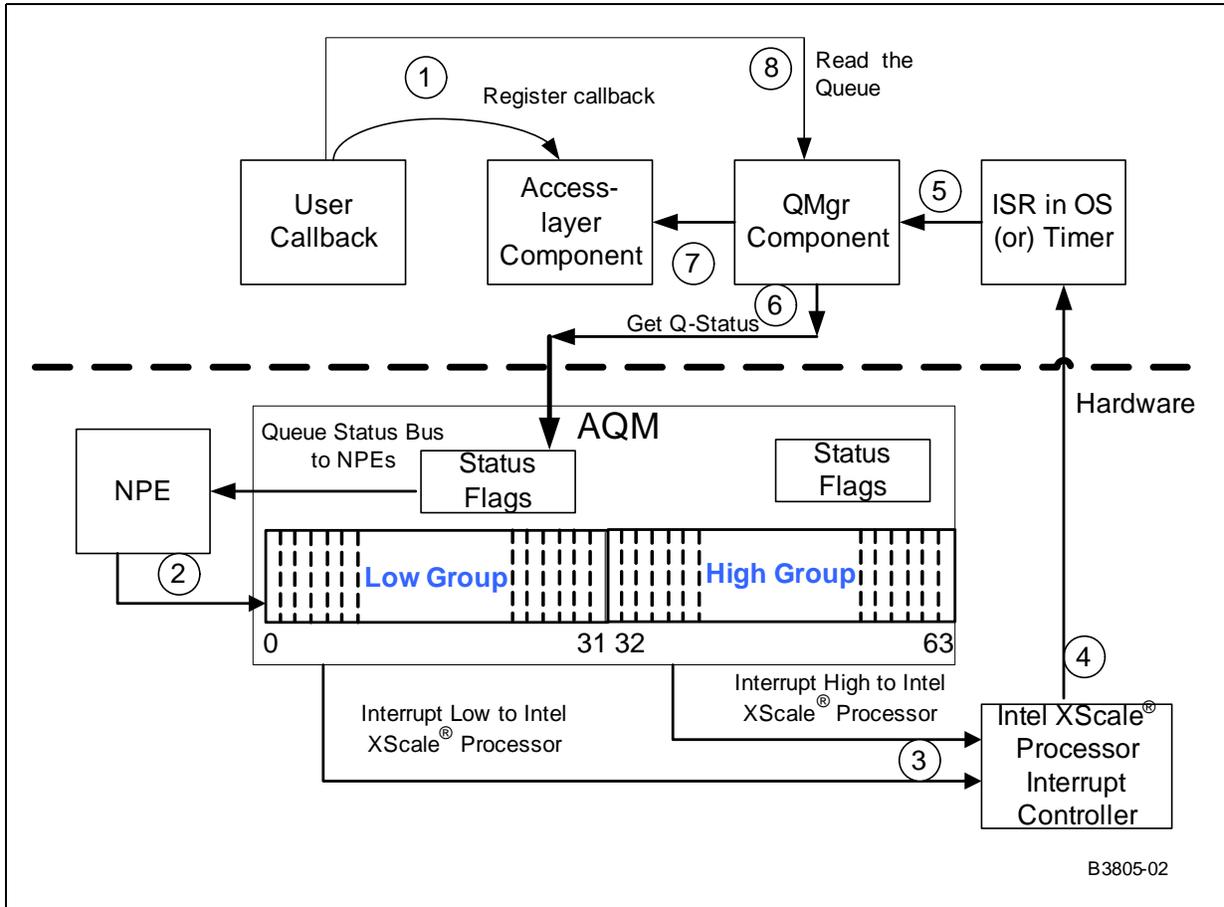


Figure 96 shows the sequence of events that occurs when a dispatcher is run in the context of a polling mechanism.

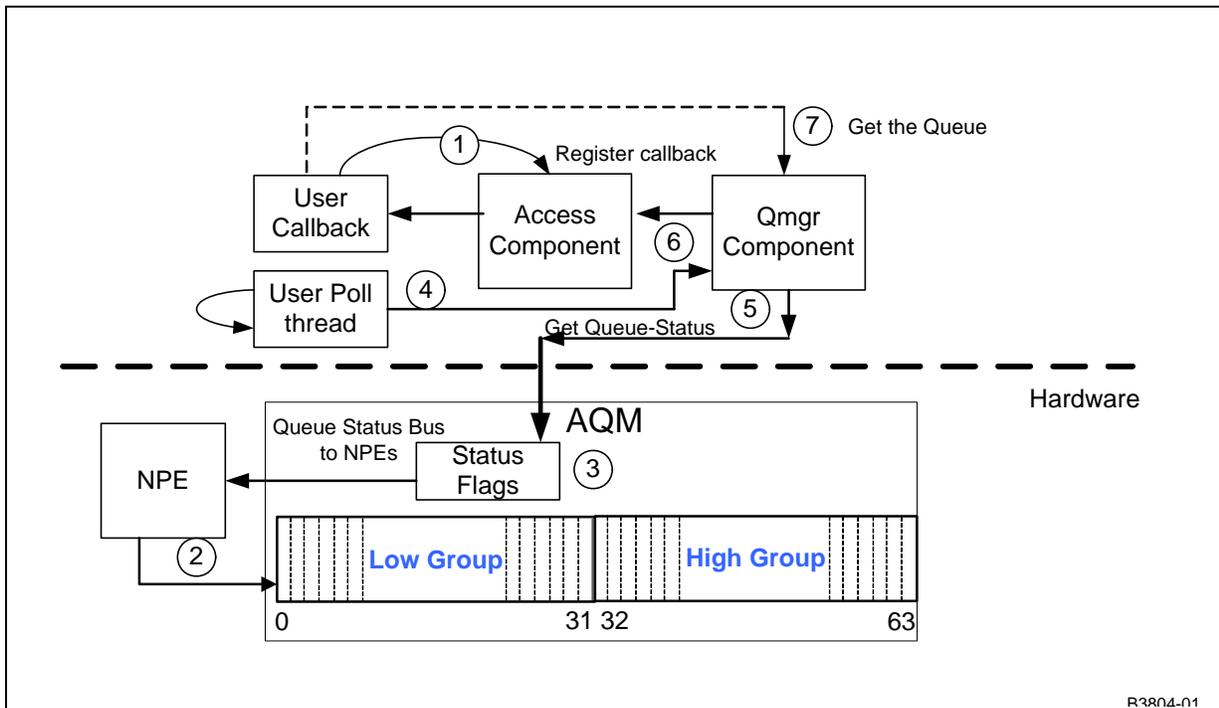
At the start of the dispatcher a call is made to read the status of the status flag to check if the queue watermark threshold has been crossed. It then immediately clears the status flag. In case of livelock prevention feature, the status flag is not cleared immediately because of the sticky interrupt implementation.

1. The user registers a callback function with the access-layer component (for example, EthAcc). The dispatcher invokes callback in the access-layer component, and the access-layer component then invokes the user callback.
2. When the NPE receives a packet, it updates the Rx queue with location of the buffer.
3. When the watermark is crossed the status flag gets updated corresponding to that queue.

4. The polling thread calls the dispatcher.
5. The dispatcher loop gets the status of the updated flag and resets it.
6. The dispatcher invokes the registered access component.
7. The access-layer components re-routes the call back to the client and the client gets the buffer pointer through the callback on the Rx queue through the access-layer.

Following this, the Intel XScale® Processor may allocate a free buffer from the memory pool to the RxFree queue for the next incoming packet from the NPE.

**Figure 96. Dispatcher in Context of a Polling Mechanism**



## 18.10 Livelock Prevention

Livelock occurs when a task cannot finish in an expected time due to it being interrupted. The livelock prevention feature allows the critical task as in case of voice processing, being serviced by a particular queue, to run for a given set of time without it being interrupted in event of a system overload. For this to happen, a periodic queue is assigned to the critical task. Periodic queues are defined as queues which generate an interrupt at a regular interval leading to a task that runs for a set length of time (periodic task). Sporadic queues are queues that can generate an interrupt at any time. Livelock prevention is used to ensure that a periodic task is not interrupted by servicing for queues set as sporadic. This is achieved by disabling notifications for sporadic queues while the periodic task is running. When the periodic task is completed the sporadic queues have their notifications re-enable. Any servicing required for sporadic queues will occur at this time. The co-existence of HSS service and Ethernet service is an example that uses livelock prevention for better efficiency.

Live lock prevention is available for IXP42X B0 and IXP46X silicon and is disabled by default. To use livelock prevention, only one queue can be set as type periodic. One or more queues may be set as type sporadic using the *ixQMGrCallbackTypeSet()* function.



By default, all the other queues that are not set to be in either a periodic or a sporadic mode are set in IX\_QMGR\_TYPE\_REALTIME\_OTHER mode. The IX\_QMGR\_TYPE\_REALTIME\_OTHER represents the default behavior of the callback function associated with respective queues when the livelock prevention feature is not in use. In a Livelock implementation, these “other” queues will not have their interrupts disabled during the servicing of the periodic queue.

The ixQMGrCallbackTypeSet() function should be used to assign IX\_QMGR\_TYPE\_REALTIME\_PERIODIC to one queue and IX\_QMGR\_TYPE\_REALTIME\_SPORADIC to **all other** queue(s) by passing a Queue-ID along with the desired queue type.

Livelock prevention is disabled by default. In order to enable the livelock option the IX\_FEATURECTRL\_ORIGBO\_DISPATCHER must be disabled using the ixFeatureCtrlSwConfigurationWrite() function before the ixQMGrInit() and ixQMGrDispatcherLoopGet() functions are called.

Queue assignments are located at ixp400\_xscale\_sw\src\include\IxQueueAssignments.h. If Ethernet QoS features are used, the Rx Priority queues are assigned in ixp400\_xscale\_sw\src\include\IxEthDBQoS.h. Queue type assignments may be checked with the ixQMGrCallbackTypeGet() function.

When ixQMGrDispatcherLoopRunB0LLP() reads the interrupt register and sees that a periodic queue is to be serviced, all queues that are set to be sporadic have their notification disabled. This prevents sporadic queues from generating interrupts, which may stall a task resulting from the periodic queue callback (periodic task). The ixQMGrPeriodicDone() function should be called after the periodic task is completed to ensure that sporadic queues are re-enabled.

*Note:* Because livelock prevention enables and disables notifications for queues set as sporadic, users should not enable and disable sporadic queues notifications other than at startup / shutdown.

*Note:* Livelock prevention operates on lower interrupt register queues only. (lower queue group 0-31).

The following is an example sequence to show how livelock would be used is to set the HSS queue to periodic and the Eth Rx queue to sporadic using the ixQMGrCallbackTypeSet() function. When codec processing (the periodic task) as a result of a HSS callback is finished, the ixQMGrPeriodicDone() function is called and Eth Rx is then serviced. This will ensure that any codec processing that is done as a result of HSS notifications is not interrupted by a burst in Eth Rx.

- Use ixFeatureCtrlSwConfigurationWrite() to disable IX\_FEATURECTRL\_ORIGBO\_DISPATCHER.
- Initialize the Queue Manager by using ixQMGrInit().
- Make a call to ixQMGrDispatcherLoopGet() to get the appropriate dispatcher function for livelock functionality.
- Initialize access-layer components, register the callback functions.
- Set the callback type for the HSS queue to periodic and the Eth Rx queue to sporadic using the ixQMGrCallbackTypeSet() function.

*Note:* All other queues (Tx queues, RxFree queues and TxDone queues) will have the callback type set to the default callback type of IX\_QMGR\_TYPE\_REALTIME\_OTHER.

- Start the dispatcher by calling the ixQMGrDispatcherLoop function.
- On completion of the periodic task, make a call to the ixQMGrPeriodicDone() function to enable the sporadic task.



## 18.11 Threading

The IxQMgr does not perform any locking on accesses to the IxQMgr registers and queues. If multiple threads access the IxQMgr, the following IxQMgr functions need to be protected by locking during concurrent access to the same queue:

- *ixQMgrQWrite()*
- *ixQMgQRead()*
- *ixQMgrQReadWithChecks()*
- *ixQMgrQWriteWithChecks()*
- *ixMgrQBurstRead()*
- *ixQMgrQBurstWrite()*
- *ixQMgrQReadMWordsMinus1()*
- *ixQMgrQPeek()*
- *ixQMgrQPoke()*
- *ixQMgrQNotificationEnable()*
- *ixQMgrQNotificationDisable()*
- *ixQMgrQStatusGet()*
- *ixQMgrQWatermarkSet()*
- *ixQMgrDispatcherLoopRunB0/BOLLP()*

All IxQMgr functions can be called from any thread, with the exception of *ixQMgrInit()*, which should be called only once — before any other call.

## 18.12 Dependencies

The IxQMgr component is dependant on the OSAL and Feature Control components. IxQMgr uses OSAL in *ixQMgrDispatcherInterruptConnect()* to register AQM ISRs. QMgr also uses IxFeatureCtrl to determine whether the underlying silicon is IXP42X product line B0, or IXP46X product line silicon. If the silicon is IXP42X B0 or IXP46X, Feature Control is used to determine whether the live lock prevention enabled dispatcher should be used.

## 18.13 NPE Parity Error Handling

IxQMgr fulfills the soft-error handling requirements by providing the following new functions or new APIs.

```
ixQMgrDispatcherLoopDisable(void)
```

The Queue Manager dispatcher loop is enabled by default. The ixErrHdlAcc component relies on the call to the *ixQMgrDispatcherLoopDisable()* to stop the queue dispatcher especially if it is called by an interrupt context of an ISR binded to the Queue Manager Interrupt IRQ. In an interrupt dispatcher mode, the call to *ixQMgrDispatcherLoopDisable()* disables the interrupt IRQ of Lower Queues (for queues(0-31)) queue manager IRQ and Upper Queues (for queues 32-63) queue manager IRQ. Like wise, in a polling dispatcher mode, the call to



ixQMgrDispatcherLoopDisable() disables any QM Read or Write to any of the queues and any other operations on the QM Registers. Thus, there will not be any changes to the queue status till the soft-error recovery is completed.

```
ixQMgrDispatcherLoopEnable(void)
```

It can be re-enabled by a call to ixQMgrDispatcherEnable().

```
ixQMgrDispatcherInterruptModeSet(BOOL mode)
```

The client or application must call ixQMgrDispatcherInterruptModeSet() after the initialization of the queue manager to indicate if (BOOL True) the dispatcher is in interrupt mode or (BOOL False) the dispatcher is in poll dispatcher mode. The queue dispatcher mode is by default in non-interrupt mode or poll mode.

```
ixQMgrDispatcherLoopStatusGet(void)
```

The call to the API of ixQMgrDispatcherLoopStatusGet() which returns IX\_QMGR\_DISPATCHER\_LOOP\_FREE if the QM dispatcher has stopped and IX\_QMGR\_DISPATCHER\_LOOP\_BUSY if the dispatcher loop is still busy.

**Note:** The above APIs are not advisable to be used for other purposes than in the event of soft reset.

§ §





## 19.0 Access-Layer Components: Synchronous Serial Port (IxSspAcc) API

---

This chapter describes the Intel® IXP400 Software v2.3's "SSP Serial Port (IxSspAcc) API" access-layer component.

### 19.1 What's New

There are no changes or enhancements to this component in software release 2.3.

### 19.2 Introduction

A Synchronous Serial Port is included in the Intel® IXP46X Product Line of Network Processors. The IxSspAcc API is provided to allow the configuration of the various registers related to the SSP hardware. Once configured, the API also provides the ability to transfer data to the Tx FIFO and from the Rx FIFO. Both polling and interrupt modes are supported.

### 19.3 IxSspAcc API Details

#### 19.3.1 Features

This component provides capabilities to:

- select frame format – SSP, SPI, or Microwire\*
- select data sizes – 4 to 16 bits
- select clock source – external or on-chip
- configure serial clock rate – to drive a baud rate of 7.2 Kbps to 1.8432 Mbps (if internal clock source is selected only)
- enable/disable the receive FIFO level interrupts
- enable/disable the transmit FIFO level interrupts
- set the transmit FIFO threshold – 1 to 16 frames
- set the receive FIFO threshold – 1 to 16 frames
- select operation mode – normal or loop-back operation
- select SPI SCLK polarity – polarity of SCLK idle state is low or high (only used in SPI format)
- select SPI SCLK phase – phase of SCLK starts with one inactive cycle and ends with ½ inactive cycle or SCLK starts with ½ inactive cycle and ends with one inactive cycle (only used in SPI format)
- select Microwire control word format – 8 or 16 bits
- enable/disable the SSP serial port hardware

This component also provides status and statistics for:

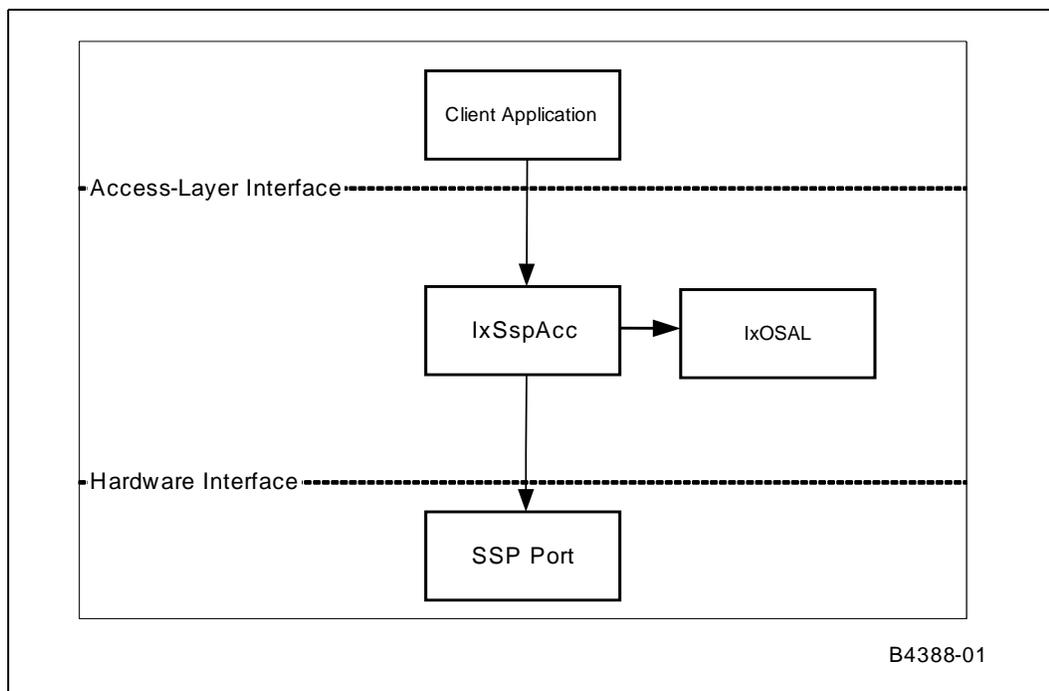


- SSP state – busy or idle
- Transmit FIFO level – 0 to 16 frames
- Receive FIFO level – 0 to 16 frames
- Transmit FIFO hit or below its threshold
- Receive FIFO hit or exceeded its threshold
- Receive FIFO overrun.
- Statistics for frames received, frames transmitted, and number of overrun occurrences.

### 19.3.2 Dependencies

IxSspAcc is dependent on the capability provided by the SSP serial port hardware. IxOSAL provides OS independency.

Figure 97. IxSspAcc Dependencies



## 19.4 IxSspAcc API Usage Models

### 19.4.1 Initialization and General Data Model

This description assumes a single client model where there is a single application-level program configuring the SSP interface and initiating I/O operations.

The client must first define the initial configuration of the SSP port by storing a number of values in the `IxSspInitVars` structure. The values include the frame format, input clock source, clock frequency, threshold values for the FIFOs, pointers to callback functions for various data scenarios, and other configuration items. After the structure is defined, `ixSspAccInit()` may be called to enable the port.



Once the port is enabled, the client will use one of the data models described later in this chapter (either Interrupt or Polling mode) to determine how and when data I/O operations need to occur. A handler (or callback) is registered for transmit and receive operations. These handlers will use the **ixSspAccFIFODataSubmit()** and **ixSspAccFIFODataReceive()** functions for transmitting and receiving data.

After the SSP port has been initialized as described above, the SSP port may be re-configured. Most of the port configuration options may be modified via available functions in the API. For example, the frame format may be changed from SPI to Microwire.

The API also provides functions to disable the SSP port, check for port activity, maintains statistics for transmitted frames, received frames and overruns, and has other debugging type functions.

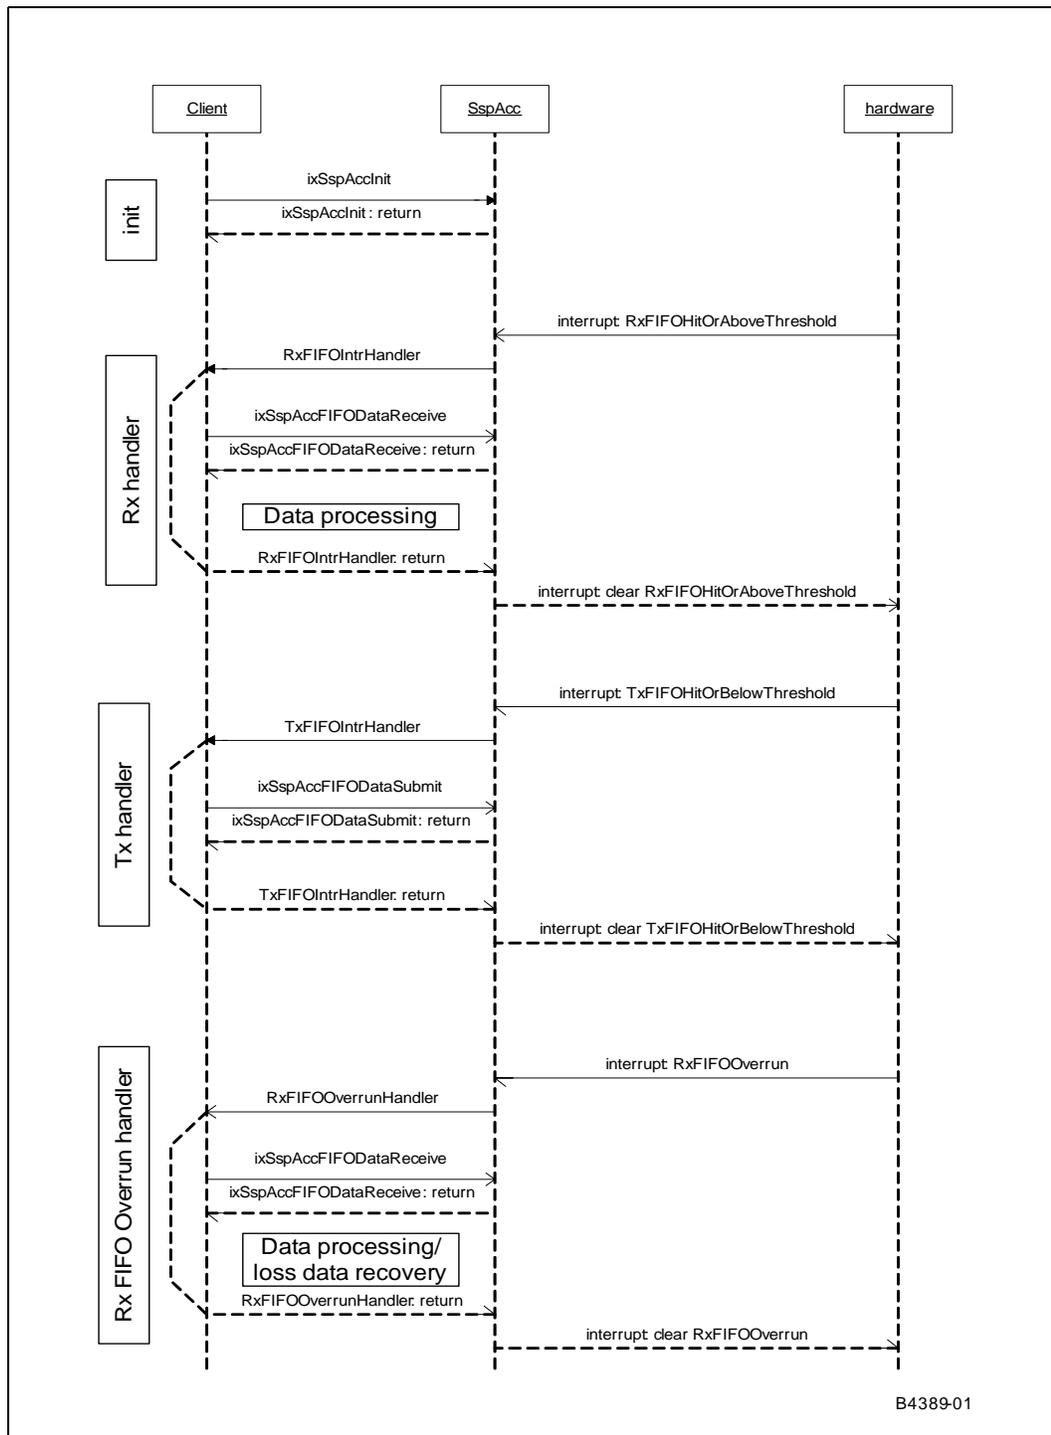
## 19.4.2 Interrupt Mode

The sequence flow for a client application using this component in interrupt mode is described below. Refer to [Figure 98](#).

1. Initialize the SSP interface with interrupts enabled.
2. For receive operations:
  - a. Interrupt is triggered due to hitting or below of threshold.
  - b. If due to Rx FIFO, Rx FIFO handler/callback is called.
  - c. Rx FIFO handler/callback extracts data from the Rx FIFO.
  - d. (handler/callback processes the extracted data)
  - e. Rx FIFO handler/callback returns.
  - f. Interrupt is cleared.
3. For transmit operations:
  - a. Interrupt is triggered due to hitting or exceeding of threshold.
  - b. If due to Tx FIFO, Tx FIFO handler/callback is called.
  - c. Tx FIFO handler/callback inserts data into the Tx FIFO.
  - d. Tx FIFO handler/callback returns.
  - e. Interrupt is cleared.
4. For an overrun:
  - a. Interrupt is triggered due to an overrun of the Rx FIFO.
  - b. Rx FIFO Overrun handler/callback is called.
  - c. Rx FIFO Overrun handler/callback extracts data from the Rx FIFO to prevent the overrun from triggering again.
  - d. (processes data extracted and perform necessary steps to recover data loss if possible)
  - e. Rx FIFO Overrun handler/callback returns.
  - f. Interrupt is cleared



Figure 98. Interrupt Scenario





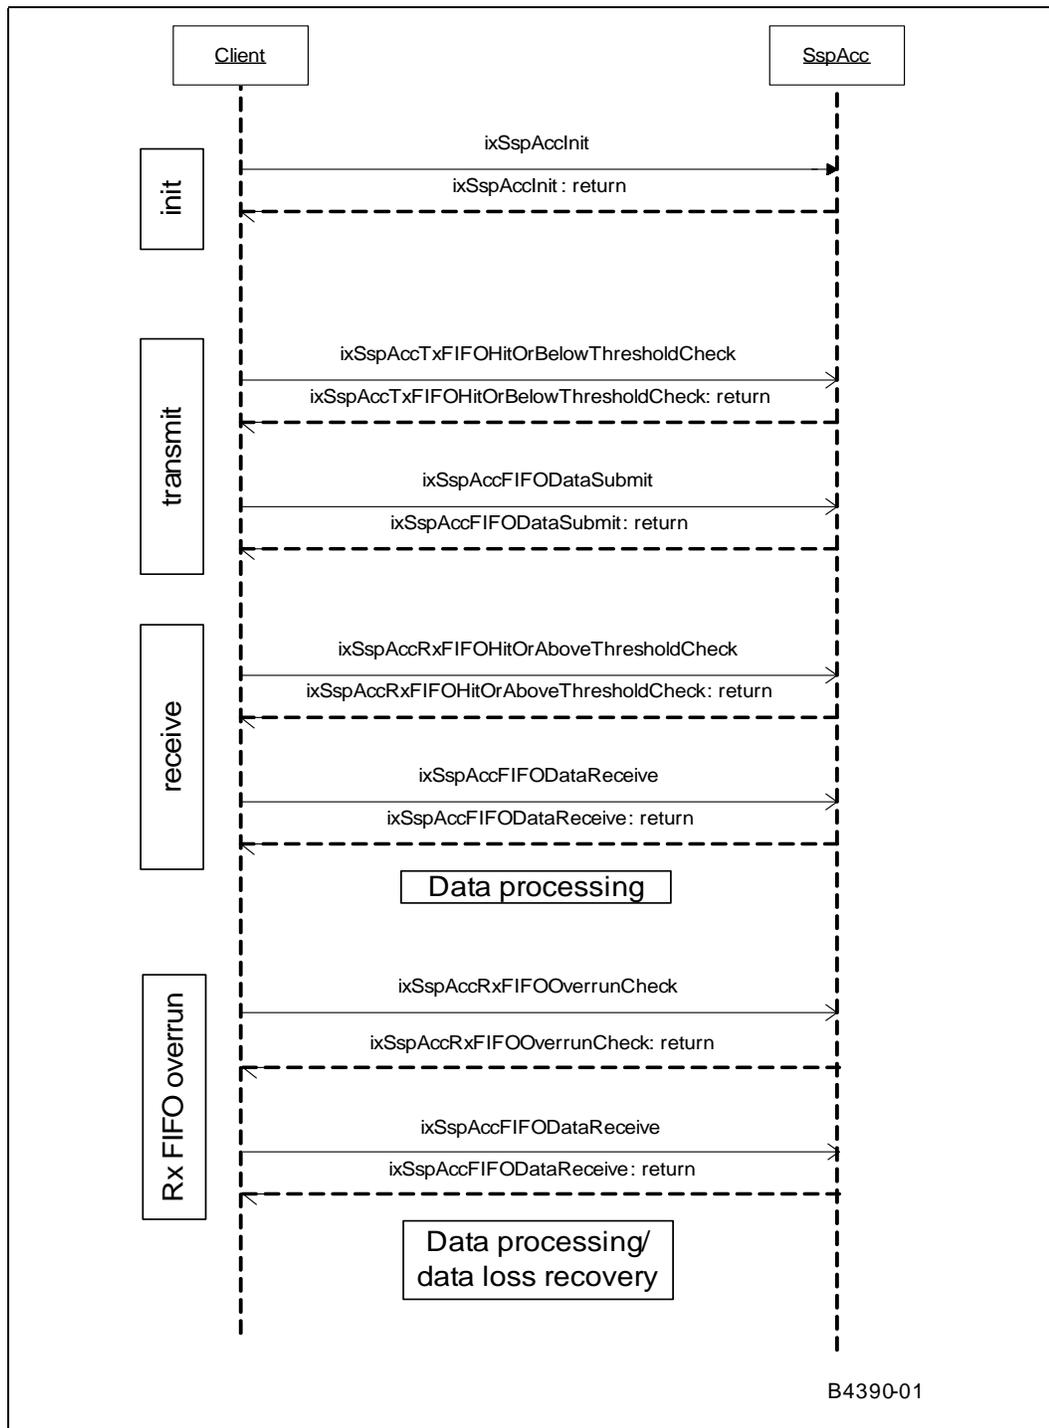
### 19.4.3 Polling Mode

The sequence flow for a client application using this component in polling mode is described below. Refer to [Figure 99](#).

1. Initialize the SSP with interrupts disabled.
2. For transmit operations:
  - a. Check if the Tx FIFO has hit or is below its threshold.
  - b. If it has, then insert data into the Tx FIFO.
3. For receive operations:
  - a. Check if the Rx FIFO has hit or exceeded its threshold.
  - b. If it has, then extract data from the Rx FIFO.
  - c. Process the data if needed.
4. For an overrun:
  - a. Check if the Rx FIFO Overrun has occurred.
  - b. If it has, then extract data from the Rx FIFO.
  - c. Process the data and recover any lost data if needed.



Figure 99. Polling Scenario





## 20.0 Access-Layer Components: Time Sync (IxTimeSyncAcc) API

---

This chapter describes the Intel® IXP400 Software v2.3's "Time Sync (IxTimeSyncAcc) API" access-layer component.

The IxTimeSyncAcc access-layer component enables a client application, which implements the IEEE 1588\* Precision Time Protocol (PTP) to configure the IEEE 1588 Hardware Assist block on the Intel® IXP46X Product Line of Network Processors.

### 20.1 What's New

There are no changes or enhancements to this component in software release 2.3.

### 20.2 Introduction

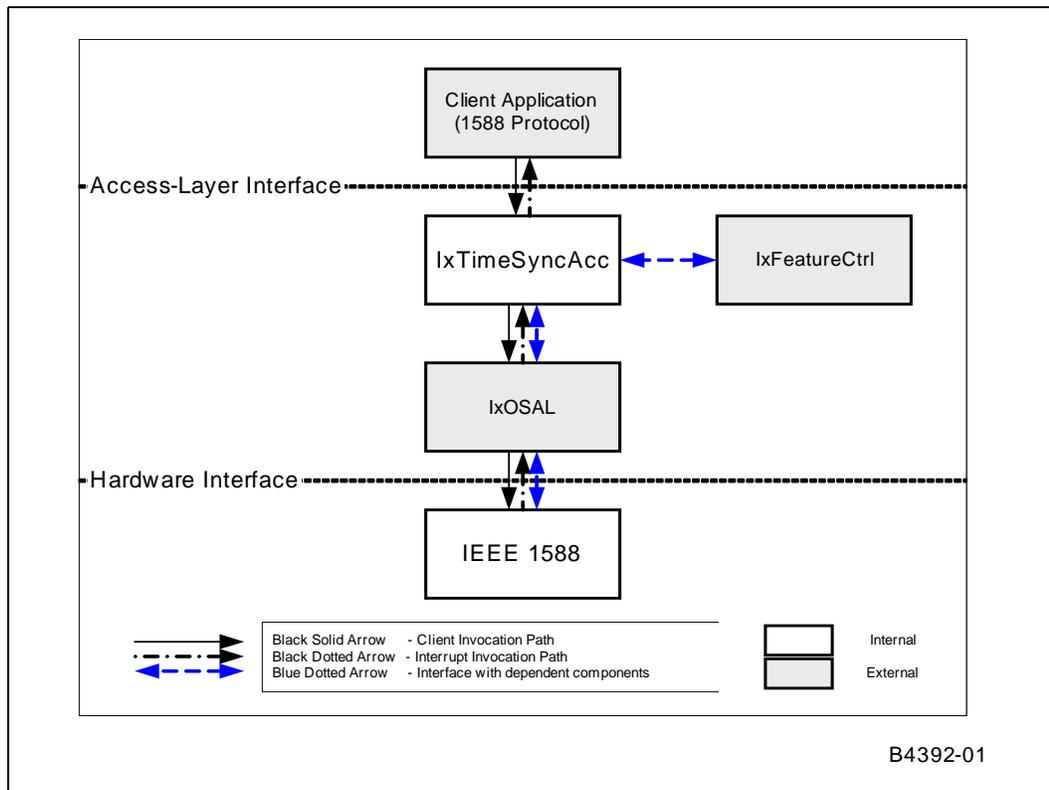
The IEEE 1588 Precision Time Protocol (PTP) is used to synchronize independent clocks running in distributed network elements/nodes to a high degree of accuracy, in the microsecond to sub-microsecond range. There are three main elements involved in supporting IEEE 1588 on the IXP46X network processors:

- IEEE 1588 Hardware Assist block, available on the IXP46X network processors. The hardware provides necessary features to allow timestamping of the IEEE 1588 PTP messages.
- IxTimeSyncAcc Access-Layer component, running on the Intel XScale® Processor. This software component provides the functionality required to enable the IEEE 1588 Hardware Assist block on various MII ports, set and receive timestamps, receive and transfer interrupt requests to client applications, and other functions.
- A IEEE 1588 PTP client application that would use the other two components to implement and use PTP messages and timestamps according to the IEEE 1588 specifications.

**Note:** This client application is not provided as part of the IXP400 software. But it does include a codelet that demonstrates the basic usage of the APIs in some IEEE 1588 scenarios. See [Chapter 23.0](#).

These three elements are depicted in [Figure 100](#).

Figure 100. IxTimeSyncAcc Component Dependencies



### 20.2.1 IEEE 1588 PTP Protocol Overview

As mentioned at the beginning of this chapter, the IEEE 1588 Precision Time Protocol (PTP) is used to synchronize independent clocks running in distributed network elements/nodes to a high degree of accuracy (in the nanosecond to sub-microsecond range). This section provides a very brief overview of the IEEE 1588 specification elements that relate to this IEEE 1588 hardware and software subsystem. For a more complete understanding of IEEE 1588, refer to *IEEE Std 1588 -2002, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, November 8, 2002 (available at <http://ieee1588.nist.gov/>).

The PTP protocol defines four timing-related messages: **Sync**, **Delay\_Req**, **Follow\_Up** and **Delay\_Resp**. Furthermore, the protocol identifies a network element/node as either a master or a slave. The sequence and usage of the protocol messages vary depending on whether the node is configured in slave or master mode. Components within the PTP messages, such as the UUID and Sequence ID fields, are used by the master and slave elements/nodes to identify themselves and relate the sequence in which the PTP messages are exchanged.

#### Synchronization Sequence

The master provides the clock source to which all the slave nodes synchronize.

The master sends a **Sync** message to the slave node, carrying in it the master node's system time as a timestamp. The master may also use **Follow\_Up** message with the timestamp of the last **Sync** message to provide more accurate timestamp details to a



slave, after accounting for the PHY, synchronization, and internal processing delays. The slave element/node, after detecting the **Sync** or **Follow\_Up** message, will begin the process to synchronize its system clock based on the master clock timestamp.

The slave may also initiate a synchronization request by sending a **Delay\_Req** message with its local system time as the timestamp to the master. The master will then respond with **Delay\_Resp**, carrying both the timestamp at which the **Delay\_Req** was received and the timestamp included by the slave in the **Delay\_Req** message. This allows the slave to determine the transit delay and accordingly to update its system time.

## 20.2.2 IEEE 1588 Hardware Assist Block

### Overview

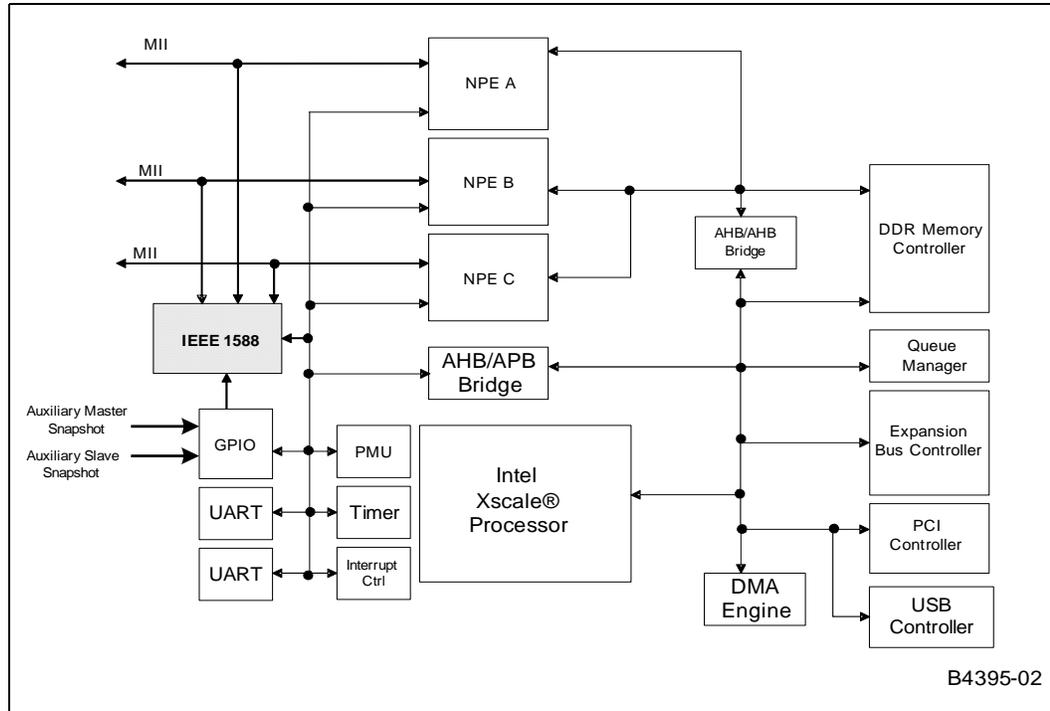
The hardware provides necessary features to allow timestamping of the IEEE 1588 PTP messages. The IEEE 1588 Hardware Assist block internally snoops the MII interfaces that extend from the NPE components on the processor to Ethernet PHYs populated on the development or customer board. This provides the IEEE 1588 Hardware Assist block with the capability to detect the transversal of PTP protocol messages between the PHY and the MAC, and set internal timestamp registers with the appropriate data from these messages. When the timestamps of inbound or outbound messages are read by the hardware, the hardware block stores this information in a register.

The IEEE 1588 Hardware Assist block maintains a system time, which can be adjusted via API by the client application. Additionally, the block can be configured to interrupt the client application if the system time exceeds a specified target value.

Although IEEE 1588 PTP can be used for time synchronization of network elements/nodes over various communication media, this IEEE 1588 Hardware Assist block is designed to detect PTP messages over the NPE Ethernet interfaces only (not over the PCI interface).

Figure 101 shows the location of the IEEE 1588 Hardware Assist block and its main interconnects to other components in the IXP46X network processors.

Figure 101. Block Diagram of Intel® IXP46X Network Processor



### Detailed Information

The IEEE 1588 Hardware Assist block implements a 64-bit register to keep track of the system time, which is used to provide timestamp references for PTP messages. The register is incremented based on a frequency scaling value, as supplied by the client application. The frequency scaling value is accumulated on every clock cycle in the system into a 32-bit register, and an overflow condition will cause the system time to increment. Thus, the slave will make use of the system time to synchronize with that of the master by adjusting the frequency scaling value based on the difference between the local system time and the master system time.

The IEEE 1588 Hardware Assist block also implements a mechanism whereby the system timer can be verified against a predefined target time for equals or exceeds conditions. Upon these conditions, the hardware block can interrupt the Intel XScale® Processor, unless the interrupt is masked off. If the interrupt is masked off, the said condition is flagged. This interrupt or event may be used by client applications to update the frequency scaling and/or to set new system time and target time values. However, it is not mandatory to make use of this hardware feature to enable timestamping.

A timestamp may be generated for each of the channels (for example, on both incoming and outgoing MII ports of an NPE) whenever the **Sync** and **Delay\_Req** messages are detected (for example, sent or received). These timestamps are captured into respective transmit or receive snapshot registers. Corresponding event flags are set and is locked unless no errors are encountered. They can be reset by clearing their corresponding events.



The IEEE 1588 Hardware Assist block can also be set explicitly to handle timestamping for all messages detected on a channel, as determined by the detection of an Ethernet Start of Frame Delimiter (SFD). In this scenario, the snapshot registers containing the timestamps will not be locked. This usage model is useful for network traffic analysis applications.

Besides the timestamps, the hardware will also capture the UUID and Sequence ID for the Delay\_Req and Sync messages received in Master and Slaves modes, respectively.

An auxiliary timestamp feature is also provided in the IEEE 1588 Hardware Assist block, allowing for the capture of system time to be trigger via the GPIO pins. The slave or master timestamp is captured when the appropriate GPIO pins (8 and 7, respectively) are triggered by the Intel XScale® Processor or an external device. When these timestamps are captured, the Intel XScale® Processor is notified through interrupts or sets event flags, depending on whether the interrupts are masked off or not.

*Note:* On the IXDP465 platform, the Auxiliary Timestamp signal for slave mode is tied to GPIO pin 8. This signal is software routed by default to PCI for backwards compatibility with the IXDP425 / IXCDP1100 platform. This routing must be disabled for the auxiliary slave time stamp register to work properly. Refer to the *Intel® IXDP465 Development Platform User's Guide* or the BSP/LSP documentation for more specific information.

The hardware assist can be reset by software and will reflect the same state as can be observed on power-on reset. Table 68 summarizes the default behavior of certain hardware features upon power-on reset or software-initiated reset.

Upon reset, the system time, frequency scaling value and target time are all set to zero. Thus, at the time of power-on reset and software-initiated reset, the frequency scaling value will not increment. This value must be set to a non-zero value to allow the system time to increment. The UUID and Sequence ID are also cleared to zeros. A UUID with value zero is treated as invalid.

**Table 68. Default IEEE 1588 Hardware Assist Block States upon Hardware/Software Reset**

Hardware Feature	Options	Default State
Channel Mode	- Master - Slave	Each channel operates in slave mode.
TimeStamp	- <b>Sync</b> and <b>Delay_Req</b> messages only - All IPv4 packets	Timestamp is taken for valid <b>Sync</b> and <b>Delay_Req</b> messages and locked in the receive and transmit snapshot registers, respectively, since the default channel mode of operation is slave.
Auxiliary Master Mode Snapshot Interrupt Mask	- Enabled - Disabled	Disabled
Auxiliary Slave Mode Snapshot Interrupt Mask	- Enabled - Disabled	Disabled
Target Time Interrupt Mask	- Enabled - Disabled	Disabled

### IPv6 and VLAN-Tagged Ethernet Frames

The IEEE 1588 Hardware Assist block does not support the IPv6 protocol. It verifies that the Ethernet frame contains an IPv4 packet by checking for a value of 0x45 in the first byte of the IP datagram header. 0x45 represents a value of 4 in the **Version** field and a 20-byte IP header length.



VLAN-tagged Ethernet frames include an additional four bytes prior to the beginning of the original Ethernet **Type/Length** field. The IP header immediately follows the **Type/Length** field. VLAN-tagged Ethernet frames can be identified by the value of 0x8100 at offset 12 and 13 of the Ethernet frame. If the IEEE 1588 Hardware Assist block identifies a value of 0x8100 (for example, **VLAN TPID** field) at this offset, it will adjust the offsets it uses to support PTP messages by four bytes.

*Note:* Some popular Ethernet switch PHY chips use the same bytes in VLAN-tagged frames to encode the port through which a frame is received. These devices encode the physical port from which a frame is received in the least-significant four bits of offset 13. The IEEE 1588 Hardware Assist block is unable to detect **Sync** and **Delay\_Req** messages in this scenario.

### Additional Hardware Information

For more information on the IEEE 1588 Hardware Assist block, refer to the Intel hardware documentation for the Intel® IXP46X Product Line.

## 20.2.3 IxTimeSyncAcc

The IxTimeSyncAcc access-layer component provides a software interface to configure the IEEE 1588 Hardware Assist block, and provide access to the snapshot register data. More details are provided in “IxTimeSyncAcc API Details” on page 312.

## 20.2.4 IEEE 1588 PTP Client Application

A IEEE 1588 PTP client application is application code running on the Intel XScale® Processor that utilizes the IxTimeSyncAcc API (and other APIs in the IXP400 software) to implement and use PTP messages and timestamps according to the IEEE 1588 specifications.

The IXP400 software does not provide this client application, although it does include a codelet that demonstrates the basic usage of the APIs in some IEEE 1588 scenarios. Refer to [Chapter 23.0](#).

A common scenario would involve a IEEE 1588 client application implementing a slave, master, or boundary clock on the target hardware platform. When transmitting PTP protocol messages, the client application would need to obtain the appropriate timestamp information from IxTimeSyncAcc, construct the appropriate PTP protocol messages, and transmit the messages using the Ethernet subsystem of the IXP400 software. When receiving PTP protocol messages, the client application may poll via the IxTimeSyncAcc API for the existence of new timestamp and other related PTP message information. If the remainder of the PTP message content is of interest to the client application, it will need to receive the Ethernet frame via the Ethernet subsystem of the IXP400 software (for example, IxEthAcc).

When operating over Ethernet networks, these messages are carried in frames using the UDP transport-layer. UDP does not guarantee successful message transfer between sending and receiving nodes, and the IEEE 1588 client application must take this behavior into account.

## 20.3 IxTimeSyncAcc API Details

### 20.3.1 Features

IxTimeSyncAcc API provides the following features:

- Configure the PTP Ports (NPE channels) to operate in master or slave mode



- Poll for Sent Timestamp of the Sync and Delay\_Req messages in both master and slave modes
- Poll for Receive Timestamp of the Delay\_Req and Sync messages in both master and slave modes
- Poll for Timestamp of all messages Sent or Received irrespective of master or slave mode
- Set and retrieve System Time
- Set and retrieve Frequency Scaling Value, based upon which the System Time is incremented
- Enable and disable system time exceeded or equaled target time notification interrupt
- Inform when system time exceeds or equals target time through a client callback
- Poll to test whether system time exceeds or is equal to the target time
- Set and retrieve Target Time
- Inform when auxiliary master or slave timestamp captured through client callback
- Poll for auxiliary master or slave timestamp
- Enable and disable auxiliary timestamp notification interrupt
- Reset IEEE 1588 Hardware Assist block to the default state as observed upon power-on reset
- Get or clear statistics on packets transmitted and received (depending on the NPE channel mode configuration, all Ethernet or Sync & Delay\_Req messages).
- Show the configuration details of the IEEE 1588 Hardware Assist block (for example, contents of control and event registers, all snapshot registers, interrupts/ events asserted or pending).

### 20.3.2 Dependencies

Dependencies for IxTimeSyncAcc are shown in “[IxTimeSyncAcc Component Dependencies](#)” on page 308. These dependencies include:

- **IxFeatureCtrl** – This component is used to verify support for the IEEE 1588 Hardware Assist block in the Intel® IXP4XX product line processors. It also is used to confirm the availability of NPE ports.
- **IxOSAL** – This component makes use of the IxOSAL services for error logging or reporting as part of the standard error handling mechanism in the IXP400 software. IxOSAL also provide mutex locking, ISR registration, and access to hardware registers.

*Note:* Depending on the design and purpose of the client application, dependencies may exist to other access components besides IxTimeSyncAcc and the dependencies listed here.

### 20.3.3 Error Handling

IxTimeSyncAcc returns IX\_FAIL and other status values under the following circumstances:

- Inappropriate parameter values passed to an API
- Incorrect sequence of invocation of the APIs
- Polled mode request while interrupt mode is set
- Internal errors



IxTimeSyncAcc returns IX\_SUCCESS when errors are not observed. The client application is expected to handle these errors/values appropriately.

## 20.4 IxTimeSyncAcc API Usage Scenarios

The following scenarios present usage examples of the interface by a client application. They are each independent but, depending on the needs of the client application, could be intermixed.

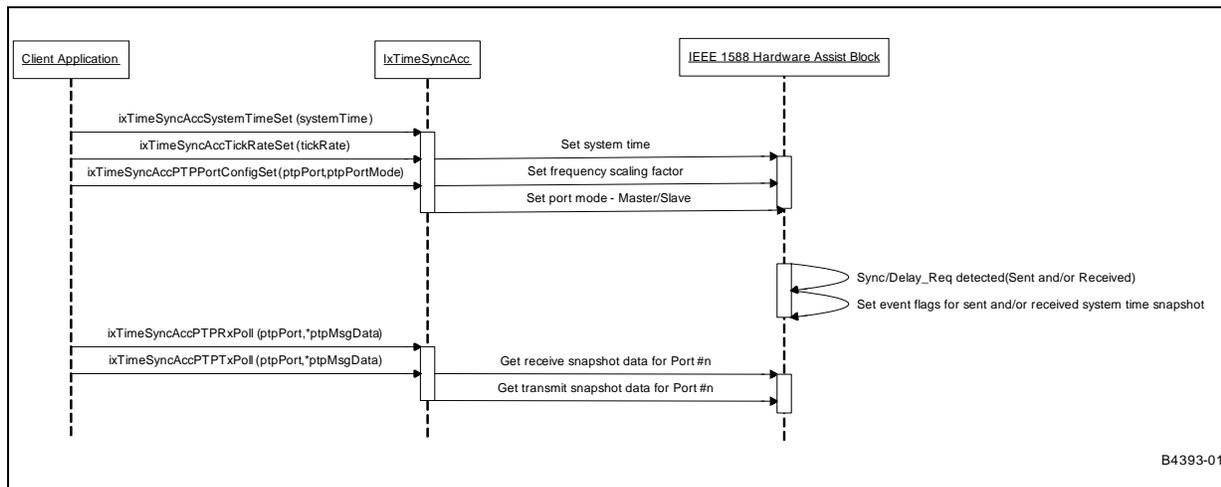
### 20.4.1 Polling for Transmit and Receive Timestamps

The IEEE 1588 Hardware Assist block detects a PTP message and then sets an event flag. The client application may poll for receive and/or transmit timestamps before or after the actual Sync/Delay\_Req message detection, which sets the event flags. The timestamps returned are valid only when the respective event flags are set. After the valid timestamps are retrieved, the event flags are cleared to allow for capturing new timestamps.

The IEEE 1588 Hardware Assist block indicates the availability of transmit and receive timestamps on the MII interfaces through events only. In other words, interrupts are not defined for these conditions (unlike the auxiliary timestamps and target time reached conditions, described later). The client application has to poll for these events to obtain the timestamps.

Figure 102 presents the timestamp polling flow.

Figure 102. Polling for Timestamps of Sync or Delay\_Req



### 20.4.2 Interrupt Mode Operations

The IxTimeSyncAcc component uses a single interrupt on IXP46X network processors to provide the client application with Target Time hit conditions or Auxiliary Master/Slave Timestamps. It implements the following priority order when the interrupt is asserted to the Intel XScale® Processor:

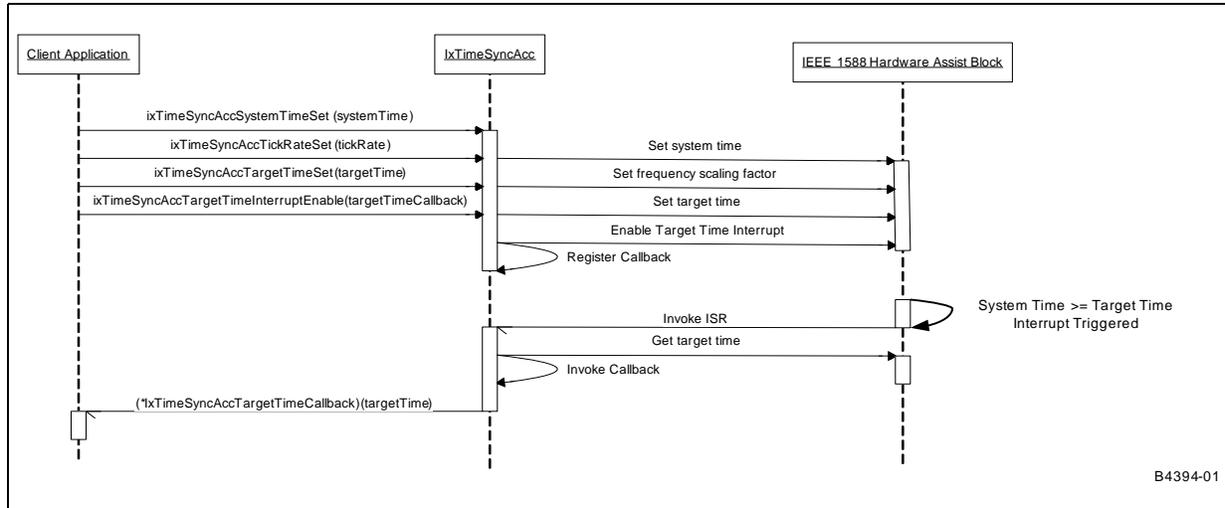
1. Target Time Reached/Hit Condition
2. Auxiliary Master Timestamp
3. Auxiliary Slave Timestamp



In order to avoid repeated invocation of the Interrupt Service Routing for the “target time reached” condition, the client application callback routine will need to either disable the interrupt handling, invoke the API to set the target time to a different value, or change the system timer value.

Figure 103 presents a scenario where the system time and target time are set, a “target time reached” condition is met, and an interrupt is used to notify the client application. A polled-mode scenario would operate similarly to what is described in “Polled Mode Operations” on page 315.

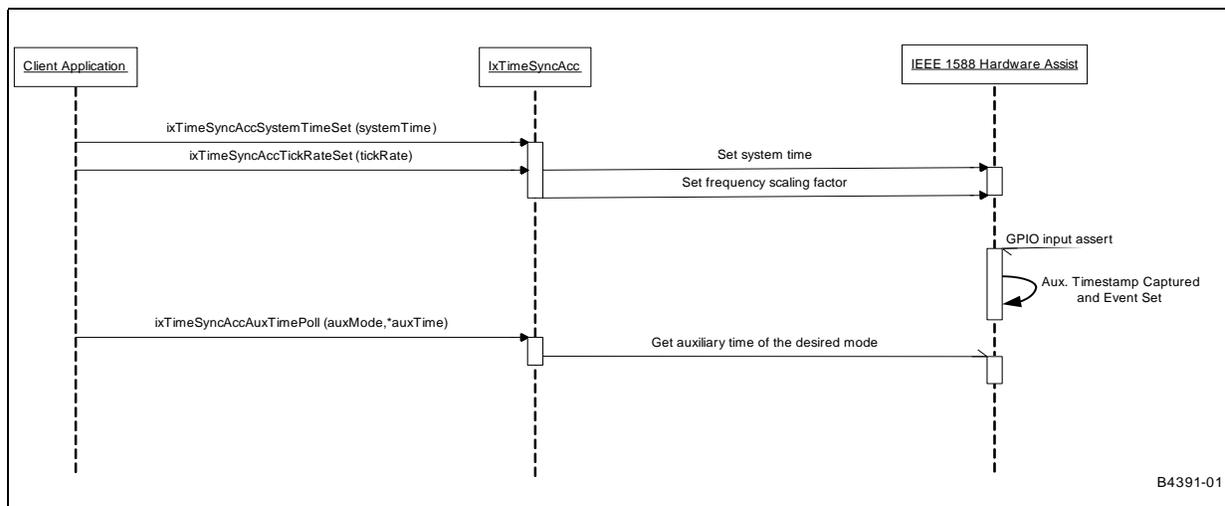
**Figure 103. Interrupt Servicing of Target Time Reached Condition**



### 20.4.3 Polled Mode Operations

Target Time events and Auxiliary snapshots can also be serviced with polling by the client application. Figure 104 shows a scenario where the client application uses polling to periodically retrieve auxiliary snapshot data.

**Figure 104. Polling for Auxiliary Snapshot Values**



§ §





## 21.0 Access-Layer Components: UART-Access (IxUARTAcc) API

---

This chapter describes the Intel® IXP400 Software v2.3's "UART-Access API" access-layer component.

### 21.1 What's New

There are no changes or enhancements to this component in software release 2.3.

### 21.2 Overview

The UARTs of the Intel® IXP4XX product line processors have been modeled on the industry standard 16550 UART. There are, however, some differences between them which prevents the unmodified use of 16550-based UART drivers. They support baud rates between 9,600 bps and 912.6 Kbps.

The higher data rates allow the possibility of using the UART as a connection to a data path module, such as Bluetooth\*. While the UART is instantiated twice on the IXP4XX product line processors, the same low-level routines is used by both. The default configuration for the processor is:

- UART0 — Debug Port (console)
- UART1 — Fast UART (for example, Bluetooth)

Any combination of debug or high-speed UART, however, could be used.

A generic reference implementation is provided that can be used as an example for other implementations/operating systems. These routines are meant to be stand-alone, such that they do not require an operating system to execute. If a new operating system is later added to those supported, these routines can be easily modified to link in to that platform, without the need for extensive rework.

The UART driver provides generic support for polled and loop back mode only.

### 21.3 Interface Description

The API covers the following functions:

- Device initialization
- UART char output
- UART char input
- UART IOCTL
- Baud rate set/get
- Parity
- Number of stop bits



- Character length 5, 6, 7, 8
- Enable/disable hardware flow control for Clear to Send (CTS) and Request to Send (RTS) signals

## 21.4 UART / OS Dependencies

The UART device driver is an API that can be used to transmit/receive data from either of the two UART ports on the processor. However, it is expected that an RTOS will provide standard UART services independent from the *IxUartAcc* device driver. That is, the RTOS UART services will configure and utilize the UART registers and FIFOs directly.

Users of the *IxUartAcc* component should ensure that the use of this device driver does not conflict with any UART services provided by the RTOS.

### 21.4.1 FIFO Versus Polled Mode

The UART supports both FIFO and polled mode operation. Polled mode is the simpler of the two to implement, but is also the most processor-intensive since it relies on the Intel XScale® Processor to check for data.

The device's Receive Buffer Register (RBR) must be polled at frequent intervals to ascertain if data is available. This must be done frequently to avoid the possibility of buffer overrun. Similarly, it checks the Transmit Buffer Register (TBR) for when it can send another character.

The FIFO on the processor's UART is 64 bytes deep in both directions. The transmit FIFO is 8 bits wide and the receive FIFO is 11 bits wide. The receive FIFO is wider to accommodate the potentially largest data word (for example, including optional stop bits and parity  $8+2+1 = 11$ ).

Interrupts can occur in one of two ways. One is when the FIFO has reached its programmed trigger level (set by the FIFO Control Register [FCR]). The other is when a character timeout has occurred (also set in the FCR). The driver will implement both modes of operation.

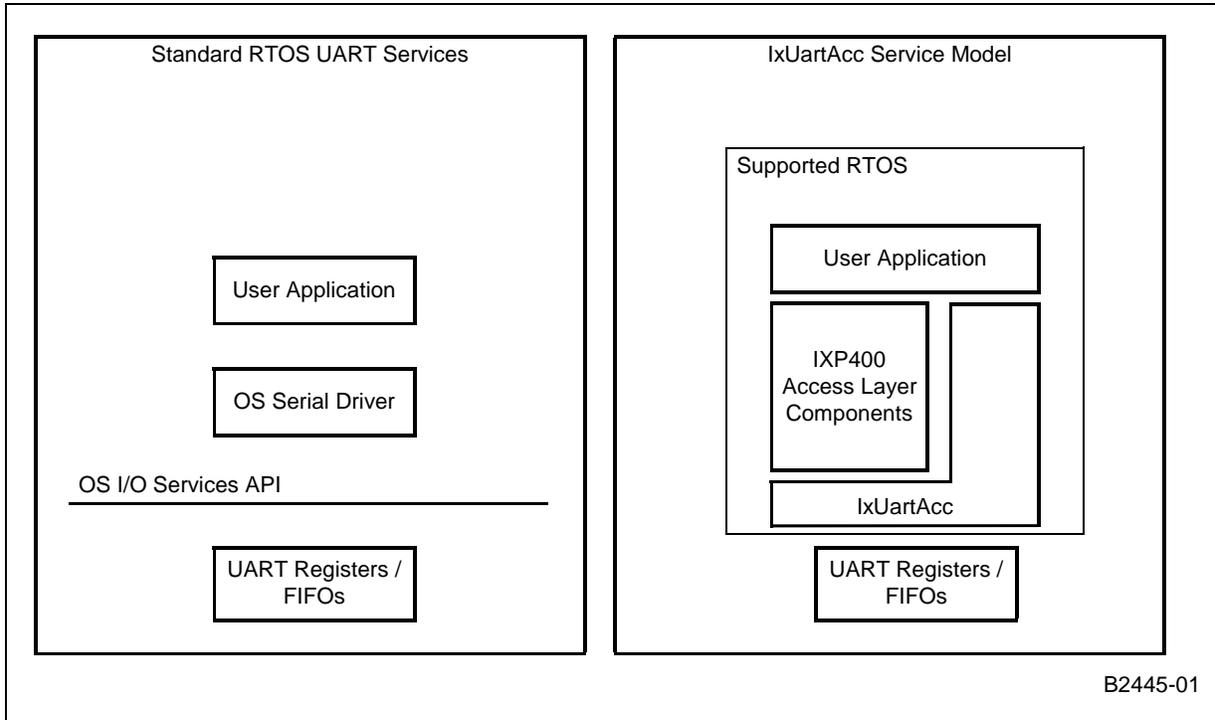
The default setup for the UART is:

- 9,600 bps baud rate
- 8-bit data word
- One stop bit
- No parity
- No flow control
- Interrupt mode (Polled for generic interface)



## 21.5 Dependencies

Figure 105. UART Services Models



§ §





## 22.0 Access-Layer Components: USB Access (ixUSB) API

---

This chapter describes the Intel® IXP400 Software v2.3's "USB Access API" access-layer component.

### 22.1 What's New

There are no changes or enhancements to this component in software release 2.3.

### 22.2 Overview

The Intel® IXP4XX Product Line of Network Processors' USB hardware components comply with the 1.1 version of the Universal Serial Bus (USB) standard.

### 22.3 USB Controller Background

The IXP4XX product line processors Universal Serial Bus Device Controller (UDC) supports 16 endpoints and can operate half-duplex at a baud rate of 12 Mbps (slave only, not a host or hub controller).

The serial information transmitted by the UDC contains layers of communication protocols, the most basic of which are fields. UDC fields include:

- Sync
- Packet identifier
- Address
- Endpoint
- Frame number
- Data
- CRC

Fields are used to produce packets. Depending on the function of a packet, a different combination and number of fields are used. Packet types include:

- Token
- Data
- Handshake
- Special

Packets are then assembled into groups to produce frames. These frames or transactions fall into four groups:

- Bulk
- Control
- Interrupt
- Isochronous

Endpoint 0, by default, is used only to communicate control transactions to configure the UDC after it is reset or hooked up (physically connected to an active USB host or hub). Endpoint 0's responsibilities include:



- Connection
- Address assignment
- Endpoint configuration
- Bus enumeration
- Disconnect

The USB protocol uses differential signaling between the two pins for half-duplex data transmission. A 1.5-KΩ pull-up resistor is required to be connected to the USB cable's D+ signal to pull the UDC+ pin high when polarity for data transmission is needed.

Using differential signaling allows multiple states to be transmitted on the serial bus. These states are combined to transmit data, as well as various bus conditions, including: idle, resume, start of packet, end of packet, disconnect, connect, and reset.

USB transmissions are scheduled in 1-ms frames. A frame starts with a SOF (Start-Of-Frame) packet and contains USB packets. All USB transmissions are regarded from the host's point of view: IN means towards the host and OUT means towards the device.

### 22.3.1 Packet Formats

USB supports four packet types:

- Token
- Data
- Handshake
- Special

A token packet is placed at the beginning of a frame, and is used to identify OUT, IN, SOF, and SETUP transactions. OUT and IN frames are used to transfer data., SOF packets are used to time isochronous transactions, and SETUP packets are used for control transfers to configure endpoints. An IN, OUT and SETUP token packet consists of a sync, a PID, an address, an endpoint, and a CRC5 field.

For OUT and SETUP transactions, the address and endpoint fields are used to select which UDC endpoint is to receive the data, and for an IN transaction, which endpoint must transmit data. A PRE (Preamble) PID precedes a low-speed (1.5 Mbps) USB transmission. The UDC supports full-speed (12 Mbps) USB transfers only. PRE packets signifying low-speed devices are ignored as well as the low-speed data transfer that follows.

Table 69. IN, OUT, and SETUP Token Packet Format

8 Bits	8 Bits	7 Bits	4 Bits	5 Bits
Sync	PID	Address	Endpoint	CRC5

A Start Of Frame (SOF) is a special type of token packet that is issued by the host at a nominal interval of once every 1 ms +/- 0.0005 ms. SOF packets consist of a sync, a PID, a frame number (which is incremented after each frame is transmitted), and a CRC5 field, as shown in Table 70. The presence of SOF packets every 1ms prevents the UDC from going into suspend mode.

Table 70. SOF Token Packet Format

8 Bits	8 Bits	11 Bits	5 Bits
Sync	PID	Frame Number	CRC5

Data packets follow token packets, and are used to transmit data between the host and UDC. There are two types of data packets as specified by the PID: DATA0 and DATA1. These two types are used to provide a mechanism to guarantee data sequence



synchronization between the transmitter and receiver across multiple transactions.

During the handshake phase, both communicate and agree which data token type to transmit first. For each subsequent packet transmitted, the data packet type is toggled (DATA0, DATA1, DATA0, and so on). A data packet consists of a sync, a PID, from 0 to 1,023 bytes of data, and a CRC16 field, as shown in Table 71. Note that the UDC supports a maximum of 8 bytes of data for an Interrupt IN data payload, a maximum of 64 bytes of data for a Bulk data payload, and a maximum of 256 bytes of data for an Isochronous data payload.

**Table 71. Data Packet Format**

8 Bits	8 Bits	0–1,023 Bytes	16 Bits
Sync	PID	Data	CRC16

Handshake packets consist of only a sync and a PID. Handshake packets do not contain a CRC because the PID contains its own check field. They are used to report data transaction status, including whether data was successfully received, flow control, and stall conditions. Only transactions that support flow control can return handshakes.

The three types of handshake packets are: ACK, NAK, and STALL.

- ACK — Indicates that a data packet was received without bit stuffing, CRC, or PID check errors.
- NAK — Indicates that the UDC was unable to accept data from the host, or it has no data to transmit.
- STALL — Indicates that the UDC is unable to transmit or receive data, and requires host intervention to clear the stall condition.

Bit stuffing, CRC, and PID errors are signaled by the receiving unit by omitting a handshake packet. Table 72 shows the format of a handshake packet.

**Table 72. Handshake Packet Format**

8 Bits	8 Bits
Sync	PID

### 22.3.2 Transaction Formats

Packets are assembled into groups to form transactions. Four different transaction formats are used in the USB protocol. Each is specific to a particular endpoint type: bulk, control, interrupt, and isochronous. Endpoint 0, by default, is a control endpoint and receives only control transactions.

The host controller initiates all USB transactions, and transmission takes place between the host and UDC one direction at a time (half-duplex).

Bulk transactions guarantee error-free transmission of data between the host and UDC by using packet-error detection and retry. The host schedules bulk packets when there is available time on the bus. The three packet types used to construct bulk transactions are: token, data, and handshake.

The eight possible types of bulk transactions based on data direction, error, and stall conditions are shown in Table 73. (Packets sent by the UDC to the host are highlighted in boldface type. Packets sent by the host to the UDC are not boldfaced.)



**Table 73. Bulk Transaction Formats**

Action	Token Packet	Data Packet	Handshake Packet
Host successfully received data from UDC	In	<b>DATA0/DATA1</b>	ACK
UDC temporarily unable to transmit data	In	None	<b>NAK</b>
UDC endpoint needs host intervention	In	None	<b>STALL</b>
Host detected PID, CRC, or bit-stuff error	In	<b>DATA0/DATA1</b>	None
UDC successfully received data from host	Out	DATA0/DATA1	<b>ACK</b>
UDC temporarily unable to receive data	Out	DATA0/DATA1	<b>NAK</b>
UDC endpoint needs host intervention	Out	DATA0/DATA1	<b>STALL</b>
UDC detected PID, CRC, or bit stuff error	Out	DATA0/DATA1	None

**Note:** Packets from UDC to host are **boldface**.

Isochronous transactions guarantee constant rate, error-tolerant transmission of data between the host and UDC. The host schedules isochronous packets during every frame on the USB, typically 1 ms, 2 ms, or 4 ms.

USB protocol allows for isochronous transfers to take up to 90% of the USB bandwidth. Unlike bulk transactions, if corrupted data is received, the UDC will continue to process the corrupted data that corresponds to the current start of frame indicator.

Isochronous transactions do not support a handshake phase or retry capability. The two packet types used to construct isochronous transactions are token and data. The two possible types of isochronous transactions, based on data direction, are shown in [Table 74](#).

**Table 74. Isochronous Transaction Formats**

Action	Token Packet	Data Packet
Host successfully received data from UDC	In	<b>DATA0/DATA1</b>
UDC successfully received data from host	Out	DATA0/DATA1

**Note:** Packets from UDC to host are **boldface**.

Control transactions are used by the host to configure endpoints and query their status. Like bulk transactions, control transactions begin with a setup packet, followed by an optional data packet, then a handshake packet. Note that control transactions, by default, use DATA0 type transfers. [Table 75](#) shows the four possible types of control transactions.



**Table 75. Control Transaction Formats, Set-Up Stage**

Action	Token Packet	Data Packet	Handshake Packet
UDC successfully received control from host	Setup	DATA0	<b>ACK</b>
UDC temporarily unable to receive data	Setup	DATA0	<b>NAK</b>
UDC endpoint needs host intervention	Setup	DATA0	<b>STALL</b>
UDC detected PID, CRC, or bit stuff error	Setup	DATA0	None

**Note:** Packets from UDC to host are **boldface**.

Control transfers are assembled by the host by sending a control transaction to tell the UDC what type of control transfer is taking place (control read or control write), followed by two or more bulk data transactions. The first stage of the control transfer is the setup. The device must either respond with an ACK; or if the data is corrupted, it sends no handshake.

The control transaction, by default, uses a DATA0 transfer, and each subsequent bulk data transaction toggles between DATA1 and DATA0 transfers. For a control write to an endpoint, OUT transactions are used. For control reads, IN transactions are used.

The transfer direction of the last bulk data transaction is reversed. It is used to report status and functions as a handshake. The last bulk data transaction always uses a DATA1 transfer by default (even if the previous bulk transaction used DATA1). For a control write, the last transaction is an IN from the UDC to the host, and for a control read, the last transaction is an OUT from the host to the UDC.

**Table 76. Control Transaction Formats**

Control Write	Setup	DATA (BULK OUT)	STATUS (BULK IN)
Control read	Setup	<b>DATA (BULK IN)*</b>	STATUS (BULK OUT)

**Note:** Packets from UDC to host are **boldface**.

Interrupt transactions are used by the host to query the status of the device. Like bulk transactions, interrupt transactions begin with a setup packet, followed by an optional data packet, then a handshake packet. [Table 77](#) shows the eight possible types of interrupt transactions.

**Table 77. Interrupt Transaction Formats**

Action	Token Packet	Data Packet	Handshake Packet
Host successfully received data from UDC	In	<b>DATA0/DATA1</b>	<b>ACK</b>
UDC temporarily unable to transmit data	In	None	<b>NAK</b>
UDC endpoint needs host intervention	In	None	<b>STALL</b>
Host detected PID, CRC, or bit stuff error	In	<b>DATA0/DATA1</b>	None
UDC successfully received data from host	Out	DATA0/DATA1	<b>ACK</b>
UDC temporarily unable to receive data	Out	DATA0/DATA1	<b>NAK</b>
UDC endpoint needs host intervention	Out	DATA0/DATA1	<b>STALL</b>
UDC detected PID, CRC, or bit stuff error	Out	DATA0/DATA1	None

**Note:** Packets from UDC to host are **boldface**.



## 22.4 ixUSB API Interfaces

Table 78. API interfaces Available for Access Layer

API	Description
ixUSBDriverInit	Initialize driver and USB Device Controller.
ixUSBDeviceEnable	Enable or disable the device.
ixUSBEndpointStall	Enable or disable endpoint stall.
ixUSBEndpointClear	Free all Rx/Tx buffers associated with an endpoint.
ixUSBSignalResume	Trigger signal resuming on the bus.
ixUSBFrameCounterGet	Retrieve the 11-bit frame counter.
ixUSBReceiveCallbackRegister	Register a data-receive callback.
ixUSBSetupCallbackRegister	Register a setup-receive callback.
ixUSBBufferSubmit	Submit a buffer for transmit.
ixUSBBufferCancel	Cancel a buffer previously submitted for transmitting.
ixUSBEventCallbackRegister	Register an event callback.
ixUSBIsEndpointStalled	Retrieve an endpoint's stall status.
ixUSBStatisticsShow	Display device state and statistics.
ixUSBErrorStringGet	Convert an error code into a human-readable string error message.
ixUSBEndpointInfoShow	Display endpoint information table.

The ixUSB API components operate within a callback architecture. Initial device setup and configuration is controlled through the callback registered during the ixUSBSetupCallbackRegister function. Data reception occurs through the callback registered during the ixUSBReceiveCallbackRegister function. Special events are signalled to the callback registered during the ixUSBEventCallbackRegister function.

Prior to using any other ixUSB API, the ixUSB client must initialize the controller with the ixUSBDriverInit API call. After this call the driver is in a disabled state. The call to ixUSBDeviceEnable allows data, setup, and configuration transmissions to flow.

### 22.4.1 ixUSB Setup Requests

The UDC's control, status, and data registers are used only to control and monitor the transmit and receive FIFOs for endpoints 1 - 15. All other UDC configuration and status reporting are controlled by the host, via the USB, using device requests that are sent as control transactions to endpoint 0. Each data packet of a setup stage to endpoint 0 is 8 bytes long and specifies:

- Data transfer direction
  - Host to device
  - Device to host
- Data transfer type
  - Standard
  - Class
  - Vendor
- Data recipient
  - Device
  - Interface



- Endpoint
- Other
- Number of bytes to transfer
- Index or offset
- Value: Used to pass a variable-sized data parameter
- Device request

The UDC decodes most commands with no intervention required by the ixUSB client. Other setup requests occur through the setup callback. The following data structure in [Figure 106](#) is passed to the setup-callback function so the software can be configured properly.

**Figure 106. USBSetupPacket**

```
typedef struct /* USBSetupPacket */
{
    UCHAR bmRequestType;
    UCHAR bRequest;
    UINT16 wValue;
    UINT16 wIndex;
    UINT16 wLength;
} USBSetupPacket;
```

[Table 79](#) shows a summary of the setup device requests.



Table 79. Host-Device Request Summary

Request	Name
SET_FEATURE	Enables a specific feature, such as device remote wake-up and endpoint stalls.
CLEAR_FEATURE	Clears or disables a specific feature.
SET_CONFIGURATION	Configures the UDC for operation. Used following a reset of the controller or after a reset has been signalled via the USB.
GET_CONFIGURATION	Returns the current UDC configuration to the host.
SET_DESCRIPTOR	Sets existing descriptors or adds new descriptors. Existing descriptors include: † <ul style="list-style-type: none"> <li>• Device Configuration</li> <li>• String</li> <li>• Interface</li> <li>• Endpoint</li> </ul>
GET_DESCRIPTOR	Returns the specified descriptor, if it exists.
SET_INTERFACE	Selects an alternate setting for the UDC's interface.
GET_INTERFACE	Returns the selected alternate setting for the specified interface.
GET_STATUS	Returns the UDC's status including: <ul style="list-style-type: none"> <li>• Remote wake-up</li> <li>• Self-powered</li> <li>• Data direction</li> <li>• Endpoint number</li> <li>• Stall status</li> </ul>
SET_ADDRESS	Sets the UDC's 7-bit address value for all future device accesses.
SYNCH_FRAME	Sets an endpoint's synchronization frame.

† Interface and endpoint descriptors cannot be retrieved or set individually. They exist only embedded within configuration descriptors.

Via control endpoint 0, the user must decode and respond to the GET\_DESCRIPTOR command.

Refer to the *Universal Serial Bus Specification Revision 1.1* for a full description of host-device requests.

### 22.4.1.1 Configuration

In response to the GET\_DESCRIPTOR command, the user sends back a description of the UDC configuration. *The UDC can physically support more data-channel bandwidth than the USB will allow.* When responding to the host, the user must be careful to specify a legal USB configuration.

For example, if the user specifies a configuration of six isochronous endpoints of 256 bytes each, the host will not be able to schedule the proper bandwidth and will not take the UDC out of Configuration 0. The user must determine which endpoints to not tell the host about, so that they will not get used.

Another option, especially attractive for isochronous endpoints, is to describe a configuration of less than 256 bytes maximum packet to the host. The direction of the endpoints is fixed and the UDC will physically support only the following maximum packet sizes:

- Interrupt endpoints — 8 bytes
- Bulk endpoints — 64 bytes
- Isochronous endpoints — 256 bytes

In order to increase flexibility, the UDC supports a total of four configurations. While each of these configurations is identical within the UDC, the software can be used to make three distinct configurations. Configuration 0 is a default configuration of



endpoint 0 only.

For a detailed description of the configuration descriptor, see the USB 1.1 specification.

### 22.4.1.2 Frame Synchronization

The SYNCH\_FRAME request is used by isochronous endpoints that use implicit-pattern synchronization. The isochronous endpoints may need to track frame numbers in order to maintain synchronization.

Isochronous-endpoint transactions may vary in size, according to a specific repeating pattern. The host and endpoint must agree on which frame begins the repeating pattern. The host uses this request to specify the exact frame on which the repeating pattern begins.

The data stage of the SYNCH\_FRAME request contains the frame number in which the pattern begins. Having received the frame number, the device can start monitoring each frame number sent during the SOF. This is recorded in the frame counter and made available through specific driver functions (see ixUSBFrameCounterGet).

### 22.4.2 ixUSB Send and Receive Requests

The USB access layer encodes and decodes data frames sending and receiving buffers to and from the client in the same format as IX\_MBUF.

Buffers are sent from the UDC to the host with the ixUSBBufferSubmit API.

Data buffers are received from the host through the callback function registered with the access layer during the ixUSBReceiveCallbackRegister API call.

### 22.4.3 ixUSB Endpoint Stall Feature

A device uses the STALL handshake in one of two distinct occasions.

The first case — known as “functional stall” — is when the *Halt* feature, associated the endpoint, is set. A special case of the functional stall is the “commanded stall.” Commanded stall occurs when the host explicitly sets the endpoint’s *Halt* feature using the SET\_FEATURE command.

Once a function’s endpoint is halted, the function must continue returning STALL packets until the condition causing the halt has been cleared through host intervention (using SET\_FEATURE). This can happen both for IN and OUT endpoints. In the case of IN endpoints, the endpoint sends a STALL handshake immediately after receiving an IN token. For OUT endpoints the STALL handshake is sent as soon as the data packet after the OUT token is received.

Figure 107. STALL on IN Transactions

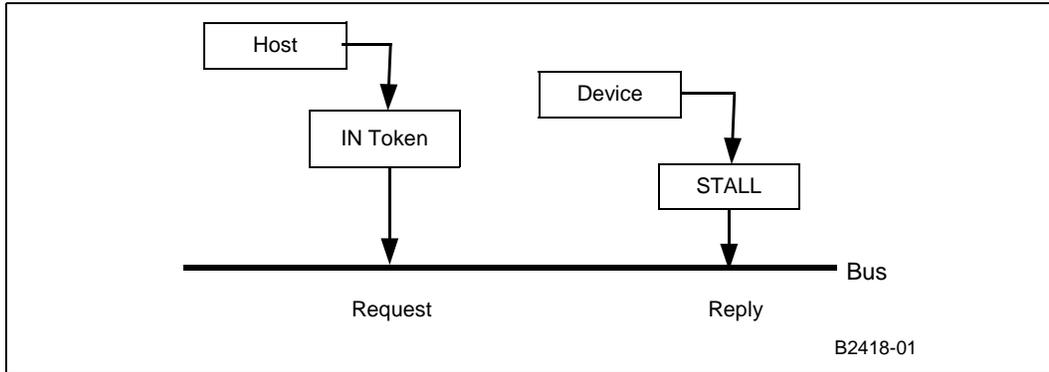
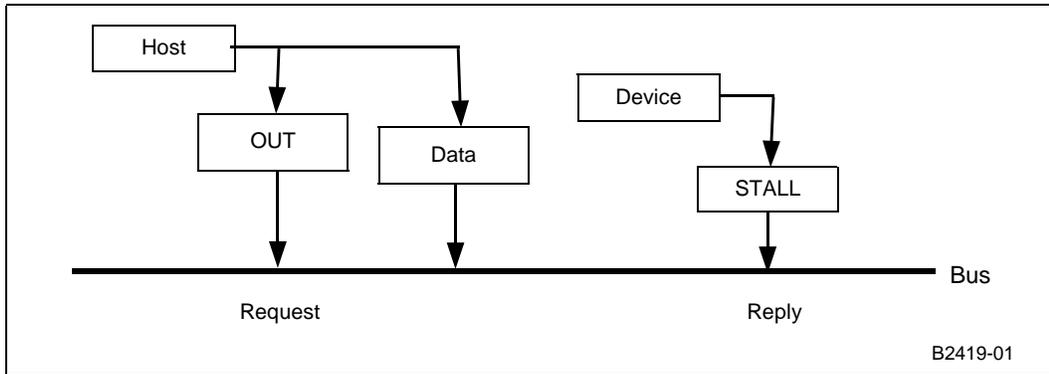


Figure 108. STALL on OUT Transactions



The second case of a STALL handshake is known as a “protocol stall” and is unique to control pipes. Protocol stall differs from functional stall in meaning and duration.

A protocol STALL is returned during the Data or Status stage of a control transfer, and the STALL condition terminates at the beginning of the next control transfer (Setup). Protocol stalls are usually sent to notify the host that a particular USB command is malformed or not implemented.



## 22.4.4 ixUSB Error Handling

The USB API calls return the IX\_FAIL error code after detecting errors. It is the responsibility of the user to implement appropriate error handling. Detailed error codes are used to report USB Driver errors. They are provided in the *lastError* field of the *USBDevice* structure that must be passed by the user in every API call. When the API calls are successful the *lastError* field is assigned the IX\_SUCCESS value.

**Table 80.** Detailed Error Codes

<pre> #ifndef IX_USB_ERROR_BASE #define IX_USB_ERROR_BASE 4096 #endif /* IX_USB_ERROR_BASE */  /* error due to unknown reasons */  #define IX_USB_ERROR(IX_USB_ERROR_BASE + 0)  /* invalid USBDevice structure passed as parameter or no device present */ #define IX_USB_INVALID_DEVICE (IX_USB_ERROR_BASE + 1)  /* no permission for attempted operation */ #define IX_USB_NO_PERMISSION(IX_USB_ERROR_BASE + 2)  /* redundant operation */ #define IX_USB_REDUNDANT(IX_USB_ERROR_BASE + 3)  /* send queue full */ #define IX_USB_SEND_QUEUE_FULL(IX_USB_ERROR_BASE + 4)  /* invalid endpoint */ #define IX_USB_NO_ENDPOINT(IX_USB_ERROR_BASE + 5)  /* no IN capability on endpoint */ #define IX_USB_NO_IN_CAPABILITY(IX_USB_ERROR_BASE + 6)  /* no OUT capability on endpoint */ #define IX_USB_NO_OUT_CAPABILITY(IX_USB_ERROR_BASE + 7)  /* transfer type incompatible with endpoint */ #define IX_USB_NO_TRANSFER_CAPABILITY(IX_USB_ERROR_BASE + 8)  /* endpoint stalled */ #define IX_USB_ENDPOINT_STALLED(IX_USB_ERROR_BASE + 9)  /* invalid parameter(s) */ #define IX_USB_INVALID_PARAMS(IX_USB_ERROR_BASE + 10) </pre>
<p><b>Note:</b> “Error due to unknown reasons” — This code is also used when there is only one possible error reason and the error was already signaled by the IX_FAIL return code.</p>

## 22.5 USB Data Flow

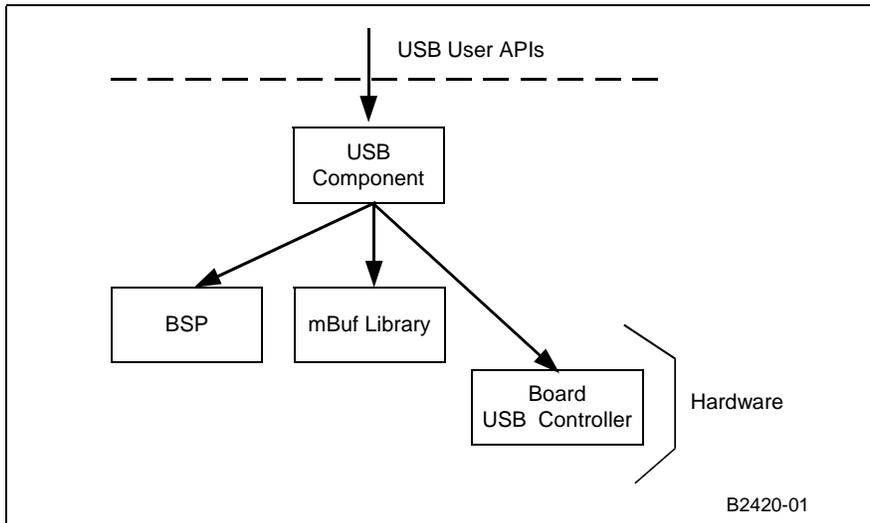
The USB device is a memory mapped device on the processor’s peripheral bus. It will not interact directly with the NPEs. Any data path between USB and other components must be performed via the Intel XScale® Processor.

## 22.6 USB Dependencies

The USB device driver is a self-contained component with no interactions with other data components. [Figure 109](#) shows the dependencies for this USD component.



Figure 109. USB Dependencies



§ §



## 23.0 Codelets

---

This chapter describes the Intel® IXP400 Software v2.3 codelets.

### 23.1 What's New

The Crypto codelet has been enhanced in software release 2.3 to use the RNG, SHA1, and EAU coprocessors in the PKE crypto engine.

### 23.2 Overview

The codelets are example code that utilize the access-layer components and operating system abstraction layers discussed in the preceding chapters. Codelets, while not exhaustive examples of the functionality available to the developer, provide a good basis from which to begin their own code development for test harnesses, performance analysis code, or even functional applications to take to market.

This chapter describes the major features of the available in each codelet. For detailed information, see the header and source files provided with software release 2.3 in the `xscale_sw/src/codelets` directory.

### 23.3 ATM Codelet (IxAtmCodelet)

This codelet demonstrates an example implementation of a working ATM driver that makes use of the `AtmdAcc` component, as well as demonstrating how the lower layer `IxAtmdAcc` component can be used for configuration and control.

This codelet also demonstrates an example implementation of OAM F4 Segment, F4 End-To-End (ETE), F5 Segment and F5 ETE loopback. `Aal5` or `Aal0` (48 or 52 bytes) traffic types are available in this codelet, as well as the display of transmit and receive statistics.

`IxAtmCodelet` makes use of the following access-layer components:

- `IxAtmdAcc`
- `IxAtmm`
- `IxAtmSch`

### 23.4 Crypto Access Codelet (IxCryptoAccCodelet)

This codelet demonstrates how to use the `IxCrypto` access-layer component and the underlying security features in the Intel® IXP4XX Product Line of Network Processors. `IxCryptoAccCodelet` runs through the scenarios of initializing the NPEs and Queue Manager, context registration, and performing a variety of encryption (3DES, AES, ARC4), decryption, and authentication (SHA1, MD5) operations. This codelet demonstrates both IPsec and WEP service types.



The codelets demonstrate how to use the PKE APIs. The codelet also performs performance measurements of cryptographic operations and PKE Crypto functions.

### 23.5 DMA Access Codelet (IxDmaAccCodelet)

The DMA Access Codelet executes DMA transfer for various DMA transfer modes, addressing modes and transfer widths. The block sizes used in this codelet are 8; 1,024; 16,384; 32,768; and 65,528 bytes. For each DMA configuration, the performance is measured and the average rate (in Mbps) is displayed.

This codelet is not supported in little-endian mode.

### 23.6 Ethernet AAL-5 Codelet (IxEthAal5App)

IxEthAal5App codelet is a mini-bridge application which bridges traffic between Ethernet and UTOPIA ports or Ethernet and an ADSL port. Two Ethernet ports and up to eight UTOPIA ports are supported, which are initialized by default at the start of application.

Ethernet frames are transferred across ATM link (through Utopia interface) using AAL-5 protocol and Ethernet frame encapsulation described by RFC 1483. MAC address learning is performed on Ethernet frames, received by Ethernet ports and ATM interface (encapsulated). IxEthAal5App filters packets based on destination MAC addresses.

IxEthAal5App makes use of the following access-layer components:

- IxEthAcc
- IxAtmdAcc
- IxAtmm
- IxAtmSch
- IxQMgr

### 23.7 Ethernet Access Codelet (IxEthAccCodelet)

This codelet demonstrates both Ethernet data and control plane services and Ethernet management services. The features can be selectively executed at run-time via the menu interface of the codelet.

- Ethernet data and control plane services:
  - Configuring both ports as a receiver sink from an external source (such as Smartbits)
  - Configuring Port-1 to automatically transmit frames and Port-2 to receive frames. Frames generated and transmitted in Port-1 are looped back into Port-2 by using cross-over cable.
  - Configuring and performing a software loopback on each of the two Ethernet ports.
  - Configuring both ports to act as a bridge so that frames received on one port are retransmitted on the other.
- Ethernet management services:
  - Adding and removing static/dynamic entries.
  - Calling the maintenance interface (run as a separate background task)
  - Calling the show routine to display the MAC address filtering tables.



IxEthAccCodelet demonstrates the use of many of the access-layer components.

## 23.8 HSS Access Codelet (IxHssAccCodelet)

IxHssAccCodelet tests packetized and channelized services, with the codelet acting as data source/sink and HSS as loopback. The codelet will transmit data and will optionally verify that data received is the same as that transmitted.

Codelet runs for a user selectable amount of time. This codelet provides a good example of different Intel XScale® Processor to NPE data transfer techniques, by using mbuf pools for packetized services and circular buffers for channelized services.

## 23.9 Parity Error Notifier Codelet (IxParityENAccCodelet)

The IxParityENAccCodelet shows how to integrate parity error detection and error handling routines into a client application, using IxParityENAcc. The API is based upon capabilities available on the IXP46X network processors. This codelet demonstrates the following:

- How to initialize IxParityENAcc.
- How to configure IxParityENAcc or modify IxParityENAcc configuration.
- How to register callback with IxParityENAcc.
- How to register data abort handler with kernel (only for VxWorks\*).
- How to inject ECC error.
- How to spawn a task to initiate SDRAM memory scan.
- How to scrub memory to correct single bit ECC error.
- How to handle various parity errors reported by IxParityENAcc.
- How to determine whether the data abort is due to multi bit ECC.
- Error initiated when Intel XScale® Processor accesses SDRAM.

## 23.10 Performance Profiling Codelet (IxPerfProfAccCodelet)

IxPerfProfAccCodelet is a useful utility that demonstrates how to access performance related data provided by IxPerfProfAcc. The codelet provides an interface to view north, south, and SDRAM bus activity, event counting and idle cycles from the Intel XScale® Processor PMU and other performance attributes of the processor.

*Note:* IxPerfProfAccCodelet has not been modified to support the IXP46X network processors at this time.

## 23.11 Time Sync Codelet (IxTimeSyncAccCodelet)

This codelet shows how to use some of the IxTimeSyncAcc API functions to utilize the following features of the IEEE 1588 unit available on the IXP46X network processors:

- How to configure a channel to operate in master or slave mode.
- How to set the frequency scaling value.
- How to set and get system time.
- How to setup target time in interrupt mode.
- How to enable and disable the target time interrupt.



- How to make use of polled mode Rx and Tx PTP message timestamps for several NPE configurations.

An external device, such as a SmartBits\*, may be used to generate PTP messages and transmit to the NPE channels.

## 23.12 USB RNDIS Codelet (IxUSBRNDIS)

The IxUSBRNDIS codelet is a sample driver implementation of an RNDIS client.

RNDIS (Remote Network Driver Interface Specification) is a specification for Ethernet-like interface compatible with Microsoft\* operating systems. This codelet allows a properly configured platform based upon Intel® IXP4XX Product Line of Network Processors, running VxWorks\* or Linux\* to communicate IP traffic over USB to a Microsoft\* Windows\* system.

§ §



## 24.0 Operating System Abstraction Layer (OSAL)

---

### 24.1 What's New

The following changes and enhancements were made to this component in software release 2.3:

- New API being added, refer to “New APIs”

### 24.2 New APIs

As mentioned above, the following new APIs have been added. More details regarding the input parameters, description, and return parameters can be found in the API reference document file, *APIReference.pdf*. This document is found in the doc directory of the software release.

- *BOOL ixOsalThreadStopCheck()*
  - This function is used within the thread to check if someone is trying to kill it. When this API returns TRUE, the thread should perform cleanup and exit gracefully. As a safe programming practice, internal thread implementation should call this API to ensure no one is killing it before it calls any APIs that might cause sleep.

### 24.3 Overview

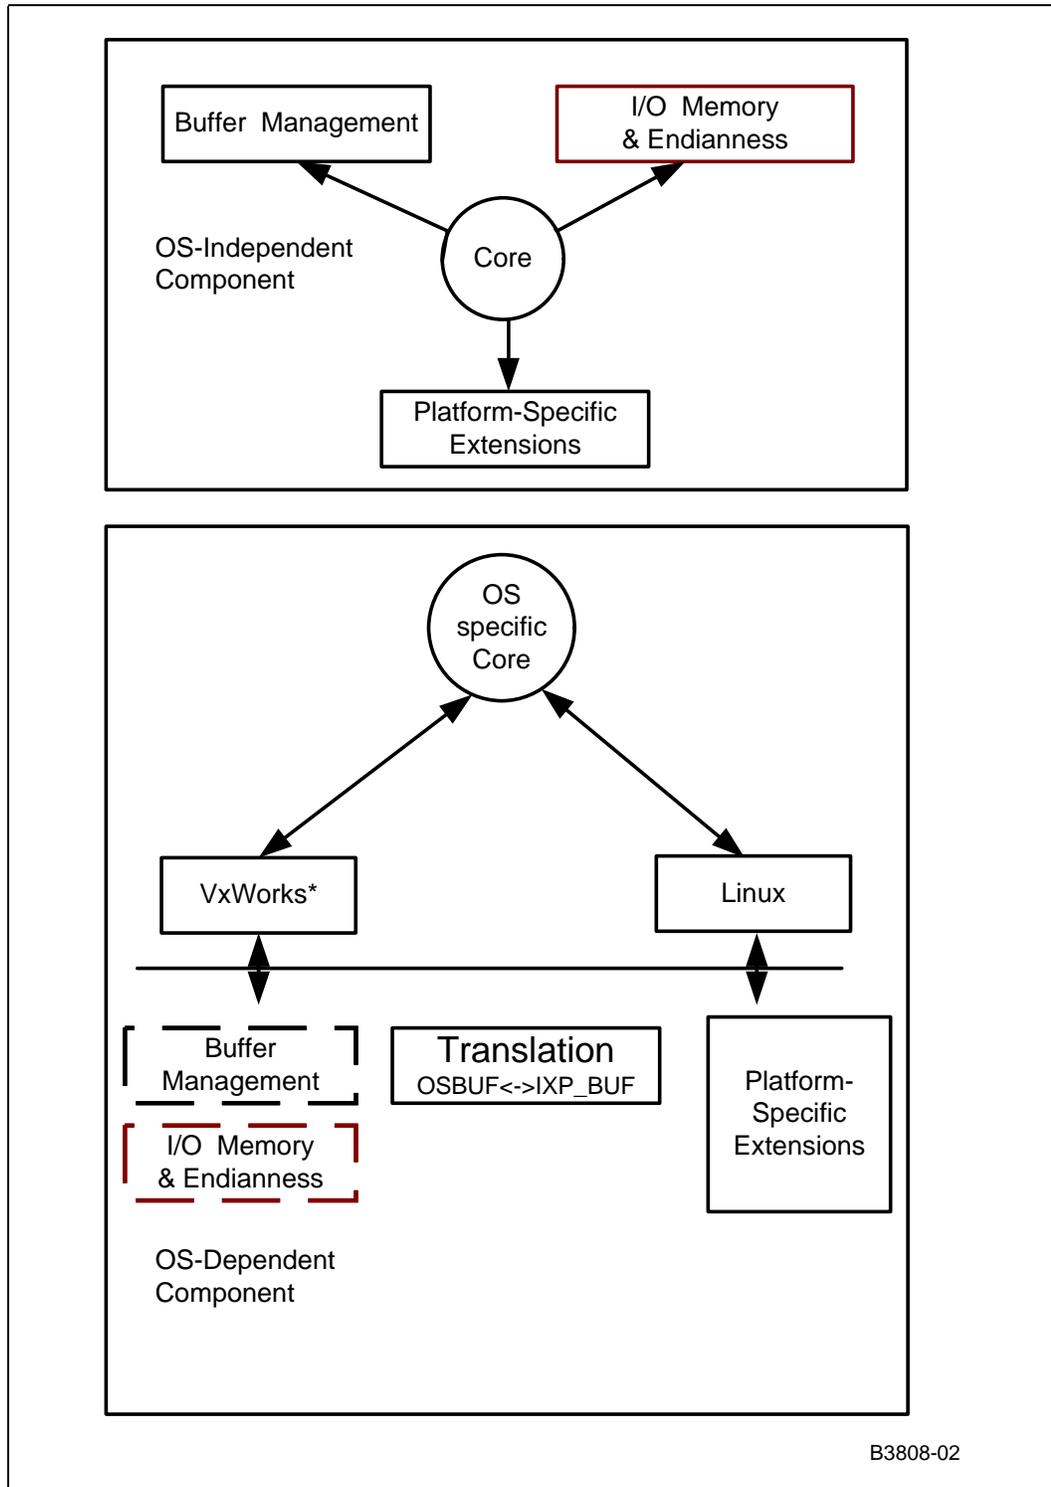
An Operating System Services Abstraction Layer (OSAL) is provided as part of the software release 2.3 architecture. [Figure 110](#) shows the OSAL architecture.

The OSAL provides a very thin set of abstracted operating-system services. All other access-layer components abstract their OS dependencies to this layer. Though primarily intended for use by the software release 2.3 access-layer component, these services are also available to the codelets and to application-layer software. The OSAL also defines an extended, more fully featured interface for different operating system services, and for different target platforms.

The OSAL layer can be categorized into two modules:

- The OS-independent core module
- The OS-dependent module

Figure 110. OSAL Architecture





## 24.4 OS-Independent Core Module

As shown in [Figure 110](#), the OS-independent component includes all the core functionality such as buffer management, platform- and module-specific OS-independent implementations, I/O memory map function implementations, and OSAL core services implementations. The Buffer Management module defines a memory buffer structure and functions for creating and managing buffer pools. The I/O Memory and Endianness module includes support for I/O memory mapping under different operating systems as well as big and little endian support.

### Core Module

The OSAL core module defines the following functionality:

- Memory allocation
- Threading
- Interrupt handling
- Thread synchronization
- Delay functions
- Time-related functions and macros
- Inter-thread communication
- Logging

The core module is a non-optional module containing fundamental types, constants, functions and macros provided by the OSAL. Many of these items are used in the other modules as well.

## 24.5 OS-Dependent Module

The OS-dependent component for a respective OS gets selected by the build system at build time. This component provides operating system services like timers, mutex, semaphores, and thread management. The OS translation functions are implemented for respective operating systems to translate the header fields of the OS buffers to IXP buffer format and vice versa. The core module is a non-optional module containing fundamental types, constants, functions and macros provided by the OSAL, and many of these items are used in the other modules as well.

## 24.6 Optional Modules

### Buffer Management Module

The OSAL Buffer Management Module implements the following functionality:

- Buffer pool management (pool initialization and allocation)
- Buffer management (buffer allocation and freeing)

### I/O Memory and Endianness Module

The I/O memory management defines a set of macros allowing the user to gain and release access to memory-mapped I/O in an operating-system-independent fashion. Depending on the target platform and OS, gaining access can vary between no special behavior (statically mapped I/O), to dynamically mapped I/O through OS-specific functions (for example, `ioremap()` in Linux\*). The Endianness module supports big and little endian.



I/O Memory and Endianness modules are OS-independent common modules. Buffer Management is common to both VxWorks\* and Linux\*, but is OS-specific for Windows CE.

### 24.6.1 Buffer Translation Module

OSAL provides buffer translation macros for users to translate OS-specific buffer formats to OSAL IXP buffer format and vice versa. The buffer translations is usually done in the driver component. However, for ease of use, the OSAL layer provides generic macros for the VxWorks\*, and Linux\* operating systems. Depending upon the build, the OSAL layer will translate the macros to its OS-specific implementation. The general syntax for using these macros is as follows:

- `IX_OSAL_CONVERT_OSBUF_TO_IXPBUF(osBufPtr,ixpBufPtr)`
- `IX_OSAL_CONVERT_IXPBUF_TO_OS_BUF(ixpBufPtr,osBufPtr)`

These macros are intended to replace Linux\* skbuf conversion, and VxWorks\* mbuf conversions. Users can also define their own conversion utilities in their package to translate their buffers to the OSAL IXP\_BUF (IX\_OSAL\_MBUF). As an option to using the translation functions, the user can choose to implement their own definitions for the ix\_mbuf structure field within the IXP\_BUF structure format.

## 24.7 OSAL Library Structure

As shown in [Figure 111](#), the OSAL library is contained in the following directories along with a “doc” folder that includes API references in HTML and PDF format.

- The “include” directory

The Include directory contains the main OSAL header files for core module and subdirectories for module-specific header files (for example, header files for the Buffer Management module grouped under the **include/modules/bufferMgt** subdirectory). It also contains subdirectories for platform-specific headers (for example, header for the IXP400 software platform grouped under “include/platforms/ixp400” subdirectory). The OSAL library is accessed via a single header file — IxOsal.h. The main header file will automatically include the core API and the OSAL configuration header file. The OSAL configuration header file (IxOsalConfig.h) contains user-editable fields for module inclusion, and it automatically includes the module-specific header files for optional modules, such as the buffer management (IxOsalBufferMgt.h), I/O memory mapping and Endianness support (IxOsalIoMem.h). Platform configuration is done in IxOsalConfig.h by including the main platform header file (IxOsalOem.h).

*Note:* Platform-specific refers to all the platforms that use the same network processor variants, that is, that use the same Intel® IXP4XX Product Line of Network Processors. A change in product line refers to using the OSAL layer for a new platform.

- The “src” directory

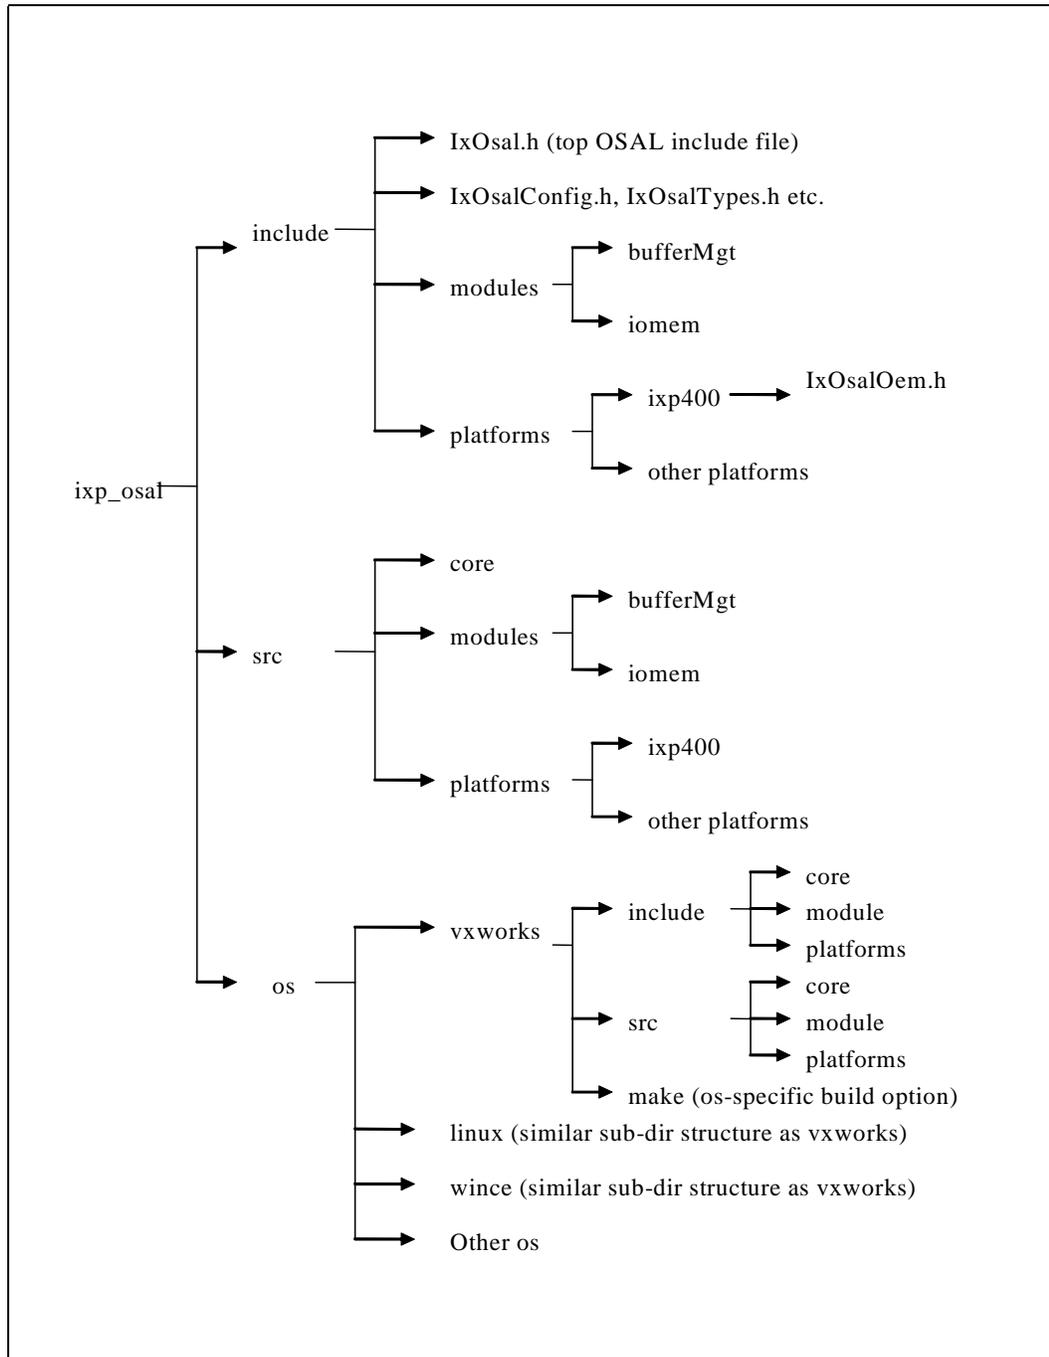
The source directory contains the actual implementation of OSAL OS-independent core module and subdirectories for OS-independent (and module-specific) implementation. Additionally, the source directory contains subdirectories for OS-independent (and platform-specific) implementation. The OSAL build system looks for all the implementations in the core module and the specified OS hierarchy. The source tree is organized in different directories, one for each target OS, and a shared directory for core implementations. Note that shared implementations do not must be common for all the possible operating systems. Instead, code that is deemed to be reusable is placed in the source directory. For each OS, the library is compiled for the OS-specific directory and the shared directory, hence it is required that each function implementation must be found either in the core directory or in the OS-specific directory.



- The “os” directory

The OS directory contains OS-dependent subdirectories with OS-specific implementations of the APIs; these directories are named after the OS they abstract (for example, “VxWorks”, “Linux”). Each “os” subdirectory has its own include directory, src directory hierarchy for OS-specific core, modules and platform implementations. The translational functions are implemented in the source subdirectory within each of the individual OS directories.

Figure 111. OSAL Directory Structure

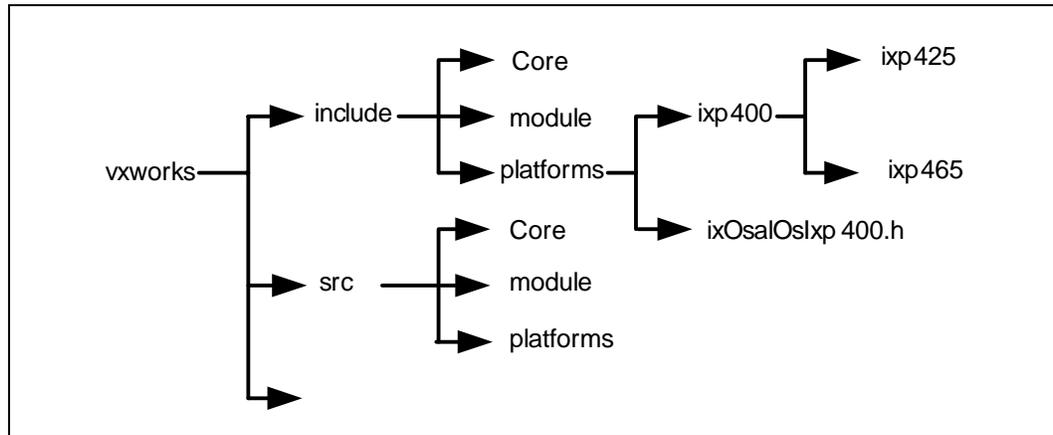


The source files for each module implementation shall reside in sub-directories in the **src** section, grouped accordingly per OS and module (**bufferMgt** and **ioMem**).

OSAL supports IXDP42X and IXDP46X development platforms. The processor variants are separated within the **ixp400** sub folders of the target OS. The directory structure of the target OS VxWorks\* is shown in [Figure 112](#).



Figure 112. OSAL Intel® IXP4XX Processor Variant Directory Structure



Processor variant configuration can be done in `ixOsAlOslxp400.h` by including either the `IxOsAlOslxp425Sys.h` or `IxOsAlOslxp465Sys.h`. The OSAL build systems builds the OSAL libraries into different release output with according to platform type (for example, IXP42X product line or IXP46X product line and target OS).

## 24.8 OSAL Modules and Related Interfaces

This section contains a summary of the types, symbols, and public functions declared by each OSAL module.

*Note:* The items shaded in light gray are subject to special platform package support, as described in the API notes for these items and the platform package requirements of each module.

### 24.8.1 Core Module

This non-optional module contains fundamental types, constants, functions and macros provided by the OSAL layer. Many of them are used in the other modules as well. Some of the common services provided by this module are:

- Thread handling
- Mutexes
- Semaphores
- Interrupt Services
- Memory allocation and translation services
- Timer services

The above-mentioned services would have its own implementation in their respective OS modules. For client purposes, the API calls will remain the same. The build system automatically switches the appropriate implementation.

Table 81 presents an overview of the OSAL core interface. Items marked in gray are specific to platform-implementation requirements.



Table 81. OSAL Core Interface (Sheet 1 of 2)

Types	IxOsalVoidFnPtr	alias for void (void) functions
	IxOsalVoidFnVoidPtr	alias for void (void *) functions
	IxOsalSemaphore	semaphore object
	IxOsalMutex	mutex object
	IxOsalFastMutex	test-and-set fast mutex object
	IxOsalThread	thread object
	IxOsalThreadAttr	thread attributes object
	IxOsalTimeval	time structure
	IxOsalTimer	timer handle
Symbols	PRIVATE	#defined as "static", except for debug builds
	PUBLIC	#defined as an empty labelling symbol
Interrupts	ixOsalIrqBind	binds interrupts to handlers
	ixOsalIrqUnbind	unbind interrupts from handlers
	ixOsalIrqLock	disables all interrupts
	ixOsalIrqUnlock	enables all interrupts
	ixOsalIrqLevelSet	selectively disables interrupts
	ixOsalIrqEnable	enables an interrupt level
	ixOsalIrqDisable	disables an interrupt level
Memory	ixOsalMemAlloc	allocates memory
	ixOsalMemFree	frees memory
	ixOsalMemCopy	copies memory zones
	ixOsalMemSet	fills a memory zone
	ixOsalCacheDmaMalloc	allocates cache-safe memory
	ixOsalCacheDmaFree	frees cache-safe memory
	IX_OSAL_MMU_PHYS_TO_VIRT	physical to virtual address translation
	IX_OSAL_MMU_VIRT_TO_PHYS	virtual to physical address translation
	IX_OSAL_CACHE_FLUSH	cache to memory flush
Threads	IX_OSAL_CACHE_INVALIDATE	cache line invalidate
	IX_OSAL_CACHE_PRELOAD	memory to preload to cache
	ixOsalThreadCreate	creates a new thread
	ixOsalThreadStart	starts a newly created thread
	ixOsalThreadStopCheck	check if the thread should stop executing
	ixOsalThreadKill	kills an existing thread
	ixOsalThreadExit	exits a running thread
	ixOsalThreadPrioritySet	sets the priority of an existing thread
	ixOsalThreadSuspend	suspends thread execution
ixOsalThreadResume	resumes thread execution	
IPC	ixOsalMessageQueueCreate	creates a message queue
	ixOsalMessageQueueDelete	deletes a message queue
	ixOsalMessageSend	sends a message to a message queue
	ixOsalMessageReceive	receives a message from a message queue



**Table 81. OSAL Core Interface (Sheet 2 of 2)**

Thread synchronization	ixOsalMutexInit	initializes a mutex
	ixOsalMutexLock	locks a mutex
	ixOsalMutexUnlock	unlocks a mutex
	ixOsalMutexTryLock	non-blocking attempt to lock a mutex
	ixOsalMutexDestroy	destroys a mutex object
	ixOsalFastMutexInit	initializes a fast mutex
	ixOsalFastMutexTryLock	non-blocking attempt to lock a fast mutex
	ixOsalFastMutexUnlock	unlocks a fast mutex
	ixOsalFastMutexDestroy	destroys a fast mutex object
	ixOsalSemaphoreInit	initializes a semaphore
	ixOsalSemaphorePost	posts to (increments) a semaphore
	ixOsalSemaphoreWait	waits on (decrements) a semaphore
	ixOsalSemaphoreTryWait	non-blocking wait on semaphore
	ixOsalSemaphoreGetValue	gets semaphore value
	ixOsalSemaphoreDestroy	destroys a semaphore object
Time	ixOsalYield	yields execution of current thread
	ixOsalSleep	yielding sleep for a number of milliseconds
	ixOsalBusySleep	busy sleep for a number of microseconds
	ixOsalTimestampGet	value of the timestamp counter
	ixOsalTimestampResolutionGet	resolution of the timestamp counter
	ixOsalSysClockRateGet	system clock rate, in ticks
	ixOsalTimeGet	current system time
	IX_OSAL_TIMEVAL_TO_TICKS	converts ixOsalTimeVal into ticks
	IX_OSAL_TICKS_TO_TIMEVAL	converts ticks into ixOsalTimeVal
	IX_OSAL_TIMEVAL_TO_MS	converts ixOsalTimeVal to milliseconds
	IX_OSAL_MS_TO_TIMEVAL	converts milliseconds to IxOsalTimeval
	IX_OSAL_TIME_EQ	"equal" comparison for IxOsalTimeval
	IX_OSAL_TIME_LT	"less than" comparison for IxOsalTimeval
	IX_OSAL_TIME_GT	"greater than" comparison for IxOsalTimeval
	IX_OSAL_TIME_ADD	"add" operator for IxOsalTimeval
IX_OSAL_TIME_SUB	"subtract" operator for IxOsalTimeval	
Logging	ixOsalLogLevelSet	sets the current logging verbosity level
	ixOsalLog	interrupt-safe logging function
Timers	ixOsalRepeatingTimerSchedule	schedules a repeating timer
	ixOsalSingleShotTimerShedule	schedules a single-shot timer
	ixOsalTimerCancel	Cancels a running timer
	ixOsalTimersShow	displays all the running timers
Misc.h	ixOsalOsNameGet	gets the running OS's name
	ixOsalOsVersionGet	gets the running OS's version



### 24.8.2 Buffer Management Module

This module defines a memory buffer structure and functions for creating and managing buffer pools.

Table 82 provides an overview of the buffer management module.

Table 82. OSAL Buffer Management Interface

Types	IX_OSAL_MBUF	memory buffer
	IX_OSAL_MBUF_POOL	memory buffer pool
Functions	ixOsalPoolInit	initializes pool with memory allocation
	ixOsalNoAllocPoolInit	initializes pool without memory allocation
	ixOsalMbufAlloc	allocates a buffer from a pool
	ixOsalMbufFree	frees a buffer into its pool
	ixOsalMbufChainFree	frees a buffer chain into its pool
	ixOsalMbufDataPtrReset	resets the buffer data pointer
	ixOsalBuffPoolFreeCountGet	gets the number of available free buffers in the pool
	IX_OSAL_MBUF_POOL_DATA_AREA_SIZE_ALIGNED	gets the pool data buffer area size required
	IX_OSAL_MBUF_POOL_MBUF_AREA_SIZE_ALIGNED	gets the buffer pool data area size required
	ixOsalPoolShow	displays pool statistics

### 24.8.3 I/O Memory and Endianness Support Module

The OSAL I/O Memory Management and Endianness Support Module implements:

- I/O memory management
- Big and little endian support

I/O memory management defines a set of macros allowing the user to gain and release access to memory-mapped I/O in an operating-system-independent fashion. Depending on the target platform and OS, gaining access can vary between statically mapped I/O to dynamically mapped I/O through OS-specific functions (for example, ioremap() in Linux\*).

Using a global memory map, which defines the specifics of each memory map cell (for example, UART registers), the access of I/O memory can be abstracted independent of operating systems, dynamic mapping, or endianness-dependent virtual memory locations. This functionality makes the code far more portable across different operating systems and platforms.

Wind River\* VxWorks\* OS maintains a 1:1 virtual to physical mapping. However, this is not the case in other OS such as Linux\*. The OSAL layer provides a portable approach that involves mapping the memory when the software is initialized to access the desired memory and unmapping the memory when the software unloads. Depending



upon the build for a particular OS (and if the memory is not statically mapped), the OSAL can create MMU entries to map the specified physical address in the usable memory range.

Additionally, the mapping automatically considers the endianness type in systems that can use mixed endian modes (such as the IXP4XX product line processors). This behavior is controlled by two defines which must be supplied by the software using these methods: `IX__OSAL_COMPONENT_MAPPING` and `IX_OSAL_MEM_MAP_TYPE`.

The OSAL layer also provides APIs for dealing with the following situations:

- Transparently accessing I/O-memory-mapped hardware in different endian modes
- Transparently accessing SDRAM memory between any endian type and big endian, for the purpose of sharing data with big endian auxiliary processing engines

The OSAL layer supports the following endianness modes:

- Big endian
- Little endian
- Little endian address coherent where
  - Core is operating in little endian mode but the bus addresses are swapped
  - 32-bit word accesses are made automatically in big endian mode
  - Byte and 16-bit half-word addresses are swapped (address XOR 3)
- Little endian, data coherent where,
  - Core is operating in little endian mode but the bus data is swapped
  - Byte accesses are made automatically in big endian mode
  - 32-bit word and 16-bit half-word values are swapped

In little endian mode, users must specify coherency modes before using the IO/Memory access macros (for example, `IX_OSAL_READ_LONG`, `IX_OSAL_WRITE_LONG`). This can be performed by declaring Little Endian Coherency mode in the customized mapping declarations under `os/vxworks/include/platforms/ixp400/IxOsalOsIxp400CustomizedMapping.h`.

[Table 83](#) provides an overview of the I/O memory and endianness support module.



**Table 83. OSAL I/O Memory and Endianness Interface**

Defines required	IX_OSAL_LE_AC	Set to Little Endian, Address Coherency
	IX_OSAL_LE_DC	Set to Little Endian, Data Coherency
	IX_OSAL_BE	Set to Big Endian mode
	IX_OSAL_MEM_MAP	map I/O memory
I/O Mapping	IX_OSAL_MEM_UNMAP	unmap I/O memory
	IX_OSAL_MMAP_PHYS_TO_VIRT	physical to virtual translation
	IX_OSAL_MMAP_VIRT_TO_PHYS	virtual to physical translation

	IX_OSAL_SWAP_LONG	32-bit word byte swap
	IX_OSAL_SWAP_SHORT	16-bit short byte swap
	IX_OSAL_SWAP_SHORT_ADDR	16-bit short address swap
	IX_OSAL_SWAP_BYTE_ADDR	byte address swap
I/O Read/Write	IX_OSAL_READ_BYTE	I/O byte read
	IX_OSAL_WRITE_BYTE	I/O byte write
	IX_OSAL_READ_SHORT	I/O 16-bit short read
	IX_OSAL_WRITE_SHORT	I/O 16-bit short write
	IX_OSAL_READ_LONG	I/O 32-bit word read
	IX_OSAL_WRITE_LONG	I/O 32-bit word write
Mixed endian systems	IX_OSAL_WRITE_BE_SHARED_BYTE	big endian byte write
	IX_OSAL_WRITE_BE_SHARED_SHORT	big endian 16-bit short write
	IX_OSAL_WRITE_BE_SHARED_LONG	big endian 32-bit word write
	IX_OSAL_READ_BE_SHARED_BYTE	big endian byte read
	IX_OSAL_READ_BE_SHARED_SHORT	big endian 16-bit short read
	IX_OSAL_READ_BE_SHARED_LONG	big endian 32-bit word read
	IX_OSAL_SWAP_BE_SHARED_SHORT	big endian 16-bit short swap
	IX_OSAL_SWAP_BE_SHARED_LONG	big endian 32-bit word swap
IX_OSAL_COPY_BE_SHARED_LONG_ARRAY	big endian 32-bit word array copy	

## 24.9 Supporting a New OS

Support for a new operating system can be added separately by creating a new OS-specific folder under the “os” directory, with necessary modification to the core module and the build system to expand the supported OS list. It is not required that a new OS be supported for all the OSAL modules. Similarly, it is not required that supporting a new OS extends to the entire API within a module. For example, the new OS might not support locking via mutexes or semaphores.

To preserve the modularity, it is recommended that any API implementation that can be reused for another OS, and that exists in an OS-specific directory, be moved into the shared directory for the other operating system.



In the process of adding support for a new OS into OSAL the designer(s) should create an OS-specific header file, placed under **osal/os/new\_os/include**, in which there should be definitions for the utility symbols used by OSAL and OSAL applications. This header file should be included by the main OSAL header file (**osal/include/IxOsal.h**).

The process of supporting a new OS should undergo the testing procedure using the OSAL automated verification kit. This test will produce a report detailing API compliance with the specifications. Not all non-compliant APIs are errors. All non-compliant or not implemented functions should be fully documented. Once it is concluded that all the required functions and macros are supported correctly, the new implementation should be submitted to the OSAL maintainers for merging the new OS support into the main OSAL source tree.

## 24.10 Supporting New Platforms

Each platform implementing the I/O memory mapping and endianness support module is required to define a global memory map array, each element in the array having the `IxOsalMemMap` type. Typically each contiguous range in the platform memory map is represented by an entry in the global memory map. To support operating systems using dynamic memory mapping, custom functions for mapping and un-mapping memory must be implemented. These functions have already been implemented for Linux\*

*Note:* Platform specific refers to all the platforms using the same network processor variants. The Intel® IXP4XX product line processors are all part of the same platform in this case. A change in product line refers to using the OSAL layer for a new platform.

The platform package must also include the definition for the global memory map using the `IX_OSAL_IO_MEM_GLOBAL_MEMORY_MAP` define, as in the following example:

```
#define IX_OSAL_IO_MEM_GLOBAL_MEMORY_MAP ixp123GlobalMemoryMap
```

The following is an example fragment of a global memory map:

### Example 1. Global Memory Map Definitions

```
/* For Linux*/

IxOsalMemoryMap ixp123GlobalMemoryMap[] =
{
/* PCI config Registers */
{
    IX_STATIC_MAP,                /* type */
    IXP123_PCI_CFG_BASE_PHYS,     /* physicalAddress */
    IXP123_PCI_CFG_REGION_SIZE,   /* size */
    IXP123_PCI_CFG_BASE_VIRT,     /* virtualAddress */
    NULL,                          /* mapFunction */
    NULL,                          /* unmapFunction */
    0,                             /* refCount */
}
```



```
IX_OSAL_BE,          /* coherency */
"pciConfig"         /* name */
},
},
```

*Note:* The definition of the memory map is very flexible in terms of what operating systems and endianness modes can share memory map cells. Typically, an OS would use only one memory map and share the same cells for big endian and little endian access types. This is exemplified above by setting the access coherency to composite types such as "IX\_OSAL\_BE or IX\_OSAL\_LE\_AC", which means the cell can be used for big endian and little endian/address coherent access. It is, however, not possible to share a cell between both little endian address coherent and data coherent, as these are fundamentally conflicting modes of operation.

### 24.10.1 Module Specific Requirements

An OSAL **platform** is a run-time environment (hardware and software) for OSAL and the software components OSAL helps abstract from the host operating system and, to a more limited extent, from the platform itself.

While the great majority of OSAL functions, constants, symbols and types are OS-dependent and not directly platform-dependent, there are certain elements which require explicit definition for most or all platforms supported by OSAL in order to abstract them from the user application; for example, functions are not covered (or incompletely covered) by the host OS, or symbols used to identify platform-specific values (such as the platform name). The set of these platform-specific definitions constitutes a platform package.

Each module definition will include a set of requirements for the target platform (if applicable). This set of requirements specifies what platform-specific functions, constants and defines are required to be implemented for the platform to be supported by the module. It is not always possible to implement all the required functionality (for example, some platforms cannot provide a timestamp), in which case the default behavior is indicated.

Platform-specific functions, constants and defines are referred to as OEM elements. For each OEM element the following set of requirements is defined:

- whether the element SHOULD or MUST be implemented
- any dependencies on other OEM elements (for example, element X MUST be implemented if element Y is implemented)
- exact type/prototype/name of the element, if applicable
  - note that OEM functions are not required to have a specific name, only a specific type
- meaning of parameters and return value, if applicable
- re-entrancy and IRQ safety requirements
- list of applicable OSes
- name of OSAL define used to link the OEM element, if applicable
  - since OEM function implementation are free to use any name, a symbol must be defined to point to the function

For example, consider the following requirement in [Figure 113](#).



**Figure 113. Requirements for the Routine TimeStamp**

Requirement	SHOULD be supported
Dependencies	Timestamp rate MUST be implemented if this function is implemented
Macro	IX_OSAL_OEM_TIMESTAMP_GET()
Description	Retrieves the timestamp
Parameters	None
Return value	the current timestamp (UINT32)
OS	Linux kernel, VxWorks, QNX, WinCE
Reentrant	Yes
IRQ safe	Yes
Symbol	IX_OSAL_OEM_TIMESTAMP_GET
Default	Default implementation always returns 0

A possible implementation — on a hypothetical IXP123 platform, follows:

```

UINT32 ixp123TimestampGet(void)
{
    return (UINT32) (* (UINT32 *) IXP123_TIMESTAMP_REG_ADDR);
}
    
```

The symbol section of the platform would include:

```
#define IX_OSAL_OEM_TIMESTAMP_GET ixp123TimestampGet
```

In this example it is not absolutely required to define the function, since the requirement specifies SHOULD instead of MUST. If this function is not implemented, the symbol MUST not be defined, which will instruct the OSAL module for which this function is implemented (in this example, the core module) to use the default implementation.

A module may require a certain data structure, constant or define to be declared and populated. These cases are similar to the one above - for example a specific type and OSAL symbol is given, but the platform implementation is free to use any name.

### 24.10.2 General Purpose Requirements

It is strongly encouraged that every platform using interrupts defines a standard set of interrupt level definitions and base I/O memory map in its platform include files. For example:

```

#define IXP123_TIMER_IRQ_LVL (1)
#define IXP123_UART_IRQ_LVL (2)
#define IXP123_PCI_IRQ_LVL(3)
...
    
```



```
#define IXP123_UART_PHYS_BASE(0xC6000000)

#define IXP123_PCI_PHYS_BASE(0xC7000000)

...
```

It is also encouraged that each supported operating system is provided with its own header file to avoid using `#ifdefs` when selecting values and data structures (for example the global memory map) between different operating systems. If the package uses a different header file for each OS, it is required to identify this use case by adding in the platform header file:

Each OS-specific header file must reside in the **include** directory of the package and be named **os/OsName/IxOsaiOs.h**, where **OsName** is replaced (case-sensitive) by the name of the operating system:

- Linux\*
- VxWorks\*
- WinCE\*

This will ensure that the OSAL build system will correctly include the specialized OS header file.

A platform package is required to declare the size of one cache line, using:

```
#define IX_OSAL_CACHE_LINE_SIZE size
```

where **size** is the size of a cache line.

A platform package is required to declare a name for the platform, using:

```
#define IX_OSAL_PLATFORM_NAME name
```

Where **name** is the name of the platform (quotes excluded).

The initial release will include support for the IXP400 software platform as an example to provide platform support across multiple OSes.

## 24.11 Testing Strategy

The OSAL library will have an integrated automated verification kit, the purpose of which is to achieve maximum coverage of the OSAL function calls. All the tests is self-contained and will not require external equipment or user input.

The user is required to include and configure a platform package prior to running the verification kit. Any installed optional module will register its test interface hooks into the main verification kit and therefore is verified automatically as part of the main test.

It is required to test all the supported OSAL modules for each supported operating system to ensure API conformity.

The output of the test procedure is a comprehensive test report detailing tests passed, failed and skipped (if the selected platform does not implement optional functionality).

Note that the OSAL testing module will contain no OS-specific elements, as the intention behind the OSAL primitives is — when implemented — to function in the same manner in all the operating systems.





## 25.0 ADSL Driver

---

This chapter describes the ADSL driver for the Intel® IXP425 Development Platform and Intel® IXP465 Development Platform that supports the STMicroelectronics\* (formally Alcatel\*) MTK-20150 ADSL chipset in the ADSL Termination Unit-Remote (ATU-R) mode of operation.

The ADSL driver is provided as a separate package along with the Intel® IXP400 Software v2.3.

### 25.1 What's New

There are no changes or enhancements to this component in software release 2.3.

### 25.2 Device Support

STMicroelectronics MTK-20150 is supported on IXP425 and IXP465 development platforms. The MTK-20150 chipset is made up of MTC-20154 integrated analog front end and the MTC-20156 DMT/ATM digital modem and ADSL transceiver controller.

### 25.3 ADSL Driver Overview

The two main interfaces to the ADSL chipset are:

- The parallel CTRL-E interface — via the processor's expansion bus
- The ATM UTOPIA data path interface — via the processor's UTOPIA interface

The ADSL driver only supports communication with the ADSL chipset via the CTRL-E interface. All data path communication (ATM UTOPIA) must be performed via the ATM Access Layer component of the software release 2.3.

The driver uses the CTRL-E interface to download the STMicroelectronics firmware, configure and monitor the status of the ADSL chipset. The advantage of downloading the firmware via the CTRL-E interface is that it removes the requirement for a separate flash for the STMicroelectronics ADSL chipset.

The driver provides an API to bring the ADSL line up in ATU-R mode. The line is configured to negotiate the best possible line rate, given the conditions of the local loop when the line is opened. The line rate is not renegotiated once the modems are in the "show-time" mode.

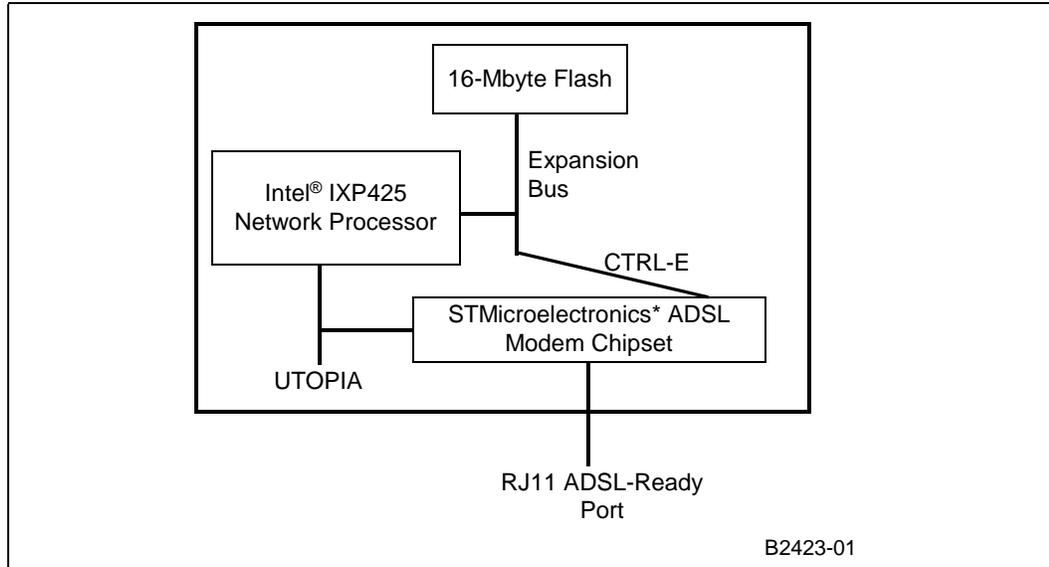
There is very little configuration information required to open an ATU-R line. Almost all line configuration parameters are supplied by the ATU-C side.

APIs are provided to take the modem off line and to check the state of the line to see if the modem is in "show-time" mode.

### 25.3.1 Controlling STMicroelectronics\* ADSL Modem Chipset Through CTRL-E

The STMicroelectronics ADSL chipset CTRL-E interface is memory-mapped into the processor's expansion bus address space. Figure 114 shows how the chipset is connected to the processor.

Figure 114. STMicroelectronics\* ADSL Chipset on the Intel® IXP425 / IXCDP1100 Development Platform



The CTRL-E interface is used for all non-data-path communication between the processor and the ADSL chipset. The ADSL driver public APIs use private driver utilities to convert client requests into CTRL-E commands to the ADSL chipset.

## 25.4 ADSL API

The ADSL driver provides a number of API that provide several general types of functionality. APIs are provided in the following areas:

- Firmware download to the ADSL chipset
- Initialization of the ADSL devices
- Opening, closing and monitoring an ADSL line.
- Soft reset

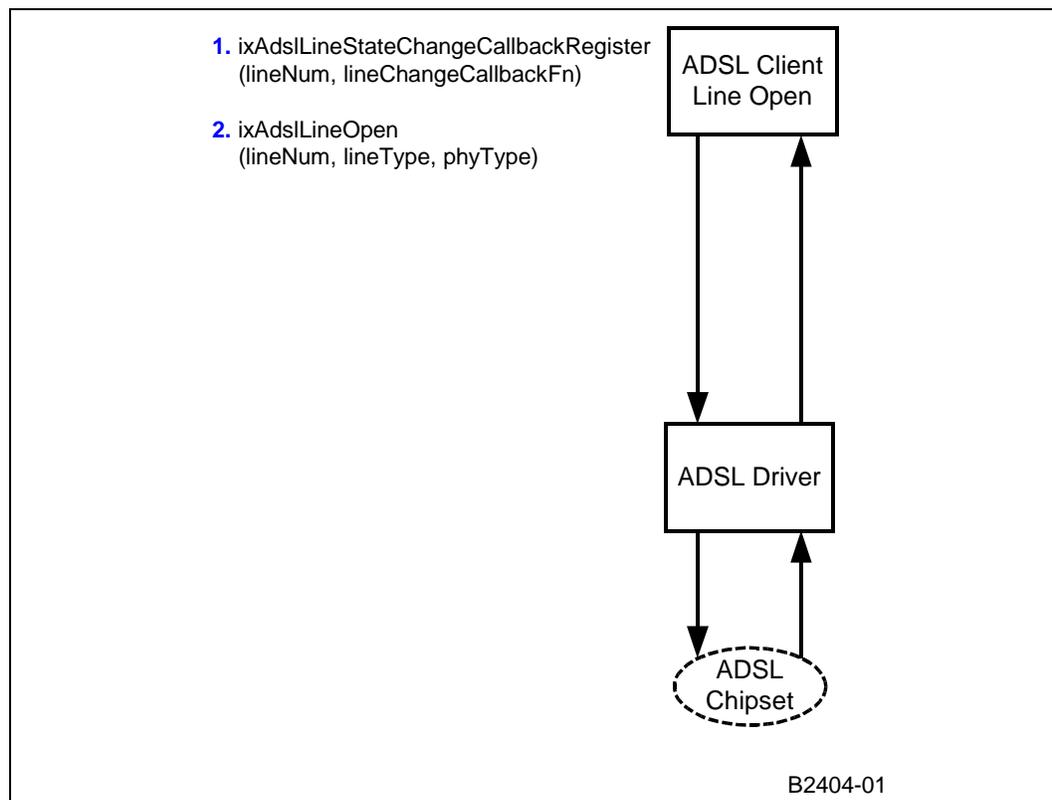
## 25.5 ADSL Line Open/Close Overview

*Note:* Before calling the ADSL driver line open function the ATM Access Layer must be started.

Figure 115 on page 355 provides an example of the ADSL driver functions that the client application code will call to open an ADSL line.



Figure 115. Example of ADSL Line Open Call Sequence



Step 1 of Figure 115 is only required if the client application wants to be notified when a line state changes occurs.

Step 2 of Figure 115 is called by the client application to establish an ATU-R ADSL connection with another modem. This function call performs the following actions within the private context of the ADSL driver:

- a. Invokes the private ixAdslDriverInit function which creates an ixAdslLineSupervisoryTask. This task invokes the ixAdslLineStateMachine.
- b. Invokes the private ixAdslUtilDeviceDownload function which downloads the STMicroelectronics\* ADSL firmware and configures the chipset.
- c. Invokes the private ixAdslCtrlEnableModem function which enables the ADSL chipset to start opening the line.

The client application can close an ADSL line by calling the ixAdslLineClose() API which will disable the modem (for example, close the line) but not kill the ixAdslLineSupervisoryTask.

## 25.6 Limitations and Constraints

- The driver only supports the ATU-R mode of operation.
- The driver can operate in single PHY mode only.

§ §





## 26.0 I<sup>2</sup>C Driver (IxI2cDrv)

---

This chapter describes the I<sup>2</sup>C Driver provided with Intel® IXP400 Software v2.3, which is for use with the Intel® IXP46X Product Line of Network Processors.

### 26.1 What's New

There are no changes or enhancements to this component in software release 2.3.

### 26.2 Introduction

The IXP46X network processors include an I<sup>2</sup>C hardware interface. This I<sup>2</sup>C driver is provided to configure and enable I<sup>2</sup>C hardware and provide a mechanism for transferring data serially through the I<sup>2</sup>C bus in both master and slave mode. Four methods of data transfer are supported by the driver: single-byte read, multi-byte read, single-byte write, and multi-byte write. The driver allows the addressing to any I<sup>2</sup>C Slave on the bus.

The capability to enable/disable the response to I<sup>2</sup>C slave address and general address calls is also provided. Transaction records/counters between the I<sup>2</sup>C hardware and other devices are tracked by the driver. The driver provides the capability to scan the bus to detect I<sup>2</sup>C slave devices and supports multiple I<sup>2</sup>C bus masters.

The driver is implemented in what is referred to as the **Algorithm Module**. This module performs the configuration and control of data transfers. This component is supported on both VxWorks\* and Linux\*.

The driver interface is compatible with the standard Linux\* I<sup>2</sup>C device driver, and is provided separately from the software release 2.3 access-layer. Since Linux\* does not allow direct user mode access to kernel driver functions, a separate "Adapter Module" is provided to accommodate direct access from user mode.

### 26.3 I<sup>2</sup>C Driver API Details

#### 26.3.1 Features

The I<sup>2</sup>C driver allows the setting of different configurations for the I<sup>2</sup>C hardware, as listed below:

- Mode select – fast mode (400 kbps) or normal mode (100 kbps). High Speed (3.4 Mbps) mode is not supported by hardware.
- Flow Selection - Interrupt or Polling modes
- Enable/disable I<sup>2</sup>C unit response to general calls
- Enable/disable I<sup>2</sup>C unit response to slave address calls
- Enable/disable the driving of the SCL line
- I<sup>2</sup>C slave address of the processor



The I<sup>2</sup>C driver features the following hardware and bus status items:

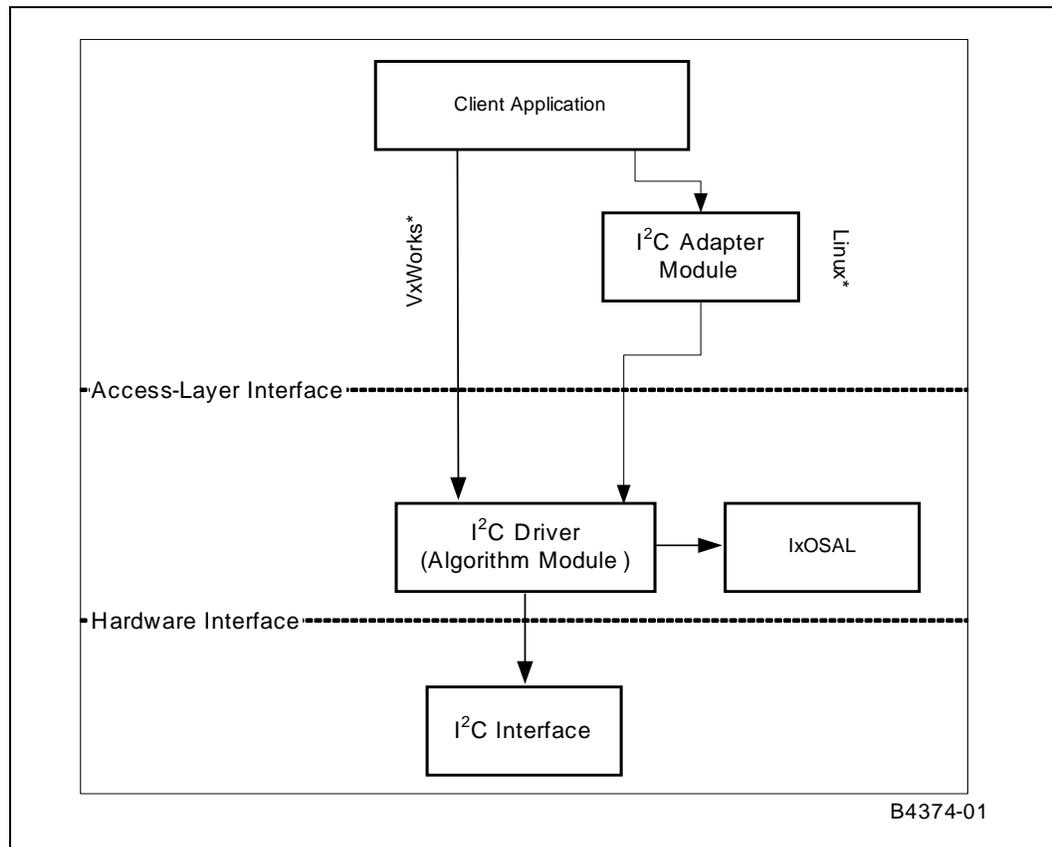
- Master transfer error
- Bus error detected
- Slave address detected
- General call address detected
- IDBR receive full
- IDBR transmit empty
- Arbitration loss detected
- Slave STOP detected
- I<sup>2</sup>C bus busy
- I<sup>2</sup>C unit busy
- Received/sent status for **ACK/NACK**
- Read/write mode (master-transmit/slave-receive or master-receive/slave-transmit)
- Selectable use of internal or OS-provided delay functions.

The I<sup>2</sup>C driver supports single and multi read, single and multi write, and repeated start data transfers for both interrupt and polled mode. A repeated start data transfer is when the master sends a start instead of a stop-start to initiate the next transfer. It is different from a multi read or multi write in that it can allow a read followed by a write or vice versa. Repeated start data transfers in slave mode are not supported.

The I<sup>2</sup>C hardware does not support extended 10-bit I<sup>2</sup>C addressing; only 7-bit slave addressing is supported. The driver will allow any 7-bit slave address (0x01 to 0x7F) except 0x00, which is reserved for general calls.

### 26.3.2 Dependencies

The I<sup>2</sup>C driver is dependent on the capability provided by the I<sup>2</sup>C hardware. Also, the driver is dependant upon IxOSAL to provide OS independency. The adapter module provides the Linux\* driver interface between the user-space applications and the kernel-space adapter module of the I<sup>2</sup>C driver. Therefore the adapter module is dependent on the I<sup>2</sup>C algorithm module. VxWorks\* uses the I<sup>2</sup>C driver directly and does not need an adapter module.

Figure 116. I<sup>2</sup>C Driver Dependencies

### 26.3.3 Error Handling

The I<sup>2</sup>C driver is capable of detecting all errors that the I<sup>2</sup>C hardware is able to provide as listed below:

- Arbitration loss error
- Bus error

Any errors that occur during data transfer which do not fall into the arbitration loss or bus error categories is classified as a master transfer error (IX\_I2C\_MASTER\_XFER\_ERROR).

#### 26.3.3.1 Arbitration Loss Error

This error occurs when the I<sup>2</sup>C hardware of the IXP46X network processors loses master control while it is acting as master. Arbitration loss happens when the unit as master sends a high signal but another master sends a low. The occurrence of two masters on the bus can happen when one I<sup>2</sup>C unit does not see another I<sup>2</sup>C unit's **START** signal to take master of the bus and then sends its own **START** signal to take master of the bus. Such an occurrence can happen when an I<sup>2</sup>C unit just exited reset and has no history of previous signals. When this occurs, the I<sup>2</sup>C status register is updated with the arbitration loss by the hardware, and if the interrupt for arbitration loss is enabled, then it will call the interrupt service routine.



Once an arbitration loss error is detected, the unit will stop transmitting. The client will need to call the transfer again and the I<sup>2</sup>C status register is checked to determine the busy status of the I<sup>2</sup>C bus. If the bus is not busy, the transfer that occurred before the bus arbitration loss error is re-submitted.

### 26.3.3.2 Bus Error

This error occurs when the I<sup>2</sup>C unit, as a master transmitter, does not receive an **ACK** in response to transmission. A bus error can also occur when the I<sup>2</sup>C unit is operating as a slave receiver, and a **NACK** pulse is generated. In master transmit mode, the hardware will abort the transaction by automatically sending a **STOP** signal. As a slave receiver, the behavior will depend on the master's action. The counters for both occurrences is updated accordingly.

## 26.4 I<sup>2</sup>C Driver API Usage Models

### 26.4.1 Initialization and General Data Model

This description assumes a single client model where there is a single application-level program configuring the I<sup>2</sup>C interface and initiating I/O operations.

#### Initialization

The client must first define the initial configuration of the I<sup>2</sup>C port by storing a number of values in the `IxI2cInitVars` structure. The values include the speed selection, data flow mode, pointers to callback functions for various data scenarios, hardware address, and behavior settings for how the I<sup>2</sup>C unit responds to general call and slave address calls. After the structure is defined, `ixI2cDrvInit()` may be called to enable the port.

Once the port is enabled, the client will use one of the data models described later in this chapter (either Interrupt or Polling mode) to determine how and when data I/O operations need to occur.

A callback or handler may be registered for interrupt transmit and receive operations in the `IxI2cInitVars` structure. There are different callbacks for when the I<sup>2</sup>C unit is operating in master or slave mode, and also for general calls.

#### Master-Interrupt Mode

The client will use the `ixI2cDrvWriteTransfer()` and `ixI2cDrvReadTransfer()` functions for transmitting and receiving data on the I<sup>2</sup>C bus in master mode. The functions will return immediately, even though the transfer has not completed. Upon function return, the callback routines registered in `IxI2cInitVars` is executed. The I<sup>2</sup>C unit will handle the appropriate arbitration and bus messaging required to support the transfer type and mode.

While the I<sup>2</sup>C unit is in Master-Interrupt mode, the use of interrupt callbacks is optional. If no callbacks are registered, the read/write transfer functions discussed above will wait until the transfer operation has completed before returning to the calling application. This method can be used if transfer status information is not needed for each transaction and simplifies the implementation of repeated start transfers. The data that is passed in the callback includes transfer mode, buffer pointer and buffer size. Since this data is already known to the client application, processing of this data via the callback would be inefficient.



### Slave-Interrupt Mode

When the processor is acting in I<sup>2</sup>C slave mode or responding to general calls in interrupt mode, the client callbacks for transmit and receive are responsible for providing a buffer used to interface with the I<sup>2</sup>C Data Buffer Register (IDBR), using the **ixI2cDrvSlaveOrGenCallBufReplenish()** function.

Examples of Slave Interrupt mode operations is provided in [“Example Sequence Flows for Slave Mode”](#) on page 362.

### Slave-Polling Mode

In polling mode, the client polling task can check for pending requests to respond to slave request or general calls using the **ixI2cDrvSlaveAddrAndGenCallDetectedCheck()** function. The client can then use the **ixI2cDrvSlaveOrGenDataReceive()** or **ixI2cDrvSlaveOrGenDataTransmit()** functions to transfer data.

### Support Functions

After the I<sup>2</sup>C unit has been initialized as described above, there are several supporting functions available in the API. These include functions that set the 7-bit Slave address to which the I<sup>2</sup>C Unit responds, scan the I<sup>2</sup>C bus for slave units, check or reset port statistics, and show the current status of the I<sup>2</sup>C unit and driver. The API can also uninitialized the I<sup>2</sup>C unit and remove the driver from memory.

### 26.4.2 Example Sequence Flows for Slave Mode

Figure 117. Sequence Flow Diagram for Slave Receive / General Call in Interrupt Mode

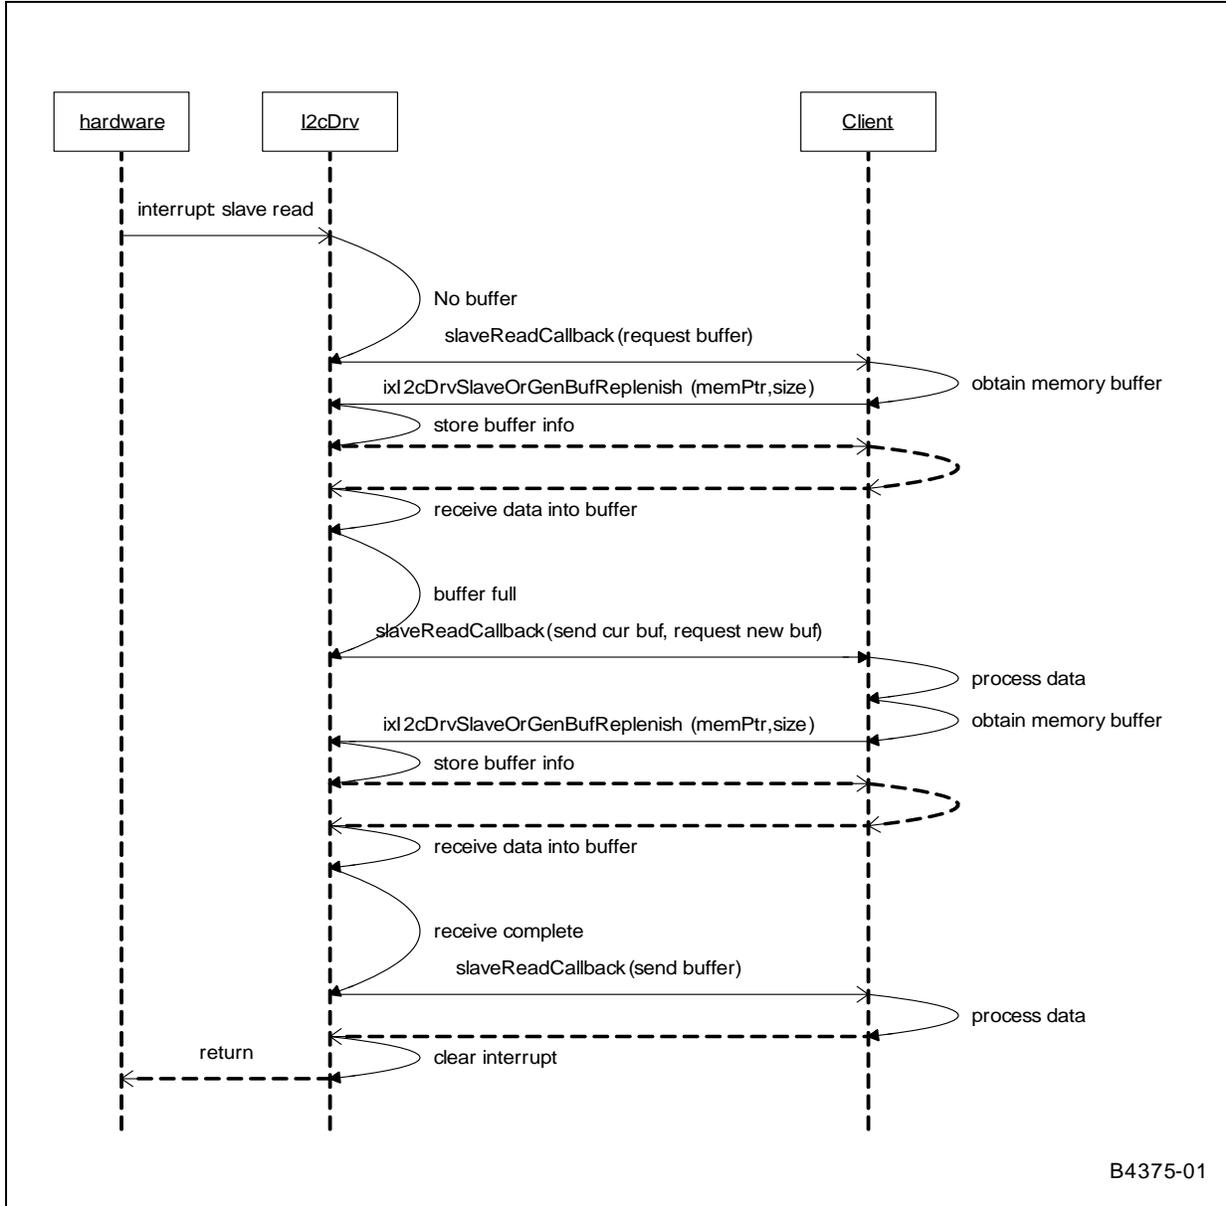




Figure 118. Sequence Flow Diagram for Slave Transmit in Interrupt Mode

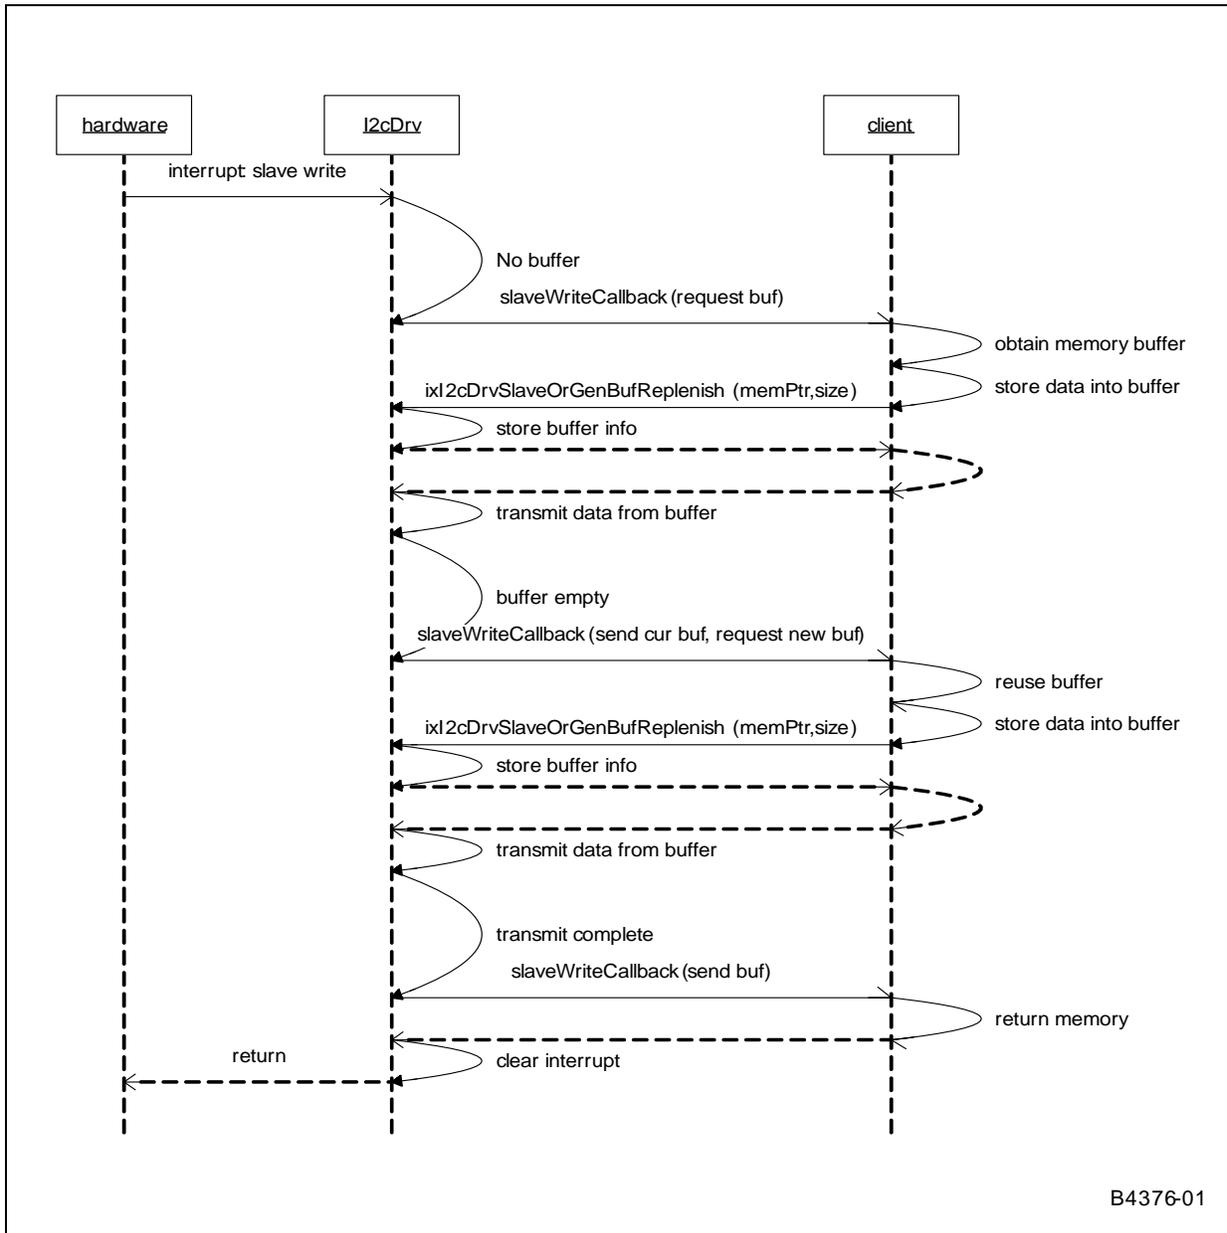


Figure 119. Sequence Flow Diagram for Slave Receive in Polling Mode

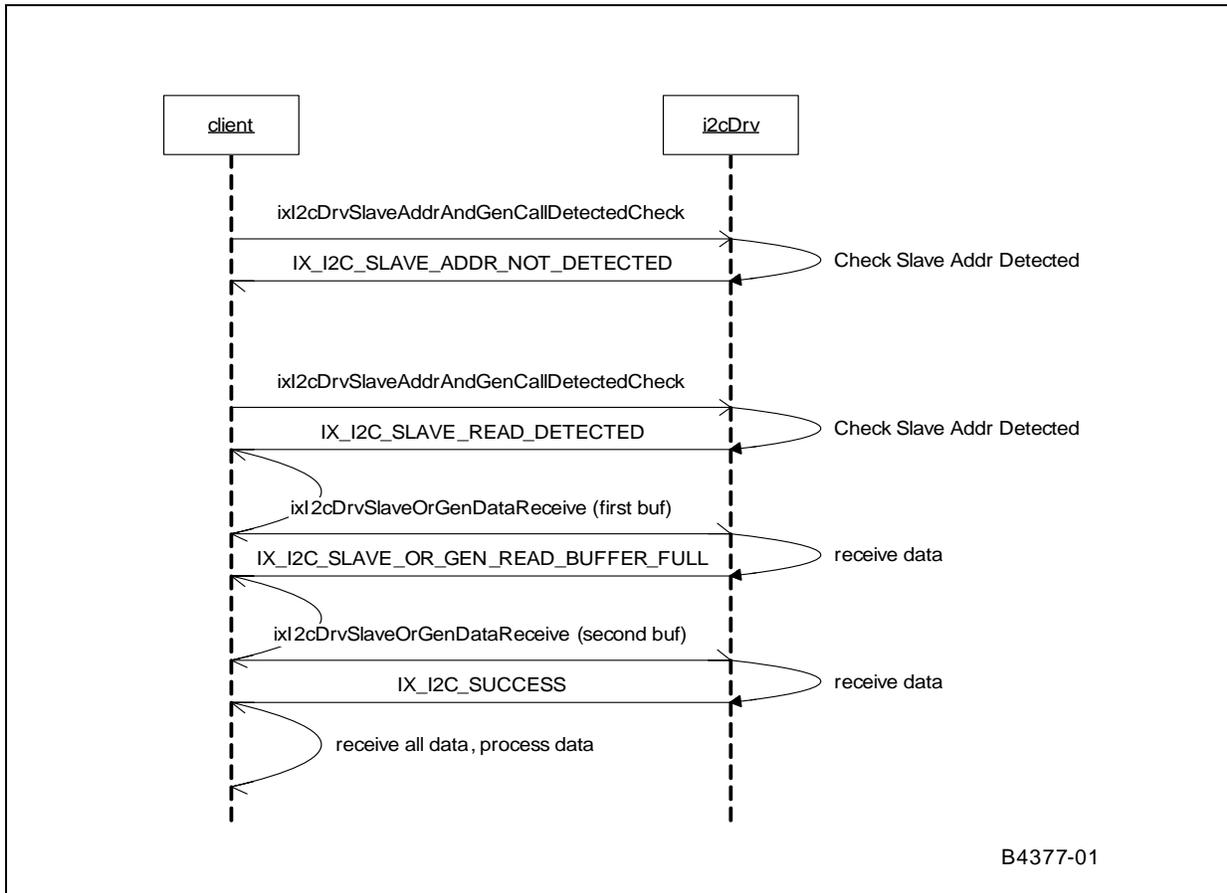
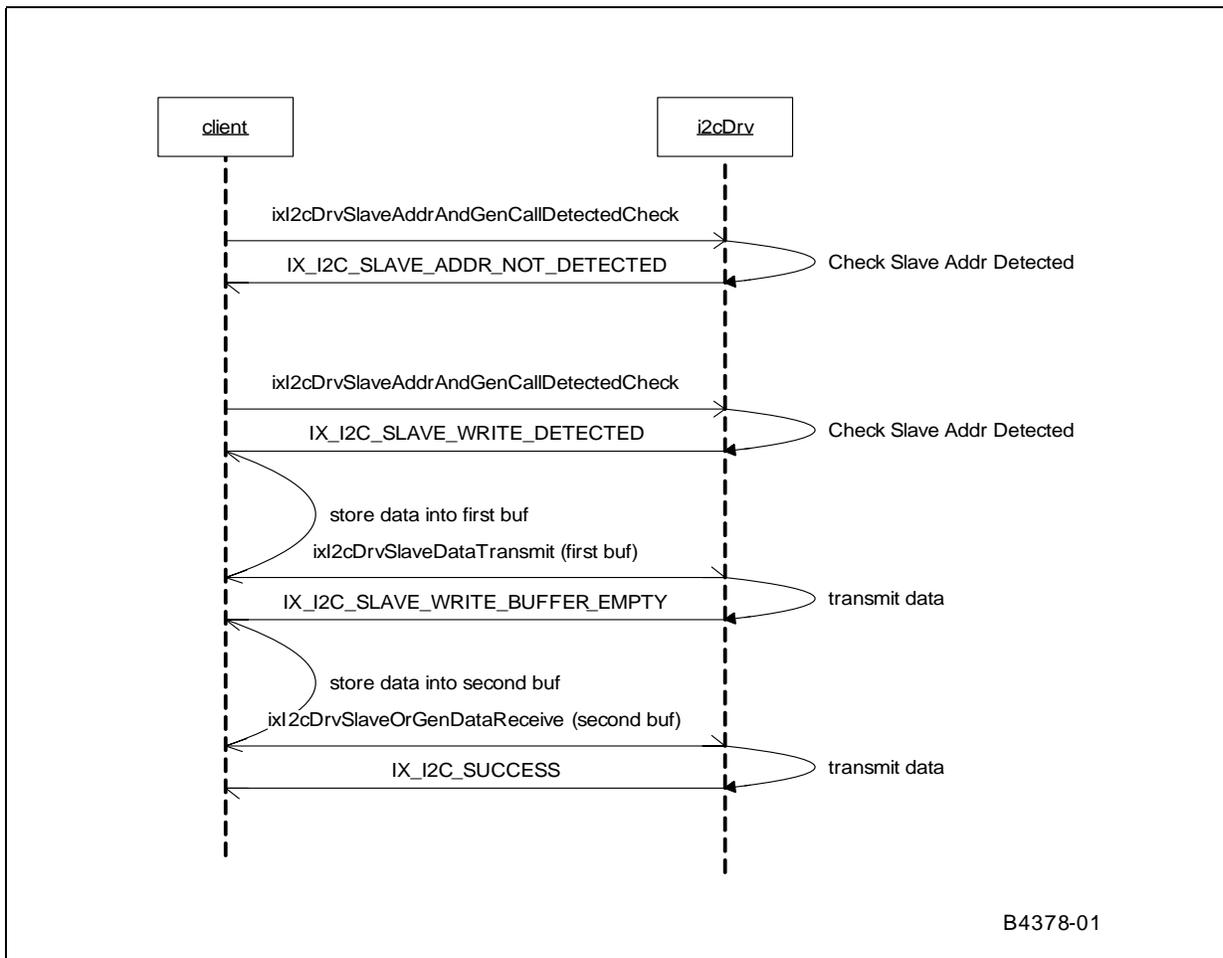




Figure 120. Sequence Flow Diagram for Slave Transmit in Polling Mode



### 26.4.3 I<sup>2</sup>C Using GPIO Versus Dedicated I<sup>2</sup>C Hardware

Some supported operating systems include support for emulating the I<sup>2</sup>C bus using GPIO lines on the processor.

The I<sup>2</sup>C driver using a dedicated I<sup>2</sup>C hardware is a totally different implementation from the driver using GPIO lines. Most of the APIs in a driver using a GPIO implementation are very low level (dedicated to controlling the SDA and SCL lines) and combine to make one transaction. The driver APIs using dedicated I<sup>2</sup>C hardware (such as with IxI2cDrv) is limited to the control provided by the hardware unit on the processor. Furthermore, the dedicated I<sup>2</sup>C hardware implementation allows more advanced features supported by the hardware, such as those to support multi-master on the bus, therefore allowing the IXP46X network processors to act as slave devices.

§ §





## 27.0 Endianness in Intel® IXP400 Software v2.3

---

### 27.1 What's New

There are no changes or enhancements to this component in software release 2.3.

Note that IXP42X product line stepping A0 processors are no longer supported.

### 27.2 Overview

The Intel® IXP4XX Product Line of Network Processors support little endian (LE) and big endian (BE) operations. This chapter discusses software release 2.3 support for LE and BE operation.

This chapter is intended for software engineers developing software or board-support packages (BSPs) that are reliant on endianness support in the processor. The chapter is intended as an introduction to the most important facts regarding endianness as it relates to the software release 2.3.

A more detailed guide to endianness in the IXP42X product line is available in the application note, *Intel® IXP42X Product Line of Network Processors and IXC1100 Control Plane Processor: Understanding Big Endian and Little Endian Modes*, which is freely available from the following Intel Developer Web site:

<http://www.intel.com/design/network/products/npfamily/docs/ixp4xx.htm>

#### Applicability to Specific Processors and Development Platforms

In general, the theories discussed in this chapter are applicable the entire Intel® IXP4XX Product Line of Network Processors. Each product generation does have some specific endianness related capabilities, as listed in “[Silicon Versions](#)” on page 379.

When discussing board-support package (BSP) issues for the Intel® IXP42X Product Line of Network Processors and IXC1100 Control Plane Processor, this chapter refers to the Intel® IXDP425 / IXCDP1100 Development Platform. For the Intel® IXP46X Product Line of Network Processors, this chapter refers to the Intel® IXDPG425 Network Gateway Development Platform.

### 27.3 The Basics of Endianness

Endianness is the numbering organization format of data representation in a computer. Endianness comes in two primary varieties: Big and Little.

Consider the 32-bit number 0x11223344 stored at address 0x1000 in a big endian system and little endian system. The memory would appear as follows:

Big-endian				Little-endian			
0x1000	0x1001	0x1002	0x1003	0x1000	0x1001	0x1002	0x1003
11	22	33	44	44	33	22	11



With big endian systems the most significant byte (11 in our example) is stored at the memory location with the lowest address. The next byte of significance (22 in our example) is stored at the next memory location, and so on. With little endian systems the least-significant byte is stored at lowest memory location.

All processors are either Big- or little endian. Some processors, such as those in the IXP4XX product line processors, have a bit in a register that allows the programmer to select the desired endianness. This is described in [Section 27.5.3.2, “Intel XScale® Processor Endianness Mode”](#) on page 377.

### 27.3.1 The Nature of Endianness: Hardware or Software?

A processor may be capable of supporting both LE and BE with the active form of endianness being dependent on bus behavior and the memory systems connected to that bus. Only correct matching between the processor's mode, bus mode (that is, how the bus and memory are connected), and the software will provide correct endian behavior.

Endianness is, in general, a hardware *and* software issue. However, a processor does not operate in a vacuum. It is part of a system. This implies that a hardware board with processors and memory components (unless specially designed to support both endians) would only support one endian mode, and software on any processor in the system must work with that same endian mode.

### 27.3.2 Endianness When Memory is Shared

Following the definition of endianness from a software point of view, and assuming a piece of hardware can be extremely complex and intelligent, can a piece of memory being shared by two processors running under different endian modes achieve all “IDEAL\_BI\_ENDIAN” objectives at the same time? The objectives for such a system are as follows:

- Share long integers correctly.  
“Correctly” could be defined as one processor ‘feeling’ that the other processor is under the same endianness mode as itself. For example, ProcessorBig writes some data starting from its view of address X. Then, if ProcessorLittle reads the same amount of data starting from its own view of address X, the data read is the same as the data written by ProcessorBig.
- Share short integers correctly.
- Share byte integers correctly.
- Each processor has its own endianness consistency.

Unfortunately, the answer is NO even with help from the most sophisticated hardware.

## 27.4 Software Considerations and Implications

Much literature is available explaining the software dependency on underlying hardware endianness.

In summary, software dependency on hardware endianness is manifested in these areas:

- Whenever a piece of software accesses a piece of memory which is treated as different sizes by manipulation of pointers in different parts of code, that code is endian-dependent. For example, IP address 0x01020304 can be treated as unsigned long. But if code must access byte 0x04 by manipulating pointers, the code becomes endian-dependent.



- If a piece of memory is accessed by other hardware or processors whose endian modes are independent of the processor on which the current software is running, then the current code becomes endian-dependent. For example, network data is always assumed to be big endian. If network data is directly moved (DMA'ed) into memory as it is, then that particular piece of memory is always big endian. As a result, the current code accessing that piece of memory becomes endian-dependent. If pointers are passed between processors, endian issues show immediately because of the fundamental difficulty, as explained in [“The Nature of Endianness: Hardware or Software?”](#) on page 368.
- The above issues can occur in many places of an operating system, a hardware driver, or even a piece of application code. Some operating systems (for example, VxWorks\*) support both endians by different compilation switches.
- Compiler, debugger, and other tools are generally endian-dependent because the translation between a high-level language (for example, C) and assembly language is endian-dependent.

Under certain application assumptions, and when programming carefully, it is possible to have a piece of code that is endian-independent.

## 27.4.1 Coding Pitfalls — Little Endian/Big Endian

The risks associated with programming in mixed endian system generally revolve around possible incompatibilities in the interpretation of data between little endian and big endian components within the system. The following examples illustrate some instances where pitfalls in coding can be interpreted differently on LE versus BE machines (and thus should be avoided). There are also examples of how to code a module in a way that permits a consistent interpretation of data structures and data accesses in general, regardless of the endianness of the processor the code may be running on. Performance can also enter into the equation, especially if byte order must be frequently shuffled by the processor.

### 27.4.1.1 Casting a Pointer Between Types of Different Sizes

The situation that this example illustrates must be avoided completely. Do not mix pointer sizes. Endianness causes different interpretation from one machine to the next, making porting problematic.

```
int J=8;

char c = *(char *) J;
```

Depending on the endianness of the processor the code is executing on, the result is:

```
Little:0x8
Big:0x0
```

The following provides another example of endianness causing the code to be interpreted differently on BE versus LE machines:

```
int myString[2] = { 0x61626364,0}; /* hex values for ascii */

Printf("%s\n", (char *)&myString);
```

Depending on the endianness of the processor the code is executing on, the result is:

```
Little:"dcba"
Big:"abcd"
```



### 27.4.1.2 Network Stacks and Protocols

**Little endian Machines:** Running a network protocol stack on a little endian processor can degrade performance due to formatting translation. If a network protocol stack is to be run on a little endian processor, at run time it will must reorder the bytes of every multi-byte data field within the various layers' headers.

**Big endian Machines:** Running a network protocol stack on a big endian processor does not degrade performance due to formatting translation. If the stack will run on a big endian processor, there is nothing to worry about; the endianness of the processor inherently matches the format of standard network data ordering.

### 27.4.1.3 Shared Data Example: LE Re-Ordering Data for BE Network Traffic

By using a macro conversion routine, the data access is re-ordered as needed to properly interpret data moving between a network (which is using big endian or network order) and a host machine, which may be little endian.

Basic Assumptions:

- TCP/IP defines the network byte order as big endian.
- Little endian machines must byte swap accesses to 16-/32-bit data types (IP address, checksum, and so forth).

Example: We want to assign the value of the IP source address field in the header of an IP packet to a 32-bit value we will call "src." Here is the code, which features a macro to translate.

```
u_long src = ntohs(ip->ip_src.s_addr);
```

Here is what the macro ntohs() looks like in actual code:

```
-ntohs()  
  
{  
  
#if (_BYTE_ORDER == _BIG_ENDIAN)  
  
        #define ntohs(x)          (x)  
  
  
#else  
  
        #define ntohs(x)((((x) & 0x000000ff) << 24) | \  
        ((x) & 0x0000ff00) << 8) | \  
        ((x) & 0x00ff0000) >> 8) | \  
        ((x) & 0xff000000) >> 24)  
  
#endif  
  
}
```

We always assume that the byte order value is set to either big endian or little endian in a define value.



## 27.4.2 Best Practices in Coding of Endian-Independence

### Avoid

- Code that assumes the ordering of data types in memory.
- Casting between different-sized types.

### Do

- Perform any endian-sensitive data accesses in macros. If the machine is big endian, the macros will not have a performance hit. A little endian machine will interpret the data correctly.

## 27.4.3 Macro Examples: Endian Conversion

A common solution to the endianness conversion problem associated with networking is to define a set of four preprocessor macros: `htons()`, `htonl()`, `ntohs()`, and `ntohl()`. These macros make the following conversions:

`htons()`: The macro name can be read “host to network short.”  
reorder the bytes of a **16-bit value** from processor order to *network order*.

`htonl()`: The macro name can be read “host to network long.”  
reorder the bytes of a **32-bit value** from processor order to *network order*.

`ntohs()`: The macro name can be read “network to host short.”  
reorder the bytes of a **16-bit value** from *network order* to processor order.

`ntohl()`: The macro name can be read “network to host long.”  
reorder the bytes of a **32-bit value** from *network order* to processor order.

### 27.4.3.1 Macro Source Code

If the processor on which the TCP/IP stack is to be run is itself also big endian, each of the four macros is defined to do nothing and there is no run-time performance impact. If the processor is little endian, the macros will reorder the bytes appropriately. These macros would be used when building and parsing network packets and when socket connections are created.

By using macros to handle any possibly sensitive data conversions, the problem of dealing with network byte order (Big endian) on a little endian machine is eliminated. Ideally all network processors would have the same endianness. Because that is not true, understand and use the following macros as needed.

#### 27.4.3.1.1 Endianness Format Conversions

```
#if defined(BIG_ENDIAN) /* the value of A will not be manipulated */
    #define htons(A) (A)
    #define htonl(A) (A)
    #define ntohs(A) (A)
    #define ntohl(A) (A)

#elif defined(LITTLE_ENDIAN) /* the value of A is byte swapped */
```



```
#define htons(A) (((A) & 0xff00) >> 8) | ((A) & 0x00ff) << 8))

#define htonl(A) (((A) & 0xff000000) >> 24) | \
    (((A) & 0x00ff0000) >> 8) | \
    (((A) & 0x0000ff00) << 8) | \
    (((A) & 0x000000ff) << 24))

#define ntohs htons
#define ntohl htohl

#else

#error "One of BIG_ENDIAN or LITTLE_ENDIAN must be #defined."

#endif
```

## 27.5 Endianness Features of the Intel® IXP4XX Product Line of Network Processors

Within the Intel® IXP4XX Product Line of Network Processors, there are several devices connected via the system bus. The system consists of the Intel XScale® Processor, network processing engines, PCI devices, APB peripherals and expansion bus peripherals. The Intel XScale® Processor may operate in either little or big endian mode. The operation of the Intel XScale® Processor in little endian mode creates a mixed-endian system.

Supporting more than one endian in a system may have two meanings:

- Case 1: Either Big or little endian in the entire system, but not mixed;
- Case 2: Some hardware components running in one endian mode while others running in the other endian mode.

The IDEAL\_BI\_ENDIAN objectives cannot be achieved in the second case but can be achieved in the first case, as explained in [“Endianness When Memory is Shared” on page 368](#). An IXP4XX processor or a system based upon such as processor belongs in the second case.

In order to support more than one endianness as implied by “Case 2”, a hardware byte-swapping or address swizzling (or munging) facility is usually employed.

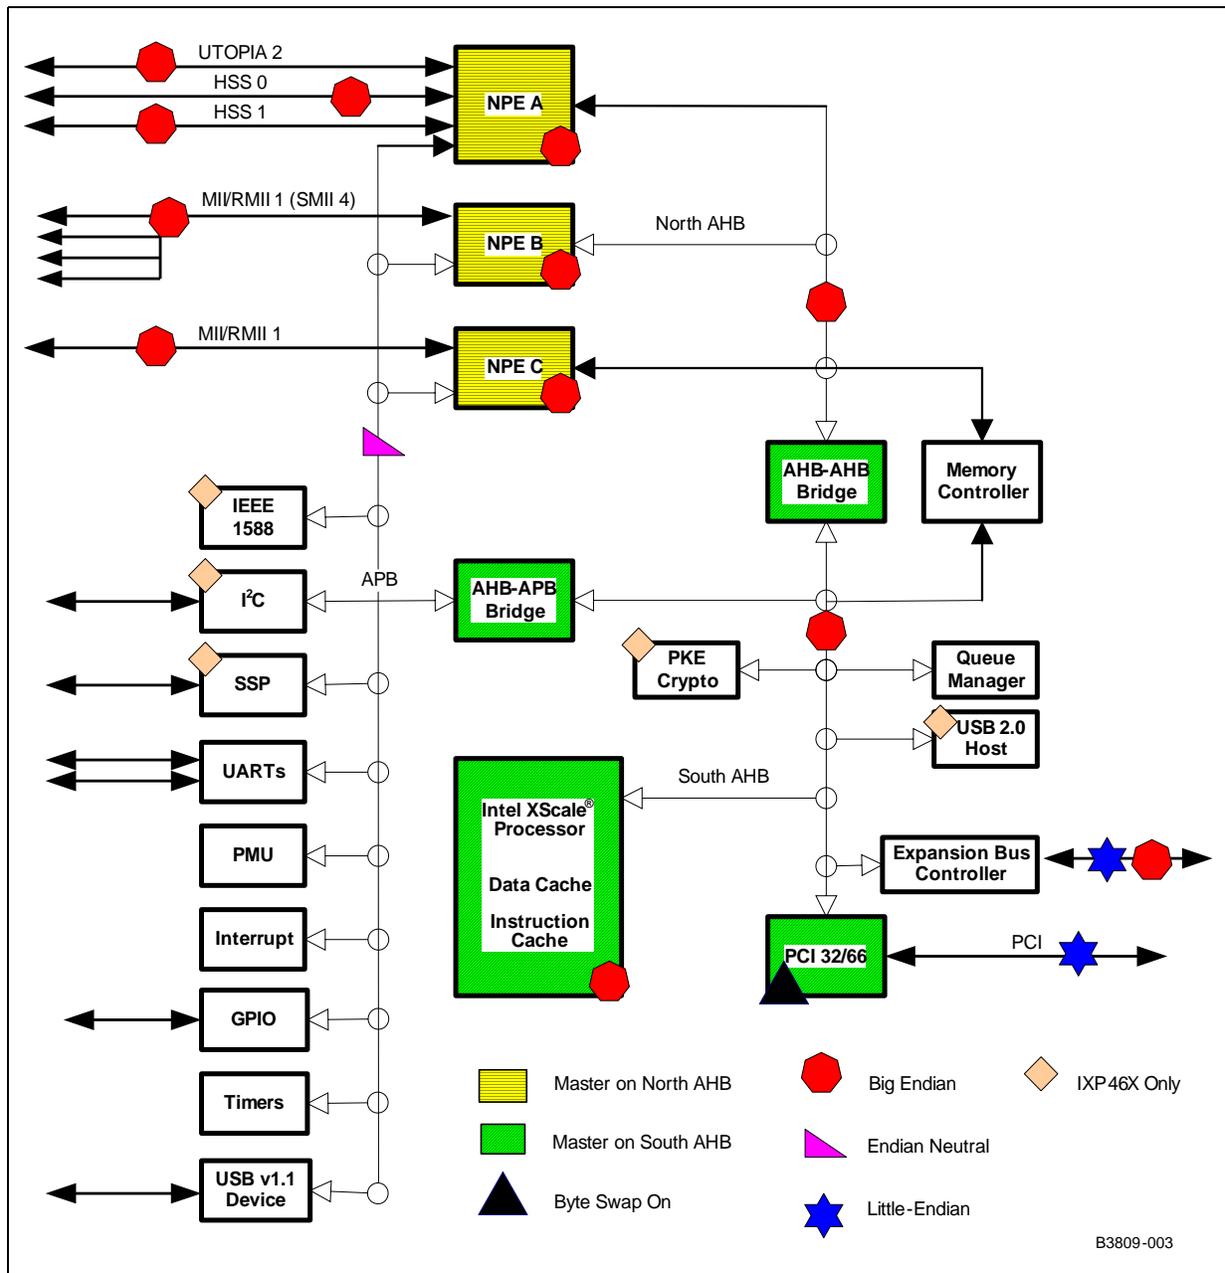
When a piece of memory is accessed by different pieces of hardware through different buses, a bus bridge is usually a good place to perform byte swapping or address swizzling. This ensures that each processor does not need to do any endian adjustments. Instead, the processor assumes the underlying hardware behaves as if it is the same endianness as the processor.



This chapter will provide an overview of the IXP4XX product line processors capabilities related to endianness. For specific detail on the various capabilities and hardware settings for the processors, refer to that processor's specific *datasheet* and *developer's manual*.

Figure 121 details the endianness of the different blocks of the IXP4XX processors when running a big endian software release.

Figure 121. Endianness in Big Endian-Only Software Release





## 27.5.1 Supporting Little Endian Mode

The following hardware items can be configured by software:

- Intel XScale® Processor running under little or big endian mode.
- The byte-swapping hardware in the PCI controller turned on or off.

The following hardware items cannot be changed by software or off-chip hardware (for example board design):

- AHB bus is running under big endian mode.
- NPEs are running in big endian mode relative to their own memory, and relative to AHB memory.

By default, the software release 2.3 is designed to operate in big endian mode and configures the Intel XScale® Processor and PCI controller as such.

Given the above hardware design, supporting little endian in the IXP4XX processors while using the IXP400 software requires the following changes in hardware:

- The Intel XScale® Processor is left to its standard default configuration, which is little endian mode.
- The byte-swapping hardware in PCI controller is turned off by setting the following register values: `pci_csr_ads=0`, `pci_csr_pds=0`, `pci_csr_abe=1`. The Intel® IXP400 Software sets the following values to support the default big endian operation: `pci_csr_ads=1`, `pci_csr_pds=1`, `pci_csr_abe=1`.

When the changes outlined above are applied, the Intel XScale® Processor will run under little endian mode while other processors in the system (for example, the NPEs) remain running under the same endian mode as defined in software release 2.3. The result is that the IXP4XX processor is running as an endian-hybrid system.

The information outlined above is a simplification of the options available in the IXP4XX product line processors, but does cover the basic concepts. Further detail is provided in following sections.

## 27.5.2 Reasons for Choosing a Particular LE Coherency Mode

Little endian mode is sub-divided into two categories:

- Intel XScale® Processor operating in **Address Coherent** mode
- Intel XScale® Processor operating in **Data Coherent** mode

Both Address and Data Coherent endian conversion are provided because there are different benefits and hazards to both approaches. If the only goal of the endian conversion was to make the Intel XScale® Processor self-consistent, meaning that the Intel XScale® Processor properly reads what it wrote, then either method would be sufficient. However, since the Intel XScale® Processor must communicate with other processors and interfaces within the IXP4XX processor, it is beneficial to provide both methods.

To understand this, consider the benefits and hazards of both approaches by examining the details of how data is stored in memory. In particular, how will the NPE read and interpret the data stored in memory? When the Intel XScale® Processor is in big endian mode, the NPE reads the data in the same format that it was written.

When the Intel XScale® Processor is in little endian **Address Coherent** mode, words written by the Intel XScale® Processor are in the same format when read by the NPE as words. However, byte accesses appear reversed and half-word accesses return the other half-word of the word. The benefit of this mode is that if the Intel XScale®



Processor is writing a 32-bit address to memory, the NPE could read that address correctly without having to do any conversion. Additionally, LE Address Coherent instructions are in the same format as they would be for big endian operation. The same program image could be used for Big- and little endian modes because instructions are the same from the point of view of the Intel XScale® Processor.

Table 84 illustrates what the NPE processor reads if the Intel XScale® Processor is in Little Endian Address coherent mode after Little Endian Writes of different sizes (byte, half-word, word).

**Table 84. Intel XScale® Processor Little Endian Writes in Address Coherent Mode and NPE Reads to/from SDRAM**

XScale: Little Endian Writes Address Coherent (Address Swizzle)			NPE: Big Endian Processor Reads		
Size	Address	Data	Size	Address	Data
Byte	0	AA	Byte	0	DD
	1	BB		1	CC
	2	CC		2	BB
	3	DD		3	AA
			Half-word	0	DDCC
				2	BBAA
			Word	0	DDCCBBAA
Half-word	0	AABB	Byte	0	CC
	2	CCDD		1	DD
				2	BB
	0			3	AA
	2		Half-word	0	CCDD
				2	AABB
			Word	0	CCDDAABB
Word	0	AABBCCDD	Byte	0	AA
				1	BB
				2	CC
				3	DD
			Half-word	0	AABB
				2	CCDD
			Word	0	AABBCCDD

When the Intel XScale® Processor is in little endian **Data Coherent** mode, bytes written by the Intel XScale® Processor are in the same format when read as bytes by the NPE. However, the bytes within a word and half-word appear reversed. This endian conversion method is beneficial when data is written and read as bytes. Additionally, many commercially available software protocol stacks were written to support both Big- and little endian modes. These stacks assume a Data Coherent endian conversion and provide all the necessary byte swapping to correct words and half-words.



Table 85 illustrates what the NPE processor reads if the Intel XScale® Processor is in Little Endian Data coherent mode after Little Endian Writes of different sizes (byte, half-word, word).

**Table 85. Intel XScale® Processor Little Endian Writes in Data Coherent Mode and NPE Reads to/from SDRAM**

XScale: Little Endian Writes Data Coherent (Byte Swizzle)			NPE: Big Endian Processor Reads		
Size	Address	Data	Size	Address	Data
Byte	0	AA	Byte	0	AA
	1	BB		1	BB
	2	CC		2	CC
	3	DD		3	DD
			Half-word	0	AABB
				2	CCDD
			Word	0	AABCCDD
Half-word	0	AABB	Byte	0	BB
	2	CCDD		1	AA
				2	DD
	0			3	CC
	2		Half-word	0	BBAA
				2	DDCC
			Word	0	BBAADDCC
Word	0	AABCCDD	Byte	0	DD
				1	CC
				2	BB
				3	AA
			Half-word	0	DDCC
				2	BBAA
			Word	0	DDCCBBAA

By providing both types of endian conversion through the use of the P-attribute bit in the MMU, the software has the flexibility to use whichever method is most convenient for the particular task.

### 27.5.3 Silicon Endianness Controls

#### 27.5.3.1 Hardware Switches

There are many hardware endianness controls available to the software. However, the following are the most important and play a significant role in the operation of software.

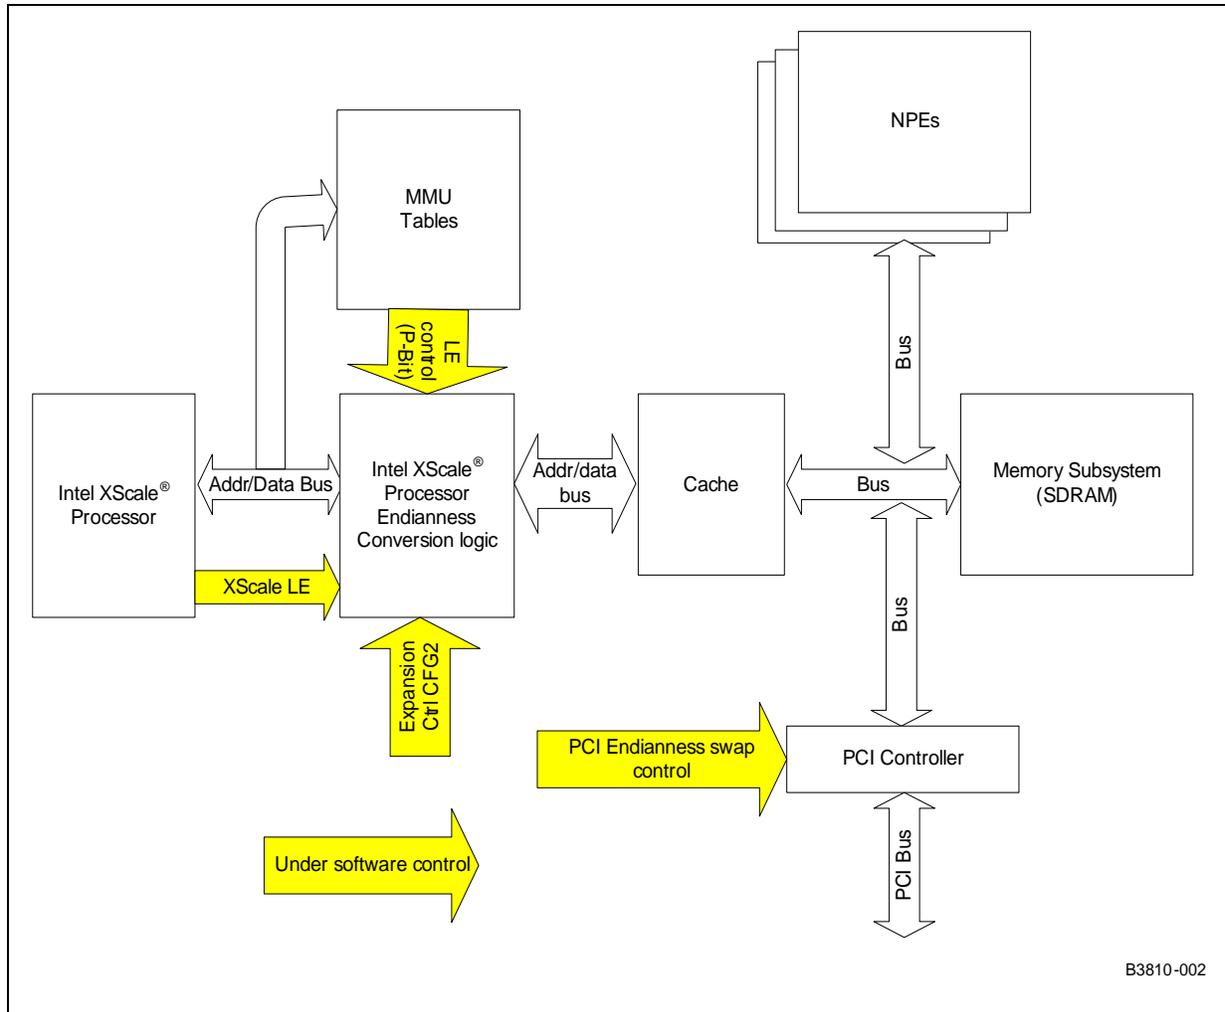
- Intel XScale® Processor BE/LE mode



- Expansion Bus Control Register 1: BYTE\_SWAP\_EN bit.
- Expansion Bus Control Register 1: FORCE\_BYTE\_SWAP bit.
- MMU Page table “P” attribute bit.
- PCI Bus swapping control.

The default operation of the IXP4XX product line processors on reset is: Intel XScale® Processor little endian, Address Coherent, MMU-disabled.

**Figure 122. Intel® IXP4XX Product Line of Network Processors Endianness Controls**



### 27.5.3.2 Intel XScale® Processor Endianness Mode

The big- and little endian modes are controlled by the B-bit, located in the “Intel StrongARM Control Register”, coprocessor 15, register 1, bit 7. The default mode at reset is little endian. To enable the big endian mode, the B bit must be set before performing any sub-word accesses to memory, or undefined results would occur. The bit takes effect even if the MMU is disabled. The following is assembly code to enable/clear the B-bit.



```
MACRO LITTLEENDIAN

MRC p15,0,a1,c1,c0,0

BIC a1,a1,#0x80 ;clear bit7 of register1 cp15

MCR p15,0,a1,c1,c0,0

ENDM
```

```
MACRO BIGENDIAN

MRC p15,0,a1,c1,c0,0

ORR a1,a1,#0x80 ;set bit7 of register1 cp15

MCR p15,0,a1,c1,c0,0

ENDM
```

The application code built to run on the system must be compiled to match the endianness. Some compilers generate code in little endian mode by default. To produce the object code that is targeted for a big endian system, the compiler must be instructed to work in big endian mode. For example, a **-mbig-endian** switch must be specified for GNU\* CC since the default operation is in little endian. For GNUPro\* assembler, **-EB** switch would assemble the code for big endian. The library being used must have been compiled in the correct endian mode.

### 27.5.3.3 Little Endian Data Coherence Enable/Disable

IXP4XX product line processors allow for MMU control of the coherence mode used on a per-MMU-page basis. These capabilities are enabled/disabled via the EXP\_CNFG1 register at physical address 0xC4000024.

#### BYTE\_SWAP\_EN (Bit 8)

This bit affects only transactions initiated by the Intel XScale® Processor. If Intel XScale® Processor endianness mode is little endian, then:

- BYTE\_SWAP\_EN = 1 - The MMU P Bit controls the selection of address or data coherency.
- ..... BYTE\_SWAP\_EN = 0 - Always address coherence mode if LE selected.

The bit has no effect if the Intel XScale® Processor is in big endian mode.

#### FORCE\_BYTE\_SWAP (Bit 9)

The IXP46X product line provides the ability to override any P-attribute bit settings in the page table. When this bit is set and the Intel XScale® Processor endianness mode is little endian, BYTE\_SWAP\_EN is ignored and Data Coherent byte swapping occurs on all transactions. This can be useful when byte-swapping is required but the MMU is disabled.

This bit is not utilized by the software release 2.3 and it not discussed further in this chapter. This bit is not available on Intel® IXP42X product line.

#### EXP\_BYTE\_SWAP\_EN (Bit 10)

The IXP46X product line provides the ability to control whether transfers initiated from master devices on the Expansion Bus should be byte swapped or not.



This bit is not utilized by the software release 2.3 and it not discussed further in this chapter. This bit is not available on Intel® IXP42X product line.

### 27.5.3.4 MMU P-Attribute Bit

The Intel XScale® Processor within the IXP4XX product line processors contains an extension to the MMU. The first level page descriptor contains an additional bit (P) which is used to control the little endian coherence mode on a per section basis (1 MB of memory).

*Note:* This bit only has effect if the Intel XScale® Processor is in LE mode and the BYTE\_SWAP\_EN(bit 8) is set and the FORCE\_BYTE\_SWAP (bit 9) is cleared.

### 27.5.3.5 PCI Bus Swap

The PCI controller has a byte lane swapping feature. The “swap” is controlled via the PCI\_CSR register’s PDS and ADS bits within the PCI controller. The swap feature must be enabled if the Intel XScale® Processor is in big endian mode or Data Coherent little endian mode. For further details, refer to the processor’s specific *datasheet* and *developer’s manual*.

*Note:* The PCI\_CSR bits on the IXP46X product line are referred to as PBS and ABS. However, they are in the same location as previous IXP4XX product line processors.

### 27.5.3.6 Summary of Silicon Controls

Table 86 summarizes the device selections and their behavior.

**Table 86. Endian Hardware Summary**

Intel XScale® Processor Endianness [1 = Big endian]	Expansion Bus Config Register [BYTE_SWAP_EN]	MMU ‘P’ Bit	Intel XScale® Processor endianness and its interaction with the AHB bus	PCI Bus Swap Enabled = PCI_CSR_PDS=1, PCI_CSR_ADS =1
1	X	X	Big endian	Enabled
0	1	1	Little endian – Data Coherent	Enabled, and PCI Bus space must be Data Coherent (0x48xx,xxxx)
0	1	0	Little endian – Address Coherent	Disabled
0	0	X	Little endian – Address Coherent	Disabled

### 27.5.4 Silicon Versions

Available hardware endianness controls vary by the stepping or product family of the processor. Identification of silicon version is indicated by markings on the devices themselves, or by accessing a register on the chip. Further details regarding this are available in the *Intel® IXP400 Software Programmer’s Guide* and the processor’s specific *datasheet*.

#### IXP42X product line B0 stepping

These processor versions support:

- Big endian
- Little endian Address Coherency
- Little endian Data Coherency



These processor part numbers are detailed in other documents, such as *Intel® IXP42X Product Line of Network Processors and IXC1100 Control Plane Processor Datasheet*.

### IXP46X network processors A0 stepping

These processor versions support:

- Big endian
- Little endian Address Coherency
- Little endian Data Coherency

These processors also add additional hardware endianness controls, including:

- Byte swapping for transactions initiated by Expansion Bus masters.
- Force byte-swapping by the Intel XScale® Processor in the event that the MMU is disabled.

## 27.6 Little endian Strategy in Intel® IXP400 Software and Associated BSPs

The little endian strategy employed is discussed in relation to two different areas:

1. The Board Support Packages (BSPs) for the supported development platforms.
2. The software release 2.3 (Access-Layer).

When adding support for little endian, there were two factors taken into account in deciding where to use Address Coherency and Data Coherency little endian modes.

1. The initial software release 2.3 releases and Board Support Packages were all big endian.
2. software release 2.3 support for little endian was required to operate on all the supported little endian operating systems.

The implications of this can be seen in two key little endian implementation decisions.

1. The little endian VxWorks\* Board Support Package uses Address Coherency. One of the properties of Address Coherency is that 32-bit accesses do not need to be swapped. Most of the processor register accesses in the BSP are 32-bit accesses, so it made sense to port the existing big endian BSP to Address Coherent little endian.
2. The software release 2.3 little endian implementation uses Data Coherency and all memory is mapped as Data Coherent. We did not want to have different little endian implementations of the software release 2.3 for the different operating systems supported, and therefore chose Data Coherency as the common implementation for all currently supported operating systems.

It should be noted that the software release 2.3 little endian implementation is designed in such a way that the coherency mode for any Access-Layer component can be changed if desired. The same is true for the memory map. There is no restriction placed on mapping memory as either Address or Data Coherent once that model is facilitated by the chosen operating-system MMU requirements. The choice of coherency mode is principally determined by the way the Operating System uses the memory management unit.

The files to consult within the software release 2.3 are:

```
\ixp_osal\include\modules\ioMem\IxOsallIoMem.h
\xp_osal\include\modules\ioMem\IxOsallMemAccess.h
\xp_osal\include\modules\ioMem\IxOsallEndianness.h
```



```
\ixp_osal\os\vxworks\include\platforms\ixp400\IxOsalOsIxp400CustomizedMapping.h
\ixp_osal\os\linux\include\platforms\ixp400\IxOsalOsIxp400CustomizedMapping.h
```

The remainder of this chapter details the processor little endian implementation. It identifies the appropriate coherency mode per hardware component and explains the implications of each selection. It also contains a detailed look at the implications of the various endianness modes and how they relate to TCP/IP stack expectations.

Details on every component are not included, but an overview of certain components is included to provide insight on which coherency modes are used. Further details on the currently supported modes of each component are available in the code comments included in the software release 2.3.

*Note:* Linux\* little endian support utilizes the existing software release 2.3 components, principally using the same VxWorks\* modifications as documented in following sections. Other changes are contained within the Linux\* board support package.

## 27.6.1 APB Peripherals

The Advanced Peripheral Bus (APB) provides access to the following peripherals:

- Blocks specific to BSP
  - UARTs
  - Performance Monitoring Unit
  - Interrupt Controller
  - GPIO Controller
  - Timer Block
  - SSP, I<sup>2</sup>C, Ethernet MAC 3 control, and IEEE 1588 units on the IXP46X product line.
- Blocks controlled by software release 2.3:
  - NPE Message Handler and Execution control registers
  - Ethernet MAC control
  - Universal Serial Bus (USB)

The APB peripherals are placed in Address Coherent mode to nullify changes from the existing big endian BSP.

## 27.6.2 IXP400 Software Little Endian Strategy on APB

In general, all access to the APB peripherals is to 32 bits (Word) memory-mapped registers (MMR). Because of this, the preferred choice is to use Address Coherence mode conversion.

By using an Address coherence conversion, the software does not need to do any byte swapping or address swizzling. These operations greatly decrease performance in an Little Endian (LE) operating system.

The existing code within the IXP400 software and BSP code base writes to all peripherals on a big endian basis. Based on the reasons stated above, all the APB peripherals should be placed in Address Coherent mode to nullify performance degradation under LE operating systems.



However, the strategy above may not be applied across all operating systems due to reasons that is explored in the following sections.

### 27.6.3 AHB Memory-Mapped Registers

There are several other memory-mapped areas within the processors:

- AHB Queue Manager. The configuration is covered in the “Queue Manager — IxQMGr” on page 382.
- PCI. Further details are provided in “PCI” on page 388.
  - Control registers. These registers are all word-wide (32 bits) and operate in Address Coherent little endian mode.
  - PCI memory (AHB mapped, 0x48xx,xxxx Phy space). This space must be mapped Data Coherent.
- Expansion Bus registers. These registers are all word-wide (32 bits) and operate in Address Coherent little endian mode.
- SDRAM control registers. These registers are all word-wide (32 bits) and operate in Address Coherent little endian mode.

In addition to the list of MMR registers presented above, the IXP46X product line has the following new components:

- USB 1.1 Host Controller
  - Control registers. These registers are all word-wide (32 bits) and operate in Address Coherent little endian mode.
  - Data buffers and USB control data structures within SDRAM. All read/write access by the Intel XScale® Processor is set up in Address Coherent little endian mode.
- AHB-RSA Bridge
  - RNG, EAU and SHA control registers are all word-wide (32 bits) and operate in Address Coherent little endian mode.
  - Internal Memory. The EAU’s internal memory can be only be access in words (32 bits). Therefore, the CryptoAcc component PKE sub-module is set up in Address Coherent mode while the main module that interfaces to the NPE is in Data Coherent mode.

### 27.6.4 Intel® IXP400 Software Core Components

Intel® IXP400 Software v2.3 contains several structural components used by all other software release 2.3 access-layer components. All of the software components are otherwise referred to as the Access-Layer and provide software interfaces for control of the various hardware blocks within the processor.

*Note:* Changes to ixEthAcc listed here are indicative of the types of changes required in other components.

#### 27.6.4.1 Queue Manager — IxQMGr

The NPE Queue Manager component provides the interface to the hardware queue manager block. All registers and hardware FIFOs are word-wide (32 bits). Data Coherent little endian mode is used.



#### 27.6.4.2 NPE Downloader — IxNpeDI

This component utilizes the NPEs' Message Handler and Execution Control registers. All registers are word-wide (32 bits). Such registers are best set up using little endian Address Coherent mode. However, this would cause the component to have differing behavior between some operating systems. As a result, the decision was made to make the NPE Execution Control registers Data Coherent.

All register reads/writes occur via the following functions, defined in npeDI/include/IxNpeDIMacros\_p.h

```
IX_NPEDL_REG_READ()
IX_NPEDL_REG_WRITE()
```

#### 27.6.4.3 NPE Message Handler — IxNpeMh

This component is dependent upon NPE Message Handler and Execution Control registers. All registers and hardware FIFOs are word-wide (32 bits).

Address Coherent little endian mode is used for messages sent via the Message Handler interface in Linux\* while Data Coherent little endian mode is used in VxWorks\*.

For example, the ixNpeMhMessageSend function is defined as follows:

```
typedef struct
{
    UINT32 data[2]; /*the actual data of the message */
} IxNpeMhMessage;
```

Although the registers would be ideally accessed in Address Coherent mode, a system-wide decision to put software release 2.3 peripherals in Data Coherent mode means the contents of the "data" within the Message Handler is modified by the underlying access-layer software.

#### 27.6.4.4 Ethernet Access Component — IxEthAcc

The decision to set up the SDRAM in Data Coherent little endian mode is driven by the primary assumption that there is more payload than control data structures exchanged between the NPEs and Intel XScale® Processor.

This approach also lends itself to using Address Coherent mode for the control structures, and, if required for a future OS porting, should be easily implemented in a particular operating system environment. Some of the information detailed is intended to facilitate use of Address Coherent mode should it be desired. It is not intended to imply that Address Coherency is used in this component in the current software from Intel.

##### 27.6.4.4.1 Data Plane

The data plane interface for IxEthAcc uses the IxQMgr component to send/receive messages between the Ethernet access and the Ethernet NPEs. All messages transferred are word-wide (32-bit) messages. These messages are modified by the



underlying access layer because the AHB Queue Manager hardware FIFOs are mapped using Data Coherent little endian (as described in “Queue Manager — IxQMgr” on page 382).

*Note:* The AHB Queue Manager can be I/O mapped into memory using either data or address Coherent conversions, and the IxQMgr software will operate correctly in either mode, transparent to the client.

The messages sent/received from the NPE contain a pointer reference to an IX\_OSAL\_MBUF, and more specifically to the NPE specific structure within the IX\_OSAL\_MBUF. See the Chapter 3.0 for more information.

The SDRAM is mapped using Data Coherency mode for all areas. This introduces two specific areas of consideration:

- NPE interpretation of the IX\_OSAL\_MBUF
- NPE interpretation of the data payload.

#### 27.6.4.4.2 IX\_OSAL\_MBUF Data Payload

The Ethernet access-layer component does not impose any alignment restrictions on the ix\_data pointer within the IX\_OSAL\_MBUF. The primary consideration in selecting the little endian coherence mode (as Data Coherent) is the expectation the standard BSD IP stack places on the data format for payloads.

The BSD IP stack makes extensive use of the htons, htonl primitives to extract IP/UDP/TCP header information within the stack. These are described in “Macro Examples: Endian Conversion” on page 371.

BSD IP Stack summary:

- Bytes can be read with a byte pointer.
- All half-word reads must be half-word-aligned and use htons/ntohs for conversions.
- All word reads must be word-aligned and use htonl/ntohl for conversions.

The issues associated with the payload is discussed in reference to an Ethernet frame. As shown in Figure 123, the frame is described in network byte order.

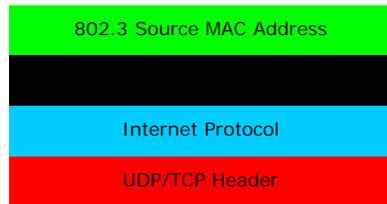
**Figure 123. Ethernet Frame (Big Endian)**

D0	D1	D2	D3
		DA[0]	DA[1]
DA[2]	DA[3]	DA[4]	DA[5]
SA[0]	SA[1]	SA[2]	SA[3]
SA[4]	SA[5]		
ver/hlen	TOS	16-bit-Len	
Identification		flag/Fragment offset	
TTL	Protocol	Header Checksum	
src-ip[0]	src-ip[1]	src-ip[2]	src-ip[3]
dst-ip[0]	dst-ip[1]	dst-ip[2]	dst-ip[3]
UDP/TCP Header			

803.2 Destination MAC Address



**Figure 123. Ethernet Frame (Continued)(Big Endian)**



The IP stack typically has an alignment restriction on the IP packet. The start of the IP packet must be word-aligned, that is, the ver/hlen field shown above must start on a 32-bit boundary. There are 14 bytes of Ethernet frame data preceding the IP header. Thus ix\_data pointers typically need to be half-word-aligned (16 bits). This is the case that is discussed in this chapter, and in the *Intel® IXP42X Product Line of Network Processors and IXC1100 Control Plane Processor: Understanding Big Endian and Little Endian Modes Application Note*.

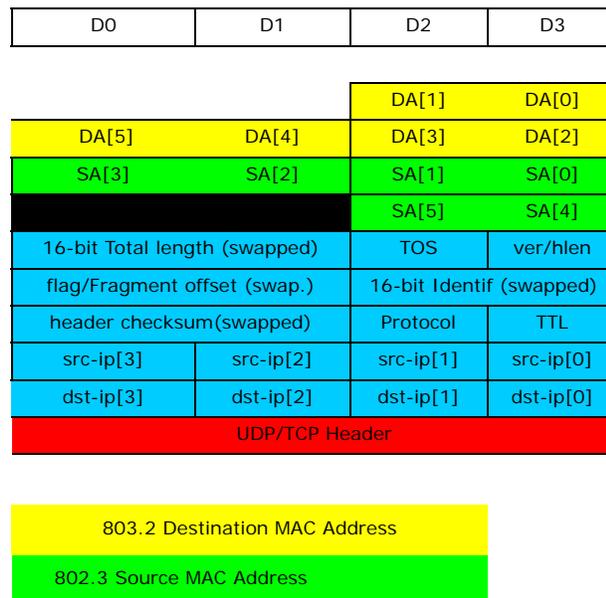
Detailed is the typical receive case for 64-byte frame (60 + CRC).

Given an IX\_OSAL\_MBUF data pointer (ix\_data) that is half-word-aligned, the NPE must transfer the frame into main memory. The transactions the NPE AHB coprocessor generates depend on the alignment and size of the transfer. For a 60-byte transfer, half-word-aligned, the NPE would generate:

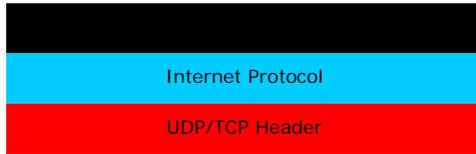
- One half-word transfer, half-word-aligned
- 14 word burst transfers, word-aligned
- One half-word transfer, half-word-aligned.

This will result in the following payload (see Figure 124) written to SDRAM from the Intel XScale® Processor (Address Coherent).

**Figure 124. One Half-Word-Aligned Ethernet Frame (LE Address Coherent)**



**Figure 124. One Half-Word-Aligned Ethernet Frame (Continued)(LE Address Coherent)**



The code below provides the read-out formation after the application of a conversion macro. Effectively, the header comes in as big endian and is then output as little endian.

The following shows the IP header structure and outlines how the payload would be read from the Intel XScale® Processor in little endian Data Coherent mode:

```

struct iphdr {
    __u8version:4,hlen:4; /* Offset 0*/
    __u8tos; /* Offset 1 byte*/
    __u16tot_len; /* Offset 2 bytes*/
    __u16id; /* Offset 4 bytes*/
    __u16frag_off; /* Offset 6 bytes*/
    __u8ttl; /* Offset 8 bytes*/
    __u8protocol; /* Offset 9 bytes*/
    __u16check; /* Offset 0xA bytes*/
    __u32saddr; /* Offset 0xC bytes*/
    __u32daddr; /* Offset 0xF bytes*/
    /*The IP options start here. */
};

```

The Header contents assume the following reads: (See [Figure 125](#))

- Half-word read at DA[1], half-word-aligned
- Word read at DA[2], word-aligned
- Word read at SA[3], word-aligned
- Half-word read type/len field, word-aligned
- Half-word read SA[5], half-word-aligned.

**Figure 125. Intel XScale® Processor Read of IP Header (LE Data Coherent)**

D0	D1	D2	D3
		DA[1]	DA[0]
DA[5]	DA[4]	DA[3]	DA[2]
SA[3]	SA[2]	SA[1]	SA[0]
		SA[5]	SA[4]
ver/hlen	TOS	16-bit Total length (swapped)	
16-bit Identif (swapped)		flag/Fragment offset (swap)	



**Figure 125. Intel XScale® Processor Read of IP Header (LE Data Coherent) (Continued)**

TTL	Protocol	header checksum(swapped)	
src-ip[3]	Src-ip[2]	src-ip[1]	src-ip[0]
dst-ip[3]	Dst-ip[2]	dst-ip[1]	dst-ip[0]
UDP/TCP Header			

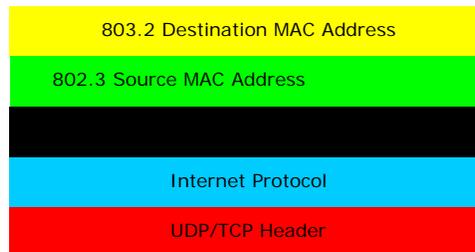


Figure 125 shows that the IP protocol stack operates correctly with the payload offered to the stack for half-word-aligned ix\_data using Data Coherent little endian mode and the IP protocol stack's use of data conversion macros.

#### 27.6.4.4.3 Learning Database Function

There are two main communication mechanisms between the Ethernet NPEs and the Intel XScale® Processor Ethernet learning function:

- NPE messages passed using the IxNpeMh interface
- Direct data structure exchanges between the IxEthDB access-layer component and NPEs

The messages passed to/from the NPE and Intel XScale® Processor are transferred via the IxNpeMh interface. Messages are written in the native endianness (BE or LE) and swapped independently by the Message Handler, before sending them to the NPEs. As mentioned in “NPE Message Handler — IxNpeMh” on page 383, messages may contain multiple word-wide data elements.

IxEthDB does not explicitly swap data when communicating with the NPEs. Data structures directly exchanged by EthDB with the NPEs, such as trees and arrays with MAC addresses and additional information, are written in a byte-oriented manner, which guarantees correct operation when the memory is accessed in big endian or Data Coherent little endian mode. Tree uploads are handled identically, using byte accesses.

#### 27.6.4.4.4 Ethernet Access MIB Statistics

The Ethernet NPEs maintain error statistics, accessible via the IxEthAcc API. The statistics are recovered from the NPE via an SDRAM buffer. The buffer is populated from the NPEs in big endian mode. As such, all words undergo a big endian to little endian (Data Coherent) conversion before the results are returned to the user.

#### 27.6.4.4.5 Ethernet MAC

All registers are mapped to the APB, and are accessed in word mode only (32 bits). The APB memory map is data coherent, and requires data swapping on all registers reads/writes.



#### 27.6.4.4.6 Intel® IXP400 Software IxEthAcc and IxEthDB Summary

This section presents a summary of the changes that were made to the IxEthAcc component, **assuming** NPE is big endian and all SDRAM is in little endian Data Coherent mode.

- The IX\_OSAL\_MBUF ix\_npe data structure's half-word and word member fields must be swapped prior to submission to the NPE. (**ixEthAccPortTxFrameSubmit()**)

**Note:** The IX\_OSAL\_MBUF chain is walked and all IX\_OSAL\_MBUFs in a chain are updated. (**ixEthAccPortRxFreeReplenish()**)

- The IX\_OSAL\_MBUF ix\_npe data structure's half-word and word member fields are swapped before usage. This procedure must be conducted after reception from the NPE and before calling:
  - User functions registered via *ixEthAccPortTxDoneCallbackRegister*.
  - User function registered via *ixEthAccTxBufferDoneCallbackRegister*.
- Ethernet Database (IxEthDB)
  - Endianness conversion of the Ethernet learning trees when ownership is transferred to/from the Intel XScale® Processor <-> Ethernet NPEs.
  - Tree Writes. **ixEthDBNPETreeWrite**
  - Tree uploads. **ixEthDBNPESyncScan**
  - Display. **ixEthELTDumpTree**
- MAC Statistics. The memory used to return statistics from the NPE is endian-converted before returning the data.
- Ethernet MAC registers are mapped in little endian Data Coherent mode.

*Note:* The coherency modes chosen for software release 2.3 little endian implementations for VxWorks\* are summarized in "[Endian Conversion Macros](#)" on page 389.

#### 27.6.4.5 ATM and HSS

Both ATM and HSS components pass descriptors between the Intel XScale® Processor and NPEs. These descriptors undergo similar changes to those described above.

#### 27.6.5 PCI

The primary consideration for PCI network drivers is the configuration of the byte swapping within the PCI controller itself (see "[Endian Hardware Summary](#)" on page 379).

The configuration is dependent on the coherency mode of the SDRAM memory area. In case of VxWorks\*, the SDRAM memory controller is in Data Coherent mode.

Importantly, the PCI memory space must be configured in little endian Data Coherent mode. This is the physical memory area 0x4800,0000.

The PCI Configuration Space Register has PCI\_CSR\_IC, PCI\_CSR\_ABE, PCI\_CSR\_PDS, PCI\_CSR\_ADS set to '1'.

#### 27.6.6 Intel® IXP400 SoftwareOS Abstraction

All little endian system configuration information is in the `ixp_osal\os`



\\vxworks\include\platforms\ixp400 \IxOsaIOs\ixp400CustomizedMappings.h. Further information on the VxWorks\* memory map is available in the VxWorks\* BSP documentation for the supported development platforms. Depending on their implementations, other operating systems may provide similar files/documents.

The macros shown in “Intel® IXP400 Software Macros” on page 389 are provided for use in the software release 2.3 components. The defines are correct for software release 2.3, but may change for other releases.

**Table 87. Intel® IXP400 Software Macros**

#defines	Description
#IX_OSAL_BE_MAPPING	Big Endian Mapping
#IX_OSAL_LE_AC_MAPPING	Little Endian address coherent byte mode
#IX_OSAL_LE_DC_MAPPING	Little Endian data coherent byte mode

Table 88 shows the endian conversion macros that need to be mapped for developer usage.

**Table 88. Endian Conversion Macros**

Macro	Behavior	Description
IX_OSAL_BE_XSTOBUSL()	No swap	Big endian XScale to Bus Long
IX_OSAL_BE_XSTOBUSS()	No swap	Big endian XScale to Bus Short
IX_OSAL_BE_BUSTOXSL()	No swap	Big endian Bus to XScale Long
IX_OSAL_BE_BUSTOXSS()	No swap	Big endian Bus to XScale Short
IX_OSAL_LE_AC_XSTOBUSL()	No swap	Little endian Address Coherent XScale to Bus Long
IX_OSAL_LE_AC_XSTOBUSS()	Address Swap	Little endian Address Coherent XScale to Bus Short
IX_OSAL_LE_AC_BUSTOXSL()	No swap	Little endian Address Coherent Bus to XScale Long
IX_OSAL_LE_AC_BUSTOXSS()	Address Swap	Little endian Address Coherent Bus to XScale Short
IX_OSAL_LE_DC_XSTOBUSL()	Data Word swap	Little endian Data Coherent XScale to Bus Long
IX_OSAL_LE_DC_XSTOBUSS()	½ Data Word swap	Little endian Data Coherent Bus to XScale Short
IX_OSAL_LE_DC_BUSTOXSL()	Data Word swap	Little endian Data Coherent Bus to XScale Long
IX_OSAL_LE_DC_BUSTOXSS()	½ Data Word swap	Little endian Data Coherent XScale to Bus Short

### 27.6.7 VxWorks\* Considerations

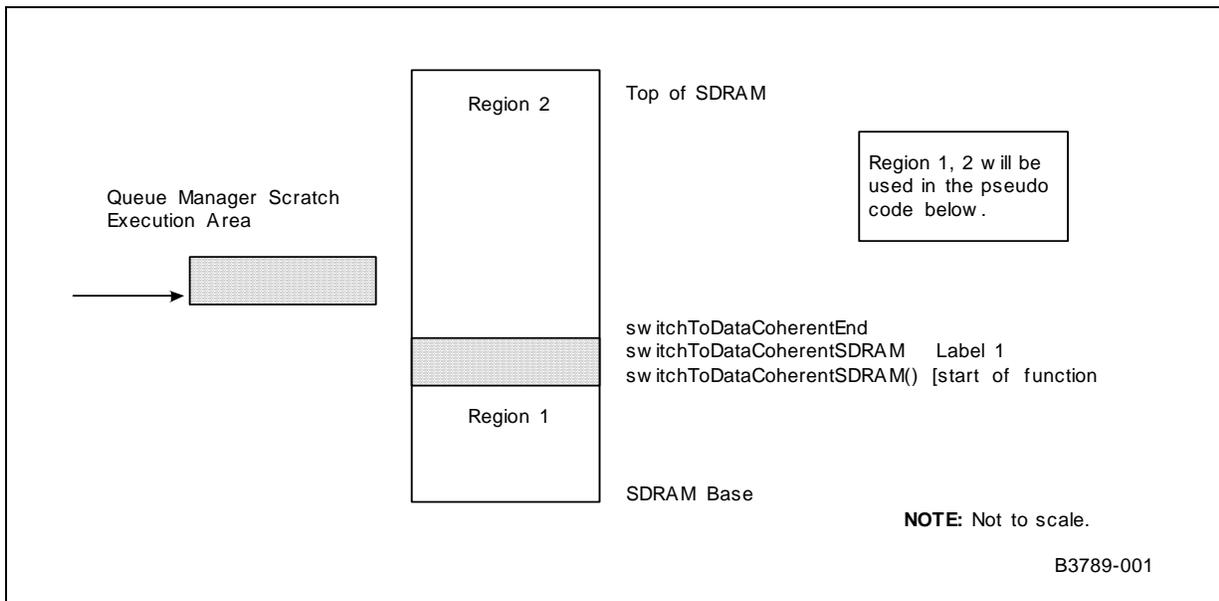
Both the AHB Queue Manager and NPE debug control registers (NPE message handler component ixNpeMh) are placed in Data Coherent little endian mode. As the NPE debug registers are in APB space, and other APB registers are mapped in Address Coherent mode, a Data Coherent alias for the APB bus is defined.

Control is transferred from the bootrom into VxWorks\* once it is downloaded via FTP. The MMU is disabled during this transition and, as such, all SDRAM is in Address Coherent mode. The SDRAM can only be converted to Data Coherent once the MMU is enabled. The MMU is enabled in usrConfig code. The first opportunity to swap the SDRAM to Data Coherent is in hardware init syshwInit0().

An example of how to place the SDRAM in Data Coherent mode while executing from this SDRAM is the function named mmuARMXScalePBitSet() in sysLib.c.

Figure 126 shows the related memory map.

Figure 126. VxWorks\* Data Coherent Swap Code



The following is example pseudo code:

```
switchToDataCoherentSDRAM:
    ; Interrupts are disabled, in hwinit2().

    Flush Cache (Instr & Data)
    Drain Write buffers
    Disable MMU
    Invalidate Instr & Data cache
    Invalidate TLB
    Walk though all MMU SDRAM Large/Section entries , setting 'P' bit for all
    entries.

    Copy MMU enable code to Q-Manager scratch.
    Perform LE endian swap on Region 1
    Perform LE endian swap on Region 2
    Set the P-Bit in MMU table walk
    Enable Byte swap in expansion bus register
    Jump to scratch memory location

    Enable MMU

    Wait for action to complete

    Jump to switchToDataCoherentSDRAM - Label1
```



Label1:

```

Enable Instr & Data cache.

Enable Branch Target buffer.

return
    
```

A similar implementation was required for execution in the VxWorks\* bootrom. The only caveat is that the SDRAM used to load the VxWorks\* image must be kept in Address Coherent mode, as execution control is transferred to that image with the MMU disabled.

*Note:* In the IXP42X product line, the only way to enable Data Coherent Byte mode (enabling byte swapping) is to set the P-bit Memory attribute and with the MMU enabled. However, in the IXP46X product line, there is a new feature that allows the Intel XScale® Processor to initiate access in LE Data Coherent mode with or without the MMU enabled or the P-bit set. This capability was made possible by the introduction of a FORCE\_BYTE\_SWAP bit (see “Little Endian Data Coherence Enable/Disable” on page 378 for details).

With this new feature, there is no need of doing LE endian swap of the region 1 and region 2 of the SDRAM as mentioned above during initial setup of the silicon. And this would greatly reduce boot-up time.

### 27.6.8 Software Versions

Table 89 provides a historical list of software releases for the IXP4XX product line processors. All versions currently support big endian operation. The table shows which versions also support little endian operation.

**Table 89. Intel® IXP400 Software Versions**

Intel® IXP400 SoftwareVersion	Little Endian Support Yes/No
IXP400 software release 1.0	No
IXP400 software release 1.1	No
IXP400 software release 1.2.1	No
IXP400 software release 1.2.2	No
IXP400 software release 1.3	Yes - VxWorks* only
Intel® IXP425 DSLAM Software	No
Intel® IXP400 DSP Software up to and including 2.5	No
IXP400 software release 1.4	Yes - VxWorks*
IXP400 software release 1.5	Yes - VxWorks* and Linux*
IXP400 software release 2.0	Yes - VxWorks* Yes - Linux* on IXDP425 Development Platform only
Intel® IXP400 Software plus Microsoft* Windows* CE.NET BSP	Yes
IXP400 software release 2.1	Yes - VxWorks* and Linux*



